COSYNE 2006

ABSTRACTS

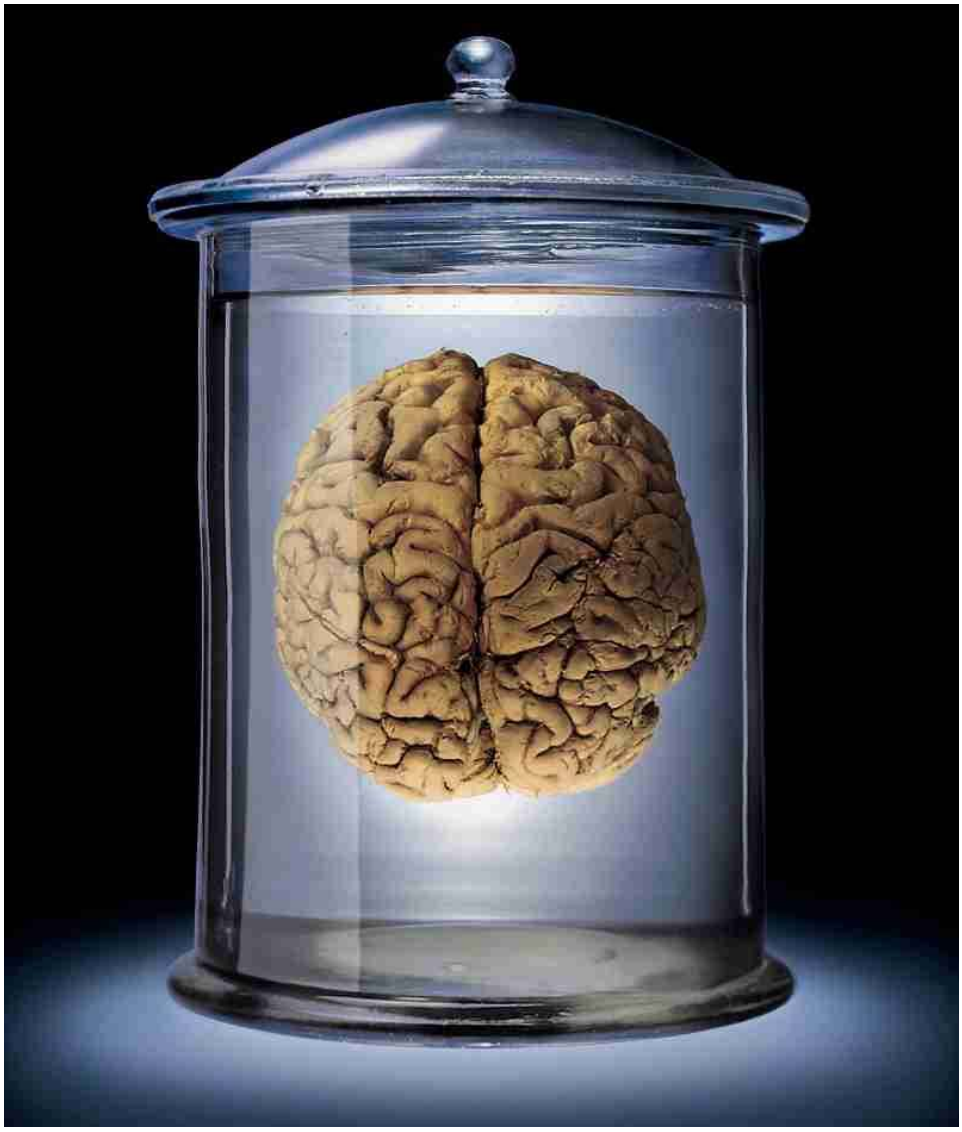

Computational and Systems Neuroscience

Main Meeting
March 5–8, 2006
Salt Lake City, Utah

4. The Metabolic Energy Cost of Action Potential Velocity

William B Levy, Patrick Crotty

University of Virginia

The action potential of the unmyelinated nerve is metabolically expensive. In the neuropil of neocortex, where axons must be unmyelinated if each one is to make several thousand sequential or neighboring synapses, the metabolic costs are surprisingly large. Attwell and Laughlin (2001) estimate that 75% of the ATP consumed by neurons in the rat brain is used for signaling. Of this 75%, the majority is used by the unmyelinated axons.

This metabolic perspective contrasts with and, as we will see, ultimately complements Hodgkin's conjectured constraint on action potential velocity. Both Hodgkin (1975) and Adrian (1975) proposed that the gating-charge movement, which inevitably accompanies rapid activation of a voltage-dependent channel, leads to an optimal density of fast Na^+ channels. This optimization occurs because the movement of charge specifically restricted to the transmembrane voltage field contributes, albeit transiently, to membrane capacitance. Since increasing capacitance slows action potential propagation, Hodgkin proposed that the Na^+ channel density has evolved to maximize velocity. Unfortunately, although the maximum velocity exists, biophysical parameter sweeps do not support the evolutionary conjecture. In a simple Hodgkin-Huxley model reparameterized to fit the action potential velocity with high precision, the conjecture misses by two-fold. When a more sophisticated model, i.e., one that takes into account that as many as six equivalent charges move across the membrane but are only sequentially available., the conjecture fails with more than a four-fold error.

Using the perspective of metabolic cost, we analyze the cost of velocity and identify an optimization that has occurred in evolution.

The energetic cost arising from an action potential is divided into three separate components (Figure 2): (1) the depolarization of the rising phase; (2) the hyperpolarization of the falling phase; and (3) the largest component, the overlapping of positive and negative currents, which has no electrical effect. Using both the Hodgkin-Huxley (HH) model and an improved version of the Hodgkin-Huxley model (denoted HHSFL after Sangrey, Friesen, and Levy (2004)), we investigate the variation of these three components as a function of two easily evolvable parameters, axon diameter and ion channel densities. Assuming conduction velocity is well-designed for each organism, the energy requirement associated with the rising phase attains a minimum near the biological values of the diameter and channel densities. This optimization is explained by another minimization, specifically the total, functional membrane capacitance per unit length. The functional capacitance is the sum of the intrinsic membrane capacitance and the gating capacitance associated with the sodium channel, and this capacitance minimizes at nearly the same values of channel density where metabolic costs are minimal (Figure 3).

Because capacitance is temperature independent and because this result is independent of the assumed velocity, the optimization result generalizes to the small, unmyelinated axons of the mammalian neuropil.

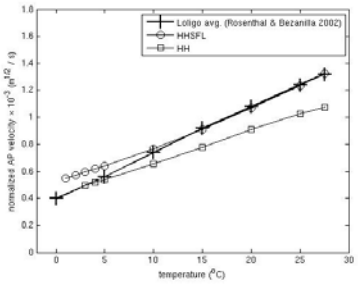

Figure 1: Action potential velocities as functions of temperature. The HHSFL model gives the best approximation to the experimental Loligo action potential velocities from 10° to 28 °C. Below 8°C, the HH action potential model is a better fit. As in Rosenthal and Bezanilla (2002), all velocities are divided by the square roots of the respective axon diameters in order to normalize for different diameter values. The experimental data points (crosses) are an average of velocity measurements for three Loligo species taken by Rosenthal and Bezanilla (2002).

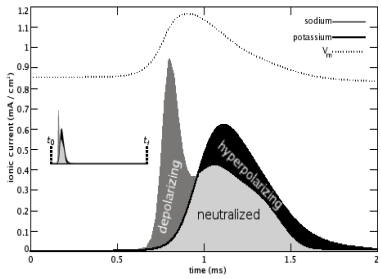

Figure 2: The Na^+ and K^+ currents during the action potential are naturally divided into three components: the depolarizing net Na^+ current in the wavefront; the net hyperpolarizing K^+ current just a little after peak depolarization; and the neutralized currents which account for the overlapping and offsetting Na^+ and K^+ fluxes. The HHSFL action potential is illustrated here. The total Na^+ current at any time is the sum of the depolarizing and neutralized currents and similarly for the K^+ current. The integral of each component (represented by the shaded areas in the figure) gives the associated flux, which is directly proportional to the energy used by the ionic pump to restore the resting concentration. Inset: A temporally expanded plot which begins at the foot of the action potential (t_0) and lasts for 10 ms ($t_f = t_0 + 10$ ms). Almost all of the ionic charge flux takes place during the first 1.5 ms.

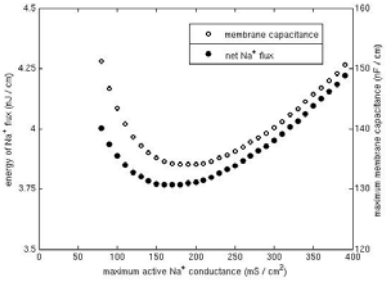

Figure 3: Total capacitance and wavefront energetic costs are convex functions at isovelocity. The metabolic energy cost per unit length incurred by the net Na^+ current (filled circles, left axis) and the total membrane capacitance per unit length (open circles, right axis) are shown for the HHSFL model at 18.5 °C. Both quantities are computed along a 21.2 m/s isovelocity curve. The similar minimization, occurring near the biological values of conductance and diameter, is sensible because the membrane capacitance is what determines the energy of the wavefront. The low-conductance cutoff of the curves is due to the difficulty of generating action potentials in this region. The overall energy normalization is derived assuming that the ATPase Na^+/K^+ pump exchanges 2 Na^+ for 2 K^+; to use the conventional 3 Na^+/2 K^+ ratio, the y-axis should be multiplied by 2/3

6. Moiré interference between grid cells: A mechanism for representing space at multiple scales

Hugh T Blair[1], Kechen Zhang[2]

[1]UCLA, [2]Johns Hopkins University

The rat hippocampus contains neurons called “place cells” which are tuned to fire only when the rat visits specific preferred locations in a familiar spatial environment (O’Keefe & Dostrovsky, 1971). Hafting et al. (2005) have recently discovered a new class of neurons called "grid cells" in the rat entorhinal cortex, an area that provides major inputs to the hippocampus. As a rat navigates in a spatial environment, grid cells fire selectively at multiple spatial locations which are geometrically arranged to form a tesselating hexagonal grid pattern that tiles the surface of the environment. Since grid cells send dense projections to hippocampal place cells, it is likely that place cells derive their location-specific firing properties by combining inputs from grid cells (Brun et al., 2002; 2005).

Here we show that two or more hexagonal grid fields can intersect to create moiré interference patterns, and these patterns can be scaled over an infinite range of sizes simply by rotating the component grids against one another (Fig. 1). When these scalable moiré interference patterns are linearly combined, they can serve as basis functions to construct complex two-dimensional images. Thus, an arbitrary image in the plane, I(x,y), can be constructed as a linear sum of moiré interference patterns, G(x,y). In the examples considered here, each G(x,y) is formed by an intersecting a pair of hexagonal grids which share the same vertex spacing, but are rotated against each other by an angle f (as in Fig. 1). Interestingly, if an output image is constructed as a linear sum of scalable moiré patterns in this way, then the output image itself is also scalable simply by rotating the component grids of the moiré basis functions against one another in unison (that is, adjusting f by a common factor simultaneously for all of the basis grids). A simple example of this image scaling phenomenon is shown in Fig. 2.

Based on these observations, we propose that moiré interference between grid fields may provide a mechanism by which hippocampal place cells can represent spatial environments at different scales. To support this argument, we first show that the spatial firing patterns of hippocampal place cells can be reproduced extremely accurately by linear combinations of several hundred grid fields with parameters that match the known firing properties of grids cells in the entorhinal cortex. We then demonstrate how moiré interference effects could explain recent data showing that hippocampal place cells are topographically organized along the dorsoventral axis to represent space at different scales (Maurer et al., 2005). Next, we present our own neurophysiological data showing that when hippocampal place cells are recorded in a chamber that can be expanded or contracted during the recording session, place cells often respond by expanding or contracting their place fields; that is, place cells rescale their preferred firing locations along with the dimensions of the environment. Finally, we show that this rescaling of hippocampal place fields can be reproduced in simulations by modeling place fields as linear combinations of several hundred moiré interference patterns, where each moiré pattern is formed by interference between a pair of hexagonal grid fields. If these grid pairs are rotated against one another, then the simulated place fields expand and contract their dimensions in a manner that resembles the behavior of place cells when the recording chamber is expanded or contracted.

We conclude that moiré interference between grid fields may provide a novel mechanism for representing space at multiple scales in the hippocampus. In the rat, this may allow hippocampal place cells to construct scale-invariant memory representations of familiar spatial environments. Similar moiré interference principles might be able to account for scale invariant memory representations in other systems, such as scale invariant recognition of familiar objects by the visual system.

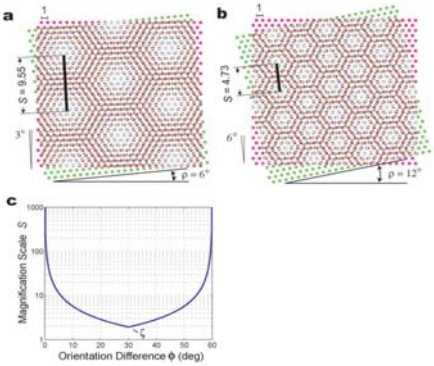

Fig. 1: Moiré interference between hexagonal grids. Two hexagonal grid fields (red & green) with unit vertex spacing are rotated against each other by an angle f so that their intersection (blue) produces a moiré interference pattern. In the cases shown, the interference pattern forms a new hexagonal grid which is a self-similar replica of the basis grids, but on a larger spatial scale.
a) Two grids are rotated against each other by an angle of f=6°, producing a moiré grid with magnified vertex spacing of S=9.55.
b) Two grids are rotated against each other by f=12° to produce a moiré grid with vertex spacing S=4.73.
c) In general, if two grids with unit vertex spacing are rotated against each other by an arbitrary angle f, then the vertex spacing of the moiré interference grid is given by S=1/[2sin(f/2)]. This function is infinite at integer multiples of 60°, and at midpoints between these infinite singularities (e.g., 30°), there are irrationally-valued singular points at where S=1.93185165.... ; we refer to this irrational constant as z.

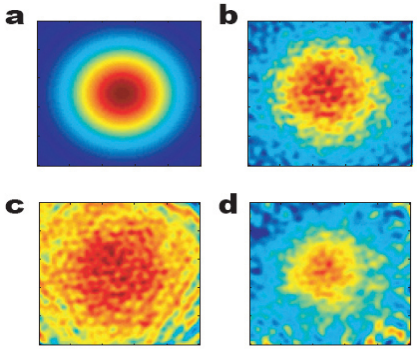

Fig. 2: Image scaling by moiré interference.
A set of 800 grids with random vertex spacing and grid orientations was grouped into 400 pairs, where each pair had the same vertex spacing and similar grid orientations that were offset from each other by a small rotation angle, f, to produce a moiré interference grid as shown in Fig. 1 above. The 400 interference grids were assigned linear coefficients using the Moore-Penrose pseudoinverse method to find the best fit to a gaussian target image.
a) The gaussian target image.
b) Least squares fit of 400 coefficients produces an output image similar to the target image.
c) Reducing the displacement angle f between grid pairs by a common factor causes the output image to expand.
d) Increasing the displacement angle f between grid pairs by a common factor causes the output image to contract.

7. Firing rates and times in the hippocampus: what are they good for?

Mate Lengyel, Peter Dayan

Gatsby Computational Neuroscience Unit, University College London

The (re)naissance of spike timing as a key primitive in neural representations owes a substantial debt to the hippocampus as a model system. Traditionally, the activity of the principal cells in the hippocampus has been characterized in terms of firing rates, and these have been shown to convey spatial [1] as well non-spatial [2] information in behaving animals. More recently, though, it has been demonstrated that the precise timing of spikes, often characterized as firing phases relative to the ongoing theta oscillation, is also central to hippocampal information processing [3,4]. Moreover, although both firing rates and phases are often tied to the spatial position of the animal [3,5,6], under some conditions they are decoupled, and apparently code for different variables [7,8]. It is still unclear, however, what the use is of such an intricate dual rate-and-phase coding scheme, and even more, how such dynamics support the key role accorded to the hippocampus in memory processing. We investigated this question in a normative theoretical framework of hippocampal memory, using data from in vitro and in vivo experiments to test our theory.

Our normative theory of hippocampal neural dynamics prescribes the way the instantaneous rates and phases of spikes in the hippocampus should interact in order to achive near-optimal memory retrieval. We suggest that the content of memory traces is specified by the firing phases of the neurons (and stored by spike timing-dependent plasticity), whereas the degree of certainty in the current memory being retrieved and/or represented is conveyed by their firing rates. The representation of uncertainty is vital in a vast range of neural computations, and behavioral performance in a wide range of tasks has indeed been shown to depend on the ability of animals and humans to correctly track uncertainty in task-relevant quantitities [9,10].

We thus explored three scenarios in which instantaneous firing rates signaled different specific forms of uncertainty: 1. uncertainty about the correctness of the retrieved memory trace represented by firing phases, 2. uncertainty about the participation of neurons in a memory trace, 3. uncertainty about both correctness and participation. The neural dynamics that implemented these computations were found to have the same qualitative characteristics in all three cases and led to competent memory retrieval performance in simulated networks. These dynamics also suggest predictions about the interactions between neural firing rates and phases in the hippocampus that are readily testable in experiments.

We tested the predictions of our theory using results from in vitro experiments (performed by Jeehyun Kwag and Ole Paulsen, Oxford University) and also preliminary results from the analysis of in vivo data from behaving animals (recorded by Francesco Battaglia, University of Amsterdam). Our earlier results confirmed that spike timing-based interactions between CA3 pyramidal cells, characterized as phase response curves, qualitatively match those predicted to be optimal [11]. We have also started testing the effect of firing rates on these interactions, and found that increased presynaptic firing rates do not change the shape of these phase response curves, consistent with our prediction that they should only scale them. The analysis of in vivo data revealed sophisticated cycle-by-cycle firing phase and rate dynamics in behaving animals, and showed that firing rates are subject to leaky dynamics, and change maximally when the firing phase of the neuron is at the point at which it changes minimally on average - again consistent with our theoretical prediction. We also find that our theory accords with a diverse set of earlier experiments [7,12] analysing the activity of the hippocampus in behaving animals. These results suggest that the neural interactions hippocampus may be optimal for retrieving memory traces encoded by firing phases augmented with an important uncertainty signal conveyed by firing rates.

This work was supported by the Gatsby Charitable Foundation and the European Bayesian-Inspired Brain and Artefacts project.

[1] O'Keefe J, Dostrovsky J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res 34:171-175, 1971.
[2] Wood ER, Dudchenko PA, Eichenbaum H. The global record of memory in hippocampal neuronal activity. Nature 397:613-616, 1999.
[3] O'Keefe J, Recce ML. Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3:317-330, 1993.
[4] Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsaki G. Organization of cell assemblies in the hippocampus. Nature 424:552-556, 2003.
[5] Harris KD, Henze DA, Hirase H, Leinekugel X, Dragoi G, Czurko A, Buzsaki G. Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. Nature 417:738-741, 2002.
[6] Mehta MR, Lee AK, Wilson MA. Role of experience and oscillations in transforming a rate code into a temporal code. Nature 417:741-6, 2002.
[7] Hirase H, Czurko A, Csicsvari J, Buzsaki G. Firing rate and theta-phase coding by hippocampal pyramidal neurons during 'space clamping'. Eur J Neurosci 11:4373-4380, 1999.
[8] Huxter J, Burgess N, O'Keefe J. Independent rate and temporal coding in hippocampal pyramidal cells. Nature 425:828-832, 2003.
[9] Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. Nature 400:233-238, 1999.
[10] Kording KP, Wolpert DM. Bayesian integration in sensorimotor learning. Nature 427:244-247, 2004.
[11] Lengyel M, Kwag J, Paulsen O, Dayan P. Matching storage and recall: hippocampal spike timing-dependent plasticity and phase response curves. Nat Neurosci 8:1677-1683, 2005
[12] Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science 309:619-623, 2005.

9. Homeostatic synaptic plasticity in dendrites – is it local or global?

Ithai Rabinowitch, Idan Segev

Hebrew University of Jerusalem

Homeostatic synaptic plasticity (HSP) slowly regulates activity in neurons. Whenever activity is chronically enhanced, HSP weakens synaptic weights and vice versa. It is still unknown whether the underlying mechanisms for HSP is local, sensing activity at the site of each synapse and modifying its strength individually, or whether HSP is global, affecting all synapses in unison. The actual behavior of HSP is determined both by the spatial resolution of its underlying biophysical mechanism as well as by the location-specificity of the dendritic activity under HSP control. Using a detailed model of a CA1 pyramidal neuron with multiple HSP-modifiable excitatory synapses distributed throughout the dendritic tree, we explored the interactions between the HSP mechanism and dendritic activity and found that: (1) A local mechanism will act indistinguishably from a global one, scaling all synapses together, following overall persistent perturbation of the entire dendritic tree such as has been done in experimental studies on HSP (e.g. Turrigiano et al.). (2) Local HSP can regulate activity at the level of the individual dendritic functional compartment and is therefore more suitable for a compartmentalized organization of neuronal activity, whereas global HSP is more in line with a centralized scheme of neuronal function at the occasional expense of overemphasizing already over-active dendritic compartments (or vice versa). For this analysis we have developed a novel method for identifying and mapping functional dendritic compartments of any kind (e.g., based on voltage or on Ca+2 concentration) under any condition (linear or nonlinear dendrites). (3) Local HSP will have an active role in selecting which spatial configurations of potentiated synapses are to persist and which are to perish, favoring spatially distributed potentiated synapses over clustered ones. Global HSP, in contrast, will uniformly scale these potentiated synapses. Thus, these two possible HSP mechanisms imply a significant functional difference and it is therefore of great importance to experimentally unveil which of the two (or perhaps both) are operating in the dendritic tree.

11. Learning and Memory in an Exactly Solvable Stochastic Spiking Network

Surya Ganguli

UCSF

A central challenge in systems neuroscience involves understanding how memories are stored as stable patterns of synaptic efficacies in recurrent spiking networks, and how learning is mediated by global changes in these synaptic patterns. At the level of the single cortical synapse, it is thought that changes in efficacy are driven by local spike timing dependent plasticity (STDP) rules. Thus to understand learning and memory, one must understand the complex connection between these local rules, operating at individual synapses, and their effects at the macroscopic level of recurrent neural circuit behavior. In situations where there is a separation of time scales between plasticity rates and spiking dynamics, the statistics of spatiotemporal correlations between spikes drive synaptic changes. In such a restricted scenario, the crux of the difficulty in understanding global network level learning under local STDP lies in the intricate relation between synaptic states and spiking statistics; the spiking statistics are a complicated global function of all synaptic states in the network and in turn these emergent statistics drive further changes in all synaptic states.

We address this difficulty by proposing a recurrent spiking model that, remarkably, is exactly solvable in the sense that the spiking statistics can be computed in closed form for arbitrary recurrent and feed forward synaptic weights and arbitrary instantaneous correlations in external inputs. In this model each synapse has a weight p that represents the probability a presynaptic spike will cause a postsynaptic spike with a fixed small latency L. If multiple input spikes arrive simultaneously on different synapses, the output spiking probability sums linearly. The weights are constrained so the ouput probability never exceeds one. This is simply the linear Poisson neuron (Kistler and van Hemmen 2002) with a delayed delta-function kernel. Furthermore all neurons are pairwise independent conditioned upon past inputs. Thus a network of N neurons is characterized by its synaptic weight matrix p_{ij}. There can also be synapses driven by external inputs with weights f_{ik}.

Given these definitions, one can calculate the cross correlation function, or mean of $s_i(t)s_j(t+t')$, where $s_i(t)$ is the spike train of neuron i, by explicitly summing over all consistent spiking histories. Naively, histories are paths through the space of all 2^N possible spike patterns at each time, but the sum over paths simplifies considerably in this particular model. The reduced sum (which is nonzero only when t'=nL for integer n) can be visualized as a sum all over paths through the network of two given types. The first type of path starts at neuron i at time t, hops from neuron to neuron and visits neuron j in precisely n steps at time t'. The second type of path originates at a third neuron in the past, and sends off two nonintersecting branches through the network. One branch must terminate at neuron i at time t and the other at neuron j at time t+t'. Each path is associated with a product of probabilites p_{ij} corresponding to links in the path, and is weighted by the firing rate of the neuron at the path's point of origin. These paths are shown in figure 1.

Surprisingly, these paths can be summed and written in closed form. The resulting expression can be substituted into any STDP rule of the the form $d/dt\ p_{ij} = F(p_{ij}, s_is_j)$. Fixed points of these equations represent the stable network memories. Changes in these fixed points in response to changes in external input correlations correspond to network learning.

An exact analysis of small 2 and 3 node networks yields the following intuitions. The existence of recurrent circuitry in addition to feed forward circuitry partitions the space of input correlations into three types: weak, strong positive, and strong negative. If correlations are weak, a large class of recurrent synaptic patterns are stable. Strong positive correlations destabilize these synaptic patterns and create a new synaptic fixed point, which reinforces the input spike correlations. If the network starts in this fixed point and is subsequently exposed to strong negative correlations, the fixed point will destabilize and the weights will move to a new fixed point that reinforces the negative input correlations. The stronger the positive correlations were, the harder it becomes to destabilize the resulting fixed point upon subsequent exposure of the network to negative correlations. Put simply, recurrent synapses, combined with local STDP rules, make it harder for a network to change its mind in the face of new observations that contradict old observations. These synapses make networks stubborn.

The analysis of small networks can be extended to larger networks, where questions about the learnability of various input correlations given initial synaptic weights arise. More generally, this work opens the door to an understanding of synaptic pattern formation in large networks with arbitrary feedbacks and sensorimotor connectivities, which is a prerequisite for a biophysical understanding of learning and memory at the spiking network level.

12. Synaptic logic of cortical neuromodulation and plasticity

Robert C Froemke, Michael M Merzenich, Christoph E Schreiner

UCSF

Receptive fields of the adult sensory cortex are plastic, capable of a high degree of reorganization. In many cases, such plasticity requires activation of one or more subcortical neuromodulatory systems. However, the synaptic consequences of neuromodulation for gating receptive field plasticity are unknown. Here we examine the contributions of excitatory and inhibitory currents to long-term changes in receptive field structure after pairing sensory stimulation with neuromodulator release.

We made whole-cell voltage-clamp recordings in vivo from neurons in the primary auditory cortex of the adult anesthetized rat (Wehr and Zador, 2003; Zhang et al., 2003). After mapping auditory receptive fields with tone pips of various frequencies for 5-10 minutes, we repetitively paired tones at a certain frequency, the conditioned stimulus (CS), with electrical stimulation of the basal forebrain, the major source of cortical acetylcholine (Bakin and Weinberger, 1996; Kilgard and Merzenich, 1998). This conditioning procedure led to large, rapid increases in excitatory currents evoked by CS. Conversely, the excitatory response to the best frequency (BF) was diminished after pairing, regardless of the spectral distance between CS and BF. We also observed that conditioning led to a dramatic reduction in inhibitory current evoked by CS. These changes in excitation and inhibition persisted for >20 minutes after conditioning, and were highly specific- aside from changes to CS and BF, responses to all other frequencies were generally unaltered. Excitatory and inhibitory plasticity were coordinated together, as both effects were prevented when NMDA receptors and muscarinic receptors were blocked. Iontophoresis of bicuculline, to block GABA-A receptors, reversed the inhibitory plasticity. Moreover, in the presence of bicuculline, acetylcholine release was no longer required for induction of excitatory plasticity, which resulted from repetitive presentation of CS alone.

Preliminary evidence, using stimulation of other subcortical neuromodulatory systems, suggests that the sign of plasticity of excitatory and inhibitory currents is specific for different neuromodulators. Noradrenalin release from the locus coeruleus enhanced both excitation and inhibition at the CS, while serotonin release from the raphe nucleus decreased excitation and inhibition.

This specific shift in the balance of excitation and inhibition might alter the firing mode of cortical neurons. Extracellular recordings revealed that, after acetylcholine pairing, there were two main changes to peristimulus time histograms of spiking responses evoked by the paired tone: an increase in peak and an increase in width. These results demonstrate that long-lasting modifications of excitatory and inhibitory inputs are rapidly induced in vivo, and can be used to control both the spectral and temporal structure of cortical receptive fields.

15. Fixational Eye Movements and the Representation of Natural Scenes

Michele Rucci, Ramon Iovin, Gaelle Desbordes

Boston University

The retina faces a challenging encoding task. Output signals need to be represented efficiently in order to be quickly and reliably transmitted to higher processing stages. Input signals, however, exhibit a high degree of redundancy, as images of natural scenes tend to change smoothly over space and time. These statistical redundancies, if not eliminated, would cause neurons to repeatedly represent the same information and would prevent the establishment of compact representations. It has been speculated that the spatial characteristics of neurons in the retina and LGN could by themselves eliminate the broad correlations of natural scenes. However, the measured contrast sensitivity functions of individual cells do not appear compatible with a decorrelation of neural responses in the presence of naturalistic stimuli.

Under natural viewing conditions, input signals to the retina depend not only on the external stimulus, but also on the way the stimulus is sampled by means of eye movements. The projection of the visual scene onto the retina is never stationary. Even during the brief periods of fixation between macroscopic saccades, small eye movements continuously shift the retinal image by an amount that should be clearly perceptible. As the eye jitters, visual receptive fields sample information from different locations, and the input signals to individual neurons change in time. Neurophysiological studies have shown that neurons in the early visual system of macaques are sensitive to the input changes produced by fixational eye movements.

In this study, we refine our proposal that the physiological instability of visual fixation contributes to the establishment of efficient representation of natural scenes. This proposal is based on two main findings. First, we show that fixational eye movements alter the statistics of the visual signals entering the eyes of human observers in a peculiar way. By reconstructing the spatiotemporal input to the eye during free-viewing of natural scenes, we found that fixational eye movements introduce uncorrelated fluctuations in the input to pairs of nearby receptors in the retina. These fluctuations depend critically on both the statistics of natural images and the characteristics of fixational eye movements. Uncorrelated fluctuations occur only in the presence of stimuli with a scale-invariant power spectrum, as in the case of natural images.

Secondly, we show that, by including temporal frequencies to which retinal and geniculate neurons are highly sensitive, the spatially uncorrelated input signal produced by fixational eye movements is amplified by the dynamics of cell response in models of the retina and LGN. In these models, fixational eye movements strongly affect the responses of magnocellular cells. On average, neighboring magnocellular cells become progressively less correlated during a period of fixation and are completely decorrelated by the end of a typical fixation.

These results have important implications regarding the neural code used by the early visual system to represent natural stimuli.

17. A theory of optimal feature selection during visual search

Vidhya Navalpakkam, Laurent Itti

USC

How does the human visual system select relevant locations and visual features (e.g., color, orientation) in order to quickly detect desired targets in distracting backgrounds? Although recent evidence suggests that humans can select relevant locations optimally, it is not yet known whether they can select visual features optimally. Several heuristics for feature selection have been proposed in the past, such as promoting the target's features in early visual areas like V1 and V2. But the correct choice of features depends on both the desired target as well as distractors in the background.

Here, we propose the first formal theory of how prior statistical knowledge of target and distractor features modulates the response gains of neurons encoding features, such that search speed is maximized. Through numerical simulations, we show that this theory successfully explains many reported behavioral and electrophysiological observations including top-down effects such as the role of priming, the role of uncertainty, target enhancement and distractor suppression, as well as bottom-up effects such as pop-out, the role of target-distractor discriminability, distractor heterogeneity, linear separabilty and others.

Contrary to most common heuristics which suggest promotion of target features in order to detect the target, the optimal theory makes surprising predictions that target features may sometimes be suppressed, or non-target features may be enhanced. We validate these counter-intuitive predictions through new psychophysics experiments. Four naive subjects performed a difficult search for a target bar tilted 55 degrees off vertical among distractor bars tilted 50 degrees. The gains thus set up were tested by randomly inserting probe trials, in which we briefly flashed (200ms) four items representing the distractor (50 degree), the target (55 degree), relevant as predicted by the theory (60 degree), and steep (80 degree) cues. As always, the task was to search for the target and report it. Although subjects searched for a 55 degree target, as predicted by the theory, there were significantly higher number of reports on the 60 degree item (paired t-test with $p < 0.05$). These results provide direct experimental evidence that humans may select visual features optimally.

This study bears implications for further research in understanding top-down attention during visual search. For instance, previous research in physiology focused on feature gain modulation during attention to a target feature, and largely ignored the role of the distracting background features. In contrast, our research suggests that the distractors play a crucial role in determining feature gains, and may even lead to suppression of target features or enhancement of non-target features. Investigating the modulatory effects during visual search call for new experiments in physiology, brain imaging and behavior.

18. Motor adaptation as Bayesian inference

Konrad P Kording[1], Josh B Tenenbaum[1], Reza Shadmehr[2]

[1]Massachusetts Institute of Technology, [2]Johns Hopkins University

During our everyday life muscles become stronger and weaker due to various causes. Training and fatigue, oxygenation and glucose levels, wounds and healing, numerous factors influence the properties of our muscles on different timescales. The nervous system must produce reliable movements in the presence of a relentlessly changing motor plant. Observing the outcome of our own movements provides noisy information about the state of the motor plant but no information about which factor caused a deviation. The nervous system thus constantly needs to estimate the state of its motor plant in the presence of uncertainty about the size and time scale of disturbances.

Here we present a model that estimates the state of the motor plant in an optimal Bayesian way, making only minimal assumptions about the sources of muscle gain perturbations. The model is based on the idea of a motor gain. When the nervous system sends a motor command that would on average lead to a perfect movement then if the motor gain has increased to a value greater than 1 the movement would overshoot and if the motor gain is smaller than 1 the movement would undershoot. The computational reason for adaptation is to compensate for the ongoing fluctuations of the motor gain. Disturbances that change the properties of the motor plant are manifested as a change in motor gain and happen simultaneously on many timescales. We model the motor gain as the sum of many random walks at different timescales. Each over which it can be expected to τdisturbance has its own timescale decay. Each disturbance is affected by Gaussian noise. The optimal solution to this problem is a Kalman filter with one observed variable, the motor gain, but many unobserved underlying variables (the disturbances). Numerically solving the resulting equations predicts how a system should adapt when presented with artificially changing motor gains.

This model predicts many properties of several motor adaptation experiments. We compared the predictions of this model with experiments that addressed saccadic gain adaptation. In those experiments a monkey saccades from one target to another. While its eyes saccade the target is moved providing the monkey with (false) information that the motor plant got stronger or weaker. Our model predicts deviations from exponential adaptation that have been found experimentally as well as novel results on temporal integration during saccadic gain adaptation. When applied to hand movement data, the model predicts adaptation to force fields in spaced versus massed training paradigms.

We provide evidence that adaptation in the motor system is not simply a response to observed errors in movement. Instead it assumes a specific causal structure of how the motor system might be affected by perturbations such as fatigue or disease. Within this framework many phenomena of motor adaptation can be described as an optimal strategy by which the brain estimates the constantly changing properties of its motor plant.

21. Population Codes for Dynamic Cue Combination

Rama Natarajan[1], Peter Dayan[2], Quentin Huys[2], Richard Zemel[1]

[1]University of Toronto, [2]Gatsby CNU

Information about the dynamic environment is conveyed to the brain via different sensory systems in the form of spikes. Each of the sensory cues provides uncertain information due to inherent noise in the sensory transmission process and variability in the environment. When the environment variables (e.g., position) change rapidly, on the same time-scale as the neuronal spiking, there is only sparse information at any given time about the current state of the variable. By efficiently integrating information from all available cues with information available from prior experience, the brain can derive robust estimates of the current state.

In our earlier work, we proposed a computational framework for representing time-varying sensory information using dynamic population codes. When the input spikes are noisy or sparse, downstream populations need to maintain a spiking history in order to access all information about the current state. Under the hypothesis that an efficient representation should enable ready access to relevant information for downstream neurons, we suggested that a recurrently connected population of neurons recodes its input by producing spikes such that each spike can be decoded independently, in a causal manner, without referring to any spike history, thus effectively embedding the spike history in the new set of spikes. This is achieved by learning relevant prior information that captures the temporal dynamics of the environment in the synaptic connections of the recurrent population.

Here, we focus on how the re-coded information can be manipulated easily and integrated efficiently in a hierarchical manner to perform neural computations through time. Particularly, we consider how a neural system might employ the proposed coding scheme recursively, to dynamically re-weight and integrate information about current state from different sensory modalities. For example, one population can encode dynamic information from one modality such as proprioception, while another represents visual information. Both utilize the recoding scheme to efficiently represent a state variable such as position. Then a third population takes these recoded spikes as inputs, combines them, and using the same representational scheme produces a spike-based representation of the posterior distribution over current position.

We apply this approach to a dynamic cue combination setting, in which information from cues must be mapped from representations specific to each cue to a target representation, such as motor output.

We show that this framework can account for recent data on probabilistic computation during sensorimotor processing (Kording \\& Wolpert, 2004), in which human subjects moved a cursor on a horizontal plane to a visual target by moving their (hidden) finger. The cursor was shown before the start of the movement, it was then hidden, except for one point of blurry visual feedback in the middle of the movement. Unbeknownst to them, at the onset of movement, the cursor was displaced by dX, drawn from a prior distribution P(dX). The subjects must estimate dX in order to compensate for the displacement and land the cursor on the target. The key result is that subjects learned and used the prior information P(dX), and integrated it with the visual information in a way that was appropriately sensitive to the degree of blur in the visual feedback.

A trial of this experiment can be viewed as a static combination problem, where the prior over visual displacements is combined with likelihoods due to visual evidence and proprioceptive evidence to determine appropriate motor outputs to compensate for the estimated displacement. Our spiking cue combination model can account for these data, yielding the same dependence of the final cursor position on the degree of blurriness in the visual feedback. In both cases, the compensation for the estimated cursor displacement dX, as determined by the finger position at the end of the trial, is well predicted by a Bayesian combination of visual and proprioceptive likelihoods with the prior distribution.

Our model allows further predictions for extended versions of this task. For example, visual feedback could be provided more than once during the trial, and the model would predict the relationship between the estimated displacement and the particular feedback (position and uncertainty) on the trials. The model yields an interesting set of predictions in the case where the prior over cursor displacements is multimodal. Whereas standard dynamic models, such as Kalman filters, are unable to represent and update a multimodal posterior, our model can maintain and generate predictions in this case.

25. Ensemble coding of visual motion in primate retina and its readout in the brain

E.J. Chichilnisky[1], Eric S Frechette[1], Alexander Sher[2], Matthew I Grivich[2], Dumitru Petrusca[2], Alan M Litke[2]

[1]The Salk Institute, [2]University of California, Santa Cruz

Sensory experience usually depends on the ensemble activity of hundreds or thousands of neurons, but little is known about how such ensemble activity faithfully encodes behaviorally important sensory information. Even less is known about how faithfully ensemble activity is decoded by downstream neurons to control perception and behavior. Together these two issues encapsulate one of the major open problems in systems neuroscience: the nature and significance of population codes. To approach this problem, we have examined the encoding of visual motion information by waves of activity in the primate retina, and the decoding of motion information using psychophysical measurements in matched conditions.

We first examined how precisely speed of movement is specified by the ensemble activity of parasol retinal ganglion cells (RGCs) in macaque monkey retina. Parasol RGCs provide the dominant inputs to the magnocellular layers of the LGN, which in turn provide major inputs to motion sensitive areas of visual cortex. We applied a new large-scale multi-electrode recording method to measure the activity of ~100 ON and OFF parasol RGCs simultaneously in isolated retinas stimulated with moving bars. To assess the fidelity of ensemble motion signals, we estimated stimulus speed directly from recorded RGC responses using an optimized algorithm that resembles modern models of motion sensing in the brain. This revealed several fundamental aspects of the ensemble code. Overall, RGC ensemble activity encoded speed with a precision of ~1%. The elementary motion signal was conveyed in ~10 ms, comparable to the inter-spike interval, and temporal structure in spike trains provided more precise speed estimates than time-varying firing rates alone. Correlated activity between RGCs, though substantial, had little effect on speed estimates. The spatial dispersion of RGC receptive fields along the axis of motion influenced speed estimates more strongly than along the orthogonal direction, as predicted by a simple model based on RGC response time variability and optimal pooling. Finally, ON and OFF cells encoded speed with similar and statistically independent variability.

These observations provide the basis for examining the fidelity of motion readout in the central visual system. To approach this, we compared the fidelity of psychophysical speed estimation in humans to the fidelity of the retinal representation. Human observers discriminated the speed of moving stimuli with intensity, contrast, size, speed, eccentricity and spatial extent matched to the retinal experiments. Over a range of conditions, speed estimate precision was ~10%, indicating that human observers exploited only a small fraction of the information about stimulus speed transmitted by the retina. Thus, surprisingly, motion sensing is limited not by the sparse retinal representation, but instead by neural processing in the brain.

However, when tested with brief, small stimuli, human performance approached the limits imposed by the retina. This suggests that sensitivity to larger patterns of neural activity in space and time is limited by the spatial scale or temporal integration of downstream computations. We therefore examined separately the dependence of human motion sensing on stimulus duration and size. Comparison to optimal retinal readout revealed that strictly local computations performed by elementary motion detectors in the brain impose a major limit on performance. Thus, readout of population codes in the brain can be highly efficient, subject to constraints imposed by the scale of the elementary neural computation.

26. Structure of interneuronal correlations in the primary visual cortex of the rhesus macaque

Andreas S Tolias[1], Alexander Ecker[1], Georgios A Keliris[1], Thanos G Siapas[2], Stelios M Smirnakis[1], Nikos K Logothetis[1]

[1]Max-Planck-Institute, [2]Caltech

Despite recent progress in systems neuroscience, basic properties of the neural code still remain obscure. For instance, the responses of single neurons are both highly variable and ambiguous (similar responses can be elicited by different types of stimuli). This variability/ambiguity has to be resolved by considering the joint pattern of firing of multiple single units responding simultaneously to a stimulus. Therefore, in order to understand the underlying principles of the neural code it is important to characterize the correlations between neurons and the impact that these correlations have on the amount of information that can be encoded by populations of neurons. Here we applied the technique of chronically implanted, multiple tetrodes to record simultaneously from a number of neurons in the primary visual cortex (V1) of the awake behaving macaque, and to measure the correlations in the trial-to-trial fluctuations of their firing rates under the same stimulation conditions (noise correlations). We find that, contrary to widespread belief, noise correlations in V1 are very small (around 0.01) and do not change systematically neither as a function of cortical distance (up to 600 um) nor as a function of the similarity in stimulus preference between the neurons (uniform correlation structure). Interestingly, a uniform correlation structure is predicted by theory to increase the achievable encoding accuracy of a neuronal population and may reflect a universal principle for population coding throughout the cortex.

28. Prefrontal neural correlates of errors in a sequential decision making task

Bruno B Averbeck, Daeyeol Lee

University of Rochester

To understand the neural mechanism for sequential decision making, we developed a task in which monkeys had to execute sequences of eye movements for a juice reward. Each sequence consisted of 3 movements selected from one of 10 possible movements, such that all of the movements occurred in multiple sequences. There were a total of 8 different sequences, and the task was structured such that one of these sequences was randomly selected for a block of trials, and the sequence had to be executed correctly 10 times before the block was complete. To initiate a trial, the animal acquired a fixation target. Following a fixation period, the animal was given a choice between two targets. If the animal made a saccade to and fixated the correct target for 500 ms, it was presented with two more targets, one of which was correct for the next movement in the sequence. If the animal selected the wrong target, it had to return to its previous fixation, and it was presented with the choice targets again and allowed to make another choice. In this manner, the animal discovered, through trial and error, which sequence was correct in the current block. After executing the sequence correctly ten times, i.e. without making any incorrect decisions in a trial, a new sequence was randomly selected. We recorded the activity of 495 neurons from the dorso-lateral prefrontal cortex of two monkeys while they performed this task.

Using ensemble decoding techniques, we compared the neural activity from error and correct trials during the intertrial interval and the initial hold period. We found that neural activity during the intertrial interval preceding correct trials indicated the sequence that was correct in the current block. The strength of this prediction increased during the hold period, becoming strongest just before the animals began executing the sequence. We also found that, on error trials, the neural activity during the intertrial interval and the hold period did not predict the sequence which the animal was supposed to execute. The divergence in neural activity between error and correct trials occurred quickly after the end of the previous trial. The fact that the neural activity on error trials did not predict what the animal was supposed to do raised the question of whether or not the error trial activity was predicting a different sequence, or was unrelated to any of the sequences. To test this hypothesis, we used the sequence that was predicted by the neural activity on error trials to predict at which point in the sequence the animal was going to make a mistake. This was possible because the sequence that the animal was supposed to execute and the sequence that the animal was planning could diverge on either the first, second or third movements. Thus, if we know which sequence the animal was supposed to execute and which sequence the neural activity predicted we could predict the movement on which the animal would first choose the wrong target. By building a contingency table of the predicted and actual movement on which the animal made its mistake, we were able to test the hypothesis that on error trials the animal was planning an incorrect sequence. In fact, we found that we were able to predict, a statistically significant number of times, the movement on which the animal would make an error, based upon the sequence predicted by the activity during the initial hold period.

We also compared perisaccadic activity between correct and incorrect movements. We found that when the animal selected the wrong target the neural activity predicted the animal's choice, and not what the animal should have done. However, when we compared the predictions of the target the animal selected incorrectly, to the predictions of the target the animal selected correctly, we found that the neural activity predicted the target towards which the animal moved more accurately for correct movements. In other words, the neural activity did not predict the wrong target as reliably as the correct target. This finding is related to the pattern of dynamic changes in neural activity we observed while the animal was discovering the correct sequence in a new block of trials. Specifically, we have previously shown that as the animals figured out the correct sequence, the neural activity predicted the correct sequence more accurately. These results together indicate that the animal's uncertainty about the correct sequence of movements is reflected in the neural activity.

30. Free choice increases synaptic interactions between frontal and parietal cortex

Bijan Pesaran[1], Matthew J Nelson[2], Richard A Andersen[2]

[1]New York University, [2]Caltech

Choices organize our daily behavior. In some situations we choose to follow instructions, such as when we are driving, but quite often our choices are made freely from multiple alternatives. We make these free choices by integrating many factors in a decision process. Experiments have shown that neural activity in areas of the frontal and parietal cortex is correlated with both movement and decision theoretic parameters. It is likely therefore that these areas are also involved in the integration of factors that comprise the decision process. How do these areas integrate information necessary for making a free choice? Areas specialized for similar movements are connected by long-range projections that go from frontal to parietal cortex as well as from parietal to frontal cortex. Therefore, synaptic interactions across these projections may support the integration of information between these areas that is necessary to make a free choice for a reach.

Despite the widespread connectivity between brain areas, identifying synaptic interactions between brain areas and their functional correlates remains a significant challenge in cognitive neuroscience. Analysis of local field potential (LFP) activity and spiking activity in different areas could provide a resolution to this problem. The LFP pools synaptic activity in the vicinity of an extracellular recording electrode, and there is increasing evidence it can be an informative measure of neural processing. Since the LFP pools synaptic activity in a region of cortex, correlating spiking activity recorded in one brain area with LFP activity simultaneously recorded in another could allow the identification of synaptic interactions between those areas and their functional correlates.

To dissociate processes related to free choice from those related to following instructions, we trained two monkeys to do two tasks, with freely chosen and instructed reaches. In both tasks, monkeys were trained to make sequences of reaches to visual targets to find juice rewards and across trials the rewards were assigned to the targets randomly with equal probability. In the choice task, all three targets were circles and the monkeys were allowed to choose reaches to find the reward. In the instructed task, targets were a circle, square and triangle and the monkeys had to reach to them in that order to find the reward. On a trial-by-trial basis the instructions given to the monkey were matched to the choices the monkey had already made in that session. Thus, the two tasks were yoked to have the same sensory, motor and reward-related properties but different choice requirements.

We identified whether there were synaptic interactions by simultaneously recording spiking activity from individual neurons and LFP activity in both the medial intraparietal area (MIP) of parietal cortex and the dorsal premotor area (PMd) of frontal cortex. To measure the strength of synaptic interaction between these areas we estimated the spike-field coherency on a 500ms window stepping 50ms between estimates for each pair of measurements while the monkeys performed the free choice task and instructed task up to the time of the first reach. We did not find significant spike-field coherency between areas during the late delay period. However, coherency for 49 PMd spike – MIP field recordings and 23 MIP spike – PMd field recordings did achieve statistical significance for the early delay period of either the free choice or instructed tasks ($p<0.05$).

Figure 1 shows the population average PMd spike – MIP field coherency for these recordings for the free choice task and instructed task. This shows a transient significant PMd spike – MIP field coherency for both tasks during the early delay period that is stronger for the free choice task compared with the instructed task ($p<0.01$; t-test). This increase is present only in a 5-15Hz frequency band during that epoch and is not present either during the baseline or late delay period. Figure 2 shows that the effect is also present in the population average MIP spike – PMd field coherency. Therefore, synaptic interactions are present between frontal and parietal cortex immediately following target onset and these interactions are stronger during free choice than instructed behaviors. Long-range coherency may reflect synaptic interactions due to cooperative computation between frontal and parietal cortex during decision making. These interactions may serve to integrate information needed to make a free choice.

Support Contributed By: NEI,ONR and Boswell Foundation to RAA and a Career Award from the Burroughs-Wellcome Fund to BP

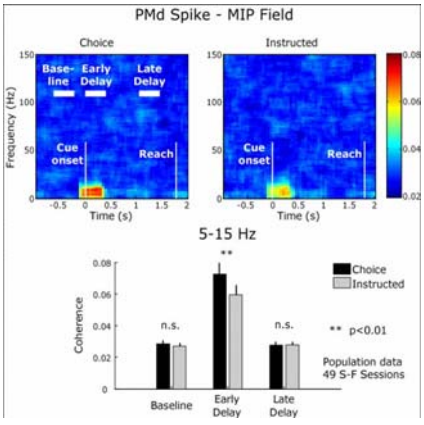

Figure 1: Population average PMd spike – MIP field coherence from cue onset through the first reach estimated on a 500 ms window moving 50 ms between estimates. Upper left panel: Free choice task. Upper right panel: Instructed task. Lower panel. Average 5 – 15Hz coherence is significantly greater during choice than instructed task in the early delay period.

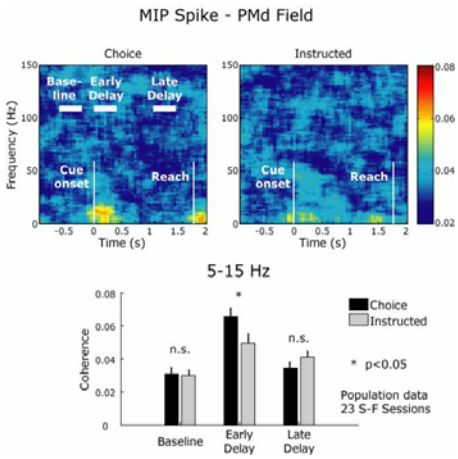

Figure 2: Population average MIP spike – PMd field coherence in similar format to Fig 1.

31. Optimal Neural Decision Boundaries for Maximal Information Transmission

Tatyana Sharpee[1], William Bialek[2]

[1]University of California, San Francisco, [2]Princeton University

We consider here how to optimally encode multidimensional signals using a spike or its absence as two possible outcomes in order to maximize the mutual information transmitted about those signals. Our goal is to understand the optimal shape of contours in the input space that would separate stimuli leading to a spike from stimuli that do not elicit spikes, and how this optimal shape changes with the average rate of firing a spike. Our assumption is that most of the variability in neural response will be caused by stimuli near the spiking decision boundary, and that the width of the uncertainty region σ is much smaller than the characteristic length scales over which the probability distribution of inputs changes. Stimuli that are far away from the decision boundary will elicit an almost certain response.

For a single neuron, the problem of maximizing information transmission given an average spike probability is equivalent to minimizing the noise entropy. The noise entropy is proportional to an integral over the probability distribution of inputs P(r) along the decision contour. Solving the variation problem with respect to the contour's shape, we derive a general equation for the decision boundary valid for arbitrary probability distributions P(r) that relates its curvature κ, scaled by the noise level σ, to the probability distribution of inputs: $\lambda + \kappa(r) + \sigma\, n \times \tilde{N} \ln P(r) = 0$. Here, n is the normal to the decision contour at a point r, and λ is the Lagrange multiplier set to satisfy the requirement from the average spike probability. Solving this equation for Gaussian inputs, we show that the optimal contours in this case are straight lines. In other words, neurons that are optimally designed to process Gaussian inputs should be sensitive to only one stimulus dimension, which is in the direction perpendicular to the decision boundary. In addition, for the case of a single neuron considered here, there are no restrictions on the orientation of this optimal contour, even if the Gaussian stimuli are correlated.

However, neurons in several sensory modalities, as well as Hodgkin-Huxley model neurons, have been shown to be sensitive to several stimulus dimensions. Is this a sign of neural non-optimally? Signals derived from natural environment can sometimes be more closely approximated by an exponential, as opposed to a Gaussian, distribution. As an example, we considered two-dimensional exponentially distributed inputs $\sim \exp(-|x|-|y|)$. The optimal decision contours in this case can be solved for exactly. They have different shapes depending on the average spike probability, but are almost always curved. At extreme spike probabilities (either close to 0 or to 1), the optimal contours are similar to squares with rounded corners. For spike probabilities near ½, the optimal contours extend to infinity, resembling a wedge with a rounded corner. For spike probability exactly equal ½, the optimal solution is a straight line, x=y, separating stimuli diagonally in two halves. The ubiquity of non-Gaussian signals in nature, particularly of the exponential distributions considered here, makes these results relevant for neurons across different sensory modalities.

Support contributed by the Swartz Foundation and NIMH K25MH068904.

32. Testing Barlow's "Efficient Coding Hypothesis": Are sensory neurons really matched to natural stimuli?

Andreas VM Herz

Bernstein Center for Computational Neuroscience Berlin

According to Barlow's seminal "Efficient Coding Hypothesis", the coding strategy of sensory neurons should be matched to the statistics of stimuli that occur in an animal's natural habitat. Many experimental findings support this hypothesis. For example, it has been demonstrated that single neurons can use their full output capacity by precisely matching their tuning curve to the statistics of natural stimuli. On the system's level, the hypothesis has helped to elucidate the organization of early visual cortex and time-frequency trade-offs of the cochlear filter bank. There is also substantial evidence that sensory systems convey far less information about simplified artificial stimuli than about signals that resemble behaviorally relevant stimuli. Recent studies even suggest that adaptation shifts the tuning curves of neurons dynamically to keep up with slow variations in the statistics of environmental stimuli.

Despite these impressive results, the Efficient Coding Hypothesis is not as assumption free as one may think at first sight. In particular, any specific prediction derived from the hypothesis crucially depends on the measured (or assumed) distribution of natural stimuli. However, even for animals living in the very same habitat, the behaviorally relevant stimuli may vary strongly from species to species. Studies that are based on a heuristic collection of natural stimuli may thus undersample or even completely miss features that are of utmost importance for the investigated species. Furthermore, a sensory system may efficiently process stimuli that are not even encountered in nature. For instance, animals often prefer artificially enhanced communication signals over those of potential mating partners.

Accordingly, sensory systems may not be optimized for natural stimuli per se; rather, their optimization may be heavily biased towards behaviorally relevant natural stimuli. Using a novel automatic search technique, we test this idea and identify stimulus ensembles that sensory neurons are optimized for. We focus on grasshopper auditory receptor neurons to study a system with well-characterized natural stimuli and rich temporal response phenomena that can be recorded for long time periods in vivo. The investigated receptor cells are therefore well suited for tackling questions of coding efficiency and read-out mechanisms on different time scales. We find that their optimal stimulus ensembles clearly differ from the natural environment, but largely overlap with a behaviorally important subensemble of natural sounds. This indicates that the receptors are optimized for peak rather than average performance. More generally, our results suggest that the coding strategies of sensory neurons are heavily influenced by differences in behavioral relevance among natural stimuli.

Apart from these findings, our approach offers an efficient on-line search algorithm for investigating high-dimensional stimulus spaces, clarifes the role of the read-out mode for optimal coding schemes, and provides a natural extension of the concept of a "tuning curve" beyond the traditional firing-rate picture. These aspects will be most helpful for studying the information processing capabilities of a wide range of sensory systems.

Joint work with C.K.Machens, T.Gollisch, and O.Kolesnikova.

Selected results have appeared in C.K.Machens et al., Neuron 47, 447-456 (2005).

35. The strategy for echolocating bats to shift attention from one target to another measured by a telemetry microphone system

Hiroshi Riquimaroux, Shizuko Hiryu, Yoshiaki Watanabe

Doshisha University

We have investigated how echolocating bats extract important information from multiple targets during flight by means of a telemetry microphone system (Telemike) mounted on the head. The system allowed us to monitor what the bat listened to during its flight. For this purpose a small and very light on-board microphone system was developed. A high-speed video camera system was also adopted together with the Telemike to trace positions in space of the flying bat for analyzing spatio-temporal pulse emission patterns and flying velocity. With those devices, how the bats executed parallel time-sharing real-time processing during their flight could be investigated. Taiwanese leaf-nosed bats (Hipposideros terasensis), CF-FM bats, were used as subjects. Measurements were done within a steal-walled flight chamber. The Doppler-shift compensation was confirmed during flight where the frequency of the dominant second harmonic of constant frequency component, CF2, of returning echoes was kept constant by adjusting CF2 frequency of emitting pulses. Estimated frequencies of CF2 of pulses from the bats' flight speed strongly agreed with actually observed values. Pulse frequencies were also estimated using echoes returning directly ahead of the bat and from its sides for two different flight conditions: landing and U-turning. Inter-pulse intervals and echo delays were also measured. Temporal patterns for Doppler-shift compensation, inter-pulse interval and echo delay showed that bats in flight alternately changed their attended direction from one to another. Bats were also observed to change the intensity and emission rate of pulses depending on the distance from the landing site.

[Work partly supported by the Innovative Cluster Creation Project promoted by MEXT and by a grant to RCAST at Doshisha University from MEXT.]

37. Temporal discounting activity in parietal neurons during intertemporal choice

Kenway Louie, Paul W. Glimcher

New York University

Decision-making requires the transformation of both current sensory information and prior knowledge of the environment into the selection of an action. An attractive mechanism for action selection is the integration of sensory and stored information into a 'value map', a topographic map of possible actions from which a single (highest valued) action can be chosen. Previous work has demonstrated decision-related neural activity of this type in the posterior parietal cortex. During visual-saccadic decision-making in macaques, the activity of saccade-generating neurons in the lateral intraparietal area (LIP) represents reward-related variables such as prior probability, reward expectation, and local reward income which control the value of the associated saccades. Consistent with an internally generated value map, neurons in area LIP encode the subjective rather than objective value of rewards and actions.

A crucial element in choice is the temporal relationship between action and reinforcement. In behavioral studies, the value of a reward is strongly influenced by the wait required to receive that reward. The subjective value of a reinforcer is a decreasing function of delay to reinforcement, a phenomenom known as temporal discounting. Such preference for smaller immediate rewards over larger later rewards is ubiquitous in animals and humans, even when such choices result in lower overall rates of reinforcement.

We investigated the activity of LIP neurons in monkey performing an intertemporal choice task. In each trial animals chose between a smaller immediate reward and a larger delayed reward, with the delay varied across trials in a block protocol. We found that monkeys exhibit consistent temporal discounting behavior, choosing the larger reward at short delays and the immediate reward at long delays. Preference data from different magnitudes of delayed reward were used to generate a discount function. Fit with a standard hyperbolic discount model, these data show rates of discounting intermediate between the faster discount rates of pigeons and rats and the slower ones reported in humans.

We found that the activity of LIP neurons is consistently modulated by the delay to reward. This modulation is strongest early in the trial, immediately after visual target onset, and weakest late in the trial, during the generation of saccades. During the visual epoch, LIP activity is a decreasing function of delay, and shows a close parallel to the behavioral discount curve derived from the monkey choice data: a psychometric-neurometric match. Furthermore, LIP neurons demonstrating decreased activity with delay also show increased firing rates with larger reward magnitude, consistent with an integrated representation of subjective value. These results suggest that early in the decision process, LIP activity encodes time-discounted subjective value, while in contrast late in the decision process, after appropriate action selection, LIP activity shifts to encode the upcoming saccade.

39. Prediction: Linear and Nonlinear Synaptic Integration zones in Basal Dendrites of Neocortical Pyramidal Cells

Bardia F Behabadi[1], Alon Polsky[2], Jackie Schiller[2], Bartlett W Mel[1]

[1]University of Southern California, [2]Technion Medical School, Haifa, Israel

The thin basal and apical oblique dendrites of pyramidal neurons receive a large fraction of these cells' excitatory inputs. Yet how the synaptic inputs to thin branches are integrated remains an important open question. A recent study of layer 5 pyramidal neurons in neocortical slices (Polsky, Mel & Schiller, 2004) found that NMDA spikes contribute to a highly nonlinear (i.e. sigmoidal) summation rule when two closely spaced synaptic inputs are activated on the same thin branch. In contrast, when a distal electrode is held fixed while a proximal electrode is moved gradually closer to the cell body, summation grows increasingly linear, just as for inputs delivered to two different branches. We used biophysically detailed compartmental models of pyramidal neurons to study this distance dependence. We found that, contrary to our expectation, the gradual decay of the nonlinear interaction between the two synaptic inputs is most likely not explained by the separation of the two electrode per se. Instead, our simulation results suggested that the nonlinear excitability of the branch, as indicated by the NMDA/AMPA peak conductance ratio, may be several times higher at distal sites compared to more proximal ones. Awating experimental confirmation, this finding leads us to the speculative conclusion that pyramidal cells maintain a relatively linear integration zone near the cell body, perhaps contributing to a neuron's "classical" receptive field, and a more nonlinear integrative zone distally, perhaps devoted to the processing of higher-order contextual inputs.

40. Response properties and synchronization of dendritic neurons: theory and experiment

Joshua A Goldberg, Chris Deister, Charles J Wilson

University of Texas at San Antonio

In this study, we consider the impact of dendritic nonlinearities on the synchronization properties of rhythmically firing neurons. We study a model of a generic Hodgkin-Huxley somatic oscillator coupled to a long dendrite. The propensity of the oscillator to synchronize to its inputs is captured by its phase response curve (PRC), which relates how small somatic voltage perturbations affect the timing of the next action potential. Combining cable theory and the theory of phase-coupled oscillators we derive an analytical expression for the relationship between the neuron's PRC and its response properties to distal dendritic inputs, expressed by its dendritic phase response density (PRD). We show that the dendrite induces a leftward shift in the PRD and attenuates it relative to the somatic PRC due to the delayed and attenuated effect of the distal perturbation on the somatic oscillator. As a result, and in agreement with the study of Crook et al. (*J Comput Neurosci.*, 5(3):315-29, 1998), we find that a synapse can induce in-phase synchronization or out-of-phase synchronization depending on its location along the dendrite. Additionally, we show that the neuron can regulate where the transition between these responses occurs along its dendrite simply by changing its firing rate. Using the quasi-active membrane approximation for a nonlinear dendrite we distinguish between amplifying currents such as the persistent sodium current and resonating currents such as the HCN current, and elucidate how these general classes of nonlinearities alter the response properties of the neuron to dendritic inputs in comparison to somatic ones. In particular, resonating currents cause a downward shift in the PRD implying that a dendritically-located fast excitatory synapse can synchronize the neuron even when it induces a phase-delay when located somatically. Finally, we used perforated-patch recordings in rat slices of autonomously firing globus pallidus neurons to compare the PRC measured in response to depolarizing somatic pulses to that measured in response to focal stimulation of the slice, which presumably activated distal synapses, while blocking $GABA_A$ and NMDA receptors. As predicted by our theoretical results, the latter is shifted leftward and is attenuated relative to the former.

41. Bayes points the way: an optimal strategy for growth cone chemotaxis

Duncan Mortimer[1], Peter Dayan[2], Kevin Burrage[1], Geoffrey J Goodhill[1]

[1]University of Queensland, [2]Gatsby Computational Neuroscience Unit

Molecular gradients provide important directional cues for growing axons in the developing nervous system. We have recently shown using a novel chemotaxis assay that axonal growth cones can detect differences as small as one molecule across their spatial extent (Rosoff et al, Nat. Neurosci, 7:678-682). Given inherent stochastic fluctuations in receptor binding, how is this remarkable sensitivity achieved? Here we develop a Bayesian account of ideal gradient detection by small sensing devices. We consider the simplest case of a one-dimensional device which is trying to determine whether a locally uniform external gradient points left or right. This sensor does not know the steepness of the gradient, the baseline concentration or the spatial distribution of its receptors. The optimal strategy in this case is to choose the direction which is most likely to have generated the observed binding statistics. Given simple prior distributions for gradient steepness and absolute concentration, we obtain a theoretical expression which characterises the extent to which signals from receptors on the growth cone surface provide information about the gradient direction. Signals originating from more peripheral positions are more informative and thus we predict the ideal form of spatial amplification by the underlying signaling network. In addition, we show that inevitable asymmetries in receptor distribution lead to bias effects which become more significant at high background concentrations. We hypothesise that receptors are actively redistributed by the growth cone in order to reduce this bias. We thus derive a sensitivity curve which fits well with our experimental data, and predict that receptor trafficking should increase at higher concentrations.

42. Studies of dendritic spike initiation and propagation in CA1 pyramidal cell models

Yael Katz, William L Kath, Nelson Spruston

Northwestern University

Dendritic excitability is determined by the complex interplay of morphology and the properties and distributions of ion channels in the neuronal membrane. In many CA1 pyramidal neuron models, terminal dendritic branches are much more excitable than their non-terminal counterparts. In our passive model, (1) for instance, a fixed synaptic conductance generates EPSPs approximately three times larger in terminal versus non-terminal branches. This implies that a much smaller synaptic conductance would be required to generate dendritic spikes in terminal branches than in non-terminal ones, and our active model (2) shows that this is indeed the case.

Although these models may be correct, from the point of view of information processing, it may be important for synaptic weights to be independent of position along the dendrite. Therefore an alternative is that neurons adjust their ion channel distributions to eliminate the privileged position of synapses on their terminal branches. Although in principle there are multiple strategies a neuron could implement to achieve this, we investigate potassium channel gradients since there is evidence that potassium channel expression is modulated in CA1 pyramidal cells. (3,4) We propose a minimal model with a uniform sodium channel density and a potassium channel distribution chosen such that dendritic spikes are initiated with the same synaptic conductance at all basal, oblique, and apical tuft dendritic locations. Constraining potassium channel density in the main apical dendrite according to the results of Hoffman and colleagues, (3) larger synaptic conductances are required to initiate spikes in this large dendritic compartment. Our model accounts for the measured voltage attenuation of backpropagating action potentials both along the primary apical branch (2) and along the oblique side branches. (4) Whether or not terminal branches are in fact hot spots for spike initiation has important consequences for dendritic integration. This work points to the need for additional experiments to determine potassium channel distributions.

We used our models to investigate under what conditions a dendritic spike, once initiated, is able to propagate into its parent branch. We found that a dendritic spike on a single daughter branch usually fails to propagate forward. Propagation into the parent branch is achieved when dendritic spikes on two daughter branches are coincident and/or when the parent branch is depolarized. A dendritic spike on one daughter branch that coincides with an EPSP on a second daughter branch is not sufficient for propagation of the dendritic spike. The implications of this finding are important for understanding synaptic integration and plasticity, processes that depend on spikes initiated in dendrites as well as backpropagating action potentials.

References:

1. Golding NL, Mickus TJ, Katz Y, Kath WL, Spruston N. (2005) Factors mediating powerful voltage attenuation along CA1 pyramidal neuron dendrites. J. Physiol. 568(1):62-82.

2. Golding NL, Kath WL, Spruston N. (2001) Dichotomy of action-potential backpropagation in CA1 pyramidal neuron dendrites. J. Neurophysiol. 86(6):2998-3010.

3. Hoffman DA, Magee JC, Colbert CM, Johnston D. (1997) K+ channel regulation of signal propagation in dendrites of hippocampal pyramidal neurons. Nature 387(6636):869-875.

4. Gasparini S, Losonczy A, Magee J. Features of back-propagating spikes in radial oblique dendrites of CA1 pyramidal neurons. (2005) Society for Neuroscience abstract.

43. Problems in Learning Efficient Nonlinear Representations: examples with quadratic codes relating to spike-triggered covariance analysis

Mark V. Albert, David J. Field

Cornell University

When applied to natural images, linear ICA and sparse coding algorithms produce filters that mimic the gabor-like receptive field structure in V1. However, these models are incapable of capturing any of the well-known nonlinearities in such neurons (such as the traditional energy model of V1 complex cells). A number of nonlinear approaches to efficient neural coding have been pursued, but these approaches must deal with a number of potential problems. 1) high dimensionality – if signals or sources combine multiplicatively. 2) signal reconstruction – which is difficult or impossible in some models, but often a requirement of an efficient representation. 3) overcomplete representation – selecting from multiple ways of representing the same signal. 4) demonstration – there is no straightforward way to display the results of most nonlinear learning algorithms.

In this poster, quadratic forms are used as an illustrative nonlinearity to approach these problems in the context of independent coding. Although the mathematics of these forms are simple, the computations quickly become intractable (O(n^4) parameters for an 'n x n' pixel patch) necessitating dimensionality reduction. This reduction invalidates certain mathematical assumptions, such as the equivalence of initial linear transforms. We will explore an independent quadratic demixing model with initial pixel, Fourier, wavelet, PCA, and ICA bases and show the bias in the hyperquadric geometries that results from the dimensionality reduction.

The difficulties in displaying, understanding, and analyzing these geometries are also present in recent electroneurophysiology approaches, such as spike-triggered covariance analysis. Quadratic form filters and spike-triggered covariance models can be displayed similarly – for instance, as a sum of squares of orthogonal filters, but we will demonstrate the interpretative bias that can result from these representations. We will include display methods with multiplicative inhibition and excitation in our analysis.

Most importantly, these results will be put in the context of current research efforts in high-dimensional efficient nonlinear coding and spike-triggered covariance analysis. Supported by NGA contract HM 1582-05-C-0007 to DJF.

44. Effects of variable inhibition on spike timing precision in the olfactory bulb

Maxime Ambard, Dominique Martinez

LORIA (France)

The olfactory bulb (OB) is a network mainly composed of excitatory neurons, the mitral cells (MCs), interconnected via local inhibitory neurons, the granule cells (GCs). Experimental evidence tends to show that spike timing in MCs plays a significant role in encoding olfactory information (Friedrich et al, 2004). A pre-requisite is that MCs fire spikes in a precise and reproducible way over repeated presentations of the same stimulus. However, GABAergic inhibition released by the GCs and received by the MCs is asynchronous and variable across repeated trials (Schoppa et al. 1998; Urban and Sakmann, 2002). In this work, we study the effects of such variable inhibition on MC spike timing precision using both computer simulations and mathematical analysis.

Simulations are first performed on a type I quadratic integrate-and-fire (QIF) model of MC. To be more in adequation with biological data (subthreshold oscillations in MCs), this MC model was then extended to a type II excitability by using an additional variable representing a recovery current. Under constant but noisy stimulation, we found that sensitivity to noise and initial conditions drives these neurons towards an unreproductible behavior. It is well known that a strong inhibitory input resets the phase of a neuron and tends to eliminate the effect of its initial conditions (Börgers and Kopell, 2003). However, when the number and the timing of the inhibitory synaptic events are variable across repeated trials, the effect of this variability on the spike timing precision is not known. To quantify this relation, we have derived an analytical expression for the spike output jitter as a function of the variability of the received inhibition. Theoretical and experimental predictions are in agreement.

This study predicts that variable inhibition is especially tolerated as the number of inhibitory cells is large, which is consistent with experimental data from early olfactory systems (antennal lobe for insects, olfactory bulb for vertebrates). In insects, where there are fewer inhibitory neurons than principal cells, the inhibitory synaptic events from different inhibitory neurons must be released synchronously in order to preserve the spike timing precision in principal cells. In the mammalian olfactory system, where the number of inhibitory neurons is higher than the one of the principal cells, the inhibitory synaptic events can be released asynchronously without deteriorating the spike reliability.

(Friedrich et al., 2004) : R.~W. Friedrich, C.~J. Habermann, G.~Laurent, Multiplexing using synchrony in the zebrafish olfactory bulb, Nature Neuroscience 7 (2004) 862–871.

(Urban and Sakmann, 2002) : N. Urban and B. Sakmann, Reciprocal intraglomerular excitation and intra- and interglomerular lateral inhibition between mouse olfactory bulb mitral cells, Journal of Physiology (2002) 355–367.

(Schoppa et al., 1998) : N. Schoppa, J.M. Kinzie, Y. Sahara, T.P. Segerson, and G.L. Westbrook, Dendrodendritic inhibition in the olfactory bulb is driven by NMDA receptors, The Journal of Neuroscience, 18 (1998) 6790–6802

(Börgers and Kopell, 2003) : C. Börgers, N. Kopell, Synchronisation in network of excitatory and inhibitory neurons with sparse, random connectivity, Neural Computation 15 (2003) 509–538.

45. Population Coding in V1

Charles H Anderson[1], Gregory C DeAngelis[1], J A Movshon[2]

[1]Washington Univ. School of Medicine, [2]New York University

With population codes, the number of neurons is proportional to the signal to noise ratio (SNR) of each degree of freedom they represent. This hypothesis allows one to use the measured distribution of the preferred spatial frequencies (SFs) of cells in V1 to quantify how SF information is transmitted down the optic nerve and represented in V1. This work describes the results of such an analysis on several hundred monkey simple and complex cells measured by J. Cavanaugh and W. Bair in the Movshon lab[1]. This SNR analysis explains a long outstanding question as to why the preferred SFs of most cells in V1 lie far below the highest SF provided by the ganglion cells, while almost all the coefficients in a wavelet decomposition of an image reside at the highest SFs. Qualitatively, the answer is that there is much more power at low SFs because of the 1/f^2 power spectrum of natural images and so more neurons are allocated to represent the high SNR available at low SFs. This analysis quantifies how the ganglion cells are organized to transmit the high SNR available at low SFs using a highly redundant coding scheme, rather than being decorrelated by "whitening" spatial filters. The redundancy of the representation within V1 is found to be on the order of 100 neurons per wavelet coefficient at low SFs relative to the redundancy at the highest SFs, in conflict with the sparse coding hypothesis. A typical simple cell in V1 gets inputs from ~1000 ganglion cells pooled over a radius of 2-3mm, which is large compared to the 1mm prototypical Hubel and Weisel cortial column (see also [1]). In addition, the distribution of the preferred SFs of complex cells is shifted to higher SFs relative to that of simple cells by almost an octave, suggesting complex cells may not get their inputs from simple cells. In summary, this engineering analysis quantifies the way V1 cells participate in population codes to represent a much lower dimensional space, which adds to the growing weight of evidence that population codes are the primary way neuronal systems represent and process information [2].

[1] Cavanaugh JR, Bair W, and Movshon JA (2002), Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons, J Neurophysiol. 88: 2547-56.

[2] Eliasmith C, and Anderson CH (2003), "Neural Engineering", MIT Press.

46. Context dependence of neural responses in rat primary auditory cortex

Hiroki Asari[1], Hysell Oviedo[2], Anthony M Zador[2]

[1]Watson School of Biological Sciences / Cold Spring Harbor Laboratory, [2]Cold Spring Harbor Laboratory

How do past events influence auditory cortical responses? Psychophysical studies have demonstrated that stimulus context strongly affects perception and auditory scene analysis in humans. Physiological studies on the effects of stimulus context (*e.g.* forward masking, background and foreground representations) suggest that the responses of auditory cortical neurons are highly dependent on stimulus history. However, the responses cannot be fully explained by linear encoding (spectro-temporal receptive field) models despite their high trial-to-trial reliability, and little is known about the details of the underlying neuronal dynamics. To examine the stimulus context dependence of neural responses in an attempt to characterize the relevant time scales and the nonlinear encoding properties of neural computations, we used whole-cell recordings *in vivo* in the anesthetized rat to examine the subthreshold responses to snippets of natural sounds, including animal vocalizations. We found that context dependence in the rat auditory cortex decays within 1 second on average, but it can last more than 4 seconds in some neurons. We are also investigating which stimulus properties, such as intensities and frequency bandwidth, contribute to the context dependence. In addition, we are trying to relate the time constants to cortical layers and/or cell types by filling the characterized neurons with fluorescent dyes.

47. Bayesian inference with probabilistic population codes: Theory

Jeffrey M Beck[1], Weiji Ma[1], Alexandre Pouget[1], Peter Latham[2]

[1]University of Rochester, [2]Gatsby Computational Neuroscience Unit

Many experiments have shown that human behavior is nearly Bayes optimal in a variety of tasks. This implies that neural activity is capable of representing both the value and uncertainty of a stimulus, if not an entire probability distribution. Here, we argue that the observed variability in neural activity is ideally suited for the representation of the uncertainty. Specifically, we note that Bayes' rule implies that a variable pattern of activity is, in fact, a natural, implicit representation of a probability distribution for the value of a stimulus. Of course, it is by no means clear that the various operations which cortical circuits may perform are capable of manipulating, combining, or decoding such a representation efficiently or otherwise.

We address this issue through the construction of a Probabilistic Population Code, or PPC. Our model consists of two elements: a neural operation, in which the activities of two or more populations of neurons are combined according to biologically plausible rule, and an associated operation on the probability distributions of the variables represented in those population codes, which are obtained through an application of Bayes' rule. Of course, the restriction that only a Bayesian decoder may be used implies that a given operation pair is not necessiarly compatible with any given stimulus dependent distribution of the neural activity. For example, suppose two independent layers of cortex coding for the same stimulus exhibit Poisson variability and have identically shaped tuning curves, although possibly with different amplitudes. Associated with each population code is a distribution on the stimulus which can be obtained using Bayes' rule. Adding the two populations yields a third population code, such that, when Bayes' rule is applied, the resulting distribution corresponds to the product of the distributions of the stimulus for the initial two populations. Thus, we say that independent, Poisson distributed populations with identically shaped tuning curves and cortical circuits which are capable of addition are compatible with a Probabilistic Population Code which performs optimal cue combination.

We generalize this result and show that linear combination of population activities is equivalent to Bayes-optimal cue combination for a large family of probability distributions (the exponential family with linear sufficiency statistics) and argue that this family is biologically relevant. Moreover, we show that in addition to optimal cue combination, this particular PPC can be made to combine tuning curves of different shapes, implement priors, increment or decrement information, and can be used to track a time varying stimulus when divisive normalization is allowed.

48. Adaptation and the role of temporal precision in the visual code

Daniel A Butts[1], Chong Weng[2], Jianzhong Jin[2], Chun-I Yeh[2], Nick A Lesica[1], Jose-Manuel Alonso[2], Garrett B Stanley[1]

[1]Harvard University, [2]SUNY - State College of Optometry

Our understanding of the visual neuron function is rooted in the receptive field, which describes the location and temporal structure of visual stimuli that evoke responses. However, there are many aspects of the neuronal response that are not explained by its receptive field, including the fine temporal precision of the spike train itself as well as how its response properties adapt to different sensory environments. To identify which features of the response are likely important in the neural code, we observe how the responses of LGN neurons change as a function of stimulus contrast using both white noise and natural movie stimuli. Since it is likely that the role of the neuron in visual processing does not change as a function of contrast, aspects of the response that are not preserved at different contrasts likely do not play a role in the neural code.

We find that changes in the response properties of LGN neurons generally act to preserve the stimuli that neuron encodes and its overall information rate at the expense of the detailed features of the spike response. In particular, spike patterns dramatically change between high and low contrast stimuli, implying that they cannot be useful in representing information in a contrast invariant manner. On the other hand, the high temporal precision of LGN responses is largely preserved, especially during natural movie stimulation.

As a result, we analyze the role of high temporal precision in the function of visual neurons using both functional and information based measures. The fine temporal structure of the neural response can account for 30% of the information in the spike train. Using stimulus reconstruction techniques, we demonstrate how such fine temporal information can encode spatial aspects of the stimulus. Thus, using contrast adaptation to identify salient features of the neural code, we demonstrate the possible significance of temporal precision in neural spike trains that implies a function beyond that of the receptive field itself.

Supported by Charles King Trust Postdoctoral Fellowship and NGIA Grant HM1582-05-C-0009 to DAB and GBS, NIH-EY05253 and SUNY Research Foundation to CW, JZJ, CIY, and JMA.

49. Using a population reference for stimulus onset time in first spike latency coding

Steven M Chase, Eric D Young

Johns Hopkins University

This work explores the role of first spike latency as a potential coding strategy in the sensory nervous system. It is well known that certain properties of a sensory stimulus, such as intensity in the auditory system or contrast in the visual system, will modulate the timing of the first peak in a peri-stimulus time histogram. This latency variation has been shown to carry information about the stimulus in different sensory systems (auditory: Furukawa and Middlebrooks, 2002; somatosensory: Panzeri et al., 2001). Most studies that estimate the information in first spike latency assume that the brain has an independent reference for computing stimulus onset. Critics of first spike latency codes argue that since the first spike latency functions of many units behave similarly, correcting first spike latencies for a population reference of stimulus onset will at the least cause first spike latency information to decrease, and may actually cause it to disappear entirely. Here we assess the influence of population estimates of stimulus onset on the information conveyed by the first spike latency of single units in the auditory system.

Using virtual space stimuli we independently manipulated several sound localization cues while recording single unit responses in the inferior colliculus of decerebrate cats. We created a pseudo-population of neurons by combining units recorded with the same stimuli, and estimated the stimulus onset time for each stimulus from this population using a biologically plausible coincidence detector model. To test the similarity of the first spike latency functions within a population, we computed pair-wise cross-correlation coefficients between the first spike latency functions of units within the population. We then directly compared the information contained in a first spike latency code by single units under the assumption of an invariant stimulus onset reference to the information computed under the assumption of a variable stimulus onset derived from the coincidence detector model.

We find that while a significant number of units have correlated first spike latency functions, a number of neurons have first spike latency functions that are uncorrelated or even anti-correlated. In accordance with this finding, the information carried by some neurons decreases when spike times are corrected for stimulus onsets calculated from the population, while the information in other units increases. On average these effects cancel, leaving little difference in the overall first spike latency information conveyed under both models.

(Supported by NIH grants DC00115, DC05742, and DC05211)

50. A feed-forward model of spatial and directional selectivity of hippocampal place cells

Ricardo Chavarriaga, Denis Sheynikhovich, Thomas Strosslin, Wulfram Gerstner

Ecole Polytechnique Fédérale de Lausanne (EPFL), School of Computer and Communication Sciences, and Brain Mind Institute, 1015 Lausanne, Switzerland

In recent years a wealth of studies have focused on the role of the Hippocampus in spatial learning and navigation, triggered by the finding of place sensitive cells in this area. The activity of these cells, termed place cells (PC), is highly correlated to the location of the animal within the environment. According to experimental data, hippocampal place cells are directional when a rat is in radial mazes, whereas most place cells are non-directional in open environments (Markus at al., 1995; Muller et al., 1994).

We have developed a model of hippocampal place cells, which reproduces both spatial and directional selectivity of these cells. In our model, place cells are intrinsically directional (responding to specific local views), and their directionality is reduced during exploration when the rat is not constrained in the direction of its movements (like in open fields or the centre of mazes). We have modelled changes in the place cell directionality as the result of experience-dependent changes (through unsupervised Hebbian learning) in feed-forward connections from a population coding for specific local views. When the animal is able to explore the same location with different head directions, it allows a single post-synaptic cell to be connected to pre-synaptic cells coding for several local views. Therefore, the directional selectivity of that cell will be reduced.

It should be noticed that this model contrast with other models addressing the issue of place cells directionality (Brunel and Trullier, 1998; Kali and Dayan, 2000), which propose that those changes are the result of synaptic changes in recurrent connections among place cells (presumably in the CA3 region). Recent data suggests that a functional CA3 recurrent network is not required to produce changes in the directionality of both entorhinal and CA1 cells (Brun et al, 2002; Fyhn et al., 2004), supporting our assumption that feed-forward projections onto the hippocampus may suffice to yield directional changes in the PC activity.

The proposed model combines two separate representations of space driven by internal and external sensory information onto a single robust representation of space. During exploration, Hebbian learning is used to update feed-forward connections in the model. A direct interaction between the allothetic and the idiothetic representation compensate for accumulating errors in the latter and allows the resetting of the path integrator upon re-entry in known environments. As in previous models developed in our laboratory (Chavarriaga et al., 2005; Strosslin et al., 2005), the place representation is built incrementally during exploration.

Results

Simulations of exploration of open environments, linear tracks and radial mazes were performed, in order to assess the spatial selectivity and directionality dependence of place cells. An information theoretic approach was used to quantify this selectivity.
Cells in the model show localised, overlapping place fields and the activity of the population assembly forms a distributed, redundant representation of space. The information content about position of the PC activity is modulated by the presence of available external cues. After exploration of open environments, place fields located near to the walls are smaller and convey more information about the agent's location than those located in the centre of the arena. Our model agrees with experimental results suggesting that the sparseness of the hippocampal place code is influenced by the availability of external sensory cues (Battaglia et al., 2004).

Moreover, The model reproduces the directional dependence of place cells depending on the exploratory behaviour of the animal (or agent). When the agent freely explores open environments, the cell's directionality decreases with time as PCs are able to integrate information corresponding to several directional sub-components (sensitive to specific local views). In contrast, when the agent follows a linear trajectory or explores a radial maze PC remain highly directional.

51. Bayesian sampling methods for the analysis of electrophysiological data

Beau D Cronin, Konrad P Kording

MIT

In the effort to draw robust and accurate conclusions from electrophysiological recordings, researchers commonly confront two problems. (1) Which model should they use to describe how neuronal outputs depend on the stimuli? Such a model (e.g., linear-nonlinear-Poisson) will, for example, specify the types of nonlinearities which exist and how spikes are generated. This model selection problem is generally difficult, as different models will often have different numbers of parameters. (2) How to estimate parameter values and report error margins? Such parameters might be the precision of orientation tuning or the nonlinearity of the neurons' firing mechanism. In many models, there are a large number of parameters and they can only be estimated jointly – making the estimation of error boundaries complicated.

Bayesian statistics allows for the solution of both problems in a coherent framework. With these methods, a statistical model specifies a probabilistic, typically hierarchical, process by which the data is generated. Such a generative model quantifies prior knowledge about the structure and properties of the data; this prior knowledge may, for example, specify plausible ranges of firing rates derived from both previous experimental results and the current experimental conditions. Importantly, the Bayesian approach naturally avoids the normal bias towards models with many parameters, and allows for principled model selection as well as parameter estimation.

At least in theory, a correct solution of the parameter estimation and model selection problems with the Bayesian approach requires integration over high-dimensional spaces. Sampling methods have recently become popular, however, because they allow for the efficient computation of approximate solutions to such problems. To illustrate, we implement a Markov Chain Monte Carlo (MCMC) algorithm, and apply it to electrophysiological data. In particular, we use spikes from orientation-selective cells from the visual cortex of cat recorded using a reverse-correlation paradigm. We choose a Poisson spiking model with circular Gaussian tuning curves, parameterized by the baseline, peak amplitude, preferred orientation, and tuning width. In addition, we apply a prior which encapsulates our expectation that these parameters will change smoothly over the time course of the response. We show that this incorporation of prior knowledge strongly increases the quality of the data analysis; specifically, we can reveal effects that are not seen before, and we can obtain much tighter error bounds on the relevant variables than are possible without the prior. Moreover, we use a variant of this method to perform model selection by comparing the Gaussian tuning curve model to a more complex difference of two Gaussians (DoG) model.

Model selection and parameter estimation, which are central to our understanding of the nervous system, are typically performed in an ad-hoc way. The Bayesian methods we describe are generally applicable to these problems as they arise in any sensory or motor modality. Put simply, these methods allow for a cleaner analysis of data using fewer recorded spikes.

52. The Effect of the Static Nonlinearity on the Efficient Coding of the Visual input.

Mohammad Dastjerdi, Dawei W Dong

Center for Complex Systems & Brain Sciences, Florida Atlantic University

It is proposed that the early visual system exploits the statistical structures of the visual environment in order to represent the visual input efficiently. In previous studies, it has been shown that the efficient representations of natural images are localized and oriented filters similar to the receptive fields of simple cells in the visual cortex. However, the orientation selective receptive fields do not emerge before visual cortex, and a simple cell receives inputs from both ON and OFF ganglion cells. Because the ON and OFF ganglion cells process visual input in a nonlinear way, one cannot simply study the efficient coding as a linear process. In the current research, the effect of static nonlinearity on the efficient coding of the visual input is investigated. Natural time varying images are preprocessed with a biologically inspired center-surround filter (CSF). Similar to the earlier studies, the efficient representations of the direct CSF output are localized and oriented filters. However, the efficient coding of the rectified CSF output (ON/OFF channels) does not result in those filters. Instead, the filters have center surround structures similar to those of ganglion cells and most of the improvement in efficiency results from the rectification. Furthermore, the kurtosis can be increased much more by using a temporal filter similar to the temporal receptive field of lateral geniculate nucleus (LGN) cells. In conclusion, the results suggest that using biologically inspired spatial/temporal filter of retina/LGN with static nonlinearity gives more efficient representation than linear processing.

53. Population coding of natural images with sensory and channel noise

Eizaburo Doi, Michael S Lewicki

Carnegie Mellon University

We propose a model that preserves signal information subject to both sensory and representation noise and demonstrate its close agreement with the characteristics of the retinal ganglion cells (RGCs) once the synaptic cost is taken into account. Sensory noise is assumed to be white with variance proportional to the illumination level (Atick & Redlich, 1992). Intrinsic or channel noise limits the precision of the neural representation which is typically a few bits per spike (Borst & Theunissen, 1999). Recently we have proposed a robust coding model that optimizes the population of receptive fields (RFs) for minimal reconstruction error subject to fixed channel noise (Doi, Balcan, & Lewicki, 2005). Here, we generalize it to allow for sensory noise by formulating the problem as an optimization problem and solving it numerically. This model minimizes the deleterious effect of both sensory and channel noise by optimally adapting the code to the statistical structure of both the signal and noise while preserving the code's robustness to channel noise by introducing the optimal level of redundancy.

Next, we examine the relevance of the proposed coding scheme to the visual sensory system. In order to accurately compare the theory to experimental data, we took into account both the optical transfer function of the primate eye (Navarro, Aral, & Williams, 1993) and the spatial density of the cone photoreceptors (Rodieck, 1998), and created a retinally projected natural image data set as a function of retinal eccentricity. Using this data set, we derived the optimal set of RFs for different retinal eccentricities.

At 0 degrees of retinal eccentricity, we assume that the ratio between the number of cone photoreceptors and the units in the next stage is 1, considering both the ratio between cones and RGCs (Orban, 1984) and the lack of half-wave rectification in our model. The optimal theoretical RFs generally have concentric center-surround organization, but the precise structure depends on the signal-to-noise ratio (SNR) of the natural image data to which the model is adapted. When the SNR is high, RF centers are generally driven by a single cone; as the SNR becomes lower, the center becomes larger. At the same time, the strength of the surround becomes weaker resulting in a transition from band-pass to low-pass filtering properties. At 10 degrees of retinal eccentricity, we assume a cell ratio of 0.5. The resulting RFs have a trend similar to those derived at 0 degrees except that RF centers are larger and are driven by multiple cones.

These results were obtained using a sparse prior on the RF weights which represents a generic resource constraint of neural hardware and energy consumption (Attwell & Laughlin, 2001; Vincent & Baddeley, 2003). Without this sparse resource prior, RFs are not spatially localized. Importantly, we observed that there is virtually no increase in reconstruction error by adding this prior, implying that the local RFs are as good as the optimal, unconstrained population code.

The model proposed here is similar to the information theoretic model of RGCs proposed by Atick & Redlich (1992). Our model shares characteristics of that model but also allows for an arbitrary cell ratio between the photoreceptors and RGCs. The essential role of using a sparse prior on synaptic weights in reproducing characteristics of RGCs was explored previously by Vincent & Baddeley (2003), but under the assumption of zero channel noise. This model also minimized coding reconstruction error, but under the restricting assumption that the synthesis filters are the transpose of the analysis filters. In the model reported here, both sets of filters are optimized and yield improved reconstruction. Furthermore, accounting for the blurring of natural images when they are projected onto the retina as well as the cone density allows our model to account for the characteristics of cells at specific retinal eccentricities. Finally, our model can be applied to systems in which the cell ratio is greater than 1, such as LGN, V1, or the cochlear nerve in the auditory system.

54. The dynamic receptive fields of the lateral geniculate nucleus (lgn) during free-viewing natural time-varying images

Dawei W Dong[1], Theodore G Weyand[2], Martin Usrey[3]

[1]Center for Complex Systems and Brain Sciences, Florida Atlantic University, [2]Department of Cell Biology and Anatomy, Louisiana State University Health Science Center, [3] Center for Neuroscience, University of California, Davis, California

Objective: To code dynamic visual input efficiently, it has been proposed that the spatio-temporal receptive fields (STRF) of LGN neurons remove the average spatio-temporal correlations of natural time-varying images (Dong & Atick 1995) and furthermore, the STRFs change according to changing statistics of scenes and saccade timing to maintain the decorrelation (Dong 2000; Truccolo & Dong 2000). Those predictions were confirmed partially: for natural time-varying images, LGN output spike trains are decorrelated on average (Dan etal 1996) and are dynamically decorrelated according to saccade timings and scenes (Dong etal 2003). Here, direct measurements of LGN STRFs are made under natural viewing conditions.

Method: The single unit activities of LGN neurons in awake and free viewing cats and monkeys are recorded along with eye movements when animals are watching natural time varying images. The recorded eye positions are important to derive the saccade timing and the actual visual input on the retina. The auto-correlation and the cross-correlation of LGN spike train and the visual input are calculated to reveal the STRFs of LGN neurons.

Results: The measured STRFs are nonstationary. In particular, the shape of the STRF for a LGN neuron depends on recent scene statistics and relative timing to saccades. As a result, the output spike train is temporally decorrelated not only on the average but also for different saccade timings and scenes. In particular, for the lower mean light intensity, the the measured filter has bigger amplitude and is also wider (slower) in time, and furthermore, the measured temporal filter is a band-passing filter right before a saccade and is less so right after a saccade (the negative part is squeezed/chopped off near the time of a saccade). It is shown that both right before/after and between saccades, the information transfer through LGN is optimized during natural viewing. Such optimization is achieved by nonstationary processing and oculomotor integration through LGN.

Conclusions: Our results are consistent with the theory of efficient coding in LGN which dynamically changes according to input statistics and eye movements.

55. Analysis of oscillatory spiking in the subthalamic nucleus of Parkinson's patients using point process models

Uri T Eden, Ramin Amirnovin, Emery N Brown, Emad N Eskandar

Massachusetts General Hospital

Patients with neurological diseases often have abnormal electrophysiological findings that can be related to particular clinical signs and symptoms. Understanding these relations and determining their statistical significance is a fundamental step toward appreciating the pathological mechanisms of the disease. In particular, abnormal oscillatory firing patterns of neurons in the subthalamic nucleus (STN) of patients with Parkinson's disease (PD) have been postulated to play a role in the pathogenesis of motor deficits such as tremor, rigidity, and akinesia.

In order to characterize the statistical firing properties of these neurons we examined neural recordings from the STN of Parkinson's patients undergoing surgery to implant a deep brain stimulating (DBS) electrode, while they performed a voluntary hand movement task in one of four cardinal directions. We defineded a subset of these neurons based on oscillatory firing patterns using a frequency domain analysis. We then characterized multiple features of the firing properties of these neurons using a statistical point process modeling framework that is well suited to address issues of spike timing.

We constructed intensity models for neurons in the STN describing the probability of spiking at each instant as a function of the time relative to the start of the movement and of the recent spiking history of the neuron in the past 150 ms. Mathematically, the model takes the following form:

$$\lambda_t=\exp\{\Sigma\alpha_i \cdot I(i,t,d)+ \Sigma\beta_j \cdot N_{j:j+1}+ \Sigma\gamma_k \cdot N_{10k:10k+9}\}, \quad (1)$$

where i indexes a discrete set of time bins for each movement trial, j indexes the past 10 ms of firing history, and k indexes the firing history from 10-150 ms in 10 ms bins. The first sum in equation 1 is the movement related component of the spiking model. Here, I(i,t,d) is an indicator function that is equal to 1 if the current time, t, is contained in the interval [i,i+Δ) and the movement direction for the current trial is d. The second and third sums in equation 1 represent the recent and long-term history components of the model, respectively. $N_{a:b}$ denotes the number of spikes observed in the interval [a,b). The functional form for this model was chosen to relate the spiking activity to the time relative to the onset of the hand movement in each direction and to past spiking history in a generalized linear model (GLM) framework. The variable t giving the time into the trial and the set of variables $N_{a:b}$ giving the past spiking activity of the neurons are the covariates of the spiking model, which, together with the set of parameters for the model given by $\theta=[\{\alpha_i\},\{\beta_j\}_{j=0\ldots9},\{\gamma_k\}_{k=1\ldots14}]$, define the instantaneous intensity of spiking. We can use this intensity to calculate the discrete time joint probability distribution function for observing any spike train,

$$p(N_{1:T};\theta)=\Pi\, \lambda_{u(k)} \exp(-\grave{o}_{[0,T]}\, \lambda_t\, dt), \quad (2)$$

where u(1),…,u(N) are the spike times and [0,T] is the observation interval.

By calculating maximum likelihood estimators for all of the parameters and their significance levels, we were able to simultaneously characterize multiple phenomena previously associated with these neurons such as increased firing as a function of movement planning and execution, directional selectivity, refractoriness, bursting, and oscillatory spiking that is attenuated during movement. We also found that in nearly all of the recorded neurons, the probability of firing a spike was significantly reduced 20-30 ms after a previous spike, suggesting that the previously described oscillatory firing of these neurons is composed of an initial period of inhibition followed by a period of increased firing probability. In describing the relative contributions of the movement task and history effects to our ability to fit the observed spiking data, we found that short-term history effects related to refractoriness and bursting are most informative about spiking, despite the fact that most analyses focus on movement and oscillatory behavior. Therefore, we have constructed a simple model that is able to capture the relative propensity of aberrant STN spiking in terms of movement associated factors, factors associated with intrinsic properties of the neurons, and factors that may be related to dysregulated network dynamics.

56. A model of multiplicative auditory responses in the midbrain of the barn owl

Brian J. Fischer[1], Charles H. Anderson[2]

[1]California Institute of Technology, [2]Washington University School of Medicine

A model was developed for the computations performed in the auditory system of the barn owl that lead to a representation of auditory space in the external nucleus of the inferior colliculus (ICx). The goal of the study was to better understand the role of nonlinear processing of interaural time difference (ITD) and interaural level difference (ILD) cues in creating a representation of auditory space. To achieve this goal, a high-level model of the barn owl's auditory localization system was developed that assigned systems-level variables and functions to each of the neurobiological substrates that process time and level difference cues. The computational model was then implemented in neurobiologically realistic model circuits that matched the currently available neurophysiological data.

At the systems level we showed that a likelihood function model fails as a model of auditory space representation in the barn owl's ICx. The likelihood function model required a multiplication of signals across frequencies that would not produce side peaks in ITD tuning curves, as is observed experimentally. We found that a model where interaural time difference and interaural level difference cues are combined nonlinearly within frequency channels and then summed across frequency channels is a neurobiologically consistent model of auditory space representation in the barn owl's ICx.

At the neuronal level we showed that AND-like sensitivity to ITD and ILD can arise in a network where inputs to each neuron are summed linearly. Our model supposes that a subset of neurons in the lateral shell of the central nucleus of the inferior colliculus encode ITD and ILD additively at the subthreshold level. Using population coding ideas we showed that subsequent neurons in ICx can have multiplicative tuning to ITD and ILD at the subthreshold level where inputs to each neuron are summed linearly. This model provides a specific implementation of the experimentally observed nonlinear selectivity for ITD and ILD in ICx neurons.

57. Selectivity, sparseness and information transmission in the inferior temporal visual cortex

Leonardo Franco[1], Edmund T Rolls[2], Jose M Jerez[1], Nick Aggelopoulos[2]

[1]University of Malaga, [2]University of Oxford

The question of how information is encoded by neuronal activity in the brain is fundamental for understanding how the brain operates. Towards the end of the primate ventral visual system, in the inferior temporal visual cortex, neurons respond with some selectivity to different faces or objects (Perrett et al., 1982; Desimone, 1991). However, each neuron does not respond to only one face or object, but is instead tuned to respond with rates that are graded, responding to a few objects with high firing rates, and with smaller and smaller firing rates for more and more objects (Baddeley et al., 1997).

We measured the selectivity and the sparseness of firing to visual stimuli of single neurons in the primate inferior temporal visual cortex in response to a set of 20 visual stimuli including objects and faces in macaques performing a visual fixation task. Neurons were analysed with significantly different responses to the stimuli. The firing rate distribution of 34% of the neurons was exponential. 31% of the neurons had too few low rates to be fitted by an exponential distribution, and were fitted by a gamma distribution.Interestingly, the raw firing rate distribution taken across of all neurons fitted an exponential distribution closely.

The sparseness aS or selectivity of the representation of the set of 20 stimuli provided by each of these neurons (which takes a maximal value of 1.0) had an average across all neurons of 0.76, indicating a rather distributed representation. The sparseness of the representation of a given stimulus by the whole population of neurons, the population sparseness also had an average value of 0.76. The similarity of the average single neuron selectivity and population sparseness for any one stimulus taken at any one time shows that the representation is weakly ergodic. For this to occur, the stimuli to which different neurons are tuned must be uncorrelated.

Another aspect of the efficiency of the representation is how rapidly the information about which stimulus was shown is made available by the firing of neurons. Delorme and Thorpe (2001) have suggested that just one spike from each neuron is sufficient, while an alternative view is that the number of spikes in a fixed time window over which a postsynaptic neurons could integrate information is more realistic, and this time might be in the order of 20 ms for a single receiving neuron, or much longer if the receiving neurons are connected by recurrent collateral associative synapses and so can integrate information over time Although the number of spikes in a short time window of e.g. 20 ms is likely to be 0, 1 or 2, the information available may be more than that from the first spike alone, and we examine this quantitatively in this work using direct measurements of neuronal activity. Using information theory methods (Franco et al., 2004) we showed that the first spike is not as efficient as all the spikes in a fixed time window of even 20 ms from a population of neurons in encoding information about which stimulus was shown.

References

1) Baddeley RJ, Abbott LF, Booth MJA, Sengpiel F, Freeman T, Wakeman EA, Rolls ET (1997) Responses of neurons in primary and inferior temporal visual cortices to natural scenes. Proc Roy Soc Lond B 264:1775-1783.

2) Delorme A, Thorpe SJ (2001) Face identification using one spike per neuron: resistance to image degradations. Neural Networks 14:795-803.

3) Desimone R (1991) Face-selective cells in the temporal cortex monkeys. J Cogn Neurosc 3:1-8.

4) Franco L, Rolls ET, Aggelopoulos NC, Treves A (2004) The use of decoding to analyze the contribution to the information of the correlations between the firing of simultaneously recorded neurons. Exp Brain Res 155:370-384.

5) Perrett DI, Rolls ET, Caan W (1982) Visual neurons responsive to faces in the monkey temporal cortex. Exp Brain Res 47:329-342.

58. Structure of the Primate Cone Mosaic and the Statistics of Color in Natural Images

Patrick Garrigan, Charles Ratliff, Jennifer M. Klein, Peter Sterling, David H. Brainard, Vijay Balasubramanian

University of Pennsylvania

The primate retina is efficiently designed for capturing grayscale spatial information from natural images (see Simoncelli & Olshausen, 2001), but it also processes color information. Here we ask whether the organization of the cone mosaics for long (L), medium (M), and short (S) wavelength cones maximizes transmission of spatial-chromatic information. We analyzed the statistics and information content of a new calibrated database of color images from the Okavango Delta in Botswana. These statistics, combined with optical properties of the eye, suggest that optimal receptor mosaics reproduce two striking features of the human mosaics: (a) high variability in the L/M cone ratio (e.g. Hagstrom, Neitz, & Neitz, 1997); (b) low proportion of S cones, which comprise less than 10% of all cones (Curcio, et al., 1991).

We calibrated a Nikon D70 digital camera and acquired a database of images from the dry season savanna habitat of a baboon troop (images acquired by Lucy Seyfarth, June-August, 2004). Raw image data was used to estimate isomerization rates (R*/s) for the L, M, and S photopigments at each pixel. Across brightly lit scenes (mean isomerisation rates > 1 x 10^6 R*/s), the average R*/s distributions in each of the L, M and S channels all peaked at low intensities (peak L = 1.4 x 10^6, M = 9.1 x 10^5, S = 7.6 x 10^4 R*/s) and showed long tails (kurtosis L = 6.77, M = 7.79, S = 13.18). Although the L and M distributions nearly coincided, the S distribution had a significantly higher peak at low isomerization rates (peak height ratios: L/M = 0.74, L/S = 0.05). The shapes of the R*/s distributions for individual scenes varied, especially in the S channel (skewness range L = (0.77 – 1.91), M = (0.87 – 2.23), S = (1.55 – 4.66)), with some images having distributions with multiple peaks in various channels. The spatial power spectra of all three cone channels, averaged over scenes, scaled in spatial frequency with an exponent of approximately –2, consistent with scale invariance and resembling the results for luminance (Simoncelli & Olshausen, 2001). Direct computation of the correlation in each channel (Pearson's product-moment correlation, r) showed that all separations yield a small, but significant, excess in spatial correlation in the L and M channels over the S channel (r(S,S) / r(L,L) = 1.035 for points separated by 40 pixels, r(S,S) / r(L,L) = 1.02 for points separated by 140 pixels). We also computed the pairwise correlations between L, M, and S channels vs. displacement in the image. All responses were highly correlated at small distances, and spatially coincident L and M channels gave nearly indistinguishable responses (r(L,M) = 0.99 for spatially coincident points). Notably, given an L or M response at some location, S responses at nearby locations were also highly predictable (r(L,S) = 0.91, r(M,S) = 0.92 for spatially coincident points). This suggests that most surfaces are broadly reflective and that the variation of color intensity from point to point is largely explained by variation of luminance.

We estimated entropy in the L, M, and S channels separately. To do this, we mimicked adaptation by quantizing the responses of each channel into a fixed number of equally probable discrete levels and then directly measured the entropy of cone arrays. For all numbers of levels tested (4 – 4096) the S-channel had highest entropy followed by M and L-channels, which were similar. The high correlation between L and M channels implies that the amount of information transmitted by a mixed L-M mosaic will depend only weakly on the L/M ratio. Thus large fluctuations in this ratio can be expected in cone mosaics across individuals, and this is observed. The higher entropy of S-channel responses implies that the amount of information transmitted by a mixed L-M-S mosaic would be maximized when S cones are in the highest proportion, in contradiction to known physiology.

This analysis has so far neglected two phenomena that might lead to overestimation of entropy in the S channel: (1) lower intensities in the S-channel could make it more susceptible to noise (despite the S opsin's apparently greater stability (Rieke & Baylor, 2000); (2) optical properties of the ocular media increase chromatic aberration in the wavelengths to which the S-channel responds. To test for noise effects, we estimated entropy in L,M,S channels with thresholds ranging between 100 and 500 R*/s (within the reported range for dark noise (Lamb, 1987) and low light (Rieke & Baylor, 2000)). In bright light (> 1 x 10^6 R*/s), the S channel still carried the most information, but at low light (< 5 x 10^4 R*/s) the S channel had about 25-30% lower entropy than L and M for dark noise > 250 R*/s. This occurs because at low light, many S channels are operating near or below the dark noise threshold, while more L and M channels are still operating within their normal dynamic range. To test for effects of chromatic aberration (human eye) on cone response entropy, we blurred images following Marimont & Wandell (1994), whose model minimizes aberration at 580 nm. After this correction, entropy was again lowest in the S-channel. One might suppose that the optics of the eye could evolve to minimize aberration at any desirable wavelength. But if the optimal wavelength were closer to peak S-cone sensitivity, then spatial acuity could be maintained only in bright light because S-channel responses are more limited by dark noise at low mean light levels. Therefore, the optics of the trichromatic eye should be optimized for acuity at longer wavelengths, further reducing the entropy of the S-cone response. This in turn implies that the amount of information transmitted by a mixed L-M-S mosaic would be maximized when S cones are in the lowest proportion.

In sum, statistical correlations across the chromatic components of natural scenes, combined with the effects of chromatic aberration and noise predict large variability in L/M cone ratios and scarcity of S-cones, precisely as observed in the primate retina.

This research was supported by NSF IBN0344678.

59. Neuroinformatic Resources for Single- and Multi-Neuron Spike Train Analysis

David H Goldberg, Jonathan D. Victor, Daniel Gardner

Weill Medical College of Cornell University

We present recent advances in the development of an evolving computational neuroinformatic resource for the sharing and analysis of spike train data (Gardner & Victor, Cosyne, 2004; Goldberg et al., Cosyne 2005). This resource integrates the neurophysiology data repository at Neurodatabase.org with a suite of complementary methods for the information theoretic analysis of spike trains. The goal of the resource is to make data and analytical techniques available to the neuroscience community and to catalyze collaborations between experimenters and theoreticians.

We now report release of the Spike Train Analysis Toolkit, implemented for use on personal desktop workstations. The source code and documentation are available at http://cortex.med.cornell.edu/toolkit. The toolkit is written in C, can be called as a Matlab toolbox, and runs on Windows, Mac OS, and Linux. It accepts data in a general, easily-generated, text-based, platform-independent format. Three methods of mutual information estimation are implemented: The direct method (Strong et al., 1998), the metric space method (Victor & Purpura, 1997), and the binless embedding method (Victor, 2002). All three methods analyze spike trains recorded from single neurons, and the implementations of the direct method and the metric space method (Aronov et al., 2003; Aronov, 2004) have been extended for the analysis of spike trains simultaneously recorded from multiple neurons. Several entropy bias and variance estimation techniques are also included.

As proof-of-concept, we used the toolkit to analyze single- and multi-neuron spike trains obtained from parietal cortex of macaque monkeys during a prehension task (Debowy et al., 2002; Vaknin et al., 2005). In each trial, subjects reached for and grasped knobs, and the corresponding spike trains were labeled according to the approach style, grasp style, and knob identity. Metric space analysis revealed: 1. significant information about the task was encoded in a 1 sec interval around the time of knob contact, 2. neurons simultaneously encoded information about grasp style and approach style, 3. information was encoded on timescales of 200 msec to 1 sec, and 4. that for some neurons, more information was conveyed when the subject could see the workspace than when the workspace was visually blocked. Preliminary multi-neuron analysis indicates that the information conveyed by pairs of nearby neurons is largely redundant, and keeping track of the source of the spike does not increase the amount of information conveyed.

These analytical tools will additionally be implemented on a dedicated parallel computational array that will be accessible via and linked to Neurodatabase.org. Such an implementation is desirable because the analytical methods can be computationally intensive, particularly when applied to multiple simultaneously recorded neurons. This implementation will also feature Java-based graphical tools for manipulating analysis inputs and visualizing analysis results.

Support by Human Brain Project/Neuroinformatics grants MH068012 from NIMH, NINDS, NIA, NIBIB, and NSF, and MH057153 from NIMH and NINDS.

60. Spike-timing effects in reverse correlation analyses

Tim Gollisch

Harvard University

Traditional reverse correlation analyses, such as spike-triggered average and spike-triggered covariance, can be confounded by variability in spike timing. This may include spike-time jitter resulting from neuronal noise as well as dynamics of spike timing that are part of the encoding process. Here, a generic data analysis framework is presented that allows us to include spike-timing effects in the reverse correlation, thus improving the estimation of receptive fields and output nonlinearities as well as the description of spike-timing dynamics. In particular, it is investigated how systematic spike latency shifts in combination with spike jitter can be recovered from simulated and experimental spike-train data. Furthermore, the set of relevant linear filters obtained from spike-triggered covariance may be reduced if spike jitter is accounted for. This is shown to be the case, for example, in spike-train data obtained from the auditory nerve of locusts.

61. A Novel Measure from Machine Learning to Describe Neural Responses

Arnulf B.A. Graf, Adam Kohn

Center for Neural Science, New York University

The response of visual cortical neurons depends on the spatio-temporal parameters of a visual stimulus. In particular, drifting gratings of different orientations elicit different responses, as illustrated by an orientation tuning curve. The latter is computed using the mean firing rate, and ignores the temporal structure of the response. The question of how to discriminate between stimuli of different orientations on the basis of the neural response can, however, be cast into a more general framework: the discrimination, or classification, of the high-dimensional neural response vector of spike times evoked by each stimulus orientation. Machine learning, and in particular pattern classification, is a field of data analysis which is ideally suited to deal with such problems since it allows one to use the entire temporal structure of the neural response, without averaging over time. We propose here a novel measure derived from machine learning—the distance between the neural responses evoked by grating stimuli of different orientations and a separating hyperplane which classifies these responses—and apply it to responses binned over a range of time scales to gain insight into the neural code.

The neural data analyzed in this study consists of the responses of single neurons to drifting sinusoidal gratings, recorded from the primary visual cortex of anesthetized and paralyzed macaque monkeys. We first determined the optimal parameters of a drifting sinusoidal grating (orientation, spatial and temporal frequency and size tuning) for each neuron. We then measured the response of neurons to stimuli of 5 different orientations while the other parameters were kept at their optimal values. These orientations were spaced equally, typically 22.5 deg, and generally included the optimal orientation and an orientation that evoked little response. All stimuli were presented at full contrast. The spike times of the neural response were sampled at 1 ms. Each stimulus was presented on average 75 times with each presentation lasting 2560 ms. For each neuron (n=36), we studied the responses to all pairs of orientations from the 5 orientations available, yielding a total of 4+3+2+1=10 classification experiments. For each trial, we considered the neural response as a vector, the encoding vector, of dimension 2560 composed of the values 0 and 1, where 1 indicates the time of occurrence of a spike. Each classification experiment was represented by a dataset composed of the encoding vectors (2x75=150 on average) combined with their class label. We then applied linear classifiers to compute a separating hyperplane between the encodings corresponding to the two orientations of each dataset. We used two types of classifiers: the popular Prototype classifier which classifies according to the closest mean-of-class prototype of each class; and the Fisher linear discriminant classifier which finds the best linear separation of the two classes according to the Fisher score, using standard linear algebra techniques. We quantified the results with three measures: the distance of each encoding to the separating hyperplane, the training error on the whole dataset, and the classification error computed using cross-validation.

We first studied the geometry of the neural response space by computing the distance of the encodings to the separating hyperplane for both classifiers (data pooled over neurons and classification experiments). The histograms of the distances had a bimodal structure, showing that the classifiers can indeed separate both classes of stimuli by creating two clusters of neural responses. This bimodal structure is more pronounced than that in the histograms of the mean firing rate of the encodings, although our distance measure was correlated with the mean firing rate. We also noticed that the bimodality of the distance histograms was more pronounced for the Fisher than the Prototype classifier. This may suggest that a classifier that reduces the entire dataset to its prototypes may be less suited to describe the neural response space than one that finds an optimal separation in this entire space.

To investigate the importance of the temporal structure of the spike train, we performed our classification analysis after temporally binning the spike trains (bin size from 1 ms to 1024 ms in octave steps). We found that the training error increased monotonically with the size of the bin, which we attribute to the dependence of this error on the dimensionality of the encoding vector (e.g. with fine binning the data can be overfit since we are classifying 150 patterns in a 2560-dimensional space). More importantly, we found the classification error showed a minimum when we used bins on the order of 10 ms. We verified that the temporal structure on this time scale was important by shuffling the order of the bins within each response, and found that both the training and classification errors increased dramatically. Shuffling the coarsely binned response (bins of 1024 ms), on the other hand, had no effect on these errors. This suggests that when temporal structure is lost through binning, shuffling bins has little additional effect.

In summary, we find that applying a novel machine learning measure—the distance of the neural responses to a separating hyperplane—provides insight into the coding of stimulus orientation in visual cortex. Specifically, our measure suggests that although neural responses corresponding to different stimuli can be reasonably discriminated by relying on the mean firing rate, discrimination is enhanced when the temporal structure (order of 10 ms) is also taken into account.

62. Modelling Adaptive Mechanisms for Motion Processing in the Macaque Visual Cortex

Nicolas Heess, Wyeth Bair

University Laboratory of Physiology; University of Oxford

Several recent studies have shown that the computations underlying the responses of cortical direction selective (DS) cells are more sophisticated than previously thought. Complex DS cells in the primary visual cortex (V1) and DS cells in the cortical motion area (V5/MT) adapt their profile of temporal integration (Bair & Movshon, 2004, J Neurosci 24:7305-7323) and their gain (Heess & Bair, 2005, Soc for Neurosci Abstr #136.7) to the statistics of moving visual stimuli. In particular, temporal integration becomes longer and response gain increases at slower speeds. It is important to determine (1) whether these two phenomena, adaptive temporal integration (ATI) and motion gain control (MGC) are related, (2) what the underlying mechanisms are, and (3) whether they are related to other adaptive phenomena, such as cortical normalization.

To answer these questions, we applied spike-triggered average (STA) and covariance (STC) analysis techniques to a set of models of DS cells as well as to neuronal data from DS cells stimulated with binary and Gaussian random motion sequences (recorded for a previous study: Bair & Movshon, 2004). Recent work suggests that STC analysis has the potential to reveal suppressive stimulus dimensions that might be related to divisive gain control. In order to guide the interpretation of our empirical results we tested models of varying levels of complexity, ranging from a simple temporal filter that operated on the one-dimensional random stimulus motion sequence to a more elaborate motion-energy (ME) model (Adelson & Bergen, 1985) that operated on a realistic three-dimensional spatio-temporal representation of the motion stimulus. Models varied with respect to the implemented filters, the spike generating mechanism and the presence of explicit adaptive features such as spike rate adaptation and divisive gain control.

We were particularly interested in the presence of suppressive kernels as potential correlates of motion gain control. Indeed, for a substantial number of cells, STC analysis identified suppressive kernels in the form of low-variance dimensions in the spike-triggered stimulus ensemble obtained by stimulation with Gaussian motion sequences. However, these kernels appeared to be consistent with those obtained from a standard ME model and from a simple 1-D linear integrate-and-fire (LIF) neuron. Thus, these kernels might arise in the neuronal data as a result of either the velocity tuning of DS neurons or non-linearities of spike generation, rather than as the result of motion gain control. Although STC was not a good indicator of gain control, it did reveal excitatory kernels beyond the STA. These kernels were very consistent in their shape and position in time (relative to the STA) across cells. Interestingly, none of the various DS model architectures that we tested was able to explain these kernels. To test the relationship between cortical normalization and motion gain control, we implemented an ME model with divisive gain control (normalization) applied to the output of the linear filters. This model accounted for motion gain control like that observed by Heess & Bair (2005). Interestingly, it also displayed adaptive changes in temporal integration across stimulus amplitude (i.e., speed) that were consistent with many features of the neural data (Bair & Movshon 2004). This suggests that motion gain control and adaptive temporal integration could have a common origin, although, consistent with the data, the strengths of the two effects were not necessarily correlated.

These results provide further insight into the computations performed by cortical DS cells. STC analysis reveals features that are not predicted by common models of DS cells. Modelling provides a possible explanation for the adaptive phenomena demonstrated earlier and indicates that the underlying mechanisms might be related to those implicated in cortical normalization. From a methodological point of view the results demonstrate that STA/STC analysis in the motion domain can be used to reveal details of the computations performed by DS cortical cells, but at the same time that results should be interpreted cautiously.

Support: Wellcome Trust

63. Single neuron computation: from dynamical system to feature detector

Sungho Hong[1], Blaise Aguera y Arcas[2], Adrienne L Fairhall[1]

[1]Department of Physiology and Biophysics, University of Washington, [2]Program in Applied Mathematics and Computation, Princeton University

White noise methods are a powerful tool for characterizing the computation performed by neural systems. These methods allow one to identify the feature or features that the neural system extracts from a complex input, and to determine how these features are combined to drive the system's spiking response. These methods have also been applied to characterize the input/output relations of single neurons driven by synaptic inputs, simulated by direct current injection. To interpret the results of white noise analysis of single neurons, we would like to understand how the obtained feature space of a single neuron maps onto the biophysical properties of the membrane, in particular the dynamics of ion channels. Here, through analysis of simple dynamical model neurons, we try to draw explicit connections between the output of a white noise analysis and the underlying dynamical system. We find that under certain conditions, the most relevant feature space of a single neuron, characterized by white noise analysis, can be also extracted from analysis of the underlying dynamical system. In making such connection, it is revealed that the nonlinear threshold from the dynamical system may play a crucial role in white noise analysis. Inspired by these findngs, we propose and analyze a multi-dimensional version of the filter-and fire model, which is analytically feasible. We also draw other interesting implications, including that under the same conditions, the feature space must be spanned by the spike-triggered average and its successive order time derivatives.

64. Burst Temporal Coding by the Retina

Toshiyuki Ishii, Toshihiko Hosoya

RIKEN Brain Science Institute

Precise understanding of neuronal information coding is critically important in analyzing brain functions. The information conveyed by neuronal responses is often inferred by counting spike-numbers in time-windows with lengths similar to or longer than that of synaptic integration, i.e. tens to hundreds of milliseconds. However, if precise relative spike timings at the scale of milliseconds would carry information, neuronal circuits could have much larger information capacity. The significance of such coding is not generally clear, due in part to the experimental difficulty in controlling all input to a particular neuronal network.

The retina is suited to analyze such fine characteristics of neuronal coding, because one can control all the input to the circuit just by projecting movies. In response to various visual inputs the retina fires spike bursts, characterized by hundreds of milliseconds of silent periods punctuated by abrupt generation of several spikes. This burst firing is highly reproducible: when repeatedly stimulated by the same visual input, the timing jitter of burst initiation is typically a few milliseconds, and the spike number fluctuation is often less than one spike per burst. Burst initiation times and spike numbers (or firing rates) can be nicely predicted by simple models, such as a linear filter followed by a nonlinear spike generator.

We asked if combinations of interspike intervals within bursts (spike patterns) carry information about the visual input. Using the retinas of salamanders and mice, we found that a subset of ganglion cells generate bursts with various spike patterns, which are unique to the preceding visual inputs. When single bursts contain three or more spikes, the multiple interspike intervals encode the input in a cooperative, non-redundant manner. This suggests that the spike patters are not sorely determined by slowly modulating instantaneous firing rates. Differences in spike patterns at the scale of milliseconds encode differences in light intensity waveforms as long as 200-300 milliseconds. We propose that the retina compresses hundreds of milliseconds of light sequences into spike patterns at the scale of milliseconds.

65. Simultaneous electrophysiology and two-photon imaging of olfactory projection neurons in intact fruit flies

Vivek Jayaraman, Gilles J. Laurent

California Institute of Technology

Genetically encoded optical indicators hold the promise of enabling non-invasive monitoring of activity in identified neurons in a behaving organism. The interpretation of images of brain activity produced using such sensors is, however, problematic, and may be the source of seemingly contradictory findings in Drosophila olfaction. At issue is the breadth of odor tuning of projection neurons (PNs), the output neurons of the antennal lobe (the insect analog of the mammalian olfactory bulb). One group performed 2-photon imaging experiments in detached antennae-brain preparations by expressing a calcium-sensitive fluorescent protein named G-CaMP[1] in PNs. They interpreted the selective G-CaMP responses they recorded to mean that PNs are rather narrowly tuned to odors, and suggested that a given PN's responses directly reflect the uniglomerular input it receives from its olfactory sensory neurons (OSNs)[2]. This is in contrast to evidence from in vivo whole-cell recordings suggesting that PNs are more broadly tuned, consistent with a transformation of OSN input by the antennal lobe circuitry[3]. A different group using G-CaMP expressed in Kenyon cells, targets of the PNs, has suggested that odors evoke sparse, stereotyped and spatially restricted responses in these neurons[4], a finding that is consistent with electrophysiological data[5]. In both sets of imaging experiments however, it is unclear what electrophysiological signal the imaging signal being measured corresponds to. The narrow odor tuning seen in these imaging experiments could, for example, be a result of the relatively high threshold of activation of G-CaMP[6,7].

We are attempting to resolve such issues by performing simultaneous 2-photon imaging and electrophysiology of G-CaMP expressing PNs in an intact, in vivo adult fly preparation. For our electrophysiological recordings in PNs, we use both loose-patch[8] and intracellular (sharp electrode) techniques. In our current experiments, we are trying to establish the correlation between physiology and G-CaMP signal. We find that G-CaMP faithfully reports sustained (e.g., >1sec) activity at rates higher than 50Hz, albeit with very slow kinetics. However, it is unable to detect sustained activity at rates below 30Hz or higher-rate activity if it is not sustained (e.g., 50Hz for 200msec). Our results suggest that while this sensor can be useful for the identification of neural targets for more refined exploration of a functional circuit, it could be misleading if used as the sole means to assess neuronal activity.

More generally, the methodology we have developed enables any genetically encoded sensors, activators or silencers to be precisely calibrated in an intact fly and in the very neurons that they are likely to be used in, something that is crucial given the likely cell-type-dependent variability of their function. Calibrated use of such fast-improving technology is likely to be of considerable value to systems neuroscience.

Support: NIDCD & NSF (G.L.) and Sloan-Swartz Foundation (V.J.). We thank the Axel, Stocker and Ito labs for their generous gifts of fly lines.

References

1. Nakai, J. et al. (2001) Nature Biotechnology 19: 137-141.

2. Wang, J.W. et al. (2003) Cell 112: 271-282.

3. Wilson, R.I., et al. (2004) Science 303: 366-370.

4. Wang, Y. et al. (2004) J. Neuroscience 24: 6507-6514.

5. Turner, G.T. and Laurent, G. (2005) Soc. Neurosci. Abstract

6. Pologruto, T. et al. (2004) J. Neuroscience 24: 9572-9.

7. Reiff, D.F. et al. (2005) J. Neuroscience 25: 4766-78.

8. Wang, J.W. et al. (2003) Soc. Neurosci. Abstract

66. Is the Cortex a Digital Computer?

Dana H Ballard, Janneke FM Jehee

University of Rochester

Decades of single-cell recordings in mammalian cortex have revealed the correlation of increased firing rate with behavioral measures, suggestive of a rate code. But is a rate code the fundamental signaling strategy of the cortex, or is it just a correlate of another quite different strategy? The two possible answers to this question both have ardent proponents. Defenders of the rate code as the basic signaling strategy point to the huge number of experimental studies that use rate code as a basic framework. In addition, they argue that the correlations between neurons that are above the number predicted by chance are too small for any synchronous strategy to be a useful phenomena. Proponents of the idea of a more basic code based on spike timing point to the increasing number of data that show spike timing effects. There have been a wide range of models suggesting very different uses of timing, such as signal binding, fast computation, signaling the result of a computation, and even consciousness itself.

We show by simulation that the cortex can adopt a signaling strategy that makes extensive use of synchrony for fast communication, but does it in a way that is consistent with the rate code indications. The motivation for the model is that the cortex has to solve a handful of basic computational problems and these can be expressed as three constraints:

1. For reliable communication the cortex uses redundancy to compensate for cell death,

2. To combat noise in the form of other processes, the cortex must use timing circuitry,

3. For fast communication the cortex needs an analog signaling strategy.

One model that satisfies all these constraints is that of a handful of simultaneously active analog/digital feedback circuits. Each circuit is phase-locked to the data and sends spikes coincident with its own phase-locked 50 Hz 'clock,' distributed throughout the cortex. Analog quantities are communicated via delays that are very small compared to the clock frequency. Every cell in the circuit is part of a clique of copies such that, at every clock cycle, one of them is chosen randomly to send the required signal. In that way the basic spike traffic through any cell appears random and approximates a Poisson distribution. Thus, the essential progress of numbers throughout the circuit randomly transits different cells from moment to moment. We have termed this Distributed Synchrony [2].

This very non-standard way of thinking about cortical computation explains experimental data. Here, we show in simulations that simple cell receptive fields can be learned using the sparse coding model. The dynamics of the circuit are similar to that of [1] except that the inner loop uses matching pursuit instead of gradient descent. Moreover, we show that the spike train observed from multiple copies has a CV distributed about unity, as seen in Poisson models.

[1] Olshausen, B.A. & Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature, 381, 607-609.

[2] Zhang, Z. & Ballard, D.H (2001). Distributed Synchrony. Journal of Neurocomputing, 44-46C, 715-720.

67. Representation of time and states in prefrontal cortex and striatum

Dezhe Z Jin[1], Naotaka Fujii[2], Ann M Graybiel[3]

[1]Department of Physics, The Pennsylvania State University, [2]Brain Science Institute, RIKEN, Japan, [3]Department of Brain and Cognitive Sciences and the McGovern Institute for Brain Research, Massachusetts Institute of Technology

Awareness of time, external and internals states is crucial for taking intelligent actions. In this poster, we show that neurons in prefrontal cortex and striatum encode external and internal states, as well as sub-second times. In our experiments, monkeys perform guided visual saccade task. We record neural responses in prefrontal cortex and striatum. This is a simple task, and the monkey needs to know only the direction of the next target to perform it correctly. Thus, only direction selective neurons with timings pre-saccade are necessary for the task. However, we find in prefrontal cortex and striatum much more diverse task related neural response profiles. There are indeed direction selective neurons, but the timings of their responses can be not only pre-saccade but also post-saccade. Many other neurons do not encode directions. Some respond only relative to the first saccade, some to every saccade; some respond exclusively to the offset of the last target; some respond to reward delivered at end of each trial. For each response types, the timings of response peaks are often distributed. For example, timings of response peaks for those responding to first saccade only can range all the way from pre-saccade to post-saccade. Event during one-second fixation period at the beginnings of each trial, we find neuron with peak responses to the task onset with latencies from 200 ms to 600 ms. The response profiles of the neurons are quite rich. These results show that prefrontal cortex and striatum encode all aspect of the task, even information not needed for simple visually guided saccades, such as completion of a saccade (internal signal) and time elapsed since last salient external (visual guide) or internal (saccade) events. We suggest that such comprehensive representation of states and time in prefrontal cortex and striatum, which are involved in formation of action rules and habits, will be important for quick formation of new associations between the states and actions that guide intelligent behavior.

68. Transmission of rapidly changing signals through a population of noisy integrate-and-fire neurons

Peyman Khorsand, Frances S Chance

University of California, Irvine

We use analytical and computational methods to compare the responses of different integrate-and-fire models to sudden changes in either mean current or noise. In the presence of colored noise (white noise filtered by a synaptic time constant), the initial response of a leaky integrate-and-fire (LIF) neuron to a step in mean and noise currents are instantaneous jumps in firing rates followed by slower approaches to the final steady-state firing rates. The extremely rapid change in firing rate occurs because the LIF model has a “sharp” boundary condition at spike threshold. As a result, action potentials are instantaneously generated when there is a change in input. Moreover, the probability of the LIF model hovering just below threshold is significantly higher than for other integrate-and -fire models. The size of the initial instantaneous response depends on the probability distribution of the membrane potential at spike threshold. This quantity directly varies with the synaptic time constant and, for most parameters, is greater for larger synaptic time constants. The instantaneous jump in firing rate does not appear in the responses of models that have “soft” spike thresholds, such as the quadratic (QIF) or exponential (EIF) integrate-and-fire models. These differences in responses to steps in mean and noise current may explain why LIF models have been shown to more accurately track rapidly oscillating current compared with QIF and EIF models (Fourcaud-Trocmé et al, 2003, J. Neurosci. 23:11628). The responses of these models to steps in mean or noise may be described by decaying oscillations around the final firing rate. We examine how the size and time course of these responses depend on the variance and filtering of background noise. For all integrate-and-fire models, we find that the addition of noise to large neuronal populations allows faster and more accurate transmission of rapid changes in input. In fact, depending on the spectral contents of the input signal, there should be an optimal level of noise that maximizes the reliability of signal transmission through a population of neurons.

69. Information Traffic on a Neural Cable

Kristin Koch[1], Ronen Segev[2], Judith McLean[1], Vijay Balasubramanian[1], Michael Freed[1], Michael J Berry[2], Peter Sterling[1]

[1]University of Pennsylvania, [2]Princeton University

The architecture of a nerve tract should match its information rate, but the rate is unknown for any nerve. Here we report such measurements for the guinea pig optic nerve (100,000 axons). We identified 7 ganglion cell types on a multi-electrode array by their reverse and autocorrelograms; other 'sluggish' cells were not further classified. We presented video images of natural scenes to mimic saccades, optic flow, object motion, and fixational eye movements, then used the spike responses (n=107 cells) to estimate information by the "direct" method.

For all measures, including mean and peak spike rates, 'firing fraction', 'burst fraction', and spike time jitter, each cell responded similarly across stimuli. However, responses did vary across cell types. Brisk-transient cells fired at the highest mean rates (8 vs. 3-6 spikes/s for other types), the highest peak rates (171 vs. 74-117 spikes/s), and the highest burst fractions (0.26 vs. 0.07-0.10). ON-OFF directionally-selective cells fired at higher peak rates than local-edge cells (117 vs. 74 spikes spikes/s); brisk-sustained cells had higher firing fractions than DS cells (0.29 vs. 0.16); local-edge cells had the highest jitter (40 vs.15-23ms). The highest information rates were for brisk-transient cells (13 vs. 6-10 bits/s). Yet for all types, information rate was ~30% of the maximum possible given the mean firing rate.

We estimated the distribution densities of each type contributing to the optic cable as (dendritic coverage factor) ÷ (dendritic field area). Then, knowing that brisk-transient cells contribute 6% of the axons, the ratios of cell densities yielded cell numbers. The brisk-transient and brisk-sustained classes comprise 2 types (ON and OFF), ON DS and ON-OFF DS cells encode, respectively, 3 and 4 cardinal directions. We multiplied cell numbers accordingly.

The traffic down the optic nerve, assuming independent cells, gives 875,000 bits/s. Surprisingly, the non-brisk types, although least studied, send the most information. If information traffic scales with axon number, the human optic nerve (one million axons) would transmit ~10 megabits/s, roughly the rate of an Ethernet connection.

cell type	bits/s	dendritic area (mm^2)	density ($cells/mm^2$)	# of cells	bits/s per array
brisk-transient	13 ± 6	0.20	30	6,000	78,000
brisk-sustained	10 ± 5	0.05	120	24,000	240,000
ON DS	6 ± 3	0.13	36	7,000	42,000
ON-OFF DS	8 ± 4	0.10	60	12,000	96,000
local-edge	7 ± 3	0.03	100	20,000	140,000
sluggish (other)	9 ± 4			31,000	279,000
total				100,000	875,000

70. Selectivity of local field potentials and spikes to the visual stimuli in the human medial temporal lobe

Alexander Kraskov[1], Rodrigo Quian Quiroga[2], Itzhak Fried[3], Christof Koch[1]

[1]Division of Biology, Caltech, [2]Department of Engineering, University of Leicester, UK, [3]Div. of Neurosurgery and Semel Institute for Neuroscience and Human Behavior UCLA, Functional Neurosurgery Unit, Tel-Aviv Medical Center and Sackler Faculty of Medicine, Tel-Aviv University

Local field potentials reflect (LFPs) an averaged dendrosomatic activity of synaptic signals of large neuronal populations within up to 3 mm. In this paper we investigate the selectivity of local field potentials to a semantic category of visual stimuli and its relation to the selectivity of the spiking activity. The amplitude of LFPs showed a strong selectivity to a category at 10% of all analyzed microelectrodes implanted in temporal lobe of nine epileptic patients. Separate analysis of the power and phase of LFPs revealed that the mean phase was category selective for lower frequencies and the power for high frequencies. Little overlap between LFP- and spike-selective microelectrodes was found. Additionally, it was also possible to read out information about the category of a stimulus presented to a patient both with spikes and with LFPs. Combining the spiking and LFPs activity enhanced decoding performance in comparison with that achieved when each one was used separately, especially for short time intervals. This property might be useful for a brain machine interface emphasizing the speed at the expense of accuracy.

71. Common-input models for multiple neural spike-train data

Liam Paninski, Jayant E Kulkarni

Columbia University

One of the greatest challenges in systems neuroscience is to understand how large networks of neurons collectively process and encode information [1]. In this work we present a statistical model for multiple simultaneously-recorded single units that includes terms corresponding to a) the dependence of the firing-rate on an experimentally controlled stimulus, b) the spiking history of the observed cells, and c) the effect of common-input terms. These common-input terms model the effect of unobserved cells which are presynaptic to two or more cells in the observed population, therefore constituting a hidden common-driving term for our observed cells.

The consideration of these common-input terms is the primary contribution of this work. It is clear that these hidden effects have the potential to play a significant role in any model, and can be expected to become more important as the size of the observed neural population increases. While there has been a great deal of work in the area of inferring population codes, most earlier models (with the important exception of [2]) lacked any common-noise terms, thus unrealistically "explaining" correlations between cells solely on the basis of direct connections between observed cells in the population.

To perform inference in this model, we consider an escape-rate approximation of the integrate-and-fire model, developing a variant of the Fokker-Planck equation for this approximation. Though the log-concavity of the likelihood of the observed data is guaranteed for the stimulus-dependent and spike-history terms, the gradients are not easy to compute, and the storage burden is intractable for the multi-variate Fokker-Planck equation. Thus we implement a faster, although approximate, algorithm based on the extended Kalman filter to perform the expectation step in the Expectation-Maximization (EM) algorithm [3], in order to fit the model parameters.

These techniques allow us to solve a variety of important inference problems in a straightforward, computationally-efficient manner; for example, we can sample spike trains from the model to predict rasters and spike-count variances given an arbitrary new stimulus [4]; compute the likelihood of any observed population spike train; infer the magnitude of the common input to a pair of cells on any given trial; and predict the responses of a cell based on the activity of other observed cells [5]. We believe these tools will find applicability in a wide variety of population coding settings.

References

[1] E. Brown, R. Kass, and P. Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. Nat. Neuro., 7: 456-461, 2004.

[2] D. Nykamp. Revealing pairwise coupling in linear-nonlinear networks. SIAM J. Appl. Math., 65: 2005-2032, 2005.

[3] A. Smith and E. Brown. Estimating a state-space model from point process observations. Neural Comp., 15: 965-991, 2003.

[4] J. Pillow et al. Accounting for timing and variability of retinal ganglion cell light responses with a stochastic integrate-and-fire model. J. Neurosci., 25: 11003-11013, 2005

[5] L. Paninski et al. Superlinear population encoding of dynamic hand trajectory in primary motor cortex. J. Neurosci., 24: 8551-8561, 2004.

72. Propagation of Synfire Activity in Locally Connected Networks with Conductance-based Synapses

Arvind Kumar[1], Stefan Rotter[2], Ad Aertsen[3]

[1]Neurobiology and Biophysics, Insti. of Biology III, Albert-Ludwigs University Freiburg, Germany, [2]Theory and Data Analysis, IGPP, Freiburg and Bernstein Center for Computational Neuroscience Freiburg, Germany, [3]Neurobiology and Biophysics, Insti. of Biology III, Albert-Ludwigs University Freiburg, Germany, Bernstein Center for Computational Neuroscience. Freiburg, Germany

The Synfire chain (SFC) model was proposed by Abeles (1991) as a minimal model to explain experimental observations of precise spike patterns in awake behaving animals (e.g Prut et al., 1998; Riehle et al., 1997), the Synfire chain (SFC) model was proposed by Abeles (1991). Computational Models based on integrate-and-fire neurons have elucidated the conditions under which millisecond precise synchronization and stable transmission of pulse packets in a SFC can be achieved (Diesmann et al., 1999). Recent in vitro studies have confirmed that cortical neurons are indeed capable of transmitting precise synchrony using a SFC (Reyes, 2003). While SFCs provided an explanation of the experimental observations, it seemed to be incompatible with the low rate asynchronous irregular network activity usually found in the cortex. In fact, the nonrandom architecture of the SFC and the transient synchrony were found to destabilize the dynamics of the network (Mehring et al., 2003). Specifically, when a group of neurons was stimulated in a short time interval mimicking a synchronized pulse packet, it induced a massive spreading synchronization which eventually destroyed the pulse packet .

Here, we studied the evolution of activity in a SFC embedded within a locally connected random network (LCRN), using conductance-based synapses (Kuhn et al., 2004). We found that the conductance-based synapses improve the stability of the LCRN as compared to the current-based case, such that the traveling activity (pulse packet) in the embedded SFC did not destabilize the activity of the embedding network. The stable propagation of the pulse packet depended on the statistics of the background activity: Under certain conditions the propagation of transient synchrony remained stable, whereas in others the transient synchrony would eventually dissolve into the background activity. High synchrony and high firing rates in the background always annihilated the synchronous activity in the SFC. On the other hand, relatively low background firing rates and low synchrony supported the traveling synchrony and made it more precise along the chain. The background activity also strongly influenced the propagation of firing rates in the SFC. We show that synchronous activity states in the network do not support propagation of firing rates or synfire activity, while with asynchronous activity in the background, it is possible to transmit information both in synfire activity and in firing rate.

Acknowledgements: We thank Markus Diesmann, Marc-Oliver Gewaltig and Carsten Mehring for helpful discussions. Partial funding by GIF, DFG-GraKo and BMBF grant 01GQ0420 to the BCCN Freiburg is gratefully acknowledged.

References:

1. Abeles, M. (1991) Corticonics: Neural Circuits of the Cerebral Cortex, Cambridge University Press.

2. Diesmann, M. et al. (1999) Stable propagation of synchronous spiking in cortical neural networks. Nature 402, 529–533.

3. Kuhn, A.et al. (2004) Neuronal integration of synaptic input in the fluctuation driven regime. J. Neurosci. 24, 2345–2356.

4. Mehring, C. et al. (2003) Activity dynamics and propagation of synchronous spiking in locally connected random networks. Biol. Cybern. 88(5), 395–408.

5. Prut, Y. et al (1998) Spatiotemporal structure of cortical activity: properties and behavioral relevance. J. Neurophysiol. 79(6), 2857–2874.

6. Reyes, A. D. (2003) Synchronydependent propagation of firing rate in iteratively constructed networks in vitro. Nat.Neurosci. 6(6), 593–599.

7. Riehle, A. et al. (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. Science 278, 950–1953.

73. Requiem for the spike?

Peter E. Latham[1], Arnd Roth[2], Michael Hausser[2], Mickey London[2]

[1]Gatsby Computational Neuroscience Unit, UCL, [2]Wolfson Institute for Biomedical Research and Department of Physiology, UCL

A major open question in neuroscience is: "what's the neural code?". The standard approach to answering this, pioneered by Richmond and Optican almost two decades ago [1], is to record spike trains and compute information under different coding models. Unfortunately, this approach requires huge amounts of data and thus, despite considerable efforts and some success, it is still not clear to what extent the neural code, in mammalian cortex, relies on precise spike timing, and in particular on spike patterns.

An alternative approach follows from the observation that if spike patterns are to carry information, they must be precisely repeatable. We may thus ask the question: does the massively recurrent connectivity that is a salient feature of cortical networks place intrinsic limits on precise repeatability, and thus on the extent to which spike patterns can carry information?

We show here that one can answer this question by measuring the mean increase in the firing rate of an average neuron in response to a single synaptic input. If the mean increase is sufficiently large, then the network must be chaotic at the microscopic level, which in turn precludes precisely repeatable spike trains. Applying this to models of biophysically realistic neurons, we find that chaos is a likely property of in-vivo networks, a result consistent with previous studies [2,3]. We then address this issue directly, using in-vivo patch-clamp recordings from cortical pyramidal neurons in anesthetized animals, and preliminary data supports the same conclusion. Finally, we connect quantitatively the mean increase in firing rate with a lower bound on the precision at which spike timing can carry information.

1. B.J. Richmond and L.M. Optican, J. Neurophysiol. 57:132-46; 57:147-61 (1987).

2. C. van Vreeswijk and H. Sompolinsky, Neural Comput. 10:1321-1371 (1998).

3. A. Banerjee, Neural Comput. 13:161-193; 13:195-225 (2001).

74. Neural Diversity and Ensemble Encoding

Aurel A. Lazar

Columbia University

We call diversity the "technique" of combining multiple replicas of the same stimulus that are (at least partially) uncorrelated in order to improve the stimulus detection performance.

Replicas are outputs of an ensemble of neurons with a common stimulus. Different neurons (or diversity branches) fade (at least partially) in an independent fashion. We distinguish two different types of gains that an ensemble of neurons might bring about: (i) gain against neuron variability, and (ii) gain in signal-to-noise ratio.

We describe a general algorithm for recovering the stimulus at the input of an ensemble of integrate-and-fire neurons from reading the spike times of the replicas. We analyze the gain against neural variability and the gain in signal-to noise ratio.

75. Bayesian inference with probabilistic population codes: Simulations in a network of conductance-based integrate-and-fire neurons

Wei Ji Ma[1], Jeffrey M Beck[1], Peter E Latham[2], Alexandre Pouget[1]

[1]University of Rochester, [2]University College London

We have recently shown how to perform optimal Bayesian inference with spikes using a form of population codes known as probabilistic population codes. Those codes do not simply represent the value of the encoded variable, but a full probability distribution over the encoded variable. This distribution can be read out by applying Bayes' rule to the population pattern of activity. When the noise is independent and Poisson, Bayesian inference involving taking the product of distributions can be implemented by simply taking the sum of the probabilistic population codes. This approach can be applied to tasks such as cue integration and decision-making.

However, in our previous work, several simplifying assumptions were made: neuronal responses obey Poisson distributions and are uncorrelated within each population, and the tuning curves are translated copies of each other. Does Bayesian inference still reduce to taking sums of population activity when these assumptions are violated in a population of biologically realistic neurons? To examine this, we simulated a network of spiking neurons consisting of two input layers, each with 252 independent neurons, and one output layer with 1008 excitatory and 252 inhibitory conductance-based integrate-and-fire neurons, recurrently connected. The two input layers had hills of activity with Poisson variability, representing two cues. Correlations were introduced through a Gaussian connectivity pattern between each input layer and the output layer. The activity in the output layer also formed a hill, which was decoded using a locally optimal linear estimator.

To simulate psychophysical experiments, we first presented one cue at a time, of which we varied the reliability by changing the value of the gain of the activated input layer. For each gain, we computed the mean and variance of the distribution encoded in the output layer when only one cue was presented. We then presented both cues together for each combination of gains. The means and variances obtained from the two-cue conditions were compared with the Bayes-optimal combinations of the means and variances obtained from the corresponding one-cue conditions. Indeed, these were approximately equal, demonstrating that spiking neurons can obtain a close approximation to optimal Bayesian inference. Similar results were obtained even when the width of the input tuning curves and the pattern of correlations differed between input layers, thus confirming our analytical results (see Beck et al, Cosyne 2006).

76. The representation of interaural time differences in human cortex

David McAlpine, Adenike O Deane-Pratt

University College London

Humans make use of small differences in the timing of a sound at the two ears to determine the location of low-frequency (<1500 Hz) sound sources. The neural representation of these interaural time differences (ITDs) has previously been assumed to be one in which sounds leading in time at one ear activate maximally neural centres in the opposite side of the brain. Using the mismatch negativity (MMN)-evoked potential and headphone presentation of interaurally-delayed sounds, we demonstrate that this simple scheme is untenable. The MMN signals automatic detection of rare deviations from a series of otherwise similar sounds, and is sensitive to ITD. We used MMN to measure the response to bursts of binaural noise, centred at 500 Hz and with bandwidth of 400 Hz, in four deviant ITD conditions: ± 500 μs (¼ of the period of 500 Hz) and ± 1500 Hz (¾ of the period). These stimuli have perceived lateral positions on the left or right compared to the standard condition with 0 μs ITD (which is perceived in the centre of the head). In the case of ¼ cycle deviants, MMN was greater over the hemisphere contralateral to the perceived position, consistent with previous data. However, a novel finding was that ¾ (far) deviants elicited greater fronto-central negativity ipsilateral to the sound-source. This suggests that stimulus pairs with ITDs equivalent to +¼ and – ¾ period of the 500-Hz centre frequency, despite being perceived on opposite sides, share a similar neural representation, and vice versa. The data cannot be predicted by place-code models of ITD encoding as these require a full-range of ITD detectors, including ITDs beyond ½ a cycle of the period of the centre frequency of any auditory channel. Rather, the data are consistent with electrophysiological recordings in animals that suggest a restricted-range – within ± ½ a cycle – model.

77. Reconstruction of speech stimuli from population of neuronal responses in primary auditory cortex

Nima Mesgarani, Stephen David, Shihab Shamma

University of Maryland College Park

We examined the responses of neurons in primary auditory cortex (A1) of awake ferrets to phonetically labeled speech stimuli. Sentences were taken from the TIMIT database and chosen to represent a diversity of male and female speakers. We presented these stimuli to awake ferrets while recording the activity of isolated A1 neurons. For analysis, we segmented the continuous speech samples into sequences of phonemes, which represent the smallest significant units of speech. We characterized the response properties of each neuron as the peri-stimulus time histogram response to each phoneme. Across a population of A1 neurons, we observed distinct patterns of phoneme selectivity that may provide a neural basis or low-level phoneme discrimination.

We investigated how features of speech are encoded in A1 using a method for reconstructing the speech stimulus from the population of neuronal responses. Stimuli were reconstructed using a linear spectro-temporal model to map the population response to the stimulus spectrogram. We compared the accuracy of reconstruction across phonemes.

One important factor involved in stimulus reconstruction is the presence of correlations which are present in complex speech stimuli. Prior knowledge of regularities in the stimulus can benefit reconstruction in the presence of noise and when spectro-temporal coverage is limited. We studied the influence of prior knowledge of stimulus correlations, noise and spectro-temporal coverage in the reconstruction accuracy using neural data and in simulation.

78. Unbiased Estimator of Shape Parameter for Spiking Irregularities under Changing Environments

Keiji Miura[1], Masato Okada[2], Shun-ichi Amari[3]

[1]Kyoto University / JST PRESTO, [2]University of Tokyo / JST PRESTO / RIKEN BSI, [3]RIKEN BSI

We considered a gamma distribution of inter-spike intervals as a statistical model for neuronal spike generation. A gamma distribution is a natural extension of Poisson process taking the effect of a refractory period into account. The model is specified by two parameters: a time-dependent firing rate and a shape parameter that characterizes spiking irregularities of individual neurons. Because the environment changes with time, observed data are generated from the time-dependent firing rate, which is an unknown function.

A statistical model with an unknown function is called a semi-parametric model, which is one of the unsolved problem in statistics and is generally very difficult to solve. We used a novel method of estimating functions in information geometry to estimate the shape parameter without estimating the unknown function. We obtained an optimal estimating function analytically for the shape parameter independent of the functional form of the firing rate. This estimation is efficient without Fisher information loss and better than maximum likelihood estimation.

We suggest a measure of spiking irregularity based on the estimating function, which can be easily applied to experimental data in changing environments. It may be possible to classify neurons into functional groups according to their spiking irregularities.

References
[1] K. Miura, M. Okada, and S. Amari. (2006). "Unbiased Estimator of Shape Parameter for Spiking Irregularities under Changing Environments", Advances in Neural Information Processing Systems, in press.
[2] K. Miura, M. Okada, and S. Amari. "Estimating Spiking Irregularities under Changing Environments", submitted.

Acknowledgments
This work was supported in part by grants from the Japan Society for the Promotion of Science (Nos. 14084212 and 16500093).

79. Formation of attractor representations of abstract rules in cortical networks

Emanuele Curti[1], Xiao-Jing Wang[2], Stefano Fusi[3]

[1]Columbia University, New York, NY, [2]Brandeis University, Waltham, MA, [3]Columbia Univ. New York, NY and ETH, Zurich, Switzerland

Animals and in particular primates have a remarkable ability to modify their interpretation of events depending on the context. The same event can have different meanings in different environments, and the executive control of behavior should take into account all the task-relevant conditions (context) in which the event occurs. Animals can learn what is the task-relevant information from their experience, and eventually abstract rules which guide the behavior. Several sets of rules can be stored simultaneously and animals are able to select rapidly the proper set depending on the context. Recently, experimental [Asaad et al 1998, Wallis et al 2001, Genovesio et al 2005] and theoretical [O'Reilly and Munakata 2000, Loh and Deco 2005] studies have begun to investigate the neural basis of context-dependent behavior. Here we propose a neural mechanism for rule abstraction which produces internal representations of the rules, as attractors of neural dynamics [Amit 1989]. In the proposed scenario every event like the presentation of a contextual cue or a sensory stimulus steers the neural activity towards a previously learned attractor representing the context and containing information about the proper interpretation of future events. In particular the pattern of activation dictated by the context can then affect the final decision of the animal about action selection, in response to a stimulus. When the context changes, a different pre-existent attractor corresponding to the rule in effect is selected, without the need to modify synaptic couplings.

We illustrate this mechanism with a simple example in which visual stimuli are associated to one of two possible motor responses (say two saccadic movements, left and right). In one context the first stimulus (A) should generate Left and the second stimulus should lead to Right in order to receive a reward. In the second context the associations are reversed: A-Right and B-Left are the rewarded associations. The rules to get reward can be expressed in words as follows: Rule 1: "when A is associated to Left, then B is associated to Right" and Rule 2:"when A is associated to Right, then B is associated to Left" There are three fundamental questions that we address: 1) how are the rule representations built? 2) how can the active representation of a context lead to the decision about the motor response? 3) how can contextual cues indicate explicitly what rule is in effect? In order to answer these three questions we first need to describe the neural network which will implement the rules. We assume that there are populations of neurons which are selective to the intented motor response, similarly to [Fusi et al. 2005]. In our case we will group together all the neurons with a preference for Left (population L) and those with a preference for Right (population R). The activation of one of the two groups would express the decision of the monkey to make a saccadic movement to a specific direction. The two populations of neurons compete through a population of inhibitory neurons, as in the decision making network introduced in [Wang 2002]. Each pattern of neural activity in which one population is active (expressing the decision) and the others are inactive, is a global attractor of the neural dynamics. We now consider the heterogeneity across cells. Within each decision population (L or R), we can identify and tag the neurons that have a preference for one specific sensory stimulus. For example, we define neurons within population L with preference to A as those that exhibit the largest response when stimulus A is presented, and we tag them with the label AL. Analogously we can define population BL, again within population L, and AR, BR within population R. This kind of heterogeneity has been observed in prefrontal and in premotor cortex [Asaad et al 1998,Wallis et al 2003]. Rule representations are created by the temporal proximity of events in rewarded trials: for example when rule 1 is in effect, A-Left trials are followed by either A-Left or by B-Right trials. Previous experimental results have shown that neural representations of events that occur in a fixed order tend to merge into a single representation linking neighboring events in a sequence [Miyashita 1988, Griniasty et al 1993, Yakovlev et al 1998]. We then expect that if AL and BR are separated attractors, after long enough sequences of trials in which rule 1 is in effect, AL and BR will merge into a single attractor. Analogously AR and BL will fuse to represent rule 2. Can then the activation of one of these attractors affect the competition between L and R and express the decision which is dictated by the rule in effect? Intuitively this should be possible because the activation of the representation of a rule can bias the competition. Indeed, if rule 1 is in effect, upon the presentation of A, AL and BR receive the recurrent inputs from AL and BR due to the fact that the network is in the attractor corresponding to rule 1. In addition to the these inputs, AL receives also a stronger activation from sensory stimulus A, which can then favor L in the competition. Finally how is the rule selected? Modification of the context can be determined by the feedback the monkey receives (e.g. when the monkey applies one rule and it is not rewarded any longer), or it can explicitly signaled by one or more contextual cues (also called occasion setters in psychology literature) [Schmajuk and Holland 1998]. We shall consider the second case below.

We show that these mechanisms can be implemented in a simple rate model of a network of neurons. We illustrate the simulated network dynamics after learning in the Figure below. A trial starts from the attractor corresponding to the representation of rule 1 (AL-BR). After one second, a contextual cue is presented to indicate that the rule in effect will be rule 2 (AR-BL). Network's activity steers toward the attractor corresponding to rule 2 (between 1.5 and 2.5s). When stimulus B is shown, the competition between L and R starts, and the previous activation of rule 2 attractor favors L, as dictated by the rule in effect. Notice that the selection of L does not disrupt the information about the rule in effect. Indeed the activity of AR remains high, surviving the decision process and its reset following the execution of the motor response.

The model proposed in this work generates several interesting predictions and provides a new way of interpreting the recorded cortical activity. In our scenario activity that is conventionally named `spontaneous' acquires functional meanings: the activity recorded in any interval between two task relevant events is expected to encode the internal state of the network which represents the rule in effect. This state of persistent activity might reflect factors which are under control in the experiment (e.g. a context dictated by a cue) or which might not be under control (e.g. the mood or the motivation of the animal). We predict that it should be possible to find a correlation between the controlled factors that specify the internal representation of a rule, and the recorded neural activity, in particular in areas like prefrontal cortex for which there is already experimental evidence for rule representation [Miller 2000]. Moreover the presence of this inter-event persistent activity should be correlated with behavior: erroneous trials can be due to a number of reasons, but those mistakes which are due to the wrong interpretation of the sensory stimuli (e.g. when the monkey believes it is in the wrong context) should be correlated with the trials in which the inter-event persistent activity is not observed or in which the representation of the wrong rule is activated. Interestingly the theoretical framework that we propose can be easily extended to other cognitive functions like attention (attending a specific feature of a stimulus to perform a task can be regarded as a rule).

80. Stochastic Multi-stability in Neural Decision-Making Systems

Gustavo Deco[1], Alexander Roxin[2], Ralph Andrzejak[2], Daniel Martí[2]

[1]ICREA/Universitat Pompeu Fabra, [2]Universitat Pompeu Fabra

Humans and animals constantly have to make decisions between alternative behavioral choices based on perceptual information and expected reward values of the different alternatives. A deeper insight into such decision-making processes at the nervous system level would help to understand the basic computational principles underlying the link between perception and action, i.e. intelligent behavior. Over the course of the last decade theoretical, neurophysiological, and psychological studies have shed light on the neural mechanisms underlying decision-making. At the behavioral level, so-called diffusion models explain a wide range of experimental results. In these models it is assumed that information that drives the decision process is accumulated continuously over time until it reaches a decision boundary. Given the success of diffusion models in explaining behavioral data it seems likely that some decision-making processes in the nervous system indeed rely on a similar accumulation of evidence. Motivated by these facts, a number of neurophysiological experiments on decision-making analyze the responses of neurons that correlate with the animal's behavior. By recording the activity of single neurons, signals that correlate with the subjects' decisions have been found in several areas of the cerebral cortex, most notably in area LIP in the parietal lobe and pre-motor areas of the frontal lobe. An important finding is that cortical areas involved in generating motor responses also show activity reflecting a gradual accumulation of evidence for choosing one or another decision. Therefore, the process of making a decision and action generation cannot be differentiated .

On the other hand, in computational and theoretical neuroscience biologically realistic neural circuits are designed to implement the diffusion process putatively underlying decision-making. These models involve two or more populations of excitatory neurons that are engaged in competitive interactions mediated by inhibition. External sensory inputs bias this competition in favor of one of the two populations leading to a gradually developing binary choice. Decision-making is thus understood as a probabilistic settling into one or another attractor state using competition biased by the stimulus values in a neuronal network with finite size noise effects, i.e it can be described by means of a system of nonlinear coupled stochastic differential equations. However, the study of such equations can be numerically time-consuming due to the need for sufficiently many trials to capture the effects of noise. This makes the use of analytical techniques appealing, when possible.

For example, the coupled equations for the neuronal populations can be studied in the presence of noise by solving the associated Fokker-Planck equation for the probability distribution of the rates. However, the nonlinear nature of the equations hinders analytical progress in this framework. An alternative approach involves the derivation of deterministic differential equations for the first and second order moments of distributions of the rates. This yields, for n populations, a $2n+n(n-1)/2$ dimensional approximation of the full dynamics, including the mean rates, variances and cross-correlations. Traditional analytical techniques for ODEs can then be applied to the reduced system. We derive such a reduced framework for a multi-stable system and show how it allows for straightforward analysis of the role of noise in decision-making.

81. Orbitofrontal cortex responses during acquisition of novel stimulus-response associations

Claudia E Feierstein, Zachary F Mainen

Watson School of Biological Sciences, Cold Spring Harbor Laboratory

Orbitofrontal cortex (OFC) is implicated in olfactory processing and learning. We previously recorded the activity of OFC neurons in rats performing a two-odor discrimination task, finding that these neurons represented the behaviorally relevant spatial locations; a subset of neurons integrated spatial information with reward value.

To better understand how these representations evolve, here we recorded during the learning of novel odor discriminations. We used tetrode recordings to monitor the activity of OFC neurons in rats during the performance of a modified two-alternative choice task, designed to compare the neuronal responses of a standard, well-learned odor pair and a novel one. In this task, rats started each session with the same well-learned two-odor discrimination (A/B). After ~100 trials, a new odor pair (C/D) was introduced and the rats had to learn the contingencies associated with that novel pair while continuing to discriminate the original A/B pair. In each recording session a different C/D pair was introduced.

Rats responded to the first one or two presentations of novel odors with a long odor sampling (> 1 s, compared to an average sampling duration of 255 ms), a phenomenon that could be considered a sign of novelty detection. Odor sampling durations returned to the average values afterwards.

Rats learned new odors associations fairly rapidly. Usually a criterion of 80% for the novel pair was achieved in 1 to 60 trials (median number of trials was 17), while maintaining intact performance for the original pair (>90%). However, there was variability across rats and odor pairs, and some odors were not learned (within a session).

A small fraction of OFC neurons showed odor selectivity for the well-learned odors. Interestingly, a large fraction of neurons were selective for novel odors. This selectivity developed extremely rapidly (1 or a few trials) after the new odors were introduced. Although many neurons responded to novel odors during the odor sampling period as might be expected, the introduction of novel stimuli was also associated with a more widespread change in response properties, including both increases and decreases in firing outside the odor presentation period (e.g. during movement or reward).

These data provide evidence that OFC circuits are engaged by novel olfactory stimuli. By integrating representations of novel stimuli with representations of actions leading to reward, OFC may participate in the learning of new stimulus-response-outcome contingencies.

Supported by WSBS and NIDCD #DC06104-01 as a part of the NSF/NIH CRCNS program.

82. Malignant Evaluation: Reinforcement Learning, Neuromodulation and Depression

Quentin JM Huys, Peter Dayan

Gatsby Computational Neuroscience Unit, University College London

Neuromodulators are critically implicated in major psychiatric complaints, including depression, schizophrenia, attention deficit hyperactivity disorder and obsessive compulsive disorder, as well as major neurological complaints such as Parkinson's disease. The predominant pharmacological approaches to treating these conditions act by manipulating the levels or effects of neuromodulators. Further, animal models of these diseases often involve neuromodulatory systems.

Ideas from reinforcement learning (RL) have come to offer relatively rich and robust (though evidently incomplete), models of the activity of certain neuromodulators during prediction and action learning tasks. These models have now advanced to the point at which the psychiatric and neurological diseases offer an important proving ground. It has become compelling to seek to understand malfunction from the perspective of normative normal function, and to benefit from a huge range of extra constraints available from the wealth of human and animal data.

Here, we consider depression from the perspective of RL. Depression is a devastatingly pervasive disease with catastrophic sequelae. Two of its key facets are anhedonia (an inability to experience previously rewarding events as pleasurable) and energetic alterations (hypo- and hyperactivity). These symptoms are generally amenable to pharmacological treatments acting predominantly on serotonin (5HT). All these characteristics resonate with issues studied in the RL literature on neuromodulation, including dopamine's long-studied role in appetitive predictions and control (blunted in anhedonia), and more recent normative accounts linking dopamine to the vigour or energy associated with actions (as in anergia and its opposite) and coupling dopamine and 5HT in affective opponency (suggesting a mode of action for the serotonergic treatments).

From a theoretical perspective, the two key aspects of the disorder are (a) what are the neuromodulatory characteristics of depressed states; and (b) what initiates and stabilizes transitions between normal and depressed or manic states.

Arguing from animal models such as learned helplessness and chronic mild stress, we explore the notion of depression as a particular form of metaplasticity of neuromodulatory systems (which are themselves sometimes considered as metaplasticizers of cortical inference and adaptation). More specifically, we consider how the availability, controllability and predictability of rewards and punishments in an environment set aspirational levels for possible outcomes, and thereby influence parameters controlling neuromodulatory responsivity. We explore the vulnerability of this system to parametric variations in neuromodulatory responsivity. Transitions between normal and depressed states may occur (initially) as normative or (later on and in disease) as oversensitive adaptations to prevailing environmental conditions, resulting in a stablized state with altered responsivity of dopaminergic and serotonergic systems.

Finally, this allows us to explore the parallels between the potential maladaptivity of these parameter settings for adaptive responding and the well-known potential maladaptivity of the competition between the responses elicited as a result of (classically conditioned) predictions of rewards and punishments, and those (specified by instrumental conditioning) that actually maximize the rewards and minimize the punishments. The equivalent of a classically conditioned response could be the form of systemic protective state, involving changed patterns of sleeping, eating, and somatic and neural energy regulation, that has long been suggested as characterizing depression.

83. Activity in the dorsolateral prefrontal cortex of macaques during an inter-temporal choice task

Jaewon Hwang[1], Daeyeol Lee[2]

[1]Brain & Cognitive Sciences, University of Rochester, [2]Center for Visual Science, University of Rochester

Inter-temporal choice refers to the process of decision making that involves a tradeoff among costs and benefits expected after variable temporal delays. In such cases, a small but immediate reward might be preferred over a large but delayed reward. This implies that an outcome expected in the future is depreciated over time, which is referred to as temporal discounting. Temporal discounting has been reported in various species including humans, but little is known about the neural mechanisms involved in temporal discounting and inter-temporal choice. To study such mechanisms, we developed a novel task for moneys, and examined the pattern of neural activity in the dorsolateral prefrontal cortex during the same task.

Two rhesus monkeys were used in the present study. In the beginning of each trial, the animal fixated a small white square presented at the center of a computer screen. One second later, two peripheral targets were presented along the horizontal meridian, and the animal indicated its choice with a saccade when the central fixation target was extinguished 1 second later. One of the targets was green and delivered a small reward when selected, whereas the other was red and delivered a large reward. The reward was delivered following a variable delay after the animal acquired one of these two peripheral targets. The amount of delay associated with each choice was signaled to the animal by a clock consisting of a series of small yellow dots (n=0 to 8) displayed around each peripheral target. After the animal made its choice, these dots were removed one by one at a constant rate (1 s/dot). The inter-trial interval was adjusted to prevent the animals from influencing the onset of the next trial.

During the behavioral experiments that were carried out prior to the neurophysiological recordings, various combinations of delays for small and large reward targets were tested to determine the shape of temporal discounting functions. The performance of various models with different temporal discounting function was evaluated with a Bayesian information criterion (BIC). The results showed that the hyperbolic discounting function with an exponent fit the behavioral data better than the exponential discounting function or a simple hyperbolic discounting function.

We then recorded the activity of 207 neurons from the dorsolateral prefrontal cortex (DLPFC) of the same monkeys tested in the behavioral studies, using a multi-electrode recording system. In addition to the inter-temporal choice task, each neuron was tested also in a control task, in which the animal's correct choice was indicated by the color of the central fixation target (red or green). Otherwise, the visual stimuli and the temporal sequence were identical for choice and control tasks. During the inter-temporal choice task, approximately 20.3% of the DLPFC neurons showed significant modulation in their activity according to the number of dots, or equivalently the amount of temporal delay, associated with different targets. In contrast, during the control task, the percentage of neurons in which the activity was significantly related to the number of dots decreased to 12.1%. This suggests that the key variable determining the choice of the animal during the inter-temporal choice was represented more strongly in the DLPFC when such information is behaviorally relevant. Such signals might be utilized within the DLPFC or in areas downstream to the DLPFC to compute temporally discounted values for alternative choices and influence the animal's choice according to an appropriate temporal discounting function.

84. Behavioral impact and neural representation of uncertainty in olfactory decision-making in rats

Adam Kepecs, Naoshige Uchida, Zachary F Mainen

Cold Spring Harbor Laboratory

Making good decisions requires the anticipation of the consequences associated with each alternative. Because outcomes are rarely certain, a core component of decision-making involves the prediction and evaluation of uncertainty about different options. Uncertainties may arise due a variety of factors but in the decision-making literature these are typically subdivided into two broad categories: risk and ambiguity. Risk defines outcomes to which objective probabilities can be assigned, such as playing dice or the lottery. Ambiguity, on the other hand, involves situations where the outcome probabilities are subjective and usually not a property of the environment but rather reflect some lack of knowledge on the part of the decision-maker.

While real life decisions typically involve ambiguity as well as risk, most of our understanding about uncertainty comes from the study of risk in both humans (Kahneman et al, 1982) and animals. For instance, Schultz and colleagues have found that in a Pavlovian task where highly distinct visual stimuli predicted different reward probabilities, monkeys could anticipate reward risk (as signaled by their licking) and the firing rate of some dopaminergic neurons in the VTA encoded this risk (Fiorillo et al, 2003). In a gambling task, McCoy and Platt (2005) found that monkeys preferred risky options and neurons in the posterior cingulate cortex encoded the risk associated with different choices. Here, we sought to understand the neural representation and behavioral impact of ambiguity, which is a type of uncertainty that is not directly set by task contingencies like risk. Ambiguity is caused by a lack of knowledge and hence may be reduced by learning unlike risk.

We trained rats to decide which of two odors, “A” or “B”, is the dominant in a mixture. These odors were mixed in different ratios (A/B at 100/0, 80/20, 68/32, 56/44, etc). Rats were rewarded at the left port for mixtures A/B < 1 and at the right port for A/B > 1. Note that the decision boundary (A/B =1) was an arbitrary assignment of the experimenter, since the pure odors did not necessarily have the same perceptual intensity. Discrimination accuracy for pure odors was close to 100% but for the most difficult mixtures (56/44) performance dropped to almost 60%. Since in this task the probability of reward is 100% for correct choices, there is no risk involved. However, the outcome may be uncertain due to ambiguity about the relationship of a given stimulus to the decision boundary. There are two main components contributing to this ambiguity: sensory variability associated with the encoding of odors (perceptual uncertainty) and the imprecision of memory recall about is the location of the decision boundary (memory uncertainty). Importantly, this task allowed us to systematically manipulate the ambiguity of individual discrimination problems by varying odor mixture ratios and hence change the distance of the stimulus from the decision boundary.

The first indication that rats actually possessed a representation of uncertainty about their decisions came from a peculiar observation about sniffing. We noticed that rats rapidly switched into a high frequency (9-12 Hz) sniffing mode before they entered a choice port and remained in this mode until water was delivered (~100-150 ms later). For ambiguous odor mixtures the pre-reward sniffing frequency was inversely correlated with the likelihood of reward: higher frequency sniffing bouts signaled low probability of reward while lower sniffing frequencies tended to precede rewarded trials. Since this sniffing mode occurred after rats made their choice and sniffing could not directly influence the outcome, the parsimonious interpretation is that sniffing frequency reflected the rat’s estimate of reward uncertainty due to the ambiguity associated with each decision.

In order to examine the neural representations underlying ambiguity, we have begun recordings of single neurons in the orbitofrontal cortex of rats performing the odor mixture task. Preliminary findings show that the activity of some neurons varies systematically with stimulus ambiguity, as controlled by the mixture ratio. In addition, during the period of reward anticipation the firing rates of many neurons are predictive of the outcome. In as much as outcome predictability cannot be accounted for by the stimulus alone, it suggests that orbitofrontal activity provides an estimate of uncertainty in a decision-task involving ambiguity.

These observations led us to wonder if uncertainty actually contributes to decisions made by the rats on a trial to trial basis. In principle ambiguity is the kind of uncertainty that can be reduced through learning. For example, in the odor mixture task, the decision-boundary is learned through reinforcement and may require continual updating. According to statistical learning theory, the rate of learning should be proportional to uncertainty. For difficult problems (e.g. 56/44 mixture) the outcome is very informative about location of the decision boundary, while the outcome of pure odor (100/0) trials reveals little about the boundary. Accordingly, the decision boundary should be adjusted more following difficult trials with high ambiguity than for trials with no ambiguity. Indeed, we found that animals were dynamically adjusting their decision strategy even after extensive training. Rats biased their decisions to the more recently rewarded direction as if their decision boundary was shifted. Moreover, the magnitude of this bias was proportional to the ambiguity of the problem, as predicted. Thus rats not only have an estimate of the uncertainty about ambiguous decisions, but this estimate actually contributes to future decisions.

In summary, we examined an elusive component of decision-uncertainty, the ambiguity due to lack of knowledge. Our results show that rats have a representation of uncertainty associated with ambiguous decisions and ongoing adjustments in behavior depend on the magnitude of this variable.

References:

Fiorillo, Tobler & Schultz (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. Science, 299:1898-902
Kahnemann, D., Slovic, P. & Tversky, A. (eds.) (1982): Judgement Under Uncertainty: Heuristics and Biases. Cambridge: Cambridge University Press
McCoy & Platt (2005) Risk-sensitive neurons in macaque posterior cingulate cortex. Nature Neuroscience,18:1220-7

85. Adaptive reinforcement learning for motivated behavior

Giancarlo La Camera, Zeng Liu, Dominique L Pritchett, Barry J Richmond

NIMH

Experimental and theoretical treatments of decision-making often assume that behavior will adapt to maximize future reward. However, in daily behavior humans often choose strategies that do not maximize future rewards, probably because motivation is adjusted based on both the value of the reward and the unpleasantness or aversiveness endured to obtain it. We have been trying to identify both experimentally and theoretically the factors that influence motivational levels.

Experimentally, we use a visually cued reward schedule task. In these tasks monkeys are taught to obtain a liquid reward by releasing a bar when a red target turned green. After learning this, the monkeys are placed in schedules requiring 1, 2 or 3 correct color discrimination trials to obtain a reward. Visual cues are presented during these trials. The cues can either be related to the schedule (valid), or chosen randomly (random). When the cues are valid the monkeys perform the rewarded trials virtually without errors, but in unrewarded trials the monkeys make errors, with the number of errors being proportional to the number of trials remaining in the schedule. To advance in the schedule the monkeys must repeat trials with errors until done correctly. Thus, the monkeys seem less motivated in trials in which no reward will be delivered, even though the trial must be completed correctly to obtain the reward in the future. For the 4 monkeys studied here, this strategy was chosen by the monkeys, was partially learned in the first session with the schedules, and was stable after 6-12 sessions of experience for hundreds of testing sessions.

The monkeys were also tested in a 'self-adjusting' version of the task, where if the monkeys completed 3 consecutive schedules without any errors, the reward size increased by one drop (max 5 drops). After an error the reward decreased by one drop (min 1 drop). The monkeys were tested in blocks with a large drop size and in blocks with a small drop. The monkeys were allowed to work until they stopped by themselves. All 4 monkeys worked longer and with many fewer total errors in the small reward sessions, making almost no errors until nearly at the time of quitting.

A standard temporal difference (TD) learning algorithm assigns larger values to trials that are more proximal to reward. However, if the model also uses the TD signal to learn which action is 'best', i.e., perform correctly vs. make an error as in actor-critic schemes, the errors eventually vanish in all trials. Thus, if the main factor in driving the behavior is to maximize reward, the monkeys should perform all trials without errors. In fact, the optimal strategy would be to ignore the cues, which the animals do not do. The experimental error pattern seen in the reward schedule task is reproduced if a fixed function is used to map the trial's value into the probability of correct performance (performance function). However, in the self-adjusting task, a fixed performance function predicts that performance driven by a large reward ought to be better than performance driven by a smaller reward, which is contrary to the experimental observations. Reward discounting cannot explain this because two out of four monkeys only rarely made errors in rewarded trials in any task.

We modified the 'static' TD model by having the performance function adapt to the satiation level. Performance is then sensitive to 'relative' motivational values (higher in the small reward sessions) rather then to absolute reward contingencies (larger vs. smaller reward). This adaptive model fits the behavior in all protocols. This work shows how existing reinforcement learning models can be modified to take into account suboptimal behavior due to motivational factors (like putting off work necessary to obtain reward) other than reward discounting.

86. A neural network model of the Eriksen task

Yuan Liu[1], Philip J Holmes[2]

[1]Department of Physics, Princeton University, [2]Department of Mechanical and Aerospace Engineering, Princeton University

We analyze a connectionist model of the two-alternative forced choice Eriksen task, in which subjects must correctly identify a central stimulus and disregard flankers that may or may not be compatible with it. We linearize and decouple the model, deriving a reduced drift diffusion process that describes the accumulation of net evidence in favor of either alternative, and we use this to produce analytical descriptions of how accuracy vs. response time data depends on the parameters of the model. We compare our results with numerical simulations of the full nonlinear model.

In this paper we analyze a connectionist model of the Eriksen flanker task and a drift-diffusion reduction of it. In this task subjects are asked to respond appropriately to a target arrow (< or >) displayed in the center of a five-symbol stimulus array on a display screen. The flanking symbols may be either compatible or incompatible with the central stimulus. In the compatible condition, the display reads <<<<< or >>>>>; in the incompatible condition, it reads >><>> or <<><<, and in each block of trials all four conditions are typically presented with equal probabilities. Experimental data show that subjects are slower and make more errors under the incompatible conditions

We have analyzed a linearized version of the connectionist model for the Eriksen two-alternative forced-choice task. We show that, provided solutions remain within the central domain of the logistic function in which it may be approximated by a linear function that matches its slope g at the bias point b, analytical solutions of a decoupled, linearized model modulated by a pre-determined attention signal provide reasonable estimates of critical times at which evidence in favor of the correct and incorrect alternatives cross over for incompatible trials. We then derive estimates of accuracy as a function of response time by interrogating a drift-diffusion (DD) process with variable drift rate, fitted to outputs from the perception layer of the fully nonlinear model. We compute the evolving probability distribution of solutions to the DD process and integrate to obtain the psychometric function (% correct) as an explicit function of response time and the parameters defining the drift rate and noise strength. The interrogation protocol assumes that the response delivered reflects the subject's current estimate, and corresponds best to a deadline task with a hard limit.

We also present results for the free-response protocol, in which subjects are allowed to respond in their own time, which we model as a first passage problem. The qualitative forms of the psychometric functions for compatible and incompatible trials in the interrogation and free response cases are similar, and our results suggest that parameter tuning based on the explicit formulae available for the interrogation protocol may be generally useful in matching connectionist models results with behavioral data.

87. Synaptic plasticity and decision making: a neural model for operant matching

Yonatan Loewenstein, Sebastian H Seung

Howard Hughes Medical Institute and Massachusetts Institute of Technology

According to Herrnstein's famous "matching law" of repeated choice behavior, the probability of choosing an alternative is proportional to the reward derived from that lternative. Although this empirical law of behavior has been studied extensively by psychologists, economists, and neuroscientists, the question of its neural basis has received little attention. Here we show that a classic neural model of decision-making, when combined with synaptic plasticity, reproduces matching behavior. In particular, we demonstrate equivalence between neural and behavioral models, by showing that two forms of synaptic plasticity cause our neural model to reproduce two well-known reinforcement learning models that lead to matching behavior. These two models can be distinguished both behaviorally and physiologically.

88. Phase space embedding of neural activities in the prefrontal cortex during a two-interval discrimination task

Christian K Machens[1], Ranulfo Romo[2], Carlos D Brody[1]

[1]Cold Spring Harbor Laboratory, [2]UNAM Mexico

Working memory is the ability to keep something in mind for a few seconds and manipulate it. This cognitive ability can be studied in a two-stimulus-interval discrimination task. Two stimuli, separated by a time delay of a few seconds, are presented to a subject, who must make a binary, categorical decision based on a comparison of the two stimuli. The task involves storing a short-term memory trace of the first stimulus and computing a decision. Based on electrophysiological data from monkeys that perform a version of this task, we have previously developed a simple neural network model, two populations of neurons coupled by mutual inhibition, that explains several features of the data [Machens et al. (2005), Science 307:1121-1124].

Here we address some of the aspects of the data that this model does not account for. The data show, for instance, that the firing rates of many persistently active neurons vary slowly with time. The variety and prevalence of such temporal dynamics is astonishing and has been described before [Brody et al. 2003, Cereb. Cortex 13:1196-1207]. As a first step towards understanding the function and mechanisms behind these dynamics, we have pooled the data from all neurons in prefrontal cortex, defined the firing rate of each neuron as a coordinate in a high-dimensional phase space, and then used simple projection techniques (such as principal component analysis) to obtain a compact description of the data.

We have found that the short-term memory component of the pooled data can be described compactly as a set of trajectories in three dimensions. Furthermore, several purely temporal components can be isolated as well. These components change over time but their value carries no information about the item held in memory. The activity of individual neurons can be approximated well by a linear combination of short-term memory and temporal components. While the short-term memory components can be modeled as simple (linear) integrators, the temporal components are mostly non-linear. Last but not least, we show how these results can be reconciled with our previously published mutual-inhibition model.

89. Performance following predicted reward postponement is well explained by temporal discounting.

Takafumi Minamimoto, Giancarlo LaCamera, Barry J Richmond

NIMH, NIH

Motivation is affected by the subjective value of a reward. It has long been observed that animals (including humans) act as though rewards become less valuable as a function of the delay between the current time and the reward. Thus, the subjective value of the reward depends on the delay between response and reward. Temporal discounting models have been very successful in quantifying the discount. We studied how motivation is affected by reward size and delay duration in an instrumental task without choice, and whether the temporal discounting models can explain the performance. We trained four monkeys to perform the cued reward postponement task. In this task, monkeys always work in one trial of a red-green sequential color discrimination. After successful completion of the trial, the monkey must wait before the reward is delivered. A visual cue presented at the beginning of the trial indicates the forthcoming of combination of delay duration (0.3, 3.6 or 7.2 s) and reward size (0.25 or 0.75 ml). Delay duration and reward size varied independently. For all four monkeys, the number of errors in this one trial increased as the postponement duration increased. When the monkeys knew that the reward would be delivered immediately, the error rates were the same for both the small and large rewards. However, when the monkeys had to wait for the reward, error rates in large reward trials were lower than those in small reward trials. Thus, the delay and reward size interact, so that motivation decreases more for small rewards as a function of postponement duration. We compared the error rates to two forms of temporal discounting: exponential and hyperbolic. Both predicted the interaction between reward size and delay. The hyperbolic form of the model fit the data more closely. This study shows that the temporal discounting models, by quantifying the subjective value of reward, can describe the performance accuracy in an instrumental task without choice, in addition to the well-known descriptions of performance in choice behavior tasks.

90. Temporally displaced STDP: Synaptic Competition and Stability

Baktash Babadi, Majid Arabgol

School of Cognitive Sceinces (SCS),IPM

Recent evidences indicate that in several brain regions, synaptic efficacies are regulated according to the relative timing of pre and post-synaptic action potentials, a phenomenon known as Spike Timing Dependent Plasticity (STDP). STDP has been shown to be present in layers II/III, IV and V of the cerebral cortex of rat, hippocampus of rat, tadpole tectum and cerebellum-like structures of electric fish (see Dan & Poo, 2004 for a review).

One of the crucial properties of STDP which have a substantial effect on its functional consequences is whether the magnitude of synaptic change depends on the present synaptic strength. When synaptic modification magnitude does not depend on the present synaptic strength (i.e. an additive STDP rule), a computationally favorable competition occurs among the synapses, but the distribution of synaptic weights has an inherent instability (Song et al, 2002). The competition results in a bimodal distribution of synaptic weights and may play a role in formation of cortical maps (Song & Abbott, 2001). On the other hand, if the magnitude of synaptic modification depends on the current synaptic strength (i.e. a multiplicative STDP rule), a stable unimodal weight distribution will be obtained but the synaptic competition is lost (Rubin et al, 2001). Similar results have been obtained in a mixed STDP rule, in which synaptic depression and not excitation depends on the current synaptic strengths (van Rossum et al, 2000). In these cases almost all the interesting computational properties of STDP will be lost.

In this study, we show that an additive STDP rule with a slightly displaced time window results in both synaptic competition and a stable weight distribution.

In order to study the effects of the STDP rule on the distribution of the synaptic weights, we simulate a single spiking neuron which receives $N_{ex} = 1000$ excitatory and $N_{in} = 250$ inhibitory poissonian presynaptic spike trains with rates r_{ex} and r_{in} respectively. The strength (weight) of the excitatory synapses changes according to STDP, while the strength of the inhibitory synapses remains constant. For the simulations, the integrate-and-fire is used. The sub-threshold potential dynamics of this model is governed by:

$$C\frac{V_m}{dt} = \frac{(V_l - V_m)}{R_m} + I_{ex}(t) - I_{in}(t) \quad (1)$$

where, V_m is the membrane potential, C is the membrane capacitance, V_l is the membrane leakage potential, R_m is the membrane resistance, $I_{ex}(t)$ and $I_{in}(t)$ are excitatory and inhibitory post synaptic currents respectively. When the membrane potential reaches the threshold V_{th}, a spike will be triggered; the membrane potential will reset to V_{reset} and starts its dynamic according to the equation 1 anew. The post synaptic currents in response to the incoming spikes are modeled as exponential functions. The synaptic plasticity affects the excitatory synaptic strengths (w_{ex}^i) values. Each pair of presynaptic and postsynaptic spikes that are t milliseconds apart adds the value $F(t)$ to the synaptic weight of the corresponding synapse excitatory synapse (w_{ex}^i). The value $F(t)$ is defined as:

$$F(t) = \begin{cases} -A_- e^{\frac{t-d}{\tau_-}} & ,t \le d \\ A_+ e^{\frac{-t+d}{\tau_+}} & ,t > d \end{cases} \quad (2)$$

where, $t = t_{post} - t_{pre}$ is the interval between the postsynaptic and presynaptic spikes, A_+ and A_- are the maximum values of potentiation and depression, $\tau_+ = 20\,msec$ and $\tau_- = 20\,msec$ are the time constant of potentiation and depression respectively, and d is the amount of shift of the plasticity window (Figure 1). The simulation results show that with $A_+ = 0.0053, A_- = 0.005$ and a displacement $d = 3.5\,msec$ a stable unimodal distribution will be obtained without any strict boundary condition (Figure 2).

When half of the excitatory input spike trains are correlated with each other (correlation coefficient = 0.1) and the rest are uncorrelated, a synaptic competition occurs (Figure 3). In this case the distribution of the synaptic weights will be bimodal, with the correlated synapses being the winners of the competition.

Although the applied displacement to the STDP window is small, it has a considerable effect on the final synaptic weights, endowing them both stability and competition. In fact it has been shown that different timing of the pre and post synaptic components of STDP may give rise to a displaced plasticity time window (Karmarkar & Buonomano 2002).

91. The time course of reward and punishment prediction error signals in the primate amygdala accounts for learning

Marina A Belova*, Joseph J Paton*, Daniel Salzman

Columbia University

The ability to predict positive and negative events is crucial for survival. This ability relies on learning how sensory cues are related to impending rewards and punishments. Neural signals that encode a mismatch between what is expected and what actually occurs - prediction error signals – represent potential teaching signals for forming accurate expectations. Such prediction error signals have been described in midbrain dopamine neurons for reward, but evidence for aversive prediction errors at the single cell level has previously been lacking. Computational models of reinforcement learning specifically posit that prediction error signals can drive learning by updating a temporal representation of value, which constitutes a prediction about upcoming reinforcement. Recently, we reported that single neurons in monkey amygdala represent the positive and negative value of visual stimuli during learning. In those experiments, we used a trace-conditioning procedure to induce learning about the value of novel visual stimuli. In order to demonstrate that neurons encoded the positive and negative value of conditioned images, and not just their sensory properties, we reversed the values of images. Across the population of neurons, the representation of value began to change within 1 or 2 trials after a reversal, and the change in the representation became asymptotic within 10 or fewer trials. The time course of changes in neural activity across the population closely matched the time course of behavioral learning.

In the current study, we asked whether amygdala neural responses to rewards and punishments could be modulated by monkeys' expectations of reinforcement. We compared neural responses to liquid rewards and aversive air-puffs in two behavioral contexts. In one context, the trace conditioning procedure described above, reinforcement was predicted by the presentation of a CS. In the other context, we delivered rewards and punishments at random times chosen from an exponential distribution so that the probability of receiving a reward or punishment remained relatively equal and constant as time progressed. 55/120 cells enhanced their responses to unexpected rewards compared to expected rewards, and 47/120 cells enhanced their responses to unexpected punishments. These enhanced responses to unexpected reinforcement may reflect potential prediction error signals for punishment and reward. Indeed, immediately after a reversal in image value, the reinforcements received are also unexpected since there has been a change in contingencies and monkeys' predictions are therefore violated. Consequently, we hypothesized that immediately after an image value reversal during trace-conditioning, the responses to rewards and punishments should also be enhanced. We conducted separate analyses on the 55 cells with enhanced responses to reward and on the 47 cells with enhanced responses to punishment. For both analyses, responses to reinforcement immediately after reversal were enhanced. This enhancement was transient, lasting fewer than 10 trials. Thus there was a tight correspondence between the time course of enhanced responses to unexpected rewards and punishments, and the development of both value signals in the amygdala and behavioral learning. Primate amygdala neural activity therefore appears to reflect both appetitive and aversive prediction error signals, and these signals could play an instrumental rule in creating a representation of positive and negative value that can be used to guide behavior during learning.

* These authors contributed equally to this work

92. Plasticity of reverberatory activity in a prototypic hebbian cell assembly

Pak-Ming Lau, Guo-Qiang Bi

University of Pittsburgh School of Medicine

The concept of cell assembly was proposed by Hebb to provide an elementary structure for thought process. The Hebbian cell assembly has two essential properties: 1. neuronal activity can reverberate in specific sequences within the assembly without sustained external drive; 2. synaptic modification resulted from the reverberatory activity further stabilizes the reverberation. Using whole-cell patch-clamp recording and simultaneous calcium imaging, we found that brief (e.g. 1-ms) stimulation of few neurons in a small network of about 100 cultured hippocampal neurons could trigger reverberatory activity in the network lasting for seconds. Such reverberatory activity consists of repeating motifs of specific firing patterns in the network. Paired-pulse stimuli with inter-pulse interval of ~200-400 ms are more effective in activating such oscillatory reverberation. Furthermore, repeated activation of reverberation with paired-pulse stimuli leads to long-term enhancement of subsequent activation by single stimuli. In addition, pairing a non-effective input (that does not activate network reverberation) into one neuron with an effective input (that activates reverberation) into another can convert the non-effective pathway into an effective one. Reverberatory circuits in culture may serve as a prototype of Hebbian cell assembly for studies on the properties and underlying cellular mechanisms. (Supported by NIMH and Burroughs Wellcome Fund)

93. A principle for learning egocentric-allocentric transformations.

Patrick A Byrne, Suzanna Becker

McMaster University

Numerous single-unit recording studies have found rodent hippocampal neurons which fire selectively for the animal's location in space, independent of its heading

direction. The population of such neurons, commonly known as "place cells", is generally argued to maintain an allocentric, or viewpoint-independent internal representation of the animal's location in space. Similar electrophysiological and other data from humans and from non-human primates suggest that such spatial representations might be ubiquitous across mammalian species. The fact that spatial information from the environment must reach the brain via sensory receptors in an inherently egocentric, or viewpoint-dependent fashion, leads to the question of how the brain learns to transform egocentric representations into allocentric ones. Additionally, the medial temporal brain regions in which allocentric representations of space are typically found have also been implicated in the intermediate to long-term storage of spatial memories. If these long-term memory representations of space are to be useful in guiding motor behaviour, then learning the reverse transformation from allocentric to egocentric co-ordinates must also be accomplished. In this work we present a biologically grounded learning principle based on the notions of "minimum reconstruction error" and "representational inertia" (described below) which leads to the learning of allocentric representations of simulated environments. This learning principle does not lead to a unique learning rule when applied to neural network models, but we have found that two different learning rules consistent with this principle lead to the learning of allocentric representations.

The first part of our learning principle is based on the fact that an animal must be able to reconstruct egocentric representations of space (embodied by the activity of a layer of "egocentric" neurons in our models) from stored allocentric representations (embodied by a layer of "allocentric" neurons in our models). This transformation must be mediated by information regarding how the animal is oriented in space relative to some fixed reference direction. We assume that this information is provided by a layer "head-direction cells". Finally, we assume that the "egocentric" and "allocentric" neurons are connected via a population of "transformation" neurons which also receive input from the head-direction system. By assuming this architecture, we are able to write an explicit cost function for reconstruction error in terms of the model weights.

The second part of our learning principle is based on a number of findings (e.g. Redish, et al., 2000) which indicate that certain medial temporal neuronal populations might exhibit a sort of inertia in their activity patterns. Thus, it might be desirable for the nature of allocentric representations created in medial temporal regions during navigation through an environment to be such that they vary as little as possible from one instant to the next. By assuming navigation occurs at a constant velocity, we are able to write an explicit cost function for this representational inertia as well.

By minimizing the combined cost function as our model navigates through a simple environment, allocentric representations of space are learned. We are also able to show that a set of weights which generates egocentric-allocentric transformations for general two-dimensional environments is a local minimum of the cost function. Finally, we show that a more biologically realistic learning rule based on the restricted Boltzmann machine and representational inertia also generates allocentric representations of space.

Reference: Redish, AD McNaughton, BL & Barnes, CA (2000) Neurocomputing 32-33, 235-241.

94. Neural Correlates of Difference in Strategy of Adaptation to Force Perturbations

Xinying Cai, Yury P Shimansky, Jiping He

Arizona State University

The goal of this study is to investigate the dependence of adaptation strategy and its cortical neural correlates on the subject's capacity to compensate for perturbation force applied to the arm during reaching movement. In the first experiment, monkeys are trained to perform 3D center-out reaching movement with high-amplitude short pulse of perturbation force delivered to the wrist through a pneumatic cylinder. The perturbation force was too strong for the monkey to compensate and it induced a constant displacement of the wrist. The main feature of the adaptation strategy was a proactive shift in the initial direction of movement opposite to the direction of perturbation. In the second experiment, same monkeys performed the same reaching task in a virtual reality environment with perturbation force generated by a torque motor. The magnitude of the force was considerably smaller and the monkey was able to exert enough force to compensate for it. As a result, perturbation induced deviation in hand trajectory gradually decreased across adaptation. Chronic multi-unit recording was performed to monitor the modification of neuronal activity in MI during the time course of adaptation. The major difference in neuronal activity is between the magnitudes of the variation of neuronal directional tuning pattern (DTP), which is characterized as the dependence of spike activity on target location. In the first experiment, we observed a substantial day to day variation of DTP. It exhibits a 'trial and error' type of variation during the early phase of adaptation and gradually converged to a certain pattern towards the end. However, in the second experiment, there is no statistically significant variation in DTP across the whole adaptation period. We attribute this differentiation in neuronal activity to different adaptation strategy adopted by the monkey based on his capacity to compensate for perturbation force.

95. Transitions in a bistable model of the calcium/calmodulin-dependent protein kinase-phosphatase system in response to STDP protocols

Michael Graupner, Nicolas Brunel

Laboratoire de Neurophysique et Physiologie, CNRS UMR 8119, Université René Descartes - Paris 5, Paris, France

The calcium/calmodulin-dependent protein kinase II (CaMKII) plays a key role in the induction of long-term post-synaptic modifications following calcium entry. Experiments suggest that long-term synaptic changes are all-or none switch-like events. The biochemical network involving CaMKII and its regulating protein signaling cascade has been hypothesized to be a bistable realization of such a switch. However, it is still unclear whether LTP/LTD protocols lead to transitions between the two states in realistic models of such a network.

A detailed biochemical model of the CaMKII autophosphorylation and the protein signaling cascade governing the CaMKII dephosphorylation is presented. Dephosphorylation is mediated by protein phosphatase 1 whose activity is indirectly regulated by the calcium-dependent balance between protein kinase A and calcineurin. All of these proteins are known to be involved in synaptic plasticity.

As reported by Zhabotinsky [Biophys J 2000 Nov; 79:2211-21], two stable states of the system exist at resting intracellular calcium concentration: a weakly-(DOWN) and a highly-phosphorylated (UP) state of the CaMKII. A transition from the DOWN to the UP state can be achieved by high calcium elevations, with an UP-shifting threshold which is determined by the competing CaMKII autophosphorylation and dephosphorylation dynamics. Intermediate Ca(2+) concentrations enhance CaMKII dephosphorylation due to a relative increase in calcineurin activity. This results in depotentiation - switching from the UP to the DOWN state - during respective Ca(2+) transients. The transitions in both directions, from the DOWN to the UP state and vice versa, are achieved in response to calcium levels which resemble those which are present during LTP and LTD induction protocols. Finally, it is shown that the CaMKII system can qualitatively reproduce results of plasticity outcomes in response to the spike-timing dependent synaptic plasticity (STDP) paradigm.

96. The Tempotron: A Neuron that Learns Spike-Timing-Based Decisions

Robert Guetig, Haim Sompolinsky

Hebrew University of Jerusalem

The timing of action potentials of sensory neurons contains substantial information about the eliciting stimuli. Although computational advantages of spike-timing-based neuronal codes have long been recognized, it is unclear whether and how neurons can learn to read out such representations. We propose a novel biologically plausible supervised synaptic learning rule enabling neurons to efficiently learn a broad range of decision rules, even when information is embedded in the spatio-temporal structure of spike patterns and not in mean firing rates. The number of categorizations of random spatio-temporal patterns that a neuron can implement is several times as large as the number of its synapses. The underlying non-linear temporal computation allows neurons to access information beyond single-neuron statistics and to discriminate between inputs also on the basis of multineuronal spike statistics. Our work demonstrates the high capacity of neural systems to learn to decode information embedded in distributed patterns of spike synchrony.

97. A computational model for self-organized learning of sparse temporal sequences in zebra finch HVC

Joseph K Jun, Dezhe Z Jin

Penn State University

Sparse and precise temporal sequences have been found to drive the timing of song features in zebra finch. Neurons in the nucleus Higher Vocal Center (HVC) that project to the premotor nucleus RA (robust nucleus of the arcopallium) burst spike with sub millisecond precision relative to song features and do so only once per motif (Hahnloser et al., 2002). The neural mechanisms that account for this behavior are as yet unknown, but are believed to be intrinsic to HVC. An important question remains: how do HVC neurons "learn" this neural activity? Is it a genetic trait? Is it trained by upstream nuclei? Or finally, is it a self-organized learning process? In this paper, we explore the third possibility.

Our neural network model starts with a random, all-to-all synaptic connectivity. We perform simulations using both conductance based neurons and, for computational efficiency, binary spiking neurons--units that spike by comparing excitatory and inhibitory synaptic strengths; results are similar for both kinds of neurons. Feedforward excitatory projections define a subset of the neurons as the first layer; this is the only structure we impose on the network. All neurons spike at low spontaneous rates due to uncorrelated synaptic noise. Feedback inhibition is mediated by interneurons in conductance based simulations and is given directly in binary simulations.

We model the following three plasticity mechanisms: 1) spike timing dependent plasticity, 2) homosynaptic depression, and 3) competitive synapse silencing. Spike timing dependent plasticity allows strengthening or weakening of synapses based on network activity. Homosynaptic depression is a constant rate of weakening on all synapses, independent of activity; it forces synapses to persistently potentiate to remain strong; therefore, it favors synapses whose postsynaptic targets receive cooperation from many presynaptic neurons. Competitive synapse silencing is a bound placed on the total number of strong--as defined by a threshold--output synapses that each neuron can have. The plasticity we model in the system have not yet been observed experimentally in HVC; they are, however, plasticity protocols that are found in other neural systems; we propose their existence in HVC.

Using the protocol above, we find that the neural network can learn a synaptic connectivity that generates sparse and precise spike timings; the learned networks are consistent with synfire chains--an architecture where neurons are organized into synchronous groups that fire successively; synfire chains have already been proposed as the neural connectivity in HVC (Jin et al. 2005). We also find that our learning algorithm is robust to "lesions" of neurons, where the synaptic input and outputs weights of the lesioned neuron are randomized; the network can retain the majority of its structure, requiring only local rearrangement to rebuild parts of the lesioned network. Neurons within HVC are constantly renewed throughout the life of the bird, including adulthood; dealing with this kind of neuronal loss must be accounted for by any model of activity learning in HVC.

References

Hahnloser, R.H.R., A.A. Kozhevnikov, & M.S. Fee (2002) Nature, 419:65-70.

Jin, D.Z., F.M. Ramazanoglu, & H.S. Seung, (2005) Submitted.

98. Basis for Training-Induced Plasticity of Auditory Localization in Adult Mammals

Andrew J King, Oliver Kacelnik, Fernando R Nodal, Carl H Parsons

Department of Physiology, Anatomy and Genetics, University of Oxford

Accurate auditory localization relies on neural computations based on spatial cues present in the sound waves at each ear. The values of these cues depend on the size, shape and separation of the two ears and can therefore vary from one individual to another. As with other perceptual skills, the neural circuits involved in spatial hearing are shaped by experience during development and retain some capacity for plasticity in later life. However, the factors that enable and promote plasticity of auditory localization in the adult brain are unknown. We have shown that mature ferrets can rapidly re-learn to localize sounds after altering their spatial cues by reversibly occluding one ear, but only if they are trained to use these cues in a behaviorally-relevant task, with greater and more rapid improvement occurring with more frequent training. Removal and subsequent re-insertion of the earplug result in relatively modest behavioral changes compared to those observed when the ear is first occluded, implying that a systematic retuning of neurons to the altered binaural cue values is unlikely to provide the basis for auditory localization plasticity. Instead, adaptation could involve learning to place greater weight on other auditory spatial cues that are less affected by monaural occlusion, such as the spectral cues provided by the open ear. To investigate this possibility, we modified these spectral cues by applying a unilateral mold that filled the various cavities and depressions of the ear while leaving a clear passage for sound to reach the external auditory meatus. This led to a significant impairment in localization accuracy by ferrets that had previously adapted to occlusion of the other ear, whereas the same unilateral manipulation of the spectral cues had a much smaller effect on the performance of control ferrets. Moreover, ferrets that had adapted to a unilateral earplug localized narrowband sounds, particularly at high center frequencies, poorly compared to control animals. Together, these findings suggest that this form of adaptive plasticity involves a shift in sensitivity away from the abnormal auditory spatial cues, specifically high-frequency interaural level differences, to other cues that are less affected by the earplug, namely spectral cues and low-frequency interaural time differences. The mature auditory system is therefore capable of adapting to abnormal spatial information by learning to re-weight different localization cues, suggesting that sound-localization training could be useful for promoting rehabilitation in patients whose perception of auditory space is compromised as a consequence of hearing disorders.

99. An integrate-and-fire model of temporal context specific episodic encoding and retrieval in the hippocampal formation

Randal A Koene, Michael E Hasselmo

Boston University Center for Memory and Brain

Making correct decisions often depends on a memory of recently experienced episodes, and the time of occurrence of a specific episode can be an essential cue during memory retrieval. Hasselmo and Eichenbaum [2] proposed mechanisms for the temporal context dependent storage and retrieval of episodic representations in the hippocampus. In that paper, threshold units were used to simulate the context sensitive responses of hippocampal neurons during a spatial alternation task, even when successive events in a familiar environment were similar. Time differences were encoded in learned episodic associations by using an exponentially decaying multiplicative factor that indicates the level of activity maintained in layer II of entorhinal cortex (ECII) for a specific representation.

In order to explore the physiological plausibility of the theoretical model, we have implemented these mechanisms using integrate-and-fire neurons and synapses that are modified by spike timing dependent plasticity (STDP). The model includes the following features:

1. The spiking neuron model includes a novel method of encoding associations with graded strength that is proportional to the ordinal distance between events in a sequence. This encoding occurs in simulated recurrent excitatory fibers of hippocampal region CA3, while a queue of recently experienced events is maintained in a simulated short-term buffer in ECII [6, 5]. The existence of the persistent firing buffer in ECII is supported by the observed intrinsic firing properties of pyramidal neurons [4, 1] in ECII and by task specific changes of the power of theta and gamma oscillations recorded in the hippocampal EEG [3].
2. In the model, layer III of entorhinal cortex stores representations of a sequence of spatial associations.
3. In the dentate gyrus of the model, unique representations are formed for each temporal context.
4. Using these temporal context specific representations, episodic associations between event representations are stored in CA3. The graded strength of these associations, established by recurrent excitatory connections between neurons in CA3, is achieved by (a) synaptic strengthening that is proportional with the number of theta cycles during which both sets of spikes representing two events are maintained in the ECII buffer, and by (b) a STDP function that modifies synaptic strength in a manner proportional to the time interval between the pre- and postsynaptic spikes of the representations that are reactivated in each cycle of ECII buffer activity.
5. The spiking of retrieved spatial associations and of retrieved episodic representations combines to elicit activity in hippocampal region CA1. Selective spiking in CA1 is therefore temporal context specific. The retrieval mechanisms in CA1 depend upon the phase of different afferent input to CA1 during the hippocampal theta rhythm.

Our simulation results give a physiologically plausible account of proposed mechanisms by which the medial temporal region may be able to retrieve episodes in a specific temporal context. Such retrieval provides the necessary information for decisions, such as those in a delayed spatial alternation task. Supported by NIMH MH60013, NIMH MH61492, NIH NIDA DA16454 and the NSF Science of Learning Center SBE 0354378 "CELEST".

[1] E. Fransen, A.A. Alonso, and M.E. Hasselmo. Simulations of the role of the muscarinic activated calcium-sensitive nonspecific cation current iNCM in entorhinal neuronal activity during delayed matching tasks. Journal of Neuroscience, 22(3):1081-1097, 2002.

[2] M.E. Hasselmo and H. Eichenbaum. Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Networks, 18(9):1172-1190, 2005.

[3] M.W. Howard, D.S. Rizzuto, J.C. Caplan, J.R. Madsen, J. Lisman, R. Aschenbrenner-Scheibe, A. Schultze-Bonhage, and M.J. Kahana. Gamma oscillations correlate with working memory load in humans. Cerebral Cortex, 13:1369-1374, 2003.

[4] R. Klink and A. Alonso. Muscarinic modulation of the oscillatory and repetitive firing properties of entorhinal cortex layer II neurons. Journal of Neurophysiology, 77(4):1813-1828, 1997.

[5] R.A. Koene, A. Gorchetchnikov, R.C. Cannon, and M.E. Hasselmo. Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. Neural Networks, 16(5-6):577-584, 2003.

[6] J.E. Lisman and M.A.P. Idiart. Storage of 7+/-2 short-term memories in oscillatory subcylces. Science, 267:1512-1515, 1995.

100. Learning of Representations in a Canonical Model of Cortical Columns

Jörg Lücke

Gatsby Computational Neuroscience Unit, UCL, UK

We present a dynamical model of processing and learning in a cortical column, which reflects microcircuit connectivity and receptive field formation in the neocortex.

In recent years efforts to analyze the wiring and functioning of small scale (order 100-10,000 neurons) neural circuits have generated a rich collection of data on axon and dendrite branching of different cortical cell types together with main synaptic target locations.
Localized cortical modules on the order of 10,000 neurons are usually referred to as cortical columns and it is widely believed that the distribution of neural activity within a column represents properties of input received from a common source, e.g., an area of the visual field for columns in the visual cortex or a patch of the skin for columns in somatosensory cortex.

The model of a cortical column presented here is described by a set of coupled differential equations describing neural activity, background oscillation, and Hebbian type synaptic plasticity within a column.
The dynamics of neural activity is motivated by recent findings about columnar connectivity. Oscillatory lateral coupling and balanced excitation and inhibition result in a representation of input by distributed neural activity within a column. A Hebbian type synaptic plasticity dynamics which is coupled to the columnar neural activity gives rise to a self-organization process of afferent fibers to the column. Self-organization is input driven and leads to the emergence of different receptive fields for different features of the presented input.

In contrast to all other systems, the suggested dynamical model captures the extra-ordinary flexibility and robustness of cortical learning, which is demonstrated by applying it, without parameter change, to hand-written digits, the bars benchmark test for feature extraction, and natural images:
For hand-written digits, images of a standard database (MNIST) are used as input and the system is shown to self-organize its receptive fields to represent typical classes of all hand-written digits in an unsupervised way.
Using overlapping bars as input (the bars test) it is shown that the system is able to represent input images distributedly. Furthermore, the bars benchmark test allows to quantitatively demonstrate that the model's robustness is superior to that of all other systems so far suggested.
Using natural images of a standard database (van Hateren) as input, the system develops Gabor wavelet like receptive fields similar to the ones of simple cells in primary visual cortex.
In all experiments the same system with the same set of parameters is used in order to demonstrate the model's flexibility.

The presented system combines bottom-up neural modeling with functional capabilities so far only reported from special-purpose systems using functional top-down approaches. In terms of flexibility and robustness it is superior to all systems suggested so far. This high flexibility together with columnar connectivity and population coding makes the presented system to a plausible model for a canonical microcircuit of the neocortex.

102. Adaptive control and the flow of information in the brain

Mohamed N Abdelghani[1], Timothy P Lillicrap[2], Douglas B Tweed[3]

[1]University of Toronto, [2]Queens University, [3]Univeristy of Toronto

Control systems in the brain are very flexible learners – more flexible, in fact, than current theories are able to explain. The problem is that any adaptive controller needs to know how its command u affects its performance error e (e.g. in reaching, e might be the vector from hand to target). Specifically, it needs to know ¶e/¶u. But how does the controller get this information? One possibility is that ¶e/¶u is learned, as in model-based adaptive control: a plant model learns the relation between e and u, and sends ¶e/¶u to the controller. In the brain, though, model-based adaptive control hits a snag: a network can learn to estimate e, but then ¶e/¶u is not coded in its firing but is distributed among its synaptic weights, and how can that information be sent to the controller? Because of this difficulty, it has been suggested that model-based adaptive control can't work in the brain, and newer theories have therefore pursued the only alternative: if ¶e/¶u can't be learned then a rough estimate of it must be known innately. But if knowledge of ¶e/¶u were solely innate, then when any element of ¶e/¶u changed sign, learning would collapse. In real sensorimotor systems, learning doesn't collapse in this way. To explain this, we show that model-based adaptive control is possible in the brain, and we propose a mechanism.

The mechanism hinges on a new type of plant model, which codes ¶e/¶u in its firing patterns. One problem has been that there is no supervisor to train such a network – no signal to tell it the correct values of ¶e/¶u. But we can solve this problem by relating ¶e/¶u to variables whose values are known. We suppose the plant model receives a vector y of signals from which e could, in principle, be estimated (e.g. in the case of the VOR, y might include head-velocity signals from the inner ear and eye-position feedback from spindles or efference copy; and always among the components of y are all components of u, because e depends on u). By the chain rule, the time-derivative e' = (¶e/¶y)y'. As e' and y' are known, it is possible, given enough samples of e' and y', to solve for ¶e/¶y, and thereby obtain ¶e/¶u. This can be done by a simple network, described by the equation x = (Wz)y', where z is an input vector, W is a third-order array of weights, and Wz is a matrix of signals. If we define E = e' – x then the learning rule $w'_{ijk} = \eta E_i y'_j z_k$ causes the signal matrix Wz to code ¶e/¶y. We call this mechanism implicit supervision because it uses e' as a sort of indirect supervisor to force Wz toward ¶e/¶y.

This mechanism explains the flexibility of neural adaptive controllers. If the genome were the sole source of knowledge about ¶e/¶u, then when the sign of any component of ¶e/¶u changed, learning would become counterproductive. But in fact, motor learning copes pretty well with such reversals, e.g. with mirror drawing, reversing spectacles, and tendon or nerve transposition. This finding shows that something like model-based adaptive control must be at work in the brain. Implicit supervision isn't the only possible mechanism, but it does have some attractive features. It is consistent with known neurophysiology. It has computational advantages (e.g. the plant model is linear in its parameters and therefore converges to a globally optimal set of weights, W; the system is also robust, as the plant model needn't learn the exact values of all components of ¶e/¶u, merely their signs). And it allows the same generic network to handle a very wide range of sensorimotor tasks.

103. Learning to learn: motor adaptive strategies change with environmental experience

Michael S. Fine, Jordan A. Taylor, Kurt A. Thoroughman

Washington University

Recent studies of trial-by-trial motor learning have shown that error in a single movement induces adaptation in the immediately subsequent movement (Thoroughman and Shadmehr, 2000; Scheidt et al., 2001). This adaptive process has been hypothesized to consist of two stages: an abstraction of a quantitative error metric from a movement, and the application of this error to generalize learning across movement space. To date, studies of trial-by-trial motor adaptation have shown that adaptation scales with sensed error (Thoroughman and Shadmehr, 2000; Scheidt et al., 2001) and generalizes broadly across the movement space (Thoroughman and Shadmehr, 2000; Donchin et al. 2003). Here we discuss two experiments that show surprising flexibility of both components of learning: sensory feedback can induce adaptation strikingly disproportional to movement error, and environments can induce narrowing of generalization across movement space.

Previous studies have shown that adaptation, across single movements, is proportional to sensed error (Thoroughman and Shadmehr, 2000; Scheidt et al., 2001). Theories presume this proportionality and have successfully mimicked human behavior when people experience novel environments throughout the extent of a movement. We found that the human response to pulsatile perturbations, in contrast, contained no proportionality to any error signal. We perturbed reaching movements using 70 ms pulses of force generated by a robotic manipulandum; pulses were distributed pseudorandomly throughout a movement set. Human subjects generated a robust adaptive response in post-pulse movements that opposed the pulse direction. All pulses yielded similar changes in predictive control regardless of the location or magnitude of the pulse. Neural models that broadly encoded movement velocity and adapted proportionally to motor error mimicked human insensitivity to pulse location, but failed to mimic human insensitivity to pulse magnitude. We conclude that single trial adaptation to force pulses reveals a categorical adaptive strategy that aims to counter the direction, rather than the magnitude, of movement error. This work strongly challenges the notion that incremental steps of motor adaptation scale with error.

Movement-by-movement generalization has been shown to be broadly tuned to movement direction (Thoroughman and Shadmehr 2000; Donchin et al. 2003), and has been suggested to remain fixed as human subjects learn novel environments (Shadmehr, 2004). However, to date, these environments have only mimicked natural environments, in which the force direction changes smoothly across movement space. To identify further the movement-by-movement generalization properties of the motor system, we trained subject in fields in which the perturbing forces changed direction as fast (low-complexity), twice (medium), or 4 times (high) as fast as the movement direction. We found that to learn the fields, subjects changed their movement-by-movement generalization of error across the movement space to match the complexity of the fields. A neural network model was capable of mimicking human adaptation, but required the underlying units of the network to narrow their tuning of movement space and reduce their gain as the environmental complexity was increased. This result challenges previous theories of motor learning that require neuronal tuning, which underlies generalization, to remain fixed during the course of learning (Thoroughman and Taylor, 2005).

The richness of human motor behavior arises from our ability to learn new movements. A successful transformation of sensed error into adaptation has been previously hypothesized to depend on real-valued error (Wolpert and Ghahramani, 2004). Here we report two novel adaptive strategies: flexible error generalization across movement direction, and categorization that defines the direction, but not magnitude of error. We now consider both components of the learning process to adapt as a function of environmental experience.

Donchin O, Francis JT, Shadmehr R (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. J Neurosci 23:9032-9045.

Scheidt RA, Dingwell JB, Mussa-Ivaldi FA (2001) Learning to move amid uncertainty. J Neurophysiol 86:971-985.

Shadmehr R (2004) Generalization as a behavioral window to the neural mechanisms of learning internal models. Hum Mov Sci 23:543-568.

Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407:742-747.

Thoroughman KA, Taylor JA (2005) Rapid reshaping of human motor generalization. J Neurosci 25:8948-8953.

Wolpert DM, Ghahramani Z (2004) Computational Motor Control. In: The Cognitive Neurosciences, 3 Edition (Gazzaniga MS, ed). Cambridge, Massachusetts: The MIT Press.

104. Learning without synaptic change: a mechanism for sensorimotor control

Kristen P Fortney, Douglas B Tweed

University of Toronto

When the brain learns, it may store new information in chemical form, by altering its synapses and other cell properties, or in reverberating loops of activity. Most theories of sensorimotor learning depend on synaptic change, but here we show that loops could also play a role, and would bring some advantages: they avoid the problem of weight transport, they learn control tasks with fewer neurons, and they account for very rapid learning, which is hard to explain based only on synaptic change.

A simple example of synaptic learning is the Widrow-Hoff rule. Suppose a neuron fires at a rate $u = w \cdot z$, where z is a vector of inputs and w are synaptic weights. If we know that the desired firing rate is u*, then the error $e = u^* - u$. Applying the learning rule $dw/dt = \eta e z$ (where η is a positive constant) makes u converge to u*. But the same effect can be achieved without synaptic change: if we compute u by running the signal ηe through an integrator, so that $du/dt = \eta e$, then u will again converge to u*. So here a reverberating loop, in the form of an integrator, does the job normally performed by adjusting the weight vector w. We will call this integrator-based mechanism weightless learning, because it avoids the need to adjust synaptic weights.

Sensorimotor learning is more complex than the above case because error signals for the neurons involved aren't readily available. In some theories this problem is solved using a forward plant model, which learns to compute appropriate error signals. With weightless learning, the same can be achieved by connecting two integrators in series. We define a vector $z = (u, y)$, where u is a vector of neural commands and y includes any other variables that affect performance error e. We also define a scalar loss function L (which might for instance be $|e|^2$) and a vector L_z, which is a neural estimate of the derivative of L with respect to z – of course we have $L_z = (L_u, L_y)$. By the chain rule, the product $L_z \cdot (dz/dt)$ is a neural estimate of dL/dt; the difference between this estimate and the true value, $dL/dt - L_z \cdot (dz/dt)$, is the model error E_m. With these variables, and two integrators, we can build an adaptive control system that minimizes L: integrating $\eta_1 E_m dz/dt$ yields L_z and then integrating $-\eta_2 L_u (L + L_y dy/dt)$ produces the motor command u. This mechanism applies to a wide range of control tasks; we show with simulations that it works for the vestibulo-ocular reflex and for two-joint reaching: it rapidly gains control of the plant (the eye or arm) and adapts when the plant changes.

One advantage of weightless learning is that it avoids the need for rapid weight transport (rapid transmission of information from inside some synapses to other, remote synapses) which is considered biologically implausible. To avoid weight transport, many weight-based learning algorithms rely on massive preprocessing, or "expansion recoding", but that calls for large numbers of neurons and synapses. The result is slowed learning. Weightlessness avoids this preprocessing, so it can learn faster and work with fewer cells.

The main disadvantage of weightless learning is that it learns on the fly. It doesn't accumulate knowledge in synapses, but instead updates commands based only on their previous values and the values of incoming signals. If these signals are changing too quickly, learning may not keep up. Consequently, weightless learning works best in conjunction with weighted learning: you can quickly acquire a new skill using weightlessness, and then consolidate it in synaptic weights. The resulting hybrid mechanism still requires far fewer neurons than purely synaptic learning. It also explains why we tend to move slowly in the early stages of learning, and why we instinctively slow down when we start to make errors.

105. Electrotaxis of C. elegans in fixed and time-varying fields

Christopher V Gabel, Aravinthan Samuel

Department of Physics. Harvard University

The nervous system of the nematode C. elegans contains only 302 neurons (in the hermaphrodite) and yet can elicit precise directed movement in response to external stimuli. In an electric field, C. elegans crawl toward the negative pole at an acute angle to the field lines. By quantifying the electrotactic movements of individual worms, we find that this preferred angle is strictly proportional to field strength. Using electric fields varied in direction and amplitude, we find that C. elegans orient toward and maintain their preferred trajectories by employing standard maneuvers including forward crawls, turns, and reversals. Finally, we identify sensory neurons, interneurons and motor neurons with critical roles in the sensorimotor pathway from electric field stimulus to behavioral output.

106. Echolocating bats use a prey intercept strategy that is time-optimal in a local, piece-wise linear sense

Kaushik Ghose[1], Timothy K Horiuchi[2], P. S. Krishnaprasad[2], Cynthia F Moss[3]

[1]Dept. Psychology, Neuroscience and Cognitive Science Program, University of Maryland, College Park, [2]Dept. Electrical and Computer Engineering, Neuroscience and Cognitive Science Program, Institute for Systems Research, University of Maryland, College Park, [3]Dept. Psychology, Neuroscience and Cognitive Science Program, Institute for Systems Research, University of Maryland, College Park

Computations in the nervous system enable animals to localize, track and capture moving targets. Studies suggest that an animal may intercept a moving target by functionally predicting its future trajectory, but the strategy employed to attain this goal is not necessarily the same for all species. In one proposed interception strategy, the animal keeps the angle between its motion and the target a constant during pursuit while it attempts to travel in a straight line. Studies of target pursuit in fish, dragonflies and humans have shown that the CB strategy is widespread across different species. Here, we show that the CB model is insufficient to explain the prey pursuit strategy of echolocating bats. The flight of echolocating bats chasing insect prey is better explained by a constant absolute target direction (CATD) strategy where the pursuer keeps the absolute direction to the target a constant while closing target distance. The CATD strategy produces a trajectory which, from the viewpoint of the target, makes the pursuer 'appear' stationary against a distant background and vice-versa. A recent observation of similar behavior in dragonfly territorial flights has been interpreted as camouflaging behavior on the part of the pursuer. Generally, insects encounter bats in darkness, with the bat emitting intense ultrasonic pulses that many insects can hear. Because motion camouflage is primarily useful for defeating visual detection and the bat is announcing its presence and direction with sonar vocalizations, the CATD strategy is unlikely to be used for camouflage. The bat's strategy is equivalent to following an intercept course to the target at every instant of time, assuming the target will continue moving at its current velocity. We argue that, for piecewise linear trajectories of the target and bat, this strategy is locally time-optimal. Interestingly, a similar strategy has been implemented in some guided missiles, and is called paralell navigation. (NSF, CRCNS-NIBIB)

107. Implications of threshold nonlinearities on mechanisms underlying persistent neural activity in a bilateral neural integrator

Itsaso Olasagasti[1], Emre Aksay[2], Guy Major[2], David W Tank[2], Mark S Goldman[1]

[1]Wellesley College, [2]Princeton University

Multilevel persistent neural activity is thought to be the neural correlate of short-term memory and has been observed in a variety of cortical and subcortical brain regions. In the oculomotor neural integrator, a bilateral brainstem region, integrator neurons maintain persistent firing rates that are threshold linearly related to eye position in the absence of external inputs. Due to the mathematical complexity of the resulting system, most previous models of the neural integrator have neglected either key nonlinearities in the firing rate vs. eye position tuning curves or known inhibitory connections. We have developed a modeling framework that directly incorporates measured neuronal tuning curves into a network with experimentally constrained inhibitory and excitatory connections. With this model, we reveal evidence for novel mechanisms underlying the functioning of the system.

Our model is based on experiments conducted in a model system, the goldfish oculomotor neural integrator. In agreement with the anatomy and physiology of this system, the model includes four populations of neurons, one excitatory and one inhibitory population on each side, and thresholds limited to the contralateral half of the eye range. Further, the model can be tuned without artificial constraints on synaptic connections such as have been used in previous models. This is accomplished by directly incorporating the known tuning curves of oculomotor neural integrator neurons. The model is implemented as a conductance-based spiking network of neurons whose firing rate versus current injection relationships qualitatively match those found experimentally.

Unilateral inactivation experiments performed in the goldfish show that stable eye positions are compromised only on one half of the oculomotor range, specifically for eye positions ipsilateral to the lesion. Most previous models of the system do not reproduce this result, instead predicting a drift in eye position over the whole oculomotor range. We find a class of networks that successfully reproduces this result. Surprisingly, although these models contain recurrent inhibitory connections, we find that disinhibitory positive feedback does not contribute to the generation of persistent neural activity in the network. Rather, we find that inhibition acts in a “feedforward” manner as a simple relay of persistent neural activity that is generated by recurrent excitation. This contrasts with longstanding assumptions about the role of recurrent inhibition as a source of positive feedback in the integrator.

By analyzing a variety of network architectures and synapse types (such as saturating, linear, and superlinear synapses), we can eliminate large classes of previously proposed models that can be shown to be inconsistent with the experimental rate vs. eye position relationships in either the lesioned or nonlesioned integrator. Our results suggest that network interactions are mediated by a high-threshold process that is turned on at elevated firing rates (for example, dendritic plateau potentials). This result is consistent with previous work suggesting that plateau potentials can enhance the robustness of persistent neural activity to perturbations in network inputs or connectivity. Our analysis may be applicable to other systems exhibiting a push-pull arrangement of inputs and nonlinearities due to variable neuronal firing rate thresholds.

108. A unified optimal control treatment of reaching and tracking

Dongsung Huh, Emanuel Todorov

UCSD

Reaching and tracking tasks have been studied extensively in motor control research. While the two are presently thought to be different behaviors, we believe that a unified treatment within the optimal control framework is possible. We view reaching as generalized tracking and vice versa. Indeed switching between the two can happen in a single motor task – such as catch-up saccades in smooth pursuit.

Here we introduce a new approach to reaching/tracking problems where a cost function suitable for both is defined, and the predicted behavior is the one which is optimal with respect to that cost function. This approach is based on two hypotheses: 1. subjects can decide when to incur a cost for being off target; 2. if they decide not to incur this accuracy cost over a certain time interval, a duration cost is automatically applied. It turns out that after some simple algebra, the hypotheses nicely reduce into a cost function which can be understood as hybrid of reaching and tracking. The resulting cost function nicely explains how reaching/tracking alternation occurs as a function of discrepancy between target and hand movement. It also suggests how duration costs might be defined in an ecologically plausible way.

We calculate the optimal control law (and corresponding optimal movement trajectories) for this hybrid cost function using a modified version of the generalized iLQG method we have recently developed (Todorov and Li, ACC 2005). The optimal solutions seem to follow realistic arm movement patterns, but a rigorous quantitative comparison remains to be done. We expect that the model will be in agreement with Fitts' law (the logarithmic relationship between target size, target distance and movement duration).

We are also training recurrent neural networks, operating in closed-loop with a simulated arm, to optimize the cost function used in the optimal control simulations. Preliminary results show that such training is possible, and that the neural responses predicted by the network are similar to what is observed in primary motor cortex.

109. Neuromechanical Modeling of Zebrafish Locomotion

Etienne Hugues[1], Donald P Knudsen[2], John A Arsenault[2], Donald M O'Malley[2], Jorge V Jose[1]

[1]SUNY at Buffalo, [2]Northeastern University

In order to understand the neural control of locomotor behaviors, one must understand both the neurodynamics of the CNS and the mechanics of the locomoting animal. In the case of piscine swimming, this requires investigation of the biomechanical properties of the fish and its hydrodynamic interactions with the aqueous environment. We are investigating the neural control of larval zebrafish locomotion because of the larva's transparent CNS and because many neurons in brainstem and spinal cord can be individually identified. These features facilitate experimental and computational modeling investigations because: (1) optical methods are efficient at identifying and systematically characterizing neurons involved in locomotor behaviors and (2) when neurons can be identified as exact individuals, one can accumulate multiple kinds of data that are relevant to its physiological functioning and subsequently incorporate such data into a comprehensive mathematical model. We report here on our modeling efforts to understand the larval locomotor system with particular emphasis on (1) interactions of the larva with its aqueous environment and (2) neurodynamical activity within the CNS that underlies a pair of conceptually important locomotor behaviors, slow and burst swims.

To understand how a zebrafish larva propels itself through the water, one needs to understand the generation of force by axial (trunk) muscles, the elastic and mechanical properties of the trunk and tail, and the hydrodynamic interactions between fish and water, within the largely viscous regime in these larvae swim. We have developed a computational model of this system that takes into account: quasi–static torques generated around the midpoint of the trunk, total drag moment acting on the body, the moment added by mass forces, lift, drag force acting on the body and fin, and the added mass forces. We assume that vortices shed from the tail fins are swept quickly away, but this will need further investigation. Based on these assumptions, we found that this model produced kinematic patterns representative of the experimentally-observed larval swimming patterns. The model also exhibited a speed/TBF (tail-beat frequency) relationship consistent with the observed faster swimming at higher TBFs. These efforts are taking us closer to an accurate representation of two distinct larval swim patterns, slow and burst swims, which are of significant interest since they may represent an early evolutionary version of distinct locomotor gaits more commonly associated with tetrapod locomotion. Determination of the axial force profiles during these swim patterns is our ultimate goal because these force patterns constitute, in effect, a transformed representation of the spatial-temporal patterns of spinal neural network activity that generated them.

In parallel, we are developing a virtual locomotor control system that will be based upon a growing body of anatomical, physiological and behavioral zebrafish data. As a first step we simulated 3 candidate spinal architectures to evaluate their suitability to serve as segmental oscillators underlying undulatory swimming. The simplest was a 6-cell model that is similar to a well-known lamprey segmental model (Grillner et al., 1988) that consists of descending excitatory and crossed inhibitory spinal interneurons. When simulated using realistic axonal projection patterns, based on the zebrafish VeMe and CoBl interneurons, we found that the 6-cell model more robustly produced the range of tail beat frequencies associated with larval burst and slow swims than did less realistic models. This was in response to varying the synaptic strengths of excitatory and inhibitory synaptic contacts. Extension of this to 8-cell models, based on additional spinal interneurons with evolutionarily conserved features, revealed that an eclectic variety of spinal architectures is able to largely support TBFs that span the larva's natural frequency range (20 to 90 Hz). Interestingly, we found that the fundamental rhythm generation mechanism varied as a function of TBF. This must be considered a first step in this effort because there is growing evidence that these larvae possess independent control systems for burst and slow swims, possibly using functionally distinct spinal CPGs. In our most recent work we have begun integrating the spiking dynamics of the neural model with our biomechanical formulation and have reproduced swimming behaviors that compare well with previous experimental results.

110. Population coding of reaction time performance in rat motor cortex and dorsal striatum

Mark Laubach, Nandakumar S Narayanan, Eyal Y Kimchi

Yale University

Variations in reaction time (RT) performance are predicted by changes in instantaneous firing rates of neurons in rat motor cortex (Laubach et al., 2000). Analysis with statistical classifiers shows that information about RT performance emerges during learning (Laubach et al., 2000) and is based on highly redundant interactions between neurons (Narayanan et al., 2005). Here, we compare properties of motor cortical neurons with neurons in the dorsal striatum, a part of the basal ganglia that is critical for RT performance (e.g., Carli et al., 1985). Using chronic, multi-electrode recording methods, we found that the percentage of neurons with task-related modulations in firing rate are roughly equal (about 60%). Inspired by Paz et al. (2005), we characterized response properties over the neuronal populations using principal component analysis and least-dependent components analysis (Stögbauer et al., 2004). This analysis found that most neurons in motor cortex were modulated around the time of response initiation. By contrast, most neurons in striatum were modulated after response initiation. Statistical classifiers were then used to study decoding and neuronal interactions were quantified using subensemble analysis (Narayanan et al. 2005), an information theoretic approach based on Schneidman et al. (2003). Two hypotheses were tested: (1) that decodings of RT performance variability (conditioned vs premature responses; short vs long RTs) should be based on pre-movement firing and therefore should involve a larger portion of the neuronal population in motor cortex than in striatum; (2) that, due to differences in anatomical connections and synaptic organization, decodings based on striatal neurons should be less redundant than those based on motor cortical neurons. We found that task-related information in motor cortex was redundant over many neurons. By contrast, task-related information in striatum was sparsely distributed over mostly statistically independent neurons, and these neurons encoded less information, on average, than cortical neurons. The consequences of these differences in population coding in the rat motor cortex and dorsal striatum for sensorimotor and conditional association learning and for the effects of aging on cortical neurons and networks will be discussed.

111. Odor identity and concentration coding in the model of the locust olfactory system

Collins G Assisi[1], Mark Stopfer[2], Gilles Laurent[3], Maxim Bazhenov[1]

[1]The Salk Institute for Biological Studies, [2]NIH-NICHD, [3]California Institute of Technology

The first olfactory relay in the locust, the antennal lobe, consists of a tightly interconnected network of excitatory projection neurons (PNs) and inhibitory interneurons (LNs). Stimulating the antennal lobe by presenting an odor to the antenna leads to field potential oscillations resulting from a distributed, coherent population response. Which PNs transiently synchronize with each other and with the local field potential depends strongly on the identity of the odor. Changing the concentration of an odor can modify the PN responses and the identity of the PNs involved while still maintaining the underlying LFP oscillation (Stopfer et al, Neuron, 2000). Both the identity and the concentration of odors, therefore, can be encoded in the spatiotemporal firing patterns of antennal lobe activity. By what mechanism, then, are these spatiotemporal patterns organized, and how can the identity and the intensity of odors be disambiguated by follower neurons?

Here, we use a realistic computational model of the locust olfactory system (Bazhenov et al., Neuron, 2001) to study its response properties for different odor concentrations. The input to the antennal lobe (an odor) was simulated as firing from olfactory receptor neurons to a set of PNs. Individual PNs were assumed to be maximally sensitive to a particular odor with a distribution of lesser sensitivities for other odors. An increase in concentration was simulated by recruiting additional PNs with input below that of the preferred PNs. We demonstrate that the simulated antennal lobe network preserves the salient properties seen in the experiment, including odor and concentration based clustering, while retaining the LFP oscillations at approximately 20 Hz. A dimension reduction analysis revealed that PN responses for different odor concentrations are grouped into clusters with different clusters corresponding to different odors. When odor representations at each time point were considered as instantaneous vectors of activity across a whole PN population, the activity elicited by different odors and concentrations could be separated easily. PN activities for different odors diverged within a few cycles of the LFP oscillation while the activity corresponding to different concentrations of a particular odor evolved along neighboring trajectories.

The representation of the odor in the antennal lobe is vastly sparsened in the mushroom body, the next stage in the olfactory relay. The Kenyon cells (KCs) of the mushroom body receive both divergent excitatory input from the PNs and feedforward inhibition from lateral horn interneurons (LHIs). Coactive PNs excite the KCs and compete with inhibition from the LHIs creating cyclic windows of time over which KCs are sensitive to direct excitatory PN input. In our study the output of the AL network was projected to the network of KCs where each neuron was modeled using a system of difference equations (maps) (Rulkov et al, JCNS, 2004). With this computationally efficient approach a network of thousands of KCs was simulated. We successfully simulated the sparsening of the odor representation in the mushroom body and examined how odor concentration and identity are represented in the mushroom body as a function of the divergence of connections from the PNs to the KCs.

112. Biologically realistic neural inhibition in arbitrary neural circuits

Christopher Parisien[1], Charles H. Anderson[2], Chris Eliasmith[1]

[1]University of Waterloo, [2]Washington University in St. Louis

In cortical neural networks, connections from a given neuron are usually either inhibitory or excitatory, but not both (this feature is often called "Dale's Principle"). Nevertheless, many current neural models ignore this constraint. Past attempts at transforming such unrealistic networks to ones respecting this principle fail in general because they are either too specific to one kind of neural circuit, or they introduce other biologically implausible assumptions. We demonstrate a constructive transformation of any neural model that solves this problem in both feedforward and dynamic recurrent networks. The resulting models are functionally equivalent to the original networks, maintain the correct temporal dynamics of the system, and are biologically plausible.

More precisely, we identify a general form for the solution to this problem. As a result, we also describe how the precise solution for a given cortical network may be determined empirically.

113. On the Dynamics of Electrically-coupled Neurons with Inhibitory Synapses

Juan Gao[1], Philp Holmes[2]

[1]Department of Mechanical and Aerospace Engineering, [2]Department of Mechanical and Aerospace Engineering & Program in Applied and Computational Mathematics. Princeton University

We study the dynamics and bifurcations of noise-free neurons coupled by gap junctions and inhibitory synapses, using Poincaré spike-to-spike maps. We focus on the case of two cells, but also show that stable synchronous splay states exist for globally coupled networks of N cells dominated by subthreshold electrical coupling.

We simplify the alpha function model of synaptic coupling to obtain the delayed delta function model and show that the explicit discontinuous maps of the latter compare well with the continuous history-dependent, implicitly-defined maps of the former.

Our results agree with those of Lewis and Rinzel (J. Comp. Neurosci. 14:283-309, 2003) in the weak coupling range, but our Poincaré map analysis yields more information about global behavior and domains of attraction. We find that increased bias currents, super-threshold electrical coupling and synapse delays promote synchrony, while sub-threshold electrical coupling and fast synapses promote asynchrony.

We compare our analytical results with simulations of an ionic current model of Hudgkin-Huxley type neurons, and briefly discuss implications of our analysis for stimulus response modes of locus coeruleus and for central pattern generators.

114. Computational consequences of lamina-specific structure in cortical microcircuit models

Stefan Haeusler, Wolfgang Maass

Institute for Theoretical Computer Science, Graz University of Technology, Austria

The neocortex is composed of neurons on different laminae that form precisely structured microcircuits. We investigate whether these stereotypical microcircuits, in particular their lamina-specific synaptic connection patterns, are distinguished by specific computational properties, which enable them to subserve the computational and cognitive capabilities of the brain in a more efficient way.

For this purpose we analyzed the information processing capabilities of detailed cortical microcircuit models with synaptic connectivity pattern between laminae chosen according to the data from [Thomson et al., 2002] that consist of Hodgkin-Huxley neurons with dynamic synapses, based on detailed data from [Markram et al., 1998] and [Gupta et al., 2000]. In addition neurons were subject to synaptic background noise based on data which are conjectured to be representative for the 'high conductance state' of cortical circuits in vivo. We studied to what extent this cortical microcircuit template supports the accumulation and fusion of information contained in generic spike inputs into layer 4 and layers 2/3, and how well it makes this information accessible to the generic "neural users", i.e. , to projection neurons in layers 2/3 and layer 5.

Our computer simulations demonstrated that these cortical microcircuit models exhibit specific computational advantages over various types of control circuits that have the same components and the same global statistics of neurons and synaptic connections, but are missing the lamina-specific structure of real cortical microcircuits. We showed that the connectivity graphs defined by this cortical microcircuit template have small-world properties, but it turned out that other properties of the connectivity graphs are more salient for the computational performance. We demonstrated that these properties of the connectivity graphs support noise suppression and thereby enhance the computational capabilities of the cortical microcircuit model.

References

[Gupta et al., 2000] Gupta, A., Wang, Y., and Markram, H., (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. Science. (article), 287(5451):273-8.

[Markram et al., 1998] Markram, H., Wang, Y., and Tsodyks, M., (1998) Differential signaling via the same axon of neocortical pyramidal neurons. Proceedings of the National Academy of Sciences of the USA 95, 5323-5328.

[Thomson et al., 2002] Thomson, A. M., West, D. C., Wang, Y., and Bannister, A. P. (2002) Synaptic connections and small circuits involving excitatory and inhibitory neurones in layers 2 to 5 of adult rat and cat neocortex: triple intracellular recordings and biocytin-labelling in vitro. Cereb. Cortex 12:936–953.

115. Pulse Packet Interaction in Associative Synfire Chain

Kazuya Ishibashi[1], Kosuke Hamaguchi[2], Masato Okada[3]

[1]Univ. of Tokyo / JST PRESTO, [2]RIKEN BSI, [3]Univ. of Tokyo / JST PRESTO / RIKEN BSI

A synfire chain is one of the networks which have stable synchronous pulse packet mode. The synfire chain which has single stable pulse packet mode has been intensively analyzed by using several neuron models. However, the synfire chains with multi-stability in pulse packet mode have not yet been investigated so thoroughly.

Here we analyze the activity of a layered associative feedforward network, which has multiple pulse packets mode corresponding to the memorized patterns. We refer to this network as an associative synfire chain. We construct the associative synfire chain by using leaky integrate-and-fire neuron model and analyze it with the Fokker-Planck equation.

First, we examine whether a memorized pattern can propagate as a pulse packet through the network. Under a certain parameter region, the layered associative network is shown to have stable pulse packet mode.

Second, we investigate the associative synfire chain with two embedded patterns to focus on the interaction of two memorized patterns. The two characteristic phenomena of the associative synfire chain are the pulse packet in mixed states and independent propagation of the two lattice patterns. Simultaneous activation of two memorized patterns with equal ratio evokes mixed state propagation, but simultaneous activation with different ratio finally develops into two temporally separated pulse packets. To understand this phenomenon, we divide neuron into four lattices by their interaction from pre-synaptic layer. (Given N binary patterns are embedded into the interaction, there are 2^N lattices in the neuron group). From this analysis, it is shown that independent neuron groups are not belong to different memorized patterns, but belong to lattices.

Finally, we activate two memorized patterns in different timing. If time difference of the activation is enough short, the second pattern is not propagated as a synchronous pulse packet, but separated into two pulse packets corresponding to the lattices. These lattice-based analysis gives clear understanding of the associative synfire chain and the dynamics of pulse packets interaction.

Acknowledgements

This work was supported in part by grants from the Japan Society for the Promotion of Science (Nos. 14084212 and 16500093).

116. A theory of object recognition:A theory of object recognition: computations and circuits in the feedforward path of the ventral stream in primate visual cortex

Thomas Serre, Minjoon Kouh, Charles Cadieu, Ulf Knoblich, Gabriel Kreiman, Tomaso Poggio

MIT

We describe a quantitative theory to account for the computations performed by the feedforward path of the ventral stream of visual cortex and the local circuits implementing them. The theory, which extends the Hubel and Wiesel hierarchical model as well as several other proposals, is qualitatively and quantitatively consistent with (and in many cases predicts) several properties of cells in V1, V2, V4, and IT, as well as results from fMRI and psychophysical experiments. The theory is a significant extension of a model introduced in [Riesenhuber & Poggio, 1999]: it now includes unsupervised learning on natural images to account for the tuning properties of neurons from V4 to IT [Serre & Poggio, Cosyne 2005]. During this learning stage, which can be regarded as an imprinting process, unitslearn from natural images by storing in their synaptic weights the specific pattern of feedforward afferent activity. We conjecture from our simulations that the resulting large number of tuned units constitutes a universal and redundant dictionary of image-features, which show a range of selectivities for image patterns, and invariances for translation and scale, and support the recognition of many different object categories. When tested on real-world natural images, the model performs at least at the level of some of the best computer vision systems on several categorization tasks [Serre & Poggio, Cosyne 2005]. We here describe recent results and show how the theory bridges several levels of understanding from computation and psychophysics through system physiology and anatomy to biophysics. The theory seems sufficiently comprehensive, detailed, and satisfactory to represent an interesting challenge for physiologists and modelers to either disprove its basic features or propose alternative theories of equivalent scope. An extended version of this abstract can be found at: http://cbcl.mit.edu/projects/cbcl/publications/ai-publications/2005/AIM-2005-036.pdf.

Predicting human performance in a rapid categorization task

We show that a model instantiating the theory is capable of performing recognition on datasets of complex images at the level of human observers. The task we use is a rapid, animal vs. non-animal recognition task [Thorpe et al., 1996] paired with a backward mask protocol [Bacon-Mace et al., 2005] that interrupts visual processing and blocks back-projections. The stimulus is presented for 20 ms, followed by a variable time interval “ISI” before the mask (1/f noise) appears for 80 ms. Results indicate that the model best predicts the level of performance of human-observers around 30 ms ISI (equivalent to 50 ms SOA). The implication is that, for this ISI, the back-projections do not play a significant role and the model provides a satisfactory description of the feedforward path. Additionally, we show that when the stimuli are rearranged into subcategories (reflecting the amount of clutter in the image and thus the level of difficulty), the model successfully predicts human performance on individual subcategories. Finally, we show that image rotation (90 and 180 deg) influences the performance of both the model and human-observers in a similar manor.

Explaining the neural tuning at the level of V4 and IT

We tested the hypothesis that units in intermediate stages of cortex are tuned to image-features that occur frequently by examining a population of intermediate layer model units that learn from natural images. We show that model units are compatible with V4 data on different stimulus datasets including Cartesian and non-Cartesian gratings [Gallant et al., 1996], boundary conformations [Pasupathy & Connor, 2001], and two-bar presentations (in the absence of attention) [Reynolds et al., 1999]. Additionally, we quantify the ability of a population of units from different stages of the model to decode information about complex objects and find good agreement with the observations obtained upon recording from populations of IT neurons. One feature of the model is its ability to quantitatively explore novel questions where experiments may be difficult or time-consuming by making quantitative predictions to guide experimental efforts. We therefore use the model to explore several new scenarios including, invariance to background changes for object recognition, invariance to the presence of multiple objects, and extrapolation to large numbers of objects and categories. We show that the model can account for object recognition under natural situations involving different backgrounds and object clutter, and that decoding is not limited to small stimulus sets.

Biophysical circuits implementing the key operations

The theory distinguishes between two main functional types of units, “simple” and “complex”, which represent the result of two key operations for invariant recognition. Simple units compute a normalized dot-product for the Gaussian-like tuning with generalization capabilities. Complex units perform a soft-max operation for the tolerance to position, scale and clutter. To support the biological plausibility of the theory, we show that these operations can be implemented by several different circuits based on known or plausible properties of neurons and synapses, involving non-spiking or spiking neurons. For non-spiking neurons, we investigate feedforward and feedback variants of these neural circuits based on shunting inhibition. In the case of spiking neurons, we assume that the information is transmitted by a set of axons, e.g., a set of redundant "wires", each wire delivering binary information over time bins on the order of 10 ms.

Summary

We describe a theory of the feedforward processing in ventral visual cortex that is compatible with and can predict experimental observations that span scales from single neuron biophysics, through neuron populations, and on up to human behavior.

117. Analysis of Cortical Microcircuits on the Systems Level

Robert Legenstein, Wolfgang Maass

Technische Universitaet Graz

The computational properties and function of cortical microcircuits are still unclear. We present in this work two approaches to analyze computational properties of randomly connected microcircuits of spiking neurons. Our analysis is based on the main assumption that the microcircuit should boost the performance of readout neurons which integrate information provided by the microcircuit about online input streams.

One approach is based on results on cellular automata and random Boolean networks, where some evidence was provided that their computational power is largest at the edge of chaos, i.e., at the transition boundary between ordered and chaotic dynamics. In the second approach, we estimate directly the properties of the neural microcircuit in terms of the diversity of computable functions and the ability of the circuit to generalize over noise.

Our results are:

1. We show that an edge of chaos analysis of computations in spiking neural circuits with dynamic synapses is feasible. We find that the edge of chaos predicts quite well those values of circuit parameters (e.g. connection strength and density) that yield maximal computational power.
2. We present theoretically founded measures to directly estimate the computational properties of microcircuits.
3. These measures enable us to predict performance of microcircuits not only for single functions, but for whole families of computations e.g., computations on spike trains or computations on firing rates.
4. We validate the proposed measures by comparing their predictions with direct evaluations of the computational performance of various quite realistic neural microcircuit models consisting of leaky integrate-and-fire neurons and dynamic synapses with parameters distributed according to emirically found data.
5. We find that best computational performance is achieved for circuits of relatively low average population firing rate (lower than 10 Hz), which points to the advantage of sparse neural coding.

The interested reader can find more detailed information in *[Legenstein R. and Maass W., Edge of chaos and prediction of computational power for neural microcircuit models. submitted for publication]*, available at http://www.igi.tugraz.at/legi/psfiles/166.pdf.

118. One-shot learning of behavioral sequences through hippocampal phase precession: A functional hypothesis on short-term synaptic plasticity

Christian Leibold[1], Kay Thurley[1], Anja Gundlfinger[2], Robert Schmidt[1], Dietmar Schmitz[2], Richard Kempter[1]

[1]Institute for Theoretical Biology HU Berlin, [2]NSRC Charite Berlin

Place cells in the CA region of the rat hippocampus typically fire a burst of spikes that exhibit a precession of their firing phases relative to field potential theta oscillations (4-12 Hz), that is, the theta phase of action potentials progressively decreases toward earlier phases. Here we show how synaptic facilitation of the hippocampal mossy fiber synapse in combination with membrane potential oscillations of pyramidal cells in CA3 can explain phase precession. This biologically plausible model reproduces experimentally observed features of phase precession such as the progressive decrease of spike phases, the nonlinear and often also bimodal relation between spike phases and the animal's place, the range of phase precession being smaller than one theta cycle, and the dependence of phase jitter on the animal's location within the place field. Phase precession of hippocampal place cells compresses behavioral sequences on a time scale of seconds to sequences of action potentials on the time scale of milliseconds. In so doing it provides a relational spike code that enables one-shot learning of associations between activity patterns in a recurrent network. One-shot online learning is achieved by means of a spike-timing dependent learning rule. We show that the synaptic weights robustly reach configurations that provide reasonably high storage capacities.

119. Mean Field Theory with Cross-Correlations for a Cortical Network Model

Alexander Lerchner[1], John Hertz[2]

[1]Laboratory of Neuropsychology, NIMH/NIH/DHHS, [2]NORDITA

Theoretical considerations, network simulations, and experimental findings all suggest that correlated neuronal activity is present in cortical networks. What network properties may underlie such correlations, and what parameters control their magnitude? Is a precise temporal code the only explanation for activity that appears to be frequently synchronized on the level of a few milliseconds?

Until now, there has been no theory for calculating cross-correlations in neural network models systematically, yet they must be present whenever two neurons share common presynaptic neurons. Here, we show for a balanced cortical network model, how our previous mean-field approach (Hertz et al., 2003, Lerchner et al., 2004 and 2006) can be extended to incorporate cross-correlations self-consistently, thereby providing a framework for investigating such questions in a quantitative way.

The model network we consider, which may be thought of describing a generic cortical column, consists of two populations of neurons, an excitatory and an inhibitory one, driven by an external excitatory input. The neurons are randomly interconnected, with on average $K>>1$ inputs per neuron. Here, we consider a situation where any two neurons share a specific fraction of their inputs. We scale the strength of the synapses as $1/\sqrt{K}$, which leads to a state of irregular firing driven by fluctuations in the neuronal input, as described previously and confirmed by network simulations. Using mean-field theory, we derive a statistical formulation of the input to pairs of neurons that includes as order parameters the firing rates, the average neuron-to-neuron rate variability, firing auto-correlations and firing cross-correlations. In our model, cross-correlations are of three different types: excitatory-excitatory, inhibitory-inhibitory, and excitatory-inhibitory. Within this framework, we also present an extension of a previously published algorithm for finding exact solutions to the mean-field equations numerically.

Relative to the autocorrelations, cross-correlations contribute with a factor K to the fluctuations in the neuronal input. This leads to two possible scenarios:

(1) The cross-correlations are extremely small (of order 1/K, to be precise). Then we show that they are in fact negative and reduce the amount of fluctuations in the membrane potential by a factor of $(1-K/N)$. We have observed this scenario in simulations of an all-inhibitory model network.

(2) The cross-correlations are of the same order as the auto-correlations. Then, the three kinds of cross-correlations must nearly cancel each other. Given the dynamic balance in the network, the spiking of a specific neuron would then be driven mostly by approximately-synchronized volleys of incoming spikes, the timing of which is nonetheless stochastic. We speculate that such a scenario may explain a reported apparent synchronization of spikes of presynaptic neurons in the neocortex.

120. Robust Propagation of Bursts in Noisy Heterogeneous Synfire Chains

Meng-Ru Li, Henry Greenside

Duke University

Recent experiments on the songbird nucleus HVC (high vocal center) show the presence of sparse brief high-frequency bursts that are believed to play an important role in sequential timing of syllables. By studying numerically a noisy heterogeneous synfire chain based on a single-compartment Hodgkin-Huxley neuronal model with five conductances, we show that bursts similar to those observed experimentally, but not individual spikes, can propagate robustly. The noise was additive Gaussian current to the voltage evolution equation with standard deviation □. The synfire chain was made heterogeneous by connecting neurons of the i-th pool and the (i+1)-th pool with probability P. We then found that the propagation of single spikes was not stable for additive Gaussian currents with □ greater than 5 pA, or for heterogeneities with $P < 0.9$. We interpret these results with a mean-field analysis, in which we average over neurons in a given synfire pool.

121. Do within modality and cross-modality sensory integration follow the same rules?

Ulrik R Beierholm[1], Steven R Quartz[1], Ladan Shams[2]

[1]Caltech, [2]UCLA - Dept. of Psych.

The human perceptual system simultaneously processes multiple streams of sensory information for most perceptual functions. Many of these simultaneous streams of sensory signals carry information about the same object/event, and this presents the nervous system with a problem of how to determine which information to combine and how to combine the information from different streams. If the signals correspond to the same source it is advantageous to integrate them, whereas if they are merely coincidental, integration would be counterproductive and the information should therefore be kept segregated [1, 2].

Here we ask whether there is a difference in how the nervous system approaches this problem when the simultaneous signals are from different modalities versus when they are within the same modality. The maximum likelihood estimation method has been previously tested in both crossmodal integration [e.g. 3, 4] as well as in visual cue combination [e.g. 5, 6] and has successfully accounted for both cross-modal and within-modality integration when the discrepancy between signals is not large and they completely get integrated. This suggests that the same rule may be at work for crossmodal and within-modality cue combination in general. On the other hand, the neurophysiological findings in superior colliculus of cat suggest that crossmodal integration is supra-additive while visual integration is not. Here, we compare auditory-visual integration-segregation with visual-visual integration-segregation within the same paradigm and task. We have previously shown that a Bayesian ideal observer that does not assume forced fusion predicts the human performance in the auditory-visual task remarkably well. We examined a visual-visual version of the same task to see whether the same model can account for the within-modality interactions.

In the auditory-visual experiment, 0 to 4 flashes were paired with 0 to 4 beeps, and observers' task was to judge the number of flashes as well as beeps in each trial. The visual stimulus was a uniform white disk presented on a black background for 10 ms at 12 degrees in the periphery below the fixation point. SOA of flashes was 70 ms. Auditory stimuli consisted of brief tones. The center of the beep trains and flash trains were aligned in all conditions. In the present experiment we replaced the beeps with flashes presented symmetrically above the fixation point. Observers' task was to judge the number of flashes in the lower visual field as well as the number of flashes in the upper visual field.

Given the subjects responses on the trials with only visual presentation in either upper or lower field, we could build the likelihood functions and we could estimate the joint priors by marginalizing the responses in all conditions from a different group of subjects. Combining the priors with the likelihoods allowed us to compute posteriors predicting the percepts of the observers in the visual-visual conditions, which we then compared with the data from these conditions. This approach is similar to the method used in refs. 1 and 2, and allows us to predict subjects' responses without any parameter tuning or fitting.

Subjects' data revealed strong interactions between the upper and lower visual stimuli. However, observers' responses deviated significantly from those predicted by the Bayesian ideal observer. This inconsistency is in contrast to the strong consistency we had obtained previously for the auditory-visual counterpart.

Further analysis shows that although the ideal observer model performs well in predicting the interaction for small discrepancies between the upper and lower visual stimuli, it severely underestimates the amount of interaction for the larger discrepancies. This suggests that the integration/segregation scheme that the nervous system uses for combining information within a modality is different from that utilized when processing information from different modalities. One possible explanation for this difference is that the noise processes that corrupt the signals in different modalities are likely to be uncorrelated justifying the assumption of conditional independence of sensory signals in different modalities (an assumption made by our Bayesian ideal observer). This assumption, however, may not hold for signals within the same modality. The noise processes that corrupt visual signals can be correlated, rendering these sensory signals statistically dependent.

References:

[1] Shams, L., Ma, W.J., Beierholm, U. Sound-induced Flash illusion as an optimal percept Neuroreport 2005; 16: 1923-1927
[2] Beierholm, U. R., Quartz, S.R., Shams, L. Bayesian inference as a unifying model of auditory-visual integration and segregation. SFN 2005 283.13.
[3] Ghahramani Z, Wolpert DM, Jordan MI. Computational models of sensorimotor integration. In: Morasso PG, Sanguineti V, editors. Selforganization, computational maps, and motor control. Amsterdam: North-Holland, Elsevier Press; 1997. pp. 117–147.
[4] Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. Nature 2002; 415:429–433
[5] Landy MS, Maloney LT, Johnston EB, Young M. Measurement and modeling of depth cue combination: in defense of weak fusion. Vision Res 1995; 35:389–412.
[6] Yuille AL, Bulthoff HH. Bayesian deicision theory and psychophysics. In: Knill DC, Richards W, editors. Perception as Bayesian inference. Cambridge: Cambridge University Press; 1996. pp. 123–161.

122. Spike timing in mechanoreceptive afferent fibers can be predicted using integrate-and-fire mechanisms.

Sliman J Bensmaia[1], Arun P Sripati[2]

[1]Johns Hopkins University, [2]Center for the Neural Basis of Cognition

SA1, RA and PC mechanoreceptive afferent fibers produce stereotyped and highly reproducible responses to mechanical oscillations. It should therefore be possible to construct quantitative models to predict the detailed timing of spike trains evoked by a vibratory stimulus. Such models can be powerful tools to understand subsequent cortical processing as well as psychophysical results. Furthermore, such models may help elucidate the mechanisms underlying mechanotransduction in cutaneous mechanoreceptive afferent fibers. For example, is transduction in fibers of a particular type (e.g. SA1) driven by a stimulus property that covaries with position, velocity, acceleration, or a combination of these variables? How is this stimulus property transformed to evoke an intracellular potential in the receptor, e.g. is it rectified or filtered? Can the evoked spikes be predicted using simple integrate-and-fire models, or do more realistic models need to be developed to achieve such predictions?

To address these issues, we recorded the responses of SA1, RA and PC afferent fibers of anaesthetized macaque monkeys to a wide variety of vibratory stimuli, including: (1) simple sinusoids that span a wide range of frequencies and amplitudes; (2) complex waves consisting of two or three superimposed sinusoids varying in frequency, amplitude, and relative phase; (3) half-wave and full-wave rectified sinusoids varying in amplitude and frequency; (4) band pass noise with varying low- and high-frequency cut-offs; and (5) ramp-and-hold stimuli varying in amplitude and ramp velocity. We then constructed an integrate-and-fire model to predict the timing of individual spikes evoked by these stimuli. Using the model, we evaluated hypotheses about four candidate quantities that might drive transduction (position, velocity acceleration, and jerk) by assessing the extent to which afferent responses can be predicted on the basis of these quantities.

The model accurately predicted individual spikes evoked by dynamic stimuli, especially in SA1 and RA fibers, which suggests that the model is a good first approximation of the chain of events leading to the neural response, and that the spiking behavior of afferent fibers can be accounted for by simple linear integrate-and-fire mechanisms. Because spikes evoked in PC fibers are less reproducible, the model was unable to predict the occurrence of individual spikes evoked by dynamic stimuli. Biophysical properties were found to vary across fiber types. For example, SA1 fibers had the lowest frequency cut-off while PC fibers had the highest frequency cut-off (as can be deduced from their threshold-frequency functions). Overall, our results suggest that, to a first approximation, SA1 fibers transduce stimulus position, RA fibers transduce stimulus velocity, PC fibers transduce stimulus acceleration. We have previously developed a model that predicts afferent responses to arbitrary spatial patterns statically indented into the skin. Furthermore, recent studies in our laboratory suggest that spatial and temporal aspects of a tactile stimulus have independent influences on SA1 and RA responses (Bensmaia SJ, Craig JC, Yoshioka T, Johnson KO. SA1 and RA afferent responses to static and vibrating gratings. J Neurophysiol). Therefore, the two models of mechanotransduction – one spatial, the other temporal – can be combined to characterize the peripheral neural representation of arbitrary tactile stimuli.

123. Brain-Inspired Neural Model of Visual Attention for Multiple Object Tracking

Roman Borisyuk[1], Yakov Kazanovich[2]

[1]University of Plymouth, United Kingdom, [2]Institute Mathematical Problems in Biology, Russian Academy of Sciences

An important experimental paradigm in the study of object-based attention is Multiple Object Tracking (MOT). In the canonical MOT experiments (Pylyshyn & Storm, 1988, Pylyshyn, 2001; Scholl, 2001), an observer views a display with simple identical objects. A certain number of the objects (*m*) are briefly flashed to mark them as **targets** . Other objects are used as **distracters** . All objects move independently and unpredictably about the screen. The subjects' task is to track **targets** with their eyes fixed at the centre of the screen. At various times during animation one of the objects is flashed and the observer should indicate whether this object is a target or a distracter. Though the number of errors increases with an increasing number of targets, even for *m* =5 targets the performance level was about 85% correct target identifications.

We present a neural attention model of MOT based on the principles of oscillation synchronization and resonance. As far as we know, this is the first attention model of MOT and the first example of oscillatory neural network application to processing scenes with moving objects (Kazanovich & Borisyuk, 2006, Neural Computation, in press).

The model design is based on our earlier published attention model with a central oscillator (AMCO) (Kazanovich & Borisyuk, 1999; 2002; 2003; Borisyuk & Kazanovich, 2003; 2004). Each element of AMCO is an oscillator whose dynamics are described by three variables: the phase of oscillations, the amplitude of oscillations, and the natural frequency of the oscillator. The interaction between oscillators is implemented in terms of phase-locking, resonance, and adaptation of the natural frequency. AMCO has a star-like architecture of connections. It contains a one-layer network of locally coupled peripheral oscillators (POs) representing the columns of cortical neurons in the primary visual cortex (areas V1-V2) that respond to specific local features of the image. Dynamics of these oscillators are controlled through global feedforward and feedback connections with a central oscillator (CO) which plays the role of the central executive of the attention system (Baddeley, 1996; Cowan, 1988). Isolated objects are represented by synchronous assemblies of POs and the focus of attention is formed by those POs which work synchronously with the CO.

The main idea of tracking m targets is to use a network that consists of m copies of interactive AMCO with each copy tracking one particular target. When implementing this idea the following problems have to be solved. Firstly, one should prevent the situation where the same target is simultaneously taken in the attention focus by several AMCO. Secondly, the model should be able to operate in the case when objects intersect during their movements. These problems were solved by finding a proper interplay of synchronizing and desynchronizing interactions in a multilayer network where each layer is responsible for tracking a single target.

The results of the model simulation are presented and compared with experimental data both in the case of non-overlapping and temporally overlapping objects. In agreement with experimental evidence, simulations with a lager number of targets have shown higher error rates. It is interesting to note that there is the similarity between structure of errors in model simulation and experiment. Studying the model has allowed to explain this phenomena as well as to suggest a new theoretical idea on the number of tracking objects. Also, the model simulations develop new hypotheses for testing by future experiments.

124. Statistics of syllable patterns in produced songs predict auditory responses in hvc of bengalese finches

Kristofer E Bouchard, Michael S Brainad

UCSF

The songbird nucleus HVC is selectively responsive to the playback of complex auditory stimuli, particularly the sound of the bird's own song. Previous studies that focused on the zebra finch have shown that HVC responds best to the single stereotyped song that the bird produces. Bengalese Finches (BFs) produce songs that have greater variability in the sequencing of syllables, making them an attractive system to study how the statistical structure of song correlates with auditory responses to patterns of syllables.

To probe this question we used synthesized songs that were composed of randomly ordered syllables from the bird's repertoire. We recorded single and multi-unit activity in HVC of sleeping adult BFs in response to strings of up to 1000 syllables. We then computed the response to a given syllable as a function of the context (i.e. the preceding syllables) in which it occurred.

As in previous studies in other species, we found that neurons in BF HVC were context sensitive; response to a given syllable was maximal only when it was randomly preceded by a naturally occurring pattern of syllables. Moreover, individual neurons were sensitive to the identity of syllables that occurred as much as four syllables earlier in a sequence.

Because there is more than one context in which a given syllable of BF song can occur, we were also able to ask whether there is a relationship between the probability of sequence production (conditioned on the last note in the sequence) and the strength of auditory response to the last note in that sequence. We found that within the family of naturally occurring sequences, responses were greatest to those patterns that were produced with the highest probability. These data are consistent with an emerging picture of HVC selectivity that responses to auditory stimuli are strongly shaped by the bird's experience of his own song.

Support Contributed By: Searle and McKnight Foundations

125. The Role of Memory in Guiding Attention

Ran Carmi, Laurent Itti

University of Southern California

What is the time frame in which perceptual memory guides attention? Current estimates range from a few hundred milliseconds to several seconds, minutes, or even days. Here we answer this question during natural vision by revealing the time course of attentional selection. First, we generated MTV-style video clips from continuous clips by using jump cuts to connect semantically unrelated clip segments. We then asked participants to visually explore either continuous or MTV-style clips, tracked their eyes, and extracted rapid gaze shifts as objective behavioral indicators of attentional selections. The utilization of perceptual memory was estimated across viewing conditions and over time by quantifying the agreement between human attentional selections and predictions made by a neurally-grounded computational model. In the critical condition, jump cuts led to sharp declines in the impact of perceptual memory on attentional selection, which then increased monotonically for up to 2.5 seconds. Our study demonstrates that previous accounts of memory utilization in simplified laboratory conditions have repeatedly led to misleading conclusions. We propose novel hypotheses and experiments with hybrid natural-artificial stimuli to further elucidate neurocomputational mechanisms of attentional selection.

126. Optimal Spatial Pooling of Neural Population Responses in the Visual Cortex

Yuzhi Chen, Zhiyong Yang, Wilson S Geisler, Eyal Seidemann

Institute for Neuroscience and Center for Perceptual Systems, University of Texas at Austin, 78712

Single-unit electrophysiological studies in behaving monkeys suggest that sensory perception is mediated by populations of cortical neurons. Despite the prevalence and importance of population coding in the brain, it is still unclear how population responses in sensory areas are read-out, or decoded, by subsequent processing stages, in order to mediate perceptually guided behavior. Optical imaging with voltage sensitive dyes (VSD) provides an ideal tool for measuring neural population responses in vivo at high spatial and temporal resolutions. In this study, we used VSD imaging in behaving monkeys to investigate possible read out mechanisms of neural population responses in the primate primary visual cortex (V1). Monkeys were trained to perform a reaction-time pattern-detection task with a small Gabor patch as the target. The contrast of the Gabor patch was varied to span the monkeys' detection thresholds. To determine the detection sensitivity of different candidate pooling models, optical imaging signals (averaged over an interval of ~160 ms) were combined over space using various pooling rules. For each trial, the pooled signals were compared to an optimally-placed criterion in order to decide whether the target was present or absent. Different spatial pooling models were evaluated in terms of their accuracy in detecting the target from the imaging data.

We find that (i) neural responses to a small Gabor patch (σ=0.33deg, eccentricity 3-4 deg) form a wide-spread activation pattern in V1 that is well fitted by a 2-D Gaussian (average σ=2.53mm), and (ii) V1 population responses at nearby sites are strongly correlated, with correlation falling off exponentially as a function of distance (average t=3.01mm). Computational analysis revealed that such high correlations at the level of neural populations are expected even if the underlying correlations between individual neurons are very weak. Because of these strong spatial correlations, most models that pool V1 signals over a large region (e.g., average rule) perform worse than a model that only considers a small (~0.25x0.25mm) but highly informative patch of cortex (best location rule). The best decoding of V1 population responses, however, was achieved with a simple pooling model having antagonistic center-surround weights (center-surround rule), because these weights reduce the effect of the large-spread spatial correlations. Surprisingly, the model that uses the center-surround rule consistently outperforms the monkeys in the detection task. These results demonstrate that (i) VSD imaging in behaving monkeys is an exceedingly sensitive technique, (ii) center-surround pooling is an efficient way of removing spatial correlations in neural population responses, and (iii) V1 responses are used sub-optimally by the monkeys. The sub-optimal performance of the monkey could be due to non-optimal spatial pooling and/or other inefficiencies at, or downstream to, V1.

Overall, our results put important constrains on the way neural population responses may be read out in the brain, and demonstrate the potential of VSD imaging in behaving monkeys for addressing fundamental questions regarding the link between neural population responses and perception. (Supported by NIH/NEI R01EY016454-01 and Sloan Fellowship)

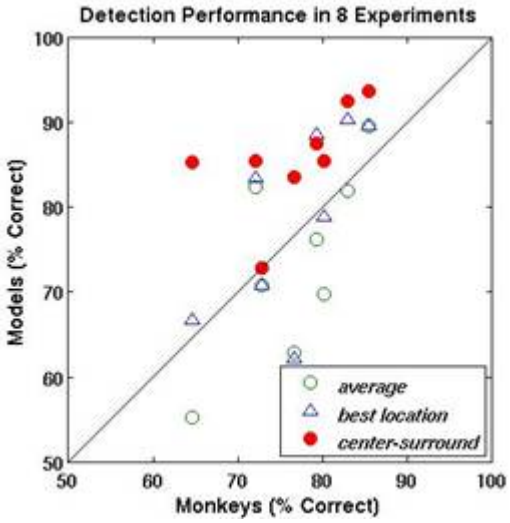

127. Characterizing contrast adaptation in a population of cat primary visual cortical neurons using Fisher information

Colin WG Clifford[1], Szonya Durant[1], Nathan A Crowder[2], Nicholas SC Price[2], Michael R Ibbotson[2]

[1]University of Sydney, [2]Australian National University

Single cell studies of contrast adaptation have shown that many individual cells change their contrast gain, their response gain, or both, after adaptation (e.g. Crowder et al., 2006). The shift in maximal contrast gain at the single cell level indicates adaptation plays a functional role, as opposed to being a result of simple fatigue - as one might deduce from a drop in response gain alone. Given that cells exhibit a wide range of behavior after adaptation, we investigated how these factors combine at the population level to affect the accuracy for detecting variations in contrast. Using the contrast response function parameters from a physiologically measured population, we use Fisher information (Dayan & Abbott, 2001) to model the population accuracy for contrast discrimination as a function of baseline contrast before and after adaptation. Our calculations are based on the measured responses of 130 cortical neurons in areas V1 and V2 of the cat visual cortex (Crowder et al., 2006), both of which are regarded as primary visual cortex in that species (Payne & Peters 2002). The effects observed here could in part be due to neural adaptation at earlier stages in the visual pathway (Solomon et al. 2004). However, this possibility in no way affects our conclusions, which are based on the adaptation of response properties of cortical neurons regardless of whether this adaptation is inherited from earlier stages of processing or generated de novo in visual cortex.

The shape of the accuracy curve over contrast of the non-adapted neural population was found to resemble the distribution of contrast levels in natural scenes (Chirimuuta et al., 2003). Adaptation at 16%, 32 % and 100% contrast caused a shift in peak accuracy. Despite an overall drop in firing rate over the whole population, accuracy was either maintained or enhanced around the adapted contrast, leading to greater efficiency of contrast coding in terms of the mean Fisher information conveyed by each spike.

Population accuracy provides a link between physiological and psychophysical contrast adaptation data. We find that the estimated contrast discrimination threshold curve becomes elevated and shifted towards higher contrasts after adaptation, as in previous psychophysical experiments. The change in the response of cells following adaptation alters the population accuracy, which in turn reflects the change in contrast levels in the visual environment.

References

Chirimuuta, M., Clatworthy, P. L., Tolhurst, D. J. (2003) Coding of the contrast in natural images by visual cortex (V1) neurons: a Bayesian approach. J. Opt. Soc. Am. A 20(7):1253-1260

Crowder, N.A., Price, N.S.C., Hietanen, M.A., Dreher, B., Clifford, C.W.G. & Ibbotson, M.R. (2006) “Relationship between contrast adaptation and directional tuning properties in areas 17 and 18 of cat visual cortex: a spectrum of adaptation effects”, Journal of Neurophysiology, 95, 271-283.

Payne B.R. and Peters A. (2002) The cat primary visual cortex. London: Academic Press.

Solomon, S. G., Pearce, J. W., Dhruv, N. T. and Lennie, P. (2004) Profound contrast adaptation early in the visual pathway. Neuron 42:155-162.

128. Multivariate Analysis of Frontal Eye Field Activity during Visual Search

Jeremiah Y Cohen[1], Pierre Pouget[2], Chenchal Rao[2], Jeffrey D Schall[2], Andrew F Rossi[2]

[1]Vanderbilt Brain Institute, [2]Vanderbilt University Department of Psychology

Simultaneous recording from multiple electrodes is becoming more prevalent in traditional experimental designs that relate the activity of neurons to behavior. The multi-electrode approach not only increases the yield in a given experiment, but presents an opportunity to observe the real-time interactions that may occur between simultaneously recorded neurons. Unfortunately most analysis methods that seek to understand the physiological relationship between neurons are limited in their capacity to pairwise comparisons (e.g. cross-correlation, JPSTHs). Because of this limitation, we have explored the use of a multivariate analysis of multiple spike train data to better understand the contribution that one or more neurons has on the firing behavior of another. We have employed a statistical framework based on the maximum likelihood estimate of the point process intensity function which relates a neuron's probability of spiking to covariates such as prior firing history, activity in other neurons, and extrinsic factors such as stimulus condition and behavior [1,2]. Modeling the firing rate of a given neuron using empirically determined covariates (e.g. activity of other neurons) yields an estimate of the relative contribution each covariate makes to the activity of the modeled neuron, from which we can infer the strength of interactions among neurons during different behavioral conditions.

Neural activity was recorded in the frontal eye field (FEF) of macaque monkeys performing a visual search for a singleton target defined by color. The stimulus array consisted of a target that was presented at one of eight iso-eccentric locations equally spaced around the fixation spot. The remaining seven locations were occupied by distracters. Reward was contingent on the monkey shifting its gaze to the oddball target after array onset. There were two levels of task difficulty, 'easy' and 'hard', determined by the degree of target-distracter similarity. Recordings were made simultaneously from four tungsten electrodes placed in the rostral bank of the arcuate sulcus. Typically, the activity of several (1-3) single units could be isolated from each electrode during a recording session. We used a point process approach to modeling the activity of FEF neurons to explore the hypothesis that interactions between simultaneously recorded neurons in FEF change in response to task difficulty.

Following the approach of Truccolo et al. [1], we modeled the activity of each neuron in the recorded ensemble using the activity of the other members of the recorded ensemble as covariates in the point process likelihood function. To obtain the model parameters, we implemented GLM fitting procedures in R (iteratively reweighted least squares algorithm). This algorithm provides a robust maximum likelihood estimate of model parameters [3]. For purposes of model validity, we applied various combinations of covariates and tested goodness-of-fit. That is, we included a set of covariates and tested the model for parsimony using Akaike's information criterion [4]. To assess the goodness-of-fit of a particular model, we employed a time-rescaling theorem [5] to transform the discrete set of events that constitutes a spike train into a continuous measure of rescaled spike times. We then applied the Kolmogorov-Smirnov test to compare the rescaled times (in increasing order) to the uniform distribution. Additionally, we examined the residual of the point process model to determine whether there was structure in the data left unexplained by the model.

When comparing the models for the 'easy' vs 'hard' search task, we found that there were greater interactions among simultaneously recorded neurons for the 'hard' task than for the 'easy'. This was revealed by the need for a fewer number of parameters to accurately model the firing rate in the 'hard' task condition. The ability to obtain a good fit with fewer parameters is indicative of greater degree of interaction among covariates (neurons) by means of more precise timing. These findings suggest a neural process in which the interactions among neurons in FEF are less variable and more structured in time for the 'hard' task when compared to the 'easy' task. We demonstrate that this point process approach to multivariate analysis accurately describes activity of FEF neurons and is well-suited to characterize the interactions between multiple neurons in the guidance of behavior.

1. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN (2005) A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. J Neurophysiol 93: 1074-1089.

2. Daley DJ, Vere-Jones D (2003) An Introduction to the Theory of Point Processes. Volume I: Elementary Theory and Methods. New York: Springer-Verlag.

3. McCullagh P, Nelder JA (1989) Generalized Linear Models (2nd ed.). Boca Raton, FL: Chapman & Hall/CRC.

4. Akaike H (1973) Information theory as an extension of the maximum likelihood principle. In: Second International Symposium on Information Theory, edited by Petrov BN and Csaki F. Budapest: Akademiai Kiado, p. 267-281.

5. Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM (2002) The time-rescaling theorem and its application to neural spike train data analysis. Neural Comput. 14: 325-346.

129. Spectral receptive field properties explain shape selectivity in V4

Stephen V David[1], Benjamin Y Hayden[2], Jack L Gallant[2]

[1]University of Maryland, [2]University of California, Berkeley

Neurons in cortical area V4 respond selectively to complex visual patterns such as curved contours and non-Cartesian gratings. Because most previous experiments in V4 have measured responses to small, idiosyncratic stimulus sets, no single functional model has been developed that generalizes across all stimuli. Thus little is known about the mechanisms by which selectivity in V4 emerges from the simpler representations in early visual areas. We hypothesized that a single model, the spectral receptive field (SRF), can explain several previous observations of selectivity in V4. The SRF describes selectivity in terms of an orientation/spatial frequency tuning surface. This model can predict the response to an arbitrary visual stimulus; thus its scope is not restricted to any limited stimulus set. We estimated SRFs for V4 neurons by linearized reverse correlation of visually-evoked responses to a large set of natural images. We found that V4 neurons have complex tuning profiles: wide orientation and spatial frequency bandwidth and often bimodal orientation tuning. For comparison, we also estimated SRFs for neurons in primary visual cortex (V1), which are known to be selective for oriented bars and Cartesian gratings. Consistent with previous observations, we found that V1 neurons have narrower orientation and spatial frequency bandwidth. To determine whether the estimated SRFs could account for selectivity observed in V1 and V4, we used them to predict responses to Cartesian gratings, non-Cartesian gratings, natural images, and curved contours. As expected, V1 neurons give the largest predicted responses to Cartesian gratings. In contrast, V4 neurons give the largest predicted responses to non-Cartesian gratings and natural images. For both V1 and V4 neurons, selectivity for a particular stimulus class is correlated with tuning properties observed in individual SRFs. These results support the hypothesis that SRFs describe general mechanisms mediating shape selectivity in area V4.

130. Robustness to reverberation of directionally-sensitive neurons in the inferior colliculus

Sasha Devore, Bertrand Delgutte

Harvard-MIT Division of Health Science and Technology

The disparate path lengths from a sound source to a listener's ears give rise to an interaural time difference (ITD), the most important cue for the localization of sounds that contain low-frequency energy e.g., speech. Previous neurophysiological experiments that examined the sensitivity of auditory neurons to ITD have typically used anechoic stimulus conditions, in which the only acoustic inputs to a listener's ears arrive directly from the sound source. This artificial situation is quite unlike that which humans encounter when listening in rooms, where the direct sound is followed by numerous acoustic reflections (echoes) from boundary surfaces. The reflections interfere with the direct wavefront, leading to a decorrelation of the signals reaching a listener's ears and pronounced temporal fluctuations in the ITD. Despite such distortion, normal-hearing listeners are generally able to accurately localize sounds in reverberant environments. A fundamental problem is that of understanding the neural mechanisms that mediate such robust perception. To date, few neurophysiological studies have addressed the effects of reverberation on neural activity, and none has examined the effects of reverberation on directional sensitivity. We investigated the effects of reverberation on the directional sensitivity of low-frequency ITD-sensitive neurons in the inferior colliculus (IC), the principal auditory nucleus in the midbrain.

Many IC neurons that are tuned to low frequencies (< 2.5 kHz) are sensitive to ITD and therefore directionally sensitive. The response of such neurons to interaurally time-delayed signals is well-described by a cross-correlation on the left- and right-ear input signals, after taking into account the frequency tuning performed by the cochlea. A goal of this study was to determine whether a physiologically-inspired cross-correlator model of binaural processing can account for neural responses in reverberation.

We recorded from single units in the central nucleus of the IC of dial-in-urethane anesthetized cats. Stimuli consisted of 400-ms bursts of exactly-reproducible Gaussian noise filtered with either anechoic or reverberant binaural room impulse responses (BRIR). The reverberant BRIRs were simulated for a typical classroom using the room-image method. The virtual spatial stimuli were presented from azimuths spanning the full frontal hemifield (–90° to +90°) at several distances through a closed acoustic system at sound pressure levels approximately 20 dB above neural thresholds.

In general, reverberation led to a demodulation of neural responses as a function of azimuth – peak firing rates were reduced and minimum firing rates were increased as compared to the anechoic condition. Reverberation also led to a decrease in the mutual information between stimulus azimuth and spike count. The degradation in azimuth sensitivity became more pronounced as the distance from the source to the listener increased, and therefore the ratio of direct to reverberant energy decreased.

To determine to what extent these degradations can be explained by interaural cross-correlation we compared our physiological results to the predictions of a cross-correlator model of IC neurons (Hancock and Delgutte, J. Neurosci. 24:7110-7). The model accurately predicted the neural responses to anechoic virtual spatial stimuli but, in general, predicted a worse degradation of directional sensitivity in reverberation than actually observed. This result suggests that low-frequency ITD-sensitive neurons in the IC are to some extent robust to reverberation, in that they perform better than predicted by current models of binaural processing. Future efforts will be directed at identifying the neural mechanisms responsible for this robustness.

Supported by NIH grants DC02258 and DC05209.

131. Extracellular Electrode Detection Range and Sampling Bias for Cat Visual Cortex

Carl Gold[1], Cyrille Girardin[2], Rodney Douglas[2], Christof Koch[1]

[1]Computation and Neural Systems, California Institute of Technology, [2]Institute of Neuroinformatics, Swiss Federal Institute of Technology (ETH)

We simulate the biophysics of intra- and extracellular current flow to calculate the distance at which an extracellular electrode could record neurons from different layers of cat visual cortex. Detailed 3-D reconstructions of neurons from identified layers of cat visual cortex (N=70; Binzegger et. al. 2004) coupled to a detailed electrophysiological membrane model, were simulated using compartmental models within NEURON (Hines and Carnevale 1997). We use the Line Source Approximation (Holt and Koch 1999) to model the extracellular action potential voltage resulting from the spiking activity of individual neurons (Gold et. al., 2005). The peak extracellular voltage resulting from an action potential was calculated in the vicinity of each model cell. A threshold determined by the background noise was used to calculate a region in which an extracellular electrode would be able to detect the action potential. For each layer of cortex we calculated an average detection range for a typical neuron located in that layer. We found that the size of the detectable region around a cell depended strongly on the size and morphology of the cells. For example, the large pyramidal cells from layer 5 were detectable in a region more than ten times larger than that of stellate cells in layer 4.

Based on the average detection volume, we calculated the number of cells that would be within range of a fixed electrode for each layer using previous measurements of the density of cells in each layer (Beaulieu and Colonnier, 1983). The number of cells within detection range varied from only 2.5 in layer 4 to moret than 30 in layer 5. (Layers 2/3 and 6A were more similar to layer 4; see Table 1) The average number of neurons within detection range, weighted by the thickness of each layer, is 8.9. Previous empirical measurements found an average of 4.4 units detectable by tetrodes in vivo (Gray et. al., 1995), suggesting that the average percentage of cells active enough to detected in cats under anesthesia is on the order of 50%. However, the fact that Gray et. al (1995) recorded from a maximum of 7 cells suggests the activity may be substantially lower in layer 5.

We also calculated the sampling bias of each layer, assuming that the fraction of cells active in each layer was uniform. This calculation suggests that the probability of sampling from any layer, given an electrode randomly placed somewhere in the grey matter, is influenced by both the size of the different layers and the detection range around the cells of that layer. For example, layer 5 cells have the largest average detection region, but layer 5 is also the thinnest. Consequently, layer 5 cells would be sampled at more than double their actual frequency but still make up only 20% of the total. Layer 4 cells would be heavily undersampled, but still make up around 22% of the total.

We are currently performing detailed measurements to calculate the fraction of neurons that are active within each layer in vivo: By combining recordings with detailed histology we can accurately estimate the number of cells within detection range of electrodes for each layer. Combining these measurements with the model calculations we will accurately estimate the fraction of cells that are active under different stimulus conditions.

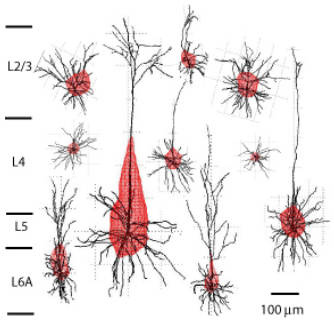

Figure 1: Illustration of detection regions around cells from different layers of Cat visual cortex.

layer	Detect Volume(mm^3)	Detect Radius (µm)	Thickness (µm)	Thick Fraction	Cell Density(/mm^3)	Cells Fraction	Cells in Detect Range	P(layer | detect)
2/3	1.2E-04	29	470	34%	5.7E+04	35%	6.1	34%
4	7.3E-05	22	510	38%	5.4E+04	37%	2.5	22%
5	1.2E-03	59	140	10%	4.3 e+04	8%	36.7	20%
6A	3.0E-04	36	250	18%	6.2 e+04	20%	11.9	24%

Table 1: Results of the computer model for typical simulation parameters. Detect Volume: Average volume of the detectable region around cells from each layer; Detect Radius: Average of the sphere radii equivalent to each cells detectable region volume; Thickness: Thickness of layer in Cat area 17; Thick. Fraction: Percent of total cortical thickness of layer; Cell Density= Cells / mm^3; Cells Fraction: Percent of total cells in layer. Cells in Detect Range: Numbers of cells that would be within detection range of a fixed electrode within each layer; P(layer|detect): Assuming equal firing rates and random sampling, the percent of recordings that would be made in each layer.

132. Hierarchical subunit model for disparity-selective complex cells in V1

Ralf M Haefner, Bruce G Cumming

LSR/NEI/NIH

Disparity-selective complex cells in primary visual cortex (V1) are considered to be instrumental in the processing of stereoscopic depth information. Many properties of these cells have been successfully described by the binocular energy model (Ohzawa et al, 1990), in which the responses of linear monocular receptive fields from each eye are summed, and the result is then squared. Some aspects of the responses of real neurons are not well described by the simplest form of this model, most notably the fact that responses to anticorrelated random dot patterns (aRDS) are weaker than responses to correlated RDS (cRDS). We have recently shown, using compound gratings, (SfN, 2005) that the model cannot be reconciled with the data by applying an additional static nonlinearity at the output. Those experiments manipulated the interocular phase difference of each component grating independently. The energy model sets stringent constraints on the relative location of maximum and minimum response in the space spanned by the two phase-disparities, a property which is not changed by any output nonlinearity. We show here that other models designed to explain the observed responses to aRDS (Read et al 2002) also fail to explain the responses to compound gratings.

We therefore present a new model, able to explain both of these data sets. This hierarchical model consists of two (or more) energy-model like subunits. The subunits differ in phase disparity and/or in position disparity. Importantly each subunit is passed through different output nonlinearities (exponents). A striking feature of the neuronal responses to compound gratings was that those neurons deviating from the energy model (45%) did so in a systematic fashion: In the responses of the energy model to compound gratings, much of the dynamic range is devoted to combinations of phase disparity that do not correspond to a simple displacement of a broadband pattern (i.e. these combinations do not correspond to any real-world disparity). Neurons that deviated from the energy model systematically dedicated more of their dynamic range to realistic combinations than the energy model. This aspect of the behavior is captured by our new model, when the subunits have different position disparities.

In order to challenge this model further, we examined the responses of these neurons to disparity in single sinusoidal luminance gratings, using at least 7 different spatial frequencies. The energy model predicts a linear relationship between frequency and preferred disparity. The slope of this relationship is determined by the positional disparity of the RF, and offset is determined by the phase disparity (Zhu & Quian 1996, Fleet et al 1996). Because the subunits differ in phase/position disparity, the hierarchical model can generate data that deviate from the straight line. We find that 20% of V1 complex cells deviate significantly from this straight line prediction (chi^2>10). Furthermore, these deviations have a characteristic shape which is in qualitative agreement with the responses of the hierarchical model.

Finally, we examine the relationship between these binocular responses to gratings and the shape of the responses to disparity in RDS. In the energy model, the responses to broadband patterns, such as RDS, can be predicted summing the responses to individual gratings. We show that for many of our cells this sum is significantly different from the observed responses to RDS. Three features of these deviations are interesting. First, these deviations are systematically larger for responses to aRDS than for cRDS (in 3/4 of our cells). This suggests that the attenuated response to aRDS is not simply explained by a gain change. Furthermore, larger deviations for aRDS are a feature of our hierarchical model. Second, it has been previously reported that the power in low frequencies is increased for RDS disparity tuning compared with monocular SF tuning (Ohzawa et al 1997, Read and Cumming 2003). Third, neurons classified as unlike the energy model on the basis of responses to compound gratings systematically showed larger deviations than those neurons that were energy model-like. All three of these features follow naturally from the hierarchical model.

Summary: we compare the responses of disparity selective V1 neurons to four different binocular stimuli: cRDS, aRDS, sinusoidal luminance gratings of several frequencies, and compound gratings. Comparing these responses reveals several failures of the energy model in a subset of V1 neurons. Each of these aspects of neuronal behavior can be reproduced by a hierarchical model that combines inputs from energy-model like subunits. We are currently investigating whether this model can give a complete account of all the observations on each cell simultaneously.

133. Sound discrimination in awake head-fixed rats

Tomas Hromadka, Anthony M Zador

Watson School of Biological Sciences, Cold Spring Harbor Laboratory

The contribution of single auditory neurons during tasks involving attention and auditory stream segregation is unknown. Most of what is known about the properties of neurons in auditory cortex has been learned using extracellular recording techniques. However, in vivo intracellular and cell-attached recordings with a patch electrode can yield insight into neural mechanisms not possible through purely extracellular techniques. For example, with cell-attached recordings it is possible to fill a neuron with dye (via electroporation) for later histological recovery, providing an opportunity to correlate structure and function. Recently, we have modified these techniques so that we can routinely apply them in the awake preparation. In order to use the power of cell-attached and whole-cell methods to study these questions, we are adapting well-controlled behavioral paradigms to the head-fixed preparation.

As a first step, we have developed a rapid and reliable method to train head-fixed rats to discriminate a target sound from a distractor. Subjects are rewarded with water for licking in response to a frequency modulated tone (warble), and punished with an air puff for licking in response to a pure tone. Individual trials are separated by a random duration silent period. Incorrect licks, and licks during inter-trial intervals, are penalized by brief air puffs. In addition, each lick during the inter-trial interval restarts the interval.

Training takes 10-14 days, after which subjects typically achieve high performance with low false positive and false negative rates. During training, subjects sit in a narrow plastic tube. Each rat that learns the task is then implanted with a headpost; the left auditory cortex is exposed and re-sealed with a plastic well. After recovery from surgery, subjects are retrained in head-fixed position. In 1-2 sessions subjects typically reach their previous behavioral performance and are ready for recording. The cell-attached recordings are stable and robust to both licking and occasional body movements.

This simple behavior demonstrates that high-resolution recording methods can be applied even in the awake head-fixed configuration. We expect that this approach can be extended to even more complex behaviors, and will allow unprecedented insight into the cortical mechanisms underlying attention.

134. A model for stimulus competition and selective visual attention in area V4

Etienne Hugues[1], Scott A Hill[2], Paul H Tiesinga[3], Jorge V José[1]

[1]SUNY at Buffalo, [2]CIRCS, Northeastern University, [3]University of North Carolina at Chapel Hill

Selective visual attention is central to human behavior, and understanding its neural basis is a fundamental problem in cognitive neuroscience. When a monkey selectively focuses attention on a visual stimulus within the receptive field of a V4 pyramidal neuron, there is an increase in the neuron's firing rate (McAdams and Maunsell, J. Neurosci. 19:431, 1999) and an increase of the local field potential (LFP) power in the gamma frequency range. There may also be an increase in coherence between the spiking and the LFP in this frequency range (Fries et al., Science 291:1560, 2001). When two stimuli are presented in the receptive field of the V4 neuron, the neuron's firing rate is intermediate between the ones observed for both stimuli alone, i.e. there is competition between them (Reynolds et al., J. Neurosci. 19:1736, 1999). Furthermore, varying the strength of one of the two stimuli, a linear dependency is found between these firing rates. When the monkey focuses its attention towards one of the stimuli, the firing rate gets closer to its value observed when this stimulus is presented alone. A computational and a theoretical approach were used here to investigate the neuronal mechanisms underlying competition and attention.

We have built a physiological model to simulate the above experiments. We consider V4 readout neurons activated by two different columns in area V2, these columns receiving independent visual stimuli. As projections onto V4 pyramidal neurons are still not fully characterized, we have considered several connectivity patterns. We have used realistic Hodgkin-Huxley type models for pyramidal neurons and interneurons and conductance-based synapses. We have found competition and a linear dependency in the firing rates only when interneurons project back to the V4 readout neurons. Then, we tested a preceding hypothesis that attention could be mediated by the increased synchrony of the interneuron population. Although this mechanism could explain the firing rate modulation for one stimulus, we found that it could not explain the experimental results for two stimuli.

To go further, we then investigated theoretically the behavior of a V4 readout neuron. As in the model, this neuron receives many synaptic inputs and is in a high-conductance state. Its firing is driven by fluctuations in the synaptic conductances. We consider that it receives stochastic excitatory and inhibitory inputs and are interested to describe the case of asynchronous and oscillatory inputs. Using the Fokker-Planck formalism and numerical simulations, we characterize the variations of its firing rate with respect to the statistical properties of the excitatory and inhibitory inputs. We first apply this analysis to unravel the competition mechanisms found in the model. From these results, we can also infer attentional mechanisms, that is how inputs to a V4 neuron should change to reproduce the experimental observations. Finally, we test the proposed mechanisms using the computational model.

135. Functional topology of attention in the pulvinar

Oliver Hulme[1], Justin Chumbley[1], Simon B Eickhoff[2], Simon Prince[1], Stuart Shipp[1]

[1]University College London, [2]Research Center Julich

From an operational perspective attention is about co-ordinating multiple brain areas to work towards a common objective. In the functional imaging literature much effort has been directed to the role of the cortex in attention, whereas the role of subcortical regions has been relatively neglected. This is mainly due to the popularity of fMRI analysis methods (e.g. spatial smoothing and group studies) that can preclude the discovery of small subcortical activations. Studies, however, that investigated subcortical correlates of attention have shown that in particular the pulvinar nucleus of the thalamus is involved in spatial shifts of attention (Kastner et al. 2004, Haynes et al 2005).

Recent anatomical findings have prompted the speculation that the pulvinar might act as a remote hub for coordinating spatial activity within multiple cortical visual maps (Shipp 2004). Central to this speculation is the topographic organisation of the pulvinar, which contains several retinotopic maps, each selectively and coherently projecting to different cortical visual areas (e.g.V1,V2,V3 etc). The pulvinar is organised such that there are axes of isorepresentation along which the different retinotopic maps are registered (analogous to the writing along a stick of candy). The registration is such that by travelling along these axes the cells encountered project to different cortical visual areas and yet are selective for the same part of visual space. Intuitively, this topographic feature affords a means by which the brain could spatially co-ordinate the activity of multiple cortical visual maps by selective amplification along the axis of isorepresentation of an attended retinotopic location within the pulvinar. If shown to exist, this amplification would provide evidence for a 'spotlight of attention' and could provide a framework with which to understand the role of the pulvinar in solving the classic 'binding' problem.

The methodological problems with studying the pulvinar at this fine scale are compounded by its small size (~10mm in diameter). Consequently no attempt has been made to map its attentionotopic topology with any level of satisfactory fidelity. Here we report a pilot mapping study using fMRI aimed at deriving an attentionotopic map of the pulvinar response to covert attentional shifts in all possible directions (360 °). To overcome the limitations associated with the size of the pulvinar we acquired significantly more data than is usual (5000+ acquisitions per subject) and no spatial smoothing was performed (data obtained from Hulme and Zeki 2006, unpublished).

Using an event related design subjects were cued to attend to one of 16 different positions (at 3° eccentricity) around a clock face, reporting the presence or absence of disc-shaped stimuli at those locations. Using standard statistical parametric mapping we estimated the BOLD response of each voxel to each of the 16 attentional directions, yielding an attentional response profiles for every voxel in the brain. Identifying the pulvinar on individual anatomical criteria we extracted a 125-voxel-cube region of interest for each hemisphere. The main effects of attention for each of the 16 directions show that there is direction selectivity in the pulvinar with distinctly different regions preferring different directions. Although in some subjects a whole field representation was found in the left dorsal pulvinar, visualising the precise 3D topography with this method is difficult.

We have therefore piloted two new analytical approaches to this kind of fMRI data:

1. Firstly, we used a mixture of factor analysers to classify/cluster voxels on the basis of their functional homology. Identifying all members of a cluster with a unique colour, the data was more easily visualised in 3-dimensional voxel space. This is a more sensitive approach than the standard main effect method, since it takes into account the multivariate nature of each voxel's response profile.

2. Secondly, we directly addressed the question of whether there are axes of isorepresentation along which the attentional response profiles are maximally similar. To archive this goal, we first computed for each voxel the amount of change in attentional profiles along each principle direction of 3D space. This allowed us to calculate a tensor field for each voxel giving the attentionotopic isotropy for each voxel in 3D space. This first eigenvector of this tensor field then indicates the direction in which the most similar multivariate response profiles are encountered. Plotting these eigenvectors finally affords the visualisation necessary to evaluate the existence of the axes of isorepresentation.

Using such methods we provide a tentative first account of the 3D functional topology of the human pulvinar. Anisotropic tensor field mapping of functional data is a new application and its potential applicability to other questions in functional neuroimaging will also be discussed.

This work was supported by the Medical Research Council and the Wellcome Trust.

136. Feedback-mediated facilitation from the "far" receptive field surround of macaque V1 neurons

Jennifer M Ichida[1], Lars Schwabe[2], Paul C Bressloff[3], Alessandra Angelucci[1]

[1]Moran Eye Center, University of Utah, Salt Lake City, UT, USA, [2]Electrical Enginering & Computer Science, TU Berlin, Germany, [3]Department of Mathematics, University of Utah, Salt Lake City, UT, USA

Cells in V1 are tuned to stimuli of optimal size within their receptive field (RF), and are usually suppressed by large stimuli involving the RF surround. Optimal stimulus size at low contrast (lsRF) is larger than at high contrast (hsRF). We previously showed that V1 horizontal connections in the near surround extend as far as the lsRF, and suggested that extrastriate feedback (FB) connections to V1 underlie suppression from the far surround (Angelucci et al. 2002). We constructed a recurrent network model (Schwabe et al. 2005) in which FB connections target excitatory cells in the RF center and the near surround (see Johnson and Burkhalter 1996), with the latter sending horizontal connections to high threshold inhibitory cells in the RF center. The model predicts that stimulation of the far surround can facilitate or suppress the RF center, depending on the amount of excitatory input driving the local inhibitor to threshold. We tested this prediction with single unit recordings from anesthetized macaque V1. Visual stimuli were designed to reduce feedforward stimulation of horizontal connections and unmask the effects of FB connections. A central grating the size of the cell's hsRF was surrounded by an annular grating with fixed outer diameter (28 deg), and an inner diameter that was decreased from 28 deg to just outside the lsRF. As the inner diameter of a high contrast surround annulus decreased, the center's response to a low contrast grating was first facilitated, then suppressed, while only suppression occurred for high contrast center gratings. Facilitation from the far surround was seen in ~ 50% of cells, and peaked at smaller inner diameters when stimulus contrast in the surround was also lowered. Suppression from the far (FB-mediated) surround was weaker than from the near (horizontal-mediated) surround. Facilitation from the far surround can be explained by a modified version of a difference of Gaussian model of center-surround interactions. Supported by: Wellcome Trust, NIH, NSF, RPB.

137. Optimal strategies for active perception

Santiago Jaramillo, Barak A Pearlmutter

National University of Ireland, Maynooth

Perception is an active process in which the body and sensory apparatus are oriented in a task-dependent fashion. We hypothesize that during active perception, sensory resources are allocated to optimize task performance. In the case of visual perception, resources consist of visual resolution in the fovea versus the periphery. Our hypothesis therefore makes predictions concerning gaze control during difficult visual tasks. To test the hypothesis, we construct a challenging task in which we can (1) measure the sensory strategies used by subjects, and (2) calculate the optimal strategies. In particular, we explore the overt strategies used by the visual system when called upon to estimate the properties of several targets simultaneously.

The task consists of a video game in which two independent targets change their binary state according to a Markov process, and the subject must estimate the state of each target. Targets are spread in the visual field in a fashion which precludes their simultaneous identification, and the subject must therefore foveate. The subject must report the state of each target at a random time, at which point the targets disappear indicating the end of the trial. Eye position is measure throughout each trial.

Measured strategies and performance are compared with predictions from a model that performs optimal state estimation under the same dynamics as the experiment. The model is composed of two channels with shared resources, each channel representing one target. The observation noise is inversely related to the resources allocated to that channel. Subject visual acuity is measured, and we assume that through practice the dynamics of the objects are learned by the subjects. The subjects thus have all the information we use to mathematically calculate the optimal visual strategy.

We are able to predict eye-movements under the assumptions of optimal estimation. In general, this model would suggest that resources are allocated according to the uncertainty of stimuli after the priors, dynamics, and current observations have been taken into account, allocating resources preferentially to features that are less certain.

Predicting overt strategies in perception is important for understanding the role of attention and the influence of distractors. We have shown here that optimal resource allocation for maximizing estimation performance is a powerful general principle in this context.

138. Information-based fMRI analysis for predefined regions of interest

Nikolaus Kriegeskorte, Peter Bandettini

NIMH

Several recent studies using functional magnetic resonance imaging (fMRI) have approached the analysis of regional multivoxel response patterns as a pattern-classification problem. The accuracy with which the experimental condition can be determined from the multivoxel activity patterns gives an indication of the information present in the neuronal response patterns in the brain region studied. Here we present methods for visualization and multivariate analysis of regional fMRI activity patterns. We progress in a stepwise fashion from conventional univariate analysis of every single voxel and the regional average to multivariate techniques. The visualizations along the way include parallel-coordinates representations of spatial activity patterns, surface plots of activity-level over a local cortical flatmap and similarity-space icons summarizing all pairwise response-pattern similarites and multivariate significances in a small 2D graph. The analyses include multivariate randomization testing as well as linear discriminant analysis and plug-in mutual information estimates.

139. Two new visual areas in human lateral occipital cortex

Jonas Larsson, David J Heeger

Dept. of Psychology & Center for Neural Science, NYU

We identified two new visual areas, LO1 and LO2, in human lateral occipital cortex between dorsal V3 and V5/MT+. Each area contained a complete map of the contralateral visual hemifield with orderly representations of both eccentricity and polar angle. The eccentricity representations were shared with V1/V2/V3. The polar angle representation in LO1 extended from the lower vertical meridian (at the boundary with dorsal V3) through the horizontal to the upper vertical meridian (at the boundary with LO2). The polar angle representation in LO2 was the mirror-reversal of that in LO1. LO1 and LO2 overlapped with the posterior part of the object-selective lateral occipital complex (pLOC) and with the kinetic occipital region (KO). The topography and stimulus selectivity of LO1 and LO2 along with their anatomical location, midway between ventral and dorsal visual pathways, suggest that they integrate shape information in retinotopic coordinates from multiple visual submodalities.

140. An information theoretic approach to detecting and discriminating mouse communication sounds

Robert C Liu[1], Christoph E Schreiner[2]

[1]Emory University, [2]UCSF

One of the ultimate goals of auditory systems research is to understand how the auditory cortex processes communication sounds to enable functions such as detection, discrimination and categorization. Most studies of mammalian auditory coding though have utilized only simple tones, frequency modulated sweeps, or synthetic complex signals like dynamic ripples and ripple noise. Since these are not naturally communicative, these sounds usually do not carry any behavioral relevance for an animal, unless it has been instrumentally trained or conditioned to respond to them. Such training, while useful for understanding how learning can induce plasticity in neural responses, may not be equivalent to natural communication in which the significance of a call is learned through social interaction. Thus, studying auditory processing in natural contexts complements more traditional approaches. On the other hand though, real communication vocalizations are often poorly characterized and thus difficult to use in parametric coding studies.

To make progress, the ideal situation would have simple sounds that are naturally communicative. This is in fact the case for mouse ultrasound communication calls. Two call types in particular are of interest: those produced by distressed infants after they have been isolated from their nest (Ehret 2005), and those emitted by adult males in the presence of females (Liu et al. 2003; Holy et al. 2005). Whistle-like in structure, these vocalizations can be straightforwardly characterized along several acoustic dimensions, such as frequency, duration, repetition period (Liu et al. 2003) and frequency sweep. Although quite variable from call to call, these signals form discrete and complementary clusters along each of these dimensions, suggesting that this intrinsic structure could be used for identification and processing of the calls (Liu et al. 2003).

In the work reported in this abstract, we introduce a computational methodology to explore how the auditory cortex of anesthetized, adult female CBA/CaJ mice (the recipient of both types of calls) codes for these communication sounds. We go beyond simply characterizing the selectivity for a particular call, and compute the information that binned, multiunit spike trains convey for communicative functions such as detection and discrimination. While the methods can in principle be applied to other communication models, the mouse model includes additional elements that make it attractive for these studies. Specifically, the impact that communicative significance has on information processing can be examined by experimenting on two groups of mice that behave differently to the pup calls: mothers and pup-naïve females. The former prefer the calls to a non-communicative sound, while the latter do not (Ehret et al. 1987). We find that for detection, a significantly larger proportion of recordings in mothers showed significant information compared to naive females, with higher average peak information. Individual recordings also consistently reached their peak information at earlier times, suggesting that the information was available faster. For discrimination, the most significant recordings from naive females provided on average much less information than the best recordings from mothers. This suggests that the communicative significance of pup calls to mothers is correlated with an improved ability of their auditory cortical neurons to discriminate these behaviorally-relevant sounds.

Ehret, G. (2005). "Infant rodent ultrasounds -- a gate to the understanding of sound communication." Behav. Genet. 35(1): 19-29.

Ehret, G., M. Koch, B. Haack and H. Markl (1987). "Sex and parental experience determine the onset of an instinctive behavior in mice." Naturwissenschaften 74(1): 47.

Holy, T. E. and Z. Guo (2005). "Ultrasonic songs of male mice." PLoS Biol 3(12): e386.

Liu, R. C., K. D. Miller, M. M. Merzenich and C. E. Schreiner (2003). "Acoustic variability and distinguishability among mouse ultrasound vocalizations." J. Acoust. Soc. Am. 114(6): 3412-22.

141. Effect of MST microstimulation on MT motion responses.

Christin H McCool, Ken Britten

University of California, Davis

Feedback connections, which are widespread between cortical areas, remain poorly understood. At the moment, many non-exclusive hypotheses have been suggested, including the creation of the suppressive receptive field surround, providing context or predictive information, spatial summation, attentional modulation, and gain control. In the dorsal stream, the motion-processing areas MT and MST have strong reciprocal connectivity. The medial superior temporal area (MST) provides feedback to the middle temporal area (MT or V5). As both areas are quite well understood, this forms an ideal system in which to explore the role of feedback in extrastriate cortex.

We microstimulated area MST, in order to perturb its feedback activity, while measuring directional responses in area MT. We used brief, local motion stimuli (translating Gabor functions or patches of moving dots), which were presented across the receptive field (RF) and surround of MT cells. We used either 3 or 5 possible locations, in a horizontal line centered on the RF center. Single stimuli and pairs of stimuli were presented, allowing us to study both the visual response and the stimulus interactions within the RF (divisive normalization or spatial summation). We measured directional tuning in each location of the RF and surround. Brief visual stimuli were presented in rapid succession, typically at 2 or 4 repetitions/sec. A brief, high-frequency pulse train was delivered immediately before half the stimuli. The microstimulation was directed to a column of MST neurons with similar tuning, and whose RFs always overlapped that of the recorded MT neuron. We measured the tuning of the MST column and related it to any changes in the MT cell's tuning during microstimulation.

In our sample of neurons, the principal effect of microstimulation was to modestly increase visual response duration, and this effect was evident for both types of visual stimuli. We have seen no evidence that MST stimulation affects overall response gain or maintained activity. Also, we saw no effect on surround strength or on divisive interactions between stimuli within MT RFs. Also, in some cells, we have seen weak and non-systematic effects on directional tuning. These results are a challenge to the view that MST feedback provides specific context information to MT, and are inconsistent with a role in divisive normalization or surround suppression.

Support Contributed By: the National Eye Institute (EY10562 and EY12576).

142. Sparse reverse correlation sequences result in improved tuning functions in V4.

Jude F Mitchell, John H Reynolds

The Salk Institute

Subspace reverse correlation has been shown to be an efficient means of estimating neuronal tuning properties in primary visual cortex and area MT. We applied this technique to estimate orientation tuning of single units in macaque area V4. At a stimulus presentation rate of 30 Hz we found that reverse correlation produced relatively weak orientation tuning as compared to conventional mapping techniques. Conventional mapping typically employs sequences that are separated by long blank intervals. This might be advantageous when estimating the relatively sluggish responses of V4 neurons. To examine this we carried out subspace reverse correlation using stimulus sequences that varied in their sparseness. Each video frame contained a stimulus with probability p, or a blank with probability (1-p). One can consider these sequences as binary sequences, where 1 corresponds to the presence of a stimulus and 0 corresponds to a blank. In this context the value p=1/2 has special significance, as it samples all binary sequences with equal probability. We varied p from 1 (standard reverse correlation) to 1/16. Increasing the temporal sparseness significantly increased response amplitude and improved the reliability of tuning estimates as assessed using cross-validation. This improvement comes at the cost of introducing a bias, as lower values of p result in a biased sample of the space of possible binary sequences. This is particularly clear in our estimates of the temporal kernel, which often became significantly larger and more peaked for low values of p. We are now examining the effect of varying sparseness on the ability of reverse correlation to recover the known kernels of model neurons, including an LNP (linear-nonlinear-poisson) model and a more realistic model that incorporates mechanisms for fast spike-rate adaptation.

143. Non-Uniform Passive Membrane Property in Dendrite Estimated by Fitting Multi-Compartment Model to Voltage Imaging Data

Toshiaki Omori[1], Toru Aonishi[2], Hiroyoshi Miyakawa[3], Masashi Inoue[3], Masato Okada[4]

[1]PRESTO, Japan Science and Technology Agency / RIKEN, [2]Tokyo Institute of Technology / RIKEN, [3]Tokyo University of Pharmacy and Life Science, [4]The University of Tokyo / PRESTO, Japan Science and Technology Agency / RIKEN

It has been reported that dendritic membrane properties and morphology play important roles in information integration in a single neuron. The specific membrane resistance, which is one of passive membrane properties, is suggested to be non-uniformly distributed. However its distribution is still unclear since it is difficult to measure the membrane properties directly over the whole single neuron. Since the specific membrane resistance is one of determinants of membrane response to synaptic inputs at each part of dendrite, revealing distribution of the specific membrane resistance is important for understanding mechanism of synaptic integration.

In this study, we used an indirect method to reveal non-uniform specific membrane resistance; we estimated distribution of specific membrane resistance by fitting a multi-compartment model to voltage imaging data. As voltage imaging data, we considered two kinds of experimental results which are suggested to be associated with distribution of the specific membrane resistance; namely propagations of EPSP (Inoue et al., 2001) and membrane responses to extracellular electric field (Bikson et al., 2004). We chose a distribution which can consistently reproduce both the propagations of EPSP and the membrane responses to extracellular electric field. Since currents induced by the extracellular electric field depend on direction of each compartment, we used a multi-compartment model with three-dimensional morphological data (not a stylized model, but a 3-D model in NEURON simulator).

In voltage imaging of hippocampal CA1 slices, an anterograde propagation of EPSP from the apical dendrite to the soma, and a retrograde propagation from the soma to the apical dendrite were observed, but it was shown by numerical simulations that uniform distribution could not reproduce the anterograde propagation and the retrograde propagation simultaneously (Inoue et al., 2001). By taking the non-uniformity of passive membrane property into account in our numerical simulations, a step function was estimated as a distribution of the specific membrane resistance, which reproduced both the anterograde and the retrograde propagations. Dependence of error between results by numerical simulation and voltage imaging showed that the specific membrane resistance might steeply decrease at distal part of the apical dendrite.

To check validity of the distribution of the specific membrane resistance, we performed numerical simulations of membrane response to the extracellular electric field. In voltage imaging of hippocampal CA1 slices, a biphasic response was observed in the apical dendrite, whereas a monophasic response was observed around the soma (Bikson et al., 2004). The step function estimated here could reproduce not only the propagations of EPSP but also the biphasic response in the apical dendrite. Furthermore, the monophasic response around the soma was also reproduced for the step function. This result shows generality of the estimated distribution of the specific membrane resistance.

In numerical simulations of multi-compartment model fitted to voltage imaging data, the two different experimental results were consistently reproduced by assuming the non-uniform specific membrane resistance. The estimated distribution suggests the specific membrane resistance steeply decreases at distal part of the apical dendrite of hippocampal CA1 pyramidal neuron.

This work was partially supported by Grant-in-Aid for Scientific Research Nos. 14084212 and 16500093 from MEXT and JSPS, Japan.

144. Variations in response sensitivity and intrinsic properties of cortical neurons

Michael J Pesavento, David J Pinto

Departments of Biomedical Engineering and Neurobiology & Anatomy, University of Rochester School of Medicine, Rochester, NY

Previous studies have demonstrated that the spiking response of cortical neurons is sensitive to rapidly changing input signals. Recently, it has been shown that neuronal subtypes exhibit different levels of sensitivity and that this can influence network processing. In this study, we used a novel approach to match the intrinsic properties of a conductance-based model neuron with experimental data and then used the model to understand how the sensitivity of responses to input timing depends on membrane properties.

We used square current pulses to quantify several standard intrinsic response measures (e.g., input resistance, membrane time constant, threshold, adaptation, etc.). We adjusted the membrane properties of the model so that simulated responses matched those of real neurons. We then varied the model's parameters to create a 6-D parameter space within which the variability of the model is bounded by the variability of the experimental data.

We assessed the sensitivity of model neurons to input timing using the same protocol as with real neurons. We simulated currents evoked by brief volleys of synaptic input varied in spike count and temporal distribution. Responses were quantified in terms of firing probability, latency to first spike, and variability. As with real neurons, firing threshold, latency, and variability of simulated responses all decreased with rapidly changing inputs. Moreover, several intrinsic measures from both real and model neurons correlate with response sensitivity (e.g., adaptation). We are currently using the model to examine and analyze the causal relationship between membrane properties and response sensitivity and to predict experimental findings induced by analogous manipulations in real neurons.

145. Nonlinear interaction between shunting and adaptation controls a switch between integration and coincidence detection in pyramidal neurons

Steven A Prescott[1], Stéphanie Ratté[2], Yves De Koninck[3], Terrence J Sejnowski[4]

[1]Computational Neurobiology Laboratory, The Salk Institute, La Jolla, CA 92037, [2]Département de physiologie, Université de Montréal, Montréal, Québec, Canada H3C 3J7, [3]Division de Neurobiologie Cellulaire, Centre de Recherche Université Laval Robert-Giffard, Québec, Québec, Canada G1J 2G3, and Department of Pharmacology and Therapeutics, McGill University, Montréal, Québec, Canada H3A 1Y6, [4]Computational Neurobiology Laboratory, The Salk Institute, La Jolla, CA 92037, and Division of Biological Sciences, University of San Diego, La Jolla, CA 92093

The total membrane conductance of pyramidal neurons in vivo is substantially increased by background synaptic input. Increased membrane conductance, or shunting, does not simply reduce neuronal excitability. Here we demonstrate that shunting interacts nonlinearly with adaptation. Recordings from hippocampal pyramidal neurons revealed that adaptation caused complete cessation of spiking in the high conductance state, whereas repetitive spiking could persist despite adaptation in the low conductance state. This behavior was reproduced in a phase plane model and was explained by an increase in voltage threshold caused by shunting. Phase plane analysis demonstrates that the increase in threshold allows greater adaptation to be induced at subthreshold potentials and reduces the minimum adaptation required to stabilize the system. Because this phenomenon depends on adaptation's activation curve, it occurs with M current-mediated adaptation but not with AHP current-mediated adaptation. The nonlinear interaction between shunting and the M current has other computationally important consequences. First, timing of spikes elicited by brief stimuli is more precise when background spikes elicited by sustained input are prohibited by adaptation, as occurs exclusively with M current-mediated adaptation in the high conductance state. Second, activation of the M current at subthreshold potentials, which is substantially increased in the high conductance state, hyperpolarizes average membrane potential away from voltage threshold thereby allowing only large, rapid fluctuations to reach threshold and elicit spikes. These results suggest that the shift from a low to high conductance state in a pyramidal neuron, as occurs in vivo, is accompanied by a switch from encoding DC input with firing rate to encoding transient inputs with precisely timed spikes, in effect, switching the operating mode from integration to coincidence detection.

146. Spatial integration of ampa- and gaba-type excitation in hypothalamic gnrh neurons

Carson B Roberts, Kelly J Suter

Emory University

Our studies using dynamic current-clamp indicate that GABAergic excitatory inputs provide a poor stimulus for action potential firing in hypothalamic GnRH neurons. However, GnRH neurons receive other synaptic inputs which might interact with GABAergic excitation. Moreover, GABAergic excitation may interact in a spatially complex manner depending on the location of inputs (i.e., on somata or on the dendrite). Such spatial interactions are difficult to study using only experimental approaches. Thus, we used our compartmental model derived from electrophysiological and morphological data in living GnRH neurons to examine spatial interactions between GABAergic and AMPA-mediated excitation. GABA synapses, with a reversal potential of -36.5 mV, were placed on the soma, and AMPA synapses, with a reversal potential of 0 mV, were placed on the soma or at varying distances from the soma on the dendrite. Consistent with earlier findings using dynamic current-clamping in living GnRH neurons, somatic AMPA synapses alone induced firing at 10 Hz for 100 Hz input frequency, whereas applying GABAergic inputs alone to soma, led to a sustained depolarization of -39 mV, but no action potentials. Effectiveness of AMPA excitation decreased with distance from the soma. When 100 Hz AMPA inputs were applied to the dendrite at 280 µm away from the soma, spike frequency was less than 0.1 Hz, despite the fact that the total length of the dendrite was over 1200 µm. With GABA synapses at the soma, firing rates induced by AMPA input on the soma were increased by ~10% for 10 Hz GABA, and by ~60% for 100 Hz GABA. With AMPA input at 100 µm from the soma, the AMPA-induced firing rate was 6.3 Hz. With the addition of GABAergic excitation at the soma, firing rate increased to 9.8 Hz. With AMPA input at 200 µm from the soma however, the somatic input from GABA did not boost the effect of distal AMPA inputs. Spike-triggered averages of conductances showed that GABA inputs immediately preceding AMPA inputs were maximally effective at inducing action potentials.

147. Dendritic Darwinism: Artificial evolution of neurons optimized for specific computations

Klaus M. Stiefel, Terrence J. Sejnowski

CNL, The Salk Institute

Genetic algorithms (GAs) are a class of optimization algorithms inspired by biological evolution. They operate by creating a class of solutions, selecting the best performers among this class, and replenishing the population from the winners. The diversity of solutions is increased by using genetic operators (mutation, crossover) in between generations. We used GAs to find optimal neuronal morphological structures for two computational tasks. The dendritic trees of the neurons were constructed using an L-system. The parameters of the L-system were encoded in the "genes" of the GA. In the first task, neuronal morphologies were selected for summing excitatory synaptic potentials (EPSPs) linearly. In the second task, structures were selected that distinguished the temporal order of EPSPs.
GAs found satisfactory solutions, sometimes resembling the morphology of real neurons. In particular, the artificial neurons discovered by the GA when selecting for linear summation resembled the neurons of the avian nucleus laminaris. The artificial neurons discovered by the GA when selecting for EPSP order detection resembled pyramidal cells in that they had thin (basal) and thick (apical) dendrites. We have therefore achieved an automated mapping between neuronal function and structure. This method allows a large catalog of computational functions to be built indexed by morphological structure.

148. Creation and reduction of a morphologically detailed model of a leech heart interneuron

Anne-Elise Tobin[1], Ronald L Calabrese[2]

[1]Brandeis University, [2]Emory University

Conductance-based neuron models aid in understanding the role intrinsic and synaptic currents play in producing neuronal activity. Incorporating morphological detail into a model allows for additional analysis of non-homogeneous distributions of intrinsic and synaptic conductances, as well as spatial segregation of electrical events. We developed a morphologically detailed "Full Model" of a leech heart interneuron that replicates reasonably well intracellular recordings from these interneurons. However, it comprises hundreds of compartments, each increasing parameter space and simulation time. To reduce the number of compartments of the Full Model, while preserving conductance densities and distributions, its compartments were grouped into functional groups that each share identical conductance densities. Each functional group was sequentially reduced to one or two compartments, preserving surface area, conductance densities and its contribution to input resistance. As a result, the input resistance and membrane time constant were preserved. The axial resistances of several compartments were rescaled to match the amplitude of synaptic currents and low-threshold calcium currents and the shape of action potentials to those in the Full Model. This reduced model, with intrinsic conductances, matched the activity of the Full Model for a variety of simulated current-clamp and voltage-clamp data. Because surface area and conductance distribution of the functional regions of the Full Model were maintained, parameter changes introduced into the reduced model can be directly translated back to the Full Model. As an example, we used a genetic algorithm automated parameter search to find conductance densities and distributions that produce endogenous bursting in the reduced morphology model. These parameters produce similar endogenous bursting when applied to the Full Model, which computes too slowly to use for automated parameter searching. Thus, our computationally efficient reduced morphology model can be used as a tool for exploring parameter space of the Full Model and in network simulations.

This work was supported by an NIH NRSA predoctoral fellowship NS 42983 to A-E. Tobin and by NINDS Grant NS 24072 to R. L. Calabrese.

149. Entorhinal Input and the Remapping of Hippocampal Place Fields

Joseph D Monaco[1], Isabel A Muzzio[2], Liat Levita[2], L F Abbott[1]

[1]Center for Theoretical Neuroscience, Columbia University, [2]Center for Neurobiology & Behavior, Columbia University

Recent experiments have begun to elucidate the spatial coding features of entorhinal cortex (EC) neurons in the rat (e.g., Fyhn et al, 2004), but their computational role in generating place-cell activity is not known. Hippocampal place fields are stable in familiar environments, but are randomly remapped in new environments. We explore how EC spatial coding could contribute to place-cell remapping, and how place fields can be stabilized by Hebbian and other forms of synaptic plasticity.

Layer II of dorsomedial EC contains "grid" cells with spatially periodic, tessellating place fields (Hafting et al, 2005). We hypothesized that hippocampal place cells form at locations where the summation of inputs from a number of EC cells with grids of different spatial frequencies and phases rises above a threshold value. Subsequent plasticity leads to the formation of singly peaked place fields through a competitive winner-take-all process. Remapping occurs when small shifts in the EC grids lead to supra-threshold input summation and subsequent place-field formation at other locations. Through this mechanism, small perturbations of the underlying EC grids can yield large and apparently random place field remappings. Completely sub-threshold summation of EC inputs allows for the possibility that no place field forms within a given environment. We explore such remapping in both 1- and 2-dimensional models of varying complexity.

Tetrode recordings in mouse CA1 show that place field stability can be modulated by attention (Kentros et al, 2004). We explore the idea that such modulation is caused by a shift in the competitive threshold for place field formation or an explicit change in the spatial metrics of the EC grid cells. We compare model results for these two possibilities to CA1 rate maps recorded in conditions requiring selective attention to either visuospatial or olfactory cues (Muzzio et al, 2005). In conclusion, we suggest that these models provide experimentally testable candidate mechanisms for the remapping of place fields in novel environments.

References

- Fyhn M, Molden S, Witter MP, Moser EI, Moser MB. (2004). Spatial representation in the entorhinal cortex. *Science*. 305: 1258-1264.
- Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*. 436(11): 801-806.
- Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER. (2004). Increased attention to spatial context increases both place field stability and patial memory. *Neuron*. 42: 283-295.
- Muzzio IA, Levita L, Kentros CG, de Boisblanc M, Kandel ER. (2005). Attention-dependent learning of modality-specific cues modulates the firing pattern of CA1 neurons differentially. *Soc for Neurosci Abstracts*. Washington, D.C.

150. Spatial and temporal organization of glomerular representation in the moth antennal lobe

Shigehiro Namiki[1], Ryohei Kanzaki[2]

[1]the University of Tsukuba, [2]the University of Tokyo

The antennal lobe (AL) is the first relay station for olfactory information in the insect brain and the anatomical equivalent of the olfactory bulb (OB) of mammals. Both systems have common structures called glomeruli in which neurons make synapses. Olfactory receptor neurons expressing the same receptor project to the same glomeruli in the AL. Projection neurons (PNs), the AL output neurons, transmit the processed information into higher order olfactory centers. To investigate spatial and temporal patterning of glomerular activity, we reconstructed olfactory representation by pooled set of single PN recordings in the moth AL.

Most of PNs innervated single glomeruli (n=126). PNs showed various slow temporal patterns to odor. PNs innervating the same glomerulus had similar response profile so that we could reconstruct odor-evoked spatial pattern of PNs spike frequency in the AL. This reconstructed spatial activity is highly distributed and dynamic. Different odor elicited different spatial pattern in each time point. The Euclidian distances between odor representations reached maximum at 200 ms after the response onset. There were no clear correlation between physical distance of glomeruli and response similarity of odor-evoked slow temporal patterns. This result is consistent with prior calcium imaging and modeling study. We conclude that olfactory information is encoded by distributed spatial and temporal pattern of PNs spiking activity and there are no clear relationship between the physical distance and response pattern in the moth AL.

151. Efficient Neural Burst Analysis

Rama Natarajan[1], Farzan Nadim[2]

[1]University of Toronto, [2]Rutgers University and New Jersey Institute of Technology

Rhythmic movements such as breathing, chewing and locomotion are regulated by neuronal networks that produce repetitive, patterned motor output. Detailed quantitative analysis of the concerted activity of neurons in these networks provides insight into the general principles of motor pattern generation and regulation. A behaviourally significant and experimentally advantageous characteristic of the extracellular recordings from these networks is the stereotypical, rhythmic bursts of activity from the individual neuron types. These recordings are typically analyzed using parameters such as the burst duration of a neuron, its cycle period, phase relationship which measures the firing delay of each neuron with respect to a reference neuron, regularity of the bursting pattern etc. A prerequisite to such quantitative analysis is accurate detection of the start and end times of bursts of each neuron type in every cycle. This is often complicated by the fact that extracellular recordings contain activities from multiple neuron types.

Spike-sorting, the paradigmatic first step in neural data analysis, separates single-unit neural activity from a multi-unit recording. But existing approaches are not always appropriate for data recorded from bursting neurons. Spike-sorting algorithms, which were primarily motivated by cortical recordings, often rely on the distinguishable, regular waveform characteristics of spikes from different neurons to classify them. When spikes invariably occur close together in time as in the case of bursting neurons, and overlap to render irregular waveforms, these approaches perform sub-optimally. Techniques to resolve the spike-overlap would prove to be highly inefficient and often ineffective, given the high degree of overlap in neural bursts.

We present a simple yet competitive solution using Conditional Random Fields (CRF) to segment a sequence of multi-unit neural recording into bursts of neural activity, and label each burst according to the putative neuron type that fired it. Each neural spike detected during computational pre-processing on the data, is treated as an observation in the input sequence. We propose a discriminatively trained log-linear sequence labeling model, that is defined on (observation, label) pairs, specified by a vector of local spike features and corresponding weight vectors. The features are represented as a binary value, with 1 indicating the presence of a feature. Each feature function is defined over a window of the observation sequence and labels at current and previous positions in the corresponding label sequence.

The encoding scheme for the feature functions is given as follows: each spike is tagged based on its position in a burst - start, middle or end (f_1). This conveys information on whether the spike marks the beginning or end of a burst of activity, or not. The second feature (f_2), is the average inter-spike interval (ISI) of each spike that is not a boundary spike, calculated as the average interval with each of its neighboring spikes. For boundary spikes, the ISI of the start spike with the spike after, and the ISI of the end spike with the spike before, are calculated. The average amplitude of a spike is also calculated, f_3. Finally, to make use of information from before and after the label of a current observation, a conjunction of features is used to encode the feature set. The values of features of current input are multiplied with those of the previous (forming features f_4, f_5 and f_6) and next (forming features f_7, f_8 and f_9). The conjunctions of the resulting feature values are also used, resulting in features f_{10}, f_{11} and f_{12}.

Training involves maximum likelihood estimation of the weight for each feature function from the training set such that the log-likelihood of the training data is maximized. The rhythmic neural activity gives rise to long-range temporal dependencies among the bursts. Under this scheme, it is possible to achieve a globally optimal labeling by way of encoding the long-range dependencies on observation values in the sequence while keeping the inference problem tractable. Non-independent features of spikes (such as significance of spike position in the sequence and label of previous spike) can be important in dealing with observations unseen during training. Such features are readily represented in this model.

The model was evaluated on extra-cellular recordings of the pyloric network of the crustacean stomatogastric nervous system. The data show triphasic motor patterns with overlapping activity from three neuron types. Precision (P) and recall (R) are used as evaluation measures; burst classification accuracy is given using the F-Measure which is 2RP/(R+P). On the development data, the model scored 0.837 on average, 0.817 on the training data, and 0.804 on the testing data. Averaging over 20 different recordings of 5 minutes each, the model detected burst segments with 81% accuracy and assigned them to one of several labels based on putative neuron types, with accuracy of 83%. Results from analysis of phase relationships between the neurons were compared with those from a hierarchical classification approach as well as experimentalist-assisted classification. Despite using a very simple graph structure for the CRF and a minimal set of features, the model outperformed other methods with minimal computational effort.

152. Dynamical contrast gain control mechanisms in a layered model of the primary visual cortex

Laurent U Perrinet, Jens Kremkow, Alexandre Reynaud, Frédéric Y Chavane

INCM/CNRS

Computations in a cortical column are characterized by the dynamical, event-based nature of neuronal signals and are structured by the layered and parallel structure of cortical layers. But they are also characterized by their efficiency in terms of rapidity and robustness. We propose and study here a model of information integration in the primary visual cortex (V1) thanks to the parallel and interconnected network of similar cortical columns. In particular, we focus on the dynamics of contrast gain control mechanisms as a function of the distribution of information relevance in a small population of cortical columns. We show that if their lateral connectivity is tuned to maximize an efficiency criteria, we observe dynamical contrast gain control mechanisms which are characteristic of V1.
This cortical area is modeled as a collection of similar cortical columns which receive input and are linked according to a specific connectivity pattern which is relevant to this area.

These columns are simulated using the NEST simulator using conductance-based Integrate-and-Fire neurons and consist vertically in 3 different layers. This system is similar in its final implementation to local micro-circuitry of the cortical column presented by Grossberg (2005). They show prototypical spontaneous dynamical behavior to different levels of noise which are relevant to the generic modeling of biological cortical columns (Kremkow, 2005).

The visual input is first transmitted from the Lateral Geniculate Nucleus (LGN) using the model of Gazeres (1998). It transforms the image flow into a stream of spikes with contrast gain control mechanisms specific to the retina and the LGN. This spiking activity converges to the pyramidal cells of layer 2/3 thanks to the specification of receptive fields in layer 4 providing a preference for oriented local contrasts in the spatio-temporal visual flow. In particular, we use in these experiments visual input organized in a center-surround spatial pattern which was optimized in size to maximize the response of a column in the center and to the modulation of this response by the surround (bipartite stimulus). This class of stimuli provide different levels of input activation and of visual ambiguity in the visual space which were present in the spatio-temporal correlations in the input spike flow optimized to the resolution of cortical columns in the visual space. It thus provides a method to reveal the dynamics of information integration and particularly of contrast gain control which are characteristic to the function of V1.

The connectivity pattern between columns implemented the lateral propagation of information in V1 and is inspired by an efficiency criteria. This connectivity was layer-specific and in particular connects specifically bundles of neighboring columns within a macro-column and across macro-columns. This architecture was inspired by neuro-physiological observations on the influence of neighboring activities on pyramidal cells activity and correlates with the lateral flow of information observed in the primary visual cortex, notably in optical imaging experiments (Jancke, 2004). The connectivity is derived from an algorithm that was used for modeling the transient spiking response of a layer of neurons to a flashed image and which was based on the Matching Pursuit algorithm (Perrinet, 2004). It extends this algorithm to a continuous flow of information by setting the specific interactions between the different columns of the layer and their sensitivity so as to maximize locally a compromise between the sparseness of the representation and the rapidity of the response. Particular care is put on defining the dynamics of the interactions and of the sensitivities, in accordance with biological data, and which specifically tune the selective interaction between the different layers of the cortical area.

Mathematically, the neurons in this layered model of cortical mini-columns explicitly represent the probability of the match of stored patterns with the continuous input knowing the short-term history of already generated events. This evaluation is based on a generative model of image generation and input noise which enables to build a probabilistic model thanks to bayesian inference. Consequently, it provides a generic approach for integrating different modalities of possibly ambiguous signals such as is the case in V1 between feed-forward input and lateral interactions.

We observe in this network-level model of V1 principal characteristic of this layer:
* Pyramidal neurons of layer 2/3 exhibit a characteristic response to different levels of contrasts to the presentation of an oriented central patch. Particularly, the shape of the orientation selectivity curves in V1 is independent of the contrast in a large range and the contrast gain functions are steeper than in the LGN. This contrast-gain control behavior, which is often modeled as a divisive normalization, is specific to V1 and was in particular dependent on the tuning of the global sensitivity between layers in a column and was complementary to the function of the contrast gain control in the LGN model.
* We replicate the influence of noisy surround activity on a localized central response. In particular, the network exhibited a suppression of the central response as a function of the context of the neighboring activity. The dynamics of this modulation may be tuned to match neuro-physiological observations and in particular in optical imaging.
* Correlated visual input from the center and the surround provided parallel incoming events correlated in time which specifically drives the network in different macroscopic states. This dynamically modulates the central response and provided a specific contrast control gain mechanism which is comparable to biological observations.

This method therefore provides a generic approach for studying the dynamics of neural computations in cortical areas such as V1. In fact, the high number of possible parameters in large scale model of V1 renders its study computationally prohibitive and this method thus provides some macroscopic parameters, e. g. the selectivity of cortical mini-columns and the lateral connectivity between columns, which may provide insights in the functional architecture of this layer. This method was applied to understanding the data observed by optical imaging of the response of V1 to dynamical visual flow such as the line-motion illusion (Jancke, 2004) and allows to predict how the information actually flows in V1 across the local feedback in the cortical column as well as in the extent of lateral interactions across columns and macro-columns. It is particularly adapted to studying the effect of the compromise between the efficiency of information transformation and the need for a rapidly available computational solution which is specifically revealed in experiments such as temporal masking. Moreover, it takes full advantage on the parallel and event-based nature of neural computations and provides an efficient computing framework which is illustrated by basic image processing tasks such as data compression or denoising. By the nature of our architecture and constraints, this architecture is especially well fitted for an implementation on emerging hardware solution such as networks of aVLSI circuits.
[Perrinet(2004)] Perrinet L. Feature detection using spikes : the greedy approach. Journal of Physiology, Paris (special issue), 98 (4-6):530Ð9, July-November 2004. [Jancke(2004)] Jancke, D. and Chavane, F. and Naaman, S. and Grinvald, A. (2004). Imaging cortical correlates of illusion in early visual cortex. Nature 428: 423-426

153. An information-theoretic generalization of spike-triggered average and covariance analysis

Jonathan W Pillow[1], Eero P Simoncelli[2]

[1]Gatsby Computational Neuroscience Unit, UCL, [2]HHMI and New York University

One of the central problems in sensory neuroscience is that of characterizing the neural code: the mapping from sensory stimuli to spike responses. Recent work has sought to address this problem with the use of statistical tools for *dimensionality reduction*. A simple intuition underlies these methods: although the space of all stimuli (e.g. the space of all images) is enormous, most attributes of these stimuli do not have any effect on the neuron's response. We can therefore attempt to model the neural code by finding a low-dimensional *feature space* within which the neuron computes its response [Bialek & de Ruyter van Steveninck 05].

Two basic methods have been developed to estimate the features that define the subspace underlying a neuron's response. The first looks for changes in the mean and/or variance of the spike-triggered stimulus ensemble (i.e., the set of stimuli that elicited a spike from the neuron), relative to those of the raw stimulus ensemble, which correspond to the spike-triggered average (STA) and the eigenvectors of the spike-triggered covariance (STC) matrix [de Ruyter van Steveninck & Bialek 88, Bialek et al 91, Simoncelli et al 04, Bialek & de Ruyter van Steveninck 05]. A second method searches directly for the feature space that preserves maximal information about the response [Paninski 03, Sharpee et al, 04].

Here, we describe a framework for dimensionality reduction in neural models that occupies a middle ground between STA/STC analysis and full information maximization. We assume that the spike-triggered ensemble is completely characterized by its mean (STA) and covariance (STC), and can thus be approximated as Gaussian. We then use an information-theoretic criterion to find the relevant feature subspace. The resulting solution has several useful properties: (1) it provides a common framework for spike-triggered average and covariance analysis, incorporating the joint effects of the mean and covariance on neural response, and allowing subspace dimensions to be ranked in order of their informativeness; (2) the Gaussian assumption leads to computationally efficient and robust information maximization, and the data requirements for recovery of a linear stage of high dimensionality are relatively modest; (3) it provides an explicit "default" model of the nonlinear stage that maps the filter responses to spike rate, even in high-dimensional feature spaces; (4) it is equivalent to maximizing the likelihood of the spike train given the stimuli under the assumed model; and (5) it can be applied to novel problems, such as the estimation of a model with space-time separable filters. We demonstrate the effectiveness of the method by applying it to the recorded extracelluar responses of macaque retinal ganglion cells and V1 cells.

154. Visual acuity in the presence of fixational eye movements

Xaq Pitkow[1], Haim Sompolinsky[2], Markus Meister[1]

[1]Harvard University, [2]Hebrew University

Humans can distinguish visual stimuli that differ by features the size of only a few photoreceptors. We can do this despite the incessant image motion due to fixational eye movements which can be many times larger than the features to be distinguished. To perform well, the brain must identify the retinal firing patterns induced by the stimulus while discounting those similar patterns caused by spontaneous retinal activity. This is a challenge since the trajectory of the eye movements, and consequently the stimulus position, are unknown. We derive the ideal observer decision rule for using retinal spike trains to discriminate between two stimuli, given that their retinal image moves with an unknown random walk trajectory. We then show how the brain could implement this computation using a simple, biologically plausible recurrent neural network with lateral interactions matched to the statistics of eye movements. This architecture performs Bayesian inference by representing likelihoods as neural activity which then diffuses through the recurrent network and modulates the influence of information acquired later.

155. Natural scene statistics predict that larger ganglion cells should have relatively smaller surrounds

Charles P Ratliff, Peter Sterling, Vijay Balasubramanian

University of Pennsylvania

The center-surround structure of retinal ganglion cell receptive fields is thought to reduce the redundancy in natural visual stimuli (e.g. Barlow 1961, Srinivasan et. al. 1982). The center measures light intensity over the cell's dendritic field; the surround measures intensity over a broader region. The cell then compares center vs. surround and transmits a measure of contrast. Physiological observations (Linsenmeier, 1982) and modeling (Smith and Sterling, 1990) have shown that the surround/center ratio decreases as center size increases. Here we show that this trend is predicted by the statistical structure of natural scenes for a ganglion cell mosaic that minimizes redundancy. The result is explained by a deviation from scale invariance in natural images. For artificial images constructed to be scale-invariant, the optimal surround-center ratio does not change with center size.

Natural visual stimuli were modeled using two sets of calibrated, monochromatic images: a standard set (van Hateren & van der Schaaf 1998) from the northern Dutch countryside, and a new set from the Okavango Delta in Botswana. Results were nearly identical. Images were convolved with ganglion cell receptive fields modeled as divisively normalized, difference-of-Gaussians filters (Shapley & Enroth-Cugell 1984). Such filters measure local percent contrast as (Ic-Is)/Is, where Ic (Is) is the center (surround) intensity. Dividing by surround intensity models adaptation to local mean luminance. The contrast responses were rectified into ON and OFF channels and binned into 10 equally probable levels consistent with measured ganglion cell contrast sensitivity (Dhingra et al. 2003). Filters with a center standard deviation s were placed at nodes 2s apart (DeVries & Baylor 1997; Borghuis et. al., in preparation), and joint distributions of the filter responses were accumulated. Information in the collective response of arrays of up to seven filters was measured from the joint distributions. For a given center size, information was calculated as a function of surround size.

For every center size, there was an optimal surround size that maximized information transfer per filter (or, equivalently, minimized redundancy in the filter array). In every case, optimal surround size was the same for filters of ON or OFF polarity. We measured the optimal surround size as a function of center standard deviation s. For small s, the optimal surround was about 5 times broader than center, and the optimal ratio of center to surround increased with center size. For the broadest filters tested, the optimal surround was only about 2 times broader than the center. These results agree with physiological measurements (Linsenmeier, 1982). Similar results were obtained for a range of surround strengths and weightings.

We wondered if the optimal surround size for a given center is determined by the two fundamental properties of natural scenes: their skewed intensity distribution and power-law, two-point correlations. To test the effects of correlations we compared the optimal surround sizes for uncorrelated Gaussian noise and for "pink noise" (artificial images constructed by convolving Gaussian noise with a 1/f filter). For Gaussian noise, information transmitted by the filter array increased steadily with the surround size for a fixed center size – hence there was no optimal surround/center ratio. However, for "pink noise" images the optimal surround/center ratio was ~12. Thus the correlations appear to explain the existence of an optimal surround size. To test the effects of the intensity distribution we compared the "pink noise" results to "natural pink noise" (pink noise with intensities remapped to the distribution of natural intensities). These two types of artificial image are equally redundant, but the optimal surround/center ratio for "natural pink noise" was ~2, i.e., 6-fold smaller than for pink noise. These results held for all center sizes because the artificial images are scale-invariant by construction. Thus the change in optimal surround/center ratio with center size apparently arises from a violation of scale invariance in natural image statistics; whereas the precise optimal ratios depend on both the intensity distribution and the correlation structure in the images.

156. Solving the stereo correspondence problem with hybrid position and phase disparity detectors

Jenny Read[1], Bruce Cumming[2]

[1]Newcastle University, [2]National Eye Institute

Neurons in early visual cortex show two distinct mechanisms of disparity tuning: position disparity and phase disparity. This is surprising, given that phase disparity – in which the displacement of each Fourier component is proportional to spatial period – does not occur in natural stimuli. This raises the question of why the visual system does not restrict itself to pure position-disparity detectors, since these are better tuned to physical disparities. We recently showed that phase disparity can be used to help solve the stereo correspondence problem (Read & Cumming, Society for Neuroscience Abstracts 583.4. 2005). A population of hybrid position/phase-disparity energy-model neurons generally shows many local maxima in its response to a constant-disparity stimulus. Identifying which maximum corresponds to the true stimulus disparity is a fundamental problem in stereo vision. Indeed, in some instances, even a constant-disparity stimulus may not even correspond to a local maximum at all: the true disparity may fall at a saddle-point, which is maximum with respect to phase disparity but minimum with respect to position disparity. Importantly, however, the true disparity is always at a local maximum with respect to phase disparity, and is always located at zero phase disparity. Thus, the response of pure position-disparity detectors does not uniquely identify the correct match, because there are in general several local extrema in the response of this population, any one of which could correspond to the stimulus disparity. However, the response of hybrid position/phase-disparity detectors does identify the correct match. If, at one of the extrema in the pure position-disparity subpopulation, the response of energy-model units tuned to phase disparities away from zero is greater than the response of the unit tuned to zero phase disparity, then this position disparity cannot be that of the stimulus. Thus, the response of a population of hybrid position/phase-disparity detectors can be used to veto false matches in the position-disparity subpopulation. Phase disparities are useful for solving the stereo correspondence problem precisely because they do not occur in real stimuli: if moving the optimal phase disparity away from zero to non-physical values increases the response, then the position disparity tuning cannot be that of the stimulus. So, false matches can be identified and ruled out. In constant-disparity stimuli, this is enough to enable the correct disparity to be uniquely identified, even within a single spatial-frequency channel where the stimulus disparity is many spatial cycles.

Here we implement a simple algorithm based on this principle, and explore its performance with test images. One is a random-dot stereogram portraying a pyramidal depth surface, the other is a natural image pair (an aerial photograph of the Pentagon). Even in these complicated, non-planar depth structures, the algorithm succeeds in producing fairly good (though noisy) disparity maps from the output of a single spatial-frequency and orientation channel. When the outputs of several spatial-frequency and orientation channels are averaged, the results are even better. This demonstrates that this algorithm is potentially useful even in realistic situations. We suggest that this method of reducing the false-match problem may be one of the strategies employed by the stereo system, and that this may why the brain contains detectors tuned to both position and phase disparities.

157. Extracting low-dimensional task representations from neural signals

James Rebesco[1], Sara A Solla[2], Lee E Miller[1]

[1]Department of Physiology, Northwestern University, [2]Department of Physics and Astronomy, Northwestern University

How does the brain encode behavior? Our ability to obtain simultaneous recordings from an increasingly large number of neurons in awake, behaving primates provides the data for explorations of this fundamental question.

The functional description that relates neural activity to behavior involves the collective activity of a neural ensemble, which needs to be characterized through properly identified degrees of freedom. A simple starting point would be to consider all neurons whose activity can be monitored during the execution of a particular task. This choice would embed the description of the ensemble's activity in a high-dimensional space, with a coordinate axis for each neuron. But such a description would be misleading, as it ignores the significant degree of correlation in the activity among these neurons. The challenge is then to identify and characterize the sub-dimensional manifold that reveals the underlying relationship between neural activity and behavioral task.

A variety of linear methods for dimensionality reduction have been used for this purpose: principal components analysis (Chapin and Nicolelis 1999) and class-conditional modeling by Gaussians (Hatsopoulos et al. 2001). These attempts have achieved a reasonable but limited degree of success. Here we propose that a nonlinear approach is crucial for unveiling the intrinsic geometry of the relevant low-dimensional neural space. We demonstrate this hypothesis through the nonlinear analysis of recordings of M1 and PMd neurons in two macaque monkeys executing center-out reaching tasks to eight distinct targets uniformly arranged on a circle. We have analyzed five datasets, each consisting of 80-90 single-unit or multi-unit signals. The analysis is based on techniques from multi-dimensional scaling (Cox and Cox 2001), extended to nonlinear manifolds through the use of the Isomap algorithm (Tenenbaum et al. 2000). The fundamental concept of the method is to use the locally linear metric among neural signals to construct manifold-embedded geodesic paths that define a global nonlinear metric.

When applied to three datasets collected during a series of straight reaches to small targets, the analysis revealed a two-dimensional manifold that captured with surprising accuracy the eight-fold organization of the targets in physical space. The corresponding class-conditional clusters were compact but nonGaussian. A two-dimensional subspace projection obtained by PCA did not capture this crucial feature of the task. The performances of these linear and nonlinear unsupervised mappings were compared quantitatively through the computation of the mutual information between the mapped points and the target labels.

In two other datasets using a different experimental setup, hand paths to targets were more curved and less reproducible. This reduced correlation between hand path and endpoint was reflected in the Isomap projection, which was less clearly reflective of target location in physical space. However, in this case the nonlinear projection provided a strong correlation ($R^2=0.85$) between neural data and direction of hand motion 200 ms later.

These results suggest that ensemble neural activity defines a low-dimensional manifold whose intrinsic coordinate system has direct relevance to the task being performed. Rather than simply examining tuning curves for individual neurons and generating estimates of task from an ensemble activity generated by linear superposition, the approach pursued here enables us to relate the intrinsic nonlinear geometry of the ensemble of neural activities to relevant features of the task.

158. Contrast gain control in the LGN optimizes information transfer, but may not require any dynamic adaptation process

Kate S Denning, Pamela Reinagel

UCSD

Sensory neurons are believed to adapt their gain to match the variance of signals along the dimension they encode. Contrast adaptation has been the subject of extensive physiological and theoretical study. Neurons in the retina, Lateral Geniculate Nucleus (LGN) and primary visual cortex (V1) exhibit contrast adaptation to the temporal and/or spatial contrast of visual stimuli. Similarly, adaptation to motion contrast has been reported in motion sensitive neurons from primates (MT) to flies (H1). We have studied contrast adaptation in the LGN. We recorded responses of single neurons in the LGN of anesthetized cats during presentation of full field flickering white noise stimuli with different temporal contrasts.

The first question we asked is whether contrast gain control has any consequences for the quality of neural coding. To quantify the extent of gain control, we fit standard linear/nonlinear models to our recorded data at each contrast. By forcing the nonlinearity to be fixed at all three contrasts, gain changes are reflected as changes in the amplitude of the linear filter. We find that all LGN cells exhibit some degree of gain control: filter amplitudes change in the direction consistent with reducing the sensitivity when contrast increases. The extent of gain control, however, varied widely across cells. Some cells had very little gain control, while others showed nearly perfect scaling (increasing sensitivity 3x when contrast drops 3x). To assess the quality of neural coding, we measured qualitative parameters of the response (sparseness, temporal precision, and trial-to-trial reliability of spike count) as well as the information transmission rate in bits/sec or bits/spike. We found that all LGN cells show some degradation of neural coding as contrast decreases, but that cells with the most gain control were the most resistant to loss of fidelity at low contrast. These results support the hypothesis that contrast gain control serves the function of preserving the efficiency of information transmission.

Our second question is whether an active process of adaptation is required to account for contrast gain control in the LGN, in light of the fact that nonlinear systems can exhibit similar “adaptation” effects without changing any internal parameters. First, we measured the dynamics of changes in neuron response properties immediately after abrupt changes in contrast. We find that cells’ response properties change gradually after a contrast change, but the time constants were consistent with the extent to which the cells’ integration time window overlapped with the new stimulus contrast. Second, we implemented models of LGN cells with a linear filter followed by a stochastic nonlinear spike generating mechanism. The model cells recapitulated most (but not all) of the contrast dependencies of LGN neurons, even when no internal parameter of the model changed with contrast. Notably, the fixed, nonlinear model cells showed gain control (changes in the amplitudes of the recovered linear filters). By exploring the parameter space in the model, we found that the gain control in the model can confer protection from information degradation, as it does in our LGN cells. We have not yet found any fixed model parameters whose gain control is as strong as some LGN neurons, so we cannot rule out that some active adaptation also occurs. We are currently exploring whether fixed nonlinear models can also explain the dynamics in our LGN data.

159. How behavioral constraints may determine optimal sensory representations

Emilio Salinas

Wake Forest University School of Medicine

The work presented here springs from two sets of observations. On one hand, there is a puzzling dichotomy in the types of response functions or tuning curves that have been documented experimentally. Many types of neurons respond to sensory stimuli such that their response profiles have a single peak at a 'preferred' point. Population codes are often thought of as composed of arrays of such units. However, many neurons show monotonic dependencies on sensory parameters. Such monotonic responses are both widespread and difficult to reconcile with current ideas about sensory coding.

On the other hand, previous theoretical studies suggest that many response characteristics of neurons are optimal for encoding information. According to this efficient coding hypothesis, neurons should encode information compactly, so that the information conveyed by one neuron about the sensory world is as independent as possible from the information carried by others. Optimality criteria based on such considerations (redundancy reduction, decorrelation, sparseness, etc.), combined with the statistics of natural stimuli, have been extremely successful at explaining the receptive field properties of neurons in early visual areas. Such theories, however, do not take into account how the encoded information will be used afterwards, if at all.

Here I propose a complementary point of view in which the responses of a neural population are evaluated in terms of the range of outputs that they are capable of generating. The idea is that optimal sensory representations should take into account not only the statistics of the sensory world, but also the statistics of the downstream activity that generates behavior. First, a a criterion is derived for quantifying the efficacy of a set of tuning curves based on the possible downstream functions that can be constructed from them. Then, curves that are optimal according to this measure are derived numerically for different classes of output or downstream functions with various properties.

According to this criterion, if the downstream responses are non-monotonic, in general the sensory responses that are optimal for driving them will be unimodal. However, if the downstream responses are monotonic, then the optimal sensory tuning curves will be predominantly monotonic. Furthermore, monotonic tuning curves should also be observed when the downstream responses are arbitrary but have a monotonic trend. The encoding of binocular disparity fits this scheme, and the types of optimal tuning curves obtained from the model match recent experimental data. Another aspect of sensory tuning curves that seems to depend strongly on the activity that is evoked downstream is their width. A dramatic example from the bat auditory system, where the width of peaked tuning curves varies systematically, is also explained by the model based on simple behavioral considerations.

In conclusion, the results suggest that the optimality of sensory responses depends not only on considerations about stimulus representation, but also on the downstream events driven by those responses. If the downstream demands change, the responses considered optimal will change as well. Thus, unimodal and monotonic curves are optimal under different circumstances. In particular, when some values of a stimulus are more likely to lead to certain actions, not all stimulus values are equal. The results suggest that monotonic curves are more efficient when there is such a trend or asymmetry in the downstream functions. These results show that knowledge about the downstream impact of their responses is crucial for understanding important properties of sensory representations.

160. The scaling of 'Winner Takes All' accuracy with the population size

Maoz Shamir

Center for Bio-dynamics, Boston University

Empirical studies appear to support conflicting hypotheses in regard to the nature of the neural code. While some studies highlight the role of a distributed population code, others emphasize the possibility of a 'single best cell' readout. One particularly interesting example of 'single best cell' readout is provided by the 'Winner Takes All' (WTA) approach. According to the WTA, every cell is characterized by one particular preferred stimulus, to which it responds maximally. The WTA estimate for the stimulus is defined as the preferred stimulus of the cell with the strongest response.

From a theoretical point of view, not much is known about efficiency of 'single best cell' readout mechanisms, in contrast to the considerable existing theoretical knowledge on the efficiency of distributed population codes. In this work, we provide a basic theoretical framework for investigating single best cell readout mechanisms. We study the accuracy of the WTA readout. In particular, we are interested in how the WTA accuracy scales with the number of cells in the population. Using this framework, we show that, for large neuronal populations, the WTA accuracy is dominated by the tail of the single cell response distribution. Furthermore, we find that although the WTA accuracy does improve when larger populations are considered, this improvement is extremely weak compared to other types of population codes. More precisely, we show that while the accuracy of a linear readout scales linearly with the population size, the accuracy of the WTA readout scales logarithmically with the number of cells in the population.

161. Temporal Invariance and Predictive Coding

Jonathan Shaw

University of Rochester

Temporal invariance and predictive coding are two principles that have been used to explain sensory perception. Both have qualitatively reproduced the receptive fields of cells in the primary visual cortex. However, no relationship has been found between them so far. We argue that such a relationship does exist: temporal invariance can be viewed as a form of predictive coding. To establish this relationship, consider a continuous time varying input, and suppose that we wish to find a continuous function of this input such that the information that this function conveys about the future input is maximized. Other than the requirements that the input and the function be continuous, this is a fairly vanilla predictive coding formalism.

This information is usually maximized with the help of a generative model of the input. However, if the input is generated by a Markov process, then the information between successive outputs is a lower bound on the information between an output and the next input. Therefore, maximizing the information between successive outputs will cause the information between an output and the next input to be increased.

The latter approach requires some way to measure the information between successive outputs of the network. One approach is to assume that the output is generated by a Markov process as well. This will always be true if the function is invertible, and may be approximately true even if the function is not. If the output is generated by a known Markov process, then the statistics of this process can be used to measure the information between successive outputs.

The simplest processes that could be used to model the output are Levy stable processes. Among these processes, the only one that generates continuous data is Brownian motion. Since the input is continuous and the function is continuous, the output is continuous. Therefore, the simplest model of the output that makes any sense to use is Brownian motion.

The information between successive values of a process generated by Brownian motion with a fixed variance is determined by the noise variance added to each sample to make the next sample. Therefore, the mutual information between successive outputs of a function of the input is maximized when the function is chosen to minimize this variance. Thus, functions that minimize temporal variance perform predictive coding.

Having shown that temporal invariance is a form of predictive coding, we now observe two major distinctions between temporal invariance and other forms of predictive coding. First, it is not a generative model. Since the goal of the network is stated purely in terms of the output, there is no need to be able to reconstruct the input from the output. In slow feature analysis, which is an implementation of temporal invariance, the functions used are highly nonlinear, and reconstruction is difficult if not impossible. Second, temporal invariance can be accomplished in closed form (again done by slow feature analysis). This makes it attractive from both a mathematical and computational standpoint. We hope this analysis will instigate the crosspollination of ideas between the fields of temporal invariance and predictive coding.

162. Probing the structure of multi-neuron firing patterns in the primate retina using maximum entropy methods

Jonathon Shlens[1], Greg D Field[1], Jeff L Gauthier[1], Matthew I Grivich[2], Dumitru Petrusca[2], Alexander Sher[2], Alan M Litke[2], EJ Chichilnisky[1]

[1]Salk Institute, [2]UC Santa Cruz

Current understanding of many neural circuits is limited by our ability to explore the vast number of potential interactions between cells. For any given circuit, a full understanding of network interactions requires (a) the ability to record from most or all neurons in the circuit, and (b) analytical methods to examine all possible interactions, a computationally daunting task. We apply powerful new experimental and analytical approaches to understand the complexity of network interactions in the primate retina.

Large scale multi-electrode recordings were used to measure electrical activity in nearly complete, regularly spaced mosaics of several hundred ON and OFF parasol (magnocellular-projecting) retinal ganglion cells (RGCs) in macaque monkey retina [Frechette et al, 2005]. The completeness of the mosaics indicates that we recorded from nearly every cell of both types in a 1x2 mm region of retina. In the presence of uniform, constant photopic illumination, pairs of RGCs fired synchronously (+/- 5 ms) much more often than expected by chance, indicating significant network interactions. Synchrony declined systematically with distance between cells, and differed for ON-ON, ON-OFF, and OFF-OFF pairs. These results provide strong constraints on the underlying circuitry, indicating that RGC connectivity is spatially localized and universal among cells of the same type.

However, this conventional pairwise analysis fails to distinguish whether synchrony arises only from pairwise interactions (e.g. reciprocal gap junctions) or instead reflects more complex interactions (e.g. diverging common input from presynaptic amacrine cells). This distinction is fundamental: diverging common input would be expected to produce very different multi-neuron firing patterns in RGCs than purely pairwise connections, and such firing patterns have been suggested as a mechanism for the retina to convey distinct, multiplexed visual messages to the brain [Meister 1995]. Furthermore, the spatial extent of synchrony may reflect signal spread through intermediate neurons, so pairwise analysis also fails to reveal the elementary spatial scale of network interactions.

To test quantitatively whether network interactions are purely pairwise, and whether they are restricted to immediately adjacent neurons, requires a candidate null statistical model to compare against the observed data. In the simpler case of pairwise synchrony, the null model is statistical independence; but in a multi-cell circuit, many patterns of connectivity are consistent with the observed pairwise interactions. Therefore, we employ the maximum entropy framework to generate a null statistical model of firing patterns constrained by the observed pairwise correlations [Jaynes 1957, Schneidman et al 2003, 2006]. Borrowed from statistical mechanics, the maximum entropy approach has been employed in many fields of machine learning to determine a complete probabilistic model from overlapping constraints, without assuming additional statistical structure. We find that purely pairwise interactions account for >95% of the departures from statistical independence in parasol RGCs. Furthermore, the predictive ability of a pairwise model approaches the upper limits given the intrinsic reproducibility of the data. We also find that the more restricted model of pairwise interactions limited to adjacent neighbors in the mosaic also accounts for >95% of the departures from statistical independence, and approaches the intrinsic reproducibility of the data. Thus, multi-neuron firing patterns from populations of parasol cells can be understood as simply arising from reciprocal interactions between adjacent cells in the mosaic, suggestive of gap junction coupling between neighbors. This provides a dramatically simpler picture of retinal organization than has been proposed, and makes it possible to understand circuit activity with a relatively small number of measurements. We have begun to extend this approach to examine larger populations of RGCs as well as interactions between RGCs in time.

163. Hybrid Discrete/Continuous Models of Brain Dynamics: Estimation from Spikes

Lakshminarayan Srinivasan[1], Uri T Eden[2], Sanjoy K Mitter[3], Emery N Brown[4]

[1]MIT Department of Electrical Engineering & Computer Science/ Harvard Medical School, Health Sciences & Technology Track/ Massachusetts General Hospital, [2]HMS-MIT HST/MGH, [3]MIT EECS Laboratory of Information & Decision Systems, [4]MIT Deptartment of Brain & Cognitive Sciences / HMS-MIT HST / MGH Anesthesia & Critical Care

Some neural systems are well modeled by purely discrete variables, such as melatonin release, which switches between two rates: high and low. Other neural systems are well described by purely continuous variables, such as circadian modulation of core body temperature. However, in several cases, the behavior of the neural system may reflect a hybrid model that mixes continuous and discrete variables. Continuous concentration gradients are spent across four-state ion channels. Hippocampal place cell receptive fields evolve through a continuous set of parameters, but switch between periods of rapid and slow change. Each of the five stages of sleep corresponds to distinct continuous oscillatory patterns in EEG.

We are interested in relating spiking activity to hybrid discrete/continuous models of neural systems. Previously, estimation procedures had been developed for relating spiking activity to continuous or discrete variable models separately. Here, we address the hybrid case with the theory of point process statistics. Applications to neural data analysis and the design of brain-driven interfaces are illustrated.

164. Quantifying The Linear And Nonlinear Components Of A Neuronal Response Using Matching Pursuit Regression With A Redundant Dictionary Of Kernels

Pramodsingh H Thakur, Paul J Fitzgerald, Sung S Kim, Steven S Hsiao

Johns Hopkins University

The sensory signals evoked at the periphery of any sensory modality undergo a number of transformations as they traverse through multiple stages towards the higher cortical areas. At every stage, neurons extract features out of the incoming signals that are critical for perception. What are these features and how are they represented in the nervous system is a central question in systems neuroscience. Earlier attempts such as reverse correlation or generalized reverse correlation/regression involved modeling a linear system between the sensory input and the neuronal responses. These methods work moderately well for the early stages but fail miserably when applied to the responses of higher processing stages, which are highly nonlinear. Attempts on nonlinear system identification of neuronal responses have been limited because of the dimensional explosion caused by adding nonlinear terms in the regression setup.

We have developed an iterative algorithm, based on the matching pursuit procedure that picks up, at every step, a kernel that best captures the variability in the neuronal response. This adaptive procedure enables us to test the neuron's response against a set of linear and nonlinear kernels larger than what would be possible in a classical regression procedure. In this procedure, at every step the current residuals are projected onto the candidate kernels, and the kernel that picks up the maximum projection is selected. At the end of every step, the selected kernel is subtracted out from the residual and the remaining kernels in the dictionary to keep the space orthogonal to the space spanned by the selected kernels in the next iteration. The iterations are terminated when the projection of the residual onto the kernels is not significantly greater than chance.

We applied this algorithm to neuronal responses obtained from the primary (SI) and secondary (SII) somatosensory cortices of the awake macaque monkey. Responses were recorded while scanning random combinations of shapes and curves across the distal pads of digits 2, 3, and 4. We used a dictionary of linear (linear weights spanning 1mm x 1mm each) and nonlinear (specific orientation and curvature weights spanning 1mm x 1mm each) kernels located over a 40mm x 30mm window surrounding the three distal pads stimulated. In area 3b most of the kernels selected were located on single finger pad, which indicates the robustness of the algorithm. Area 3b neurons were best described by a mixture of linear as well as nonlinear kernels, while area 1 and SII neurons mostly picked up nonlinear kernels.

165. Distinct roles of synaptic connectivity and refractoriness on the spike-based and rate-based population decoding

Taro Toyoizumi[1], Kazuyuki Aihara[2], Shun-ichi Amari[3]

[1]Department of Complexity Science and Engineering, University of Tokyo, [2]Institute of Industrial Science, University of Tokyo, [3]RIKEN Brain Science Institute

We evaluate the Fisher information of a population of model neurons which receive dynamical input and interact with each other via spikes. With spatially independent threshold noise, the spike-based Fisher information that characterizes the information in individual spike-timings has a particularly simple analytical form. We calculate the information loss of abandoning spike-timings and study the effect of the inter-spike interval distribution and recurrent connections on the Fisher information. For a simple spatiotemporal input, we derive the optimal recurrent connectivity that has a local excitation and global inhibition structure. The optimal synaptic connections depend on the spatial or temporal modality of the input that the system is designed to code.

The role of precise spike-timings in addition to the spike count information [1] and the correlation between neurons [2, 3] to the information transmission are fundamental problems in neuroscience. Although it is conceptually possible to evaluate the information in a real population of neurons, this is practically difficult because of the large amount of data to be required. It is also difficult, with this approach, to understand the role of neuronal parameters such as recurrent connectivity on the information coding. In this study, we calculate the Fisher information of a network model of spiking neurons. In contrast to the literature, where Fisher information is evaluated for the firing rates of neurons (rate-based Fisher information) [2, 3], we evaluate the Fisher information given that we observe individual spike-timings of all the neurons (spike-based Fisher information).

We evaluate the amount of information in the precise timings of spikes by comparing the spike-based and rate-based Fisher information, and study the role of inter-spike interval distribution and synaptic connectivity on stimulus estimation. Due to the refractory factor that depends on inter-spike intervals, the information loss of abandoning individual spike-timings is considerable for a high firing rate regime. Moreover, the information loss is also large when two Poisson neurons are connected with a strong synaptic connection.

The effect of recurrent connectivity on stimulus estimation has been studied with a local excitation and global inhibition structure, but only within the rate decoding framework [4, 5]. We apply the spike decoding paradigm by taking account of individual spike timings and calculate the optimal recurrent connections that maximize the spike-based Fisher information. Evaluating the spike-based Fisher information, we can identify the information in individual spike-timings that code temporal features of input, which are difficult to be decoded from spike rates in a wide temporal window. Even for the same spatiotemporal stimulus, the optimal recurrent connections depend on a stimulus feature which the system is optimized to code.

[1] S. Panzeri, R. Peterson, S. Schultz, M. Lebedev, and M. Diamond, Neuron 29, 769 (2001).

[2] L. F. Abbott and P. Dayan, Neural Computation 11, 91 (1999).

[3] S. Wu, S. Amari, and H. Nakahara, Neural Computation 14, 999 (2002).

[4] P. Series, P. E. Latham, and A. Pouget, Nature Neurosci. 7, 1129 (2004).

[5] M. Spiridon and W. Gerstner, Network: Computation in Neural Systems 12, 409 (2001).

166. Sniffing cycle-based odor coding in the anterior olfactory cortex

Naoshige Uchida, Zachary F Mainen

Cold Spring Harbor Laboratory

Olfaction was often regarded as a slow sense but recent psychophysical studies in rodents show highly accurate odor discrimination can be achieved with ~300 ms of odor sampling or with one sniff cycle with theta (7-8 Hz) frequency (Uchida and Mainen, 2003). Thus, olfactory information can be processed very rapidly or within one sniff cycle, but the mechanisms underlying rapid odor coding remains to be elucidated. Furthermore, little study has been done to characterize olfactory neuronal responses in awake behaving animals. In this study, we explored the nature of olfactory processing in an olfactory cortical area while the rat is engaged in an active olfactory sampling.

The anterior olfactory cortex (known as anterior olfactory nucleus, AON) is the first cortical structure which receives direct projection of the mitral/tufted cells in the olfactory bulb, and which projects further to other olfactory cortices. Using tetrodes, we performed simultaneous recordings of neurons in the anterior olfactory cortex, while the rat was performing a two-alternative olfactory discrimination task. Sniffing was simultaneously recorded using a thermocouple implanted in the nostril. Rats were trained to respond to one of the two response ports after sampling an odor at the central odor port. For a correct choice, rats were rewarded with a drop of water. During the recording sessions, four or six odors were used, and half of them are assigned to either the left or right response ports.

Firing of neurons in the AON was tightly locked to the sniffing cycles. Some neurons exhibited respiration-locked firing upon odor sampling even in the absence of an odor. Odor responses could be characterized by an increase or a decrease of firing with respect to “baseline” firing. Interestingly, these responses often occurred within a short (~50 ms) time window with a fixed latency after inhalation onset, and the latency was often odor-specific. Such transient responses typically consisted of one or two of spikes per sniff cycle and were only evident when inhalation onset was considered. Odor-specific latency differences could be obtained even when overall spike counts (or firing rates) within a longer time window (e.g. one sniff cycle) did not show a gross difference. These results suggest that sniffing cycle-based spike timing in the olfactory cortex may play a role in rapid odor coding.

167. Differences in processing of low- and high-order image statistics revealed by classification images extracted via regularized regression

Jonathan D Victor, Ana A Ashurova, Mary M Conte

Weill Medical College of Cornell University

Classification images (CI's) provide an appealing psychophysical strategy to gain insight into the neural computations underlying visual perception (Ahumada & Lovell JASA 1971; Eckstein & Ahumada, JOV 2002). Determination of CI's typically requires analysis of thousands of psychophysical trials, and usually consists of "reverse correlation" of the psychophysical responses with the individual trial images. It thus presents a computational problem similar to that of receptive field analysis and functional brain imaging, both in terms of its highly multivariate nature and the potential for inaccuracy due to chance correlations within the multivariate dataset. Here, we determine CI's in a texture-discrimination task, and adopt strategies that have proven useful in imaging and receptive-field mapping in an attempt to improve on a reverse correlation analysis.

Subjects were asked to identify the location of a 16 x 64-pixel texture-defined target, which could appear at one of four positions within a 64 x 64-pixel background array. The target and the background texture were chosen from a two-dimensional perceptual space of binary textures. One axis in this space, γ, specifies the bias in luminance statistics (γ= 1 for all white, 0 for a 50:50 mix, -1 for all black). The second axis, α, specifies the bias in local fourth-order statistics (Julesz et al. 1978), with α=1 for the "even" texture, and -1 for the "odd" texture. γ and α vary independently within a gamut, and each pair uniquely specifies a Markov random field texture. We varied the parameters along eight rays emanating from the origin, over a range for which performance varied from just above chance to near ceiling. Data from 4320 such trials were collected from each of four subjects.

As previously reported (Victor, Chubb, & Conte, Vision Res. 2005) psychophysical performance along each ray was well-described by a Weibull function with exponent near 2. Sensitivity along the luminance axis (γ) was approximately four times the sensitivity along the fourth-order axis (α). For oblique directions, cues combined in a manner consistent with probability summation, leading to elliptical isodiscrimination contours.

CI's were determined after preprocessing each stimulus to create "derived images" that represented of pixel-by-pixel estimates of γ or α, reduction to a 32 x 32 grid, and alignment to the perceived target location. Standard reverse correlation yielded CI's that identified the footprint of the target but did not reveal internal structure. These CI's showed no difference between derived images based on γ or α . While regression alone did not yield useful CI's (the number of trials was only slightly larger than the number of grid points), regression combined with regularization identified consistent structure not seen in the reverse correlation CI's. In particular, for regularized CI's based on luminance statistics, there is an accentuation at the perceived target edge and attenuation in the interior of the target -- essentially a Mach band. Regularized CI's based on local fourth-order correlation show no interior scalloping within the target. That is, the regularized CI's suggest that low-order statistics contribute to texture segmentation via local comparisons, while high-order statistics are pooled over larger areas.

These findings were present with several kinds of regularization strategies (ridge regression (Hastie et al. 2001), penalized regression (Machens et al. 2004), and a GIFA-like approach (Yokoo et al., 2001)), alone and in combination, for a wide range of regularization parameters. They demonstrate the utility of such procedures in CI analysis, and that CI's reflecting nonlinear processes may be readily obtained via analysis of appropriately derived images.

168. An Adaptive Method of Spatiotemporal Receptive Field Estimation

Michael T Wahl

UC Berkeley Department of Physics

During natural vision the statistical properties of input stimuli are constantly changing. These changes can lead to an adjustment of the response properties of neurons in the early visual system, as measured by a spatiotemporal receptive field. While receptive field adaptations to changing mean luminance and contrast have been known for some time, recent experimental evidence indicates that receptive fields also adjust to varying spatial and temporal statistics. In particular, receptive fields measured using natural and artificial stimuli seem to be different (David et al. 2004). While traditional reverse correlation algorithms can measure a spatiotemporal receptive field given an input stimulus with statistics that are stationary for the duration of the trial, they cannot recover accurate receptive field measurements if the stimulus properties vary, and they cannot track dynamic changes to the receptive field during the course of the trial.

We implement a method to track dynamic changes in the spatiotemporal receptive fields of visual neurons using a recursive least-squares (RLS) algorithm. Such a method has been previously used with stationary stimuli due to its convergence properties (Shapley et al. 2002), and has been used to track receptive field changes to varying luminance and contrast (Lesica et al. 2003), but has not been implemented to track receptive field changes to varying spatiotemporal statistics. We perform simulations with a model neuron to test the performance of the RLS method using natural image stimuli, and demonstrate its ability to recover the model neuron's receptive field, in the case of stationary stimulus statistics. We then simulate a model neuron with response properties that change with different stimulus statistics, and observe the corresponding dynamic changes in the receptive field, as calculated by the RLS algorithm. Our results indicate that the algorithm can recover adaptive receptive fields with nonstationary stimuli. We believe such a method to be instrumental to understanding visual function during natural vision.

169. Maximum likelihood decoding of moving stimuli using divisive normalization line attractor neural networks

Robert C Wilson, Leif H Finkel

University of Pennsylvania

There is a growing body of evidence that the brain computes optimally. However, despite some excellent theoretical suggestions (e.g. Gold, J.I. and Shadlen, M.N., Nature, 404: 390-394, 2000; Deneve, S., NIPS. 13. 2004; Rao, R. Neural Computation, 16(1), 1-38, 2004; Zemel, R.S., Huys, Q.J.M., Natarajan, R and Dayan, P. NIPS 2004), precisely how the brain achieves this feat is still an open question. A particularly promising approach was suggested by Deneve et al. (Deneve, S., Latham, P.E. and Pouget, A. Nature Neuroscience. 2(8):740-745. 1999) where the authors used a line attractor neural network to optimally decode a static stimulus.

In this work we extend the network of Deneve et al. to deal with moving stimuli. We do this by connecting up a line attractor network in such a way that the stable state of the system is a rotating bump of activity whose speed can be altered by changing the input to different parts of the network. In particular we use a double ring line attractor network similar to that suggested in Zhang 1996 (Zhang K., J Neurosci. 16:2112-2126, 1996) and Xie et al. 2002 (Xie, X., Hahnloser, R.H.R., and Seung, H. S. , Physical Review E 66, 041902, 2002), but with the difference that the activation rule implements divisive normalization. We use the divisive normalization rule as this is both biologically plausible (e.g. see Carradini, M., Heeger, D.J., and Movsham, J.A., J. Neurosci. 17(21), 8621-8644, 1997) and leads to robust line attractors.

The aim of our network is to decode the position of a moving stimulus (e.g. a moving dot) given a noisy input signal as a function of time. The network then receives three types of inputs: the speed of the stimulus, which could be provided by MT cells; the external input comprising a noisy hill of activity centered around the current position of the stimulus; and the recurrent inputs from the network itself.

The figure below demonstrates the basic properties of our model. The moving stimulus creates a noisy input current (shown in the top left panel) whose position is determined by the position of the stimulus. In the present case, the noise is created by simply adding uncorrelated Gaussian noise to the input current at each time step. This input is fed into the divisive normalization network and the output of this network is shown in the top right panel. Clearly this is a much cleaner signal and the position can easily be decoded using a population vector decoder. This decoded signal is shown in the bottom left panel. The green line represents population vector decoding of the input signal; the blue line is the decoding from a divisive normalization network, but with the speed set to zero ? which is equivalent to the network used in Deneve et al.; the red line is the decoding using our network; and the dashed black line is the actual motion. Clearly our network significantly outperforms the other decoding methods.

The results of a more detailed empirical investigation are shown in the bottom right panel. In these experiments we compared the network decoding (?Network?) to population vector (?PV?) decoding; and to two versions of maximum likelihood decoding (ML and ML+v). The first, ML, is the optimal decoding of the stimulus position without taking into account any velocity information, i.e. this estimator makes no assumptions that the position of the stimulus at different time points will be correlated in any way. The second estimator, ML+v, assumes that the stimulus is moving at speed v and finds the optimal straight line path of speed v through the activity. This second version is the best possible decoding of the inputs that the network receives. In the figure we plot the rms error over 20 different trials each lasting 100 time steps for these different decoding schemes. These results clearly show that our network performs close to optimally and that the ability of the line attractor network to integrate information over time gives it an advantage over the uncorrelated, ML decoding scheme.

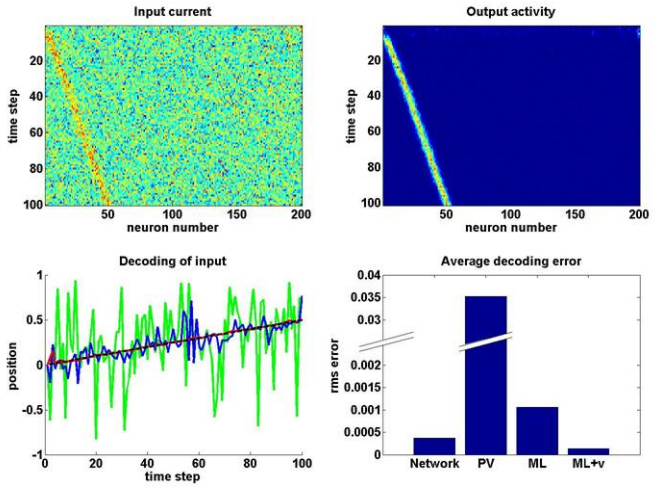

170. Behaviorally-dependent information processing in a songbird circuit required for vocal plasticity

Brian D Wright[1], Mimi H Kao[2], Allison J Doupe[3]

[1]UCSF, Sloan-Swartz Center for Theoretical Neurobiology, [2]UCSF, Dept. of Physiology, [3]UCSF, Depts. of Physiology and Psychiatry

The songbird anterior forebrain pathway (AFP) is a basal ganglia-forebrain circuit critical for vocal plasticity. The AFP has been hypothesized either to guide adaptive changes in song or to generate motor variability required for reinforcement learning. In particular, the output nucleus of the AFP, LMAN, provides input to the motor nucleus RA and modulates song on a moment by moment basis [1]. Recent work from our lab on singing-related activity in LMAN of adult zebra finches has shown a dramatic difference in firing properties depending on social context. Firing rates in LMAN are higher and activity is more variable across song renditions during undirected singing (when the bird is alone), than during directed singing to a female. Activity in LMAN also contains significantly more spike bursts in the undirected context. We seek to determine the nature of these large context-dependent differences in variability and reproducibility in LMAN activity during singing. We ask how LMAN activity encodes the ongoing song in the two conditions and how efficiently it does so. Moreover, we ask whether LMAN activity during undirected singing is simply a noisy version of activity during directed singing.

To investigate these questions, we use two representations of the LMAN activity. To understand how well the instantaneous firing rate encodes the ongoing song, we take the response to be the single spike events themselves. To see how information might be present in temporal patterns in the spike train, we also look at "words" consisting of sections of the spike train lying within a sliding time window of duration T. For each of these representations we compute the mutual information between the LMAN activity and the ongoing song in both social contexts. By comparing the information in spike train words at fixed T and the information in single spike events, we obtain an estimate of the information that may be present in temporal patterns in the spike train.

Our information-theoretic analysis leads to several answers to the above questions. First, the information carried by the instantaneous firing rate is greater during directed than undirected song. The spike train information, however, estimated for words with T=16ms, is quite similar in the two conditions. In addition, the information efficiency -- the ratio of mutual information to spike train entropy -- is nearly equal across conditions. This near equality of the efficiency of the spike trains in the two conditions occurs despite the greater variability in the undirected case. Therefore, the undirected variability does not merely add noise to the directed pattern of activity, but provides useful information about the ongoing song. Second, when we compare the mutual information in spike train words to that in single spike events, we find that the former was significantly larger for all cells in the undirected case and for 80% of cells in the directed case. This difference between words and single spike events means that temporal patterns in the spike train are contributing information about the song beyond that conveyed by the individual spikes making up the patterns. Moreover, for all cells, the proportion of information not accounted for by the instantaneous firing rate was significantly higher in the undirected condition. Together, our results suggest that the mechanism of neural encoding rapidly changes with social context from primarily a temporally precise rate code in the directed case, to more of a temporal code of extended spike patterns in the undirected case. And while LMAN switches its encoding scheme with social context, it does so in such a way that it preserves the total information it encodes about the ongoing song.

Work supported by the Swartz Foundation (B.D.W.), an HHMI predoctoral fellowship (M.H.K.), the MacArthur Foundation, the Steven and Michele Kirsch Foundation, NARSAD, the Sandler Family Supporting Foundation, and the NIH (MH55987; A.J.D.)

[1] Kao, M.H., Doupe, A.J. and Brainard, M.S. (2005). Contributions of an avian basal-ganglia forebrain circuit to real-time modulation of song. Nature 433, 638-643.

171. A Simulation of the Uniform Selection Hypothesis in a Delayed-response Task

Ahmed A. Moustafa[1], Anthony S. Maida[2]

[1]Institute of Cognitive Science, University of Louisiana at Lafayette, Lafayette, LA 70504, [2]Center for Advanced Computer Studies and Institute of Cognitive Science, University of Louisiana at Lafayette, Lafayette, LA 70504

Delayed-response tasks (DRTs) have been used to assess working memory (WM) processes in human and nonhuman animals. It is hypothesized that the basal ganglia (BG) subserve selection of both motor- and cognitive-related information to perform these tasks (Berns & Sejnowski, 1996). We report the results of simulation studies of a systems-level model of DRT performance that uses an actor/critic architecture. The matrisomes of the BG represent the actor and the striosomes the critic. The model incorporates segregated closed- and open-loop pathways between the BG and dorsolateral prefrontal cortex (DLPFC) (Middleton & Strick, 2000, 2002; Joel & Weiner, 1994). They respectively select task-relevant cognitive information to be maintained in WM and also select appropriate motor responses. Training for both functions is based on the reward-based temporal difference algorithm in the critic and a 3-factor Hebbian rule in the actor (Reynolds & Wickens, 2002; Wickens et al., 1996).

A novel feature of the model is the incorporation of delay-active neurons in the matrisomes (as well as DLPFC). Another novel feature of the model is the subdivision of the matrisomal neurons into segregated winner-take-all (WTA) networks consisting of delay- versus transiently-active units. The delay-active units gate cognitive information into WM and the transiently-active units select motor actions. This subdivision arises from functionally identical units being embedded in closed- versus open-loop pathways, respectively. Since the processing for both the closed- and open-loop subsystems is identical, we say that the model embodies the uniform selection hypothesis for both the cognitive and motoric aspects of the DRT.

By monitoring the performance of the model and tracking changes in connection strength during learning, we show that the model accounts for essential aspects of DRT performance and also explains the underlying mechanism. Particulars are given below.

First, the results show that the dopamine (DA) phasic signal shifted from the time of reward (unconditioned stimulus) to that of the trigger and then to the cue (conditioned stimuli), a phenomenon known as time shifting of DA phasic signals (Schultz et al., 1993; Braver & Cohen, 2000). Further, Fiorillo et al. (2003) recently found that DA phasic responses code for the certainty or uncertainty of receiving reward and Tobler et al. (2005) found that DA phasic responses also code for reward magnitude. Our simulation results account for these phenomena as well.

We have also conducted lesion studies of the model to simulate the effects of DA depletion in Parkinson's disease or focal damage to the striatum. The model accounts for findings that lesioning of the striatum (as well as DLPFC) leads to perseverative errors (Fung et al., 1997; Levitt et al., 2002, Picket at al., 1998) and that DA depletion may either lead to perseveration (Cooper et al., 1991; Lees & Smith, 1983) or to random responses (Taylor et al., 1986).

172. Tetrode recordings in dorsal and median raphe nuclei in awake behaving rats

Sachin P Ranade, Zachary F Mainen

CSHL

Serotonin is an important neuromodulator implicated in a range of behavioral, cognitive and psychiatric disorders. Serotonergic neurons are located in a set of nuclei in the midbrain called raphe nuclei that arborize the entire neuraxis. Neurons in the dorsal and median raphe nuclei send projections to the forebrain. While much knowledge of serotonergic function has been gained through pharmacology there are only a few recordings studies from raphe nuclei in awake animals which show that raphe neurons modulate their firing with the sleep-wake-arousal cycle and increase their activity during rhythmic behaviors like licking, whisking, breathing, locomotion, sniffing etc. However, these studies were performed in awake freely behaving animals and hence have little behavioral control, imprecise timing infrormation and minimal statistics. Therefore they may not capture fast phasic changes in firing patterns. Much insight has been gained into understanding dopamine function from dopamine neuron recordings in monkeys performing a reinforcement learning task. However, there are no existing recordings from raphe neurons during a task. In order to study response properties of raphe neurons in relation to behavior we have recorded from these neurons using multielectrode techniques in animals engaged in a controlled behavioral task.

Extensive pharmacological and behavioral data indicates that serotonin is involved in reinforcement learning. It has been noted that dopaminergic and serotonergic agents have opposing effects on behavior. While dopamine is involved in positive reinforcement, serotonin affects learning about negative reinforcers or punishment. Schultz and colleagues have shown through a series of experiments that dopamine neuronal firing can encode various reward parameters. We hypothesize that raphe neuronal firing will encode negative outcomes during reinforcement learning(Daw, Kakade and Dayan 2001).

In order to study responses of raphe neurons during goal directed behavior, rats were trained on a 2 odor olfactory discrimination task using enantiomers of Octanol. Water deprived rats were trained to poke into a center port where on a given trial one odor from the pair was delivered. Each odor was associated with availability of water at one of two choice ports (Left and Right). Depending upon odor identity the rat has to choose the correct port to receive water reward. With this task we can study correlations in firing of raphe neurons during various aspects of behavior including sensory, motor, decision and reward processes. In order to study the responses of raphe neurons to negative outcomes we introduced probabilistic reinforcement. In this modified task, we omitted reward on a small proportion of randomly interleaved correct responses. Omission of expected reward is a negative outcome and we hypothesize that raphe neurons will show firing rate modulation in response to this.

In order to record from raphe nuclei we have developed a procedure for tetrode drive implantation. This involves placing a guide cannula through the midline sagittal sinus upto a depth of 3mm and lowering the drive through it into the brain. This technique allows us to accurately localize the fine wire tetrodes to the dorsal and median raphe.

We first recorded neurons from dorsal and median raphe in awake freely behaving rats to determine their firing properties. Neurons were classified on the basis of firing rate, spike width and rhythmic modulation. A subset of high firing rate neurons showed rhythmic modulation at theta, beta and/or gamma frequencies. A very small proportion of low firing rate neurons (<10% of total population) had wide spikes (> 400 ms peak to valley), similar to serotonergic neurons reported in literature.

We have also recorded activity of raphe neurons during 2 odor discrimination task with probabilistic reinforcement. Behavioral correlations of raphe neuron firing during various phases of the task will be presented.

1. Daw ND, Kakade S, Dayan P., Neural Network 2002 Jun-Jul;15(4-6):603-16

173. A Credit Assignment Algorithm for Composite Visuo-Motor Behaviors

Constantin A Rothkopf, Dana H Ballard

University of Rochester

Mounting experimental evidence suggests that the brain may learn its behaviors through secondary reward mechanisms, using neurotransmitters such as dopamine [1]. These results are given further credence by the development of reinforcement learning algorithms such as temporal difference learning [2]. However all such algorithms suffer in that straightforward implementations scale exponentially with increasing dimensions, making them infeasible candidates for high-dimensional state spaces. A solution to this problem is to attempt to realize high dimensional state spaces with compositions of lower dimensional state spaces. Taking such an approach allows composite visuo-motor behaviors to be synthesized from simpler behaviors. For example in walking down a sidewalk, a pedestrian may have goals of staying on the sidewalk, avoiding pedestrians and picking up litter. Sprague [3], showed in a virtual reality simulation that these individual behaviors can be learned in combination by reinforcement learning. However that simulation assumed that the rewards associated with the individual behaviors were known. In practice this is an unreasonable assumption for biological systems where only the total reward for the composite behavior is likely to be available. In this case, one would like to learn the contributions from the different behaviors from the obtained total reward. This is a long-standing credit assignment problem in learning.

Kaebling et al [4] showed that an estimate for the individual rewards could be obtained by assuming that the total reward was assigned to each behavior and the variations in that reward were assumed to be noise. By using a Kalman filter, a Bayes optimal estimate of the contribution of a single agent to the total gained reward could be obtained over trials. This model made sense in their setting, which had the individual behaviors embedded in different agents, but introduced two problems in that the resultant reward estimates could have a constant bias and be suboptimal.

We show that the credit assignment problem has a solution when all the behaviors are assumed to be embedded in the same agent. This allows us to model reward not as a globally broadcasted number, but as a consumable entity. The difference is that in the latter case, the individual reward estimates must add up to the total reward estimate whereas in the former case the reward not accounted from the perspective of an individual behavior is assumed to be entirely noise. In our algorithm, each behavior needs to know which subset of other behaviors is simultaneously active. It can then keep a running estimate of its share as its current estimate adjusted by the total instantaneous reward minus the estimates of the concurrent behaviors. Our simulations show that when the order that the behaviors update is chosen randomly, the estimated reward for each behavior converges to its true value. Accordingly, the update rule for the Q-values for each agent is:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[\tilde{r} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right]$$

subject to:

$$E = \frac{1}{2} \left[\left(g - \sum_{j \neq i} \hat{r}_j \right) - \hat{r}_i \right]^2 = 0$$

Simulations on common grid world problems will demonstrate the capabilities of the algorithm. Moreover, we show how a virtual pedestrian that has to solve the three concurrent tasks mentioned above can learn the correct value functions for each behavior.

1. Schultz, W.: Nature Rev. Neurosci. 1: 199-207, 2000

2. Sutton and Barto, Reinforcement Learning: An Introduction, MIT Press 1998

3. Nathan Sprague and Dana Ballard, NIPS 2003

4. Yu-Han Chang, Tracey Ho, Leslie Pack Kaelbling, NIPS 2003

This work was supported by NIH grants EY05729 and RR09283

174. State-dependent matching law in stochastic gradient ascent

Yutaka Sakai[1], Tomoki Fukai[2]

[1]Human Informatics Course, Faculty of Engineering, Tamagawa University, Japan, [2]Brain Science Institute, RIKEN, Japan

The ability of making a correct choice behavior among various options is crucial for animals' survival. How animals or humans organize their behavior depends crucially on the expected return which may result from their choices. It is widely considered that they attempt to maximize the obtainable reward, and a similar concept underlies several efficient algorithms in the machine learning (Sutton and Barto 1998). Extensive studies have been conducted to clarify whether animal's or human's behavior does obey this conceptual rule. In fact, it is known that animal's and human's behavior often obey another rule, called as "matching law", in which the frequency of choosing each option is proportional to the amount of the past reward obtained by the choice (Herrnstein 1997; Mazur 2004). Interestingly, in many cases the matching behavior gives only a sub-optimal solution to decision making. The strategies in learning to attain the matching behavior have been proposed in literature. However, the computational implications of the matching behavior are not necessarily obvious in a generic sense. In addition, immediate rewards are assumed in the framework to concern the issue, while an actual reward is often obtained after a delay from the responsible choice in natural environments. Animal's behavior in a situation of delayed rewards and the underlying mechanisms remain unknown.

Here, we extend the matching law to incorporate state transitions and state-dependent internal values, which are employed in the reinforcement learning in order to maximize rewards which may be delivered after delays. We show that stochastic gradient ascent on the assumption of Markov decision process (MDP) always exhibits the matching behavior in non-Markov situations when the learning attains the steady state of long-term averages. As noted previously, many reinforcement learning algorithms can be derived from stochastic gradient ascent in MDP. The learning algorithms proposed to attain the matching behavior can also be derived from stochastic gradient ascent in MDP. The present result suggests that gradual maximizing of rewards on the assumption of Markovian environment develops the matching behavior in non-Markovian environments.

175. Different subregions of human striatum encode appetitive and aversive outcomes in mixed prospect predictive learning of money.

Ben Seymour[1], Nathaniel Daw[2], Peter Dayan[2], Tania Singer[3], Ray Dolan[1]

[1]Wellcome Department of Imaging Neuroscience, UCL, [2]Gatsby Compuational Neuroscience Unit, UCL, [3]Institute of Cognitive Neuroscience, UCL

There is an impressive wealth of neuro-imaging studies indicating how information about rewards and punishments is acquired, represented and processed in the human brain. Abstract rewards, such as money, elicit similar responses to those elicited by primary rewards, such as taste and olfactory stimuli, in both cases with regions of the striatum reporting on temporally sophisticated prediction errors (called TD errors). Indeed, there is by now a prevailing notion that the striatum is involved in general appetitive learning, a conclusion that fits comfortably with a wide swathe of animal data regarding the role of dopamine release in the striatum in basic reward learning.

This leads naturally to the questions of what activity is caused by abstract and concrete punishments, and how this interacts with reward processing. One main finding is that information about painful stimuli, or more specifically, prediction errors thereof, elicits very similar fMRI responses in the striatum to prediction errors for rewards. This could reflect a failing of the BOLD signal to distinguish appetitive and aversive signals that are spatially (Reynolds and Berridge, J Neurosci. 2002: 22: 7308-20) and/or neurochemically (Daw et al., Neural Networks 2002: 15:603-16) separate. Such possibilities reflect ambiguity in how the TD error theory should be mapped on the neural substrate in light of appetitive-aversive interactions.

We therefore designed a Pavlovian conditioning task involving both appetitive and aversive outcomes, in this case gains and losses of money. Twenty four healthy subjects underwent (3 second) delay conditioning of neutral arbitrary visual cues to money outcomes – indicated by a photographic image of their win or loss. We employed 5 cues, each with 50:50 probabilistic outcomes, as such: Univalent reward cue - reinforced with £1 / £0; univalent loss cue – reinforced with £-1 / £0; neutral cue – reinforced with £0 / £0; bivalent cue (£1) – reinforced with £1 / £-1; bivalent cue (50p) – reinforced with 50p/-50p.

Behavioural measures demonstrated clear conditioning to the cues: in a forced choice cue preference task performed after conditioning, subjects preferred the univalent reward cue over the neutral cue, which was preferred over the univalent loss predicting cue. Autonomic measures (pupillometry) and reaction time data also distinguished the neutral cue from the univalent gain / loss predicting cue. Bivalent cues were preferred slightly over the neutral cues, but less than the univalent gain predicting cue.

Simple trial-by-trial contrasts of brain imaging data revealed differential responses to identical outcomes, depending on the alternative potential outcome cued. Thus, contrasted between trials with bivalent and univalent cues, appetitive (£1 gain) and aversive (£1 loss) outcomes activated adjacent regions of ventral striatum. Such signals are consistent with appetitive and aversive prediction errors, respectively, and so we further explored this activity using a computational reinforcement learning model, which is known to provide an accurate account of both primary reward and punishment learning. This analysis revealed that a trial-by-trial prediction error signal correlated with BOLD activation in an area of the anterior ventral striatum (on trials with univalent and bivalent cues) and also in an area of the dorsal putamen (during bivalent trials only). The former result was suggestive of an appetitive prediction error, in that positive and negative BOLD excursions are associated with unexpectedly high or low payoffs, respectively; the latter might possibly represent an aversive prediction error. Finally, the absolute value of such an error signal (in which both unexpected gains and losses are associated with positive BOLD responses) correlated with an area of ventral striatum, slightly posterior than that associated with the purely appetitive signal.

These data provide evidence that both appetitive and aversive prediction error related activity is co-expressed in the striatum in learning about financial rewards and losses, dissociable by their expression in distinct sub-regions. It may be that more anterior regions of the ventral striatum appear to code selectively for appetitive outcomes, whereas aversive outcomes are coded more posteriorly, in keeping with animal studies. This may shed light on theories that cast striatal responses in terms of general affective salience, as opposed to valence specific, representations. The involvement of distinct appetitive and aversive systems in the representation of uncertain mixed outcome predictions has important implications for human decision making, since such basic predictive learning processes are thought to guide action selection and learning as an agent learns to exploit its environment.

176. From Chaos to Self-organization, and from Firing Fields to Place Fields. A New Hypothesis on Hippocampal Neurodynamics

Renan Vitral

NIPAN. Department of Physiology. Biological Sciences Institute. Federal University of Juiz de Fora, BR.

The past and recent studies on the development of Place Fields on the Hippocampal system sustain two factors as essentials for its development: the first one would be the presence of Theta and/or Theta/Gamma oscillations on hippocampus EEG. The second would be the attentive behavior used to explore the environment.

For the purpose of this discussion and new hypothetic proposals we raise two key questions. The first one is: Would be the called "Theta Band" computed as linear? Within this context we could ask for: 1) a dynamical pattern for definition; 2) the spike timing of single cells; 3) how would burst activity computed; and finally 4) how many brain regions working at this band would be discharging on hippocampal targets. The second question is: What are the called Place Fields? Here we argue for some definitions that we reason remain still open: 1) what the Place Fields really represent; 2) how previous conditions would interfere with the Place Fields development, like displacement for the testing chamber where the following factors certainly would be important: stress, attention and visual searching, similarities between testing chamber and cage (walls, for example). Focusing on oscillations we have also to consider factors like timing required for generate spikes from specific visual stimulus and coincidence detection, which is a difficult matter to be solved.

After these considerations we are able to ask: Can we adopt the sequence "EC → DG → CA3 → CA1 → EC + Septum" as a linear pathway for memory storage? Could progressively compressed topographic connections from neocortex be the substrate to build specific memory pathways at the hippocampal system? Could these 'memories' arrive on specific neocortical targets? Would be the hippocampus a system working through a non-dynamical pattern?

Here we call attention for the whole allocentric influences on cerebral mechanisms for update direction vectors underlying visual navigation. We reason the most important as: 1) head direction cells and systems; 2) oculo-motor and head direction continuous activities; 3) the role of vestibular and sensorimotor afferents upon the development of self-organizing networks (Place Fields); 4) the recently discovery of distinct afferent connections for hippocampal subfields and also for interneurons; and finally, 5) the role of integrated control of saccades at this context.

The Behavioral and Neurobiological data, Neurophysiologic Recording Techniques and also other Functional methods for memory evaluation sustain my hypothesis that environmental visual information is not the first one arriving on the hippocampus in a visual navigation task. I suggest that when the summation of non-visual information (head direction, vestibular, oculo-motor, sensorimotor, limbic, etc.) reaches distinct sub-fields and interneurons of the hippocampus, firing on theta band, there is the development of self-organizing networks that would build firing fields on each hippocampal subfield. Since visual signals from environment follow through I.T. region, reach E.Cx. and arrive on hippocampus, these would sustain or not the 'firing fields' already built, depending on the coherence with the intrinsic hippocampal theta oscillations. Here we would have the "Firing Fields" transformed on "Place Fields", i.e., Firing Fields on Pyramidal Cells → Place Fields on Place Cells.

This hypothesis, although yet to be tested, is coherent with the many functional attributes devoted to the HIPPOCAMPUS. Besides, it makes many predictions that certainly will be experimentally checked, using probably multi-scale techniques with a more appropriate temporal scaling and the incorporation of the mathematics of complex, chaotic systems for the understanding of the cerebral code underlying oscillations.

This study is supported by CNPq, CAPES, FAPEMIG, UFJF and NIPAN.

References - Our Proposed Model

Vitral, RWF. 2005. A model for the generation of firing fields on hippocampal CA3 place cells using opponent processing on a self-organizing neural network. Proceedings of the Cosyne05/Nature Neuroscience Conference, #272. Salt Lake City. USA.

Vitral, RWF. 2004. Hippocampal dynamics in a fixed environment studied with opponent processing added to a self organizing functional architecture. Representation of CA3 place cells firing field, hippocampal oscillations and behavioral correlates. Soc. Neurosci. Abstr. San Diego. USA.

Vitral, RWF. 2004. Toward an integrated understanding of the generation of the place fields in the different sub regions of the hippocampal region. Proceedings of the Annual Conference of Cognitive Science Society. Chicago. USA.

Vitral, RWF. 2004. CA3 place cells in a fixed environment: place fields, hippocampal oscillations, behavioral correlates and opponent processing. Proceedings of the Eighth International Conference on Cognitive and Neural Systems. Boston. USA.

Vitral, RWF & Araújo, GF. 2003. Why do the hippocampal sub-fields develop “Place Fields”? Proceedings of the Seventh International Conference on Cognitive and Neural Systems. Boston. USA.

177. Time integration in a perceptual decision task: adding and subtracting brief pulses of evidence in a recurrent cortical network model

Kong-Fatt Wong[1], Alexander C Huk[2], Michael N Shadlen[3], Xiao-Jing Wang[1]

[1]Brandeis University, [2]University of Texas at Austin, [3]University of Washington, Seattle

In a visual motion two-alternative forced-choice response time task, neural activity in the lateral intraparietal (LIP) cortex of monkey's brain is observed to ramp up for hundreds of milliseconds [1], suggesting a possible neural mechanism to temporally integrate visual motion information before a perceptual decision is made. A recent experiment [2] tested this hypothesis by considering how a brief motion pulse of 100ms duration affects the motion discrimination. In the present work, we used a neural network model to simulate this experiment. The model [3] is a reduced version of a biophysical network model of spiking neurons [4], with slow excitatory reverberation (attractor dynamics) and winner-take-all-competition. Our model reproduces the psychometric functions and the response times in the presence of a brief motion pulse during stimulus presentation. It also captures LIP activities, which, depending on whether the direction of the pulse is the same or opposite to that of motion stimulus, respectively increase or decrease persistently over a few hundred milliseconds. Furthermore, the pulse exerts a weaker influence on LIP neuronal responses when the onset of pulse is late relative to motion stimulus onset. Our attractor model accounts for this violation of time-shift invariance, but this effect is not significant in a diffusion (`noisy ramp to a fixed threshold') model [2]. This suggests that the neuronal attractor model better captures the time integration of LIP neurons. We further examine the characteristics of time integration in our neural attractor model, using two pulses of the same or opposite motion directions that are presented at different times in a single trial. Our modeling results shed insights into the microrcircuit mechanisms of decision making in LIP.

[1] J.D. Roitman, M.N. Shadlen, J. Neurosci., 22, 9475-9489 (2002)

[2] A.C. Huk, M.N. Shadlen, J. Neurosci., 25, 10420-10436 (2005)

[3] K.-F. Wong, X.-J. Wang, J. Neurosci., In Press (2006)

[4] X.-J. Wang, Neuron, 36, 955-968 (2002)

178. Bidirectional spike-timing dependent plasticity of inhibitory transmission in the hippocampus

Jake Ormond, Melanie A Woodin

University of Toronto

The mammalian hippocampus, owing to its crucial role in memory, has been the primary focus of research into synaptic plasticity. Most studies have examined plasticity at excitatory (glutamatergic) synapses, despite the fact that neuronal output is not determined by the level of excitatory transmission alone, but by the levels of coincident excitatory and inhibitory transmission. In this study, we examined plasticity of GABAA receptor mediated inhibitory transmission in area CA1 of hippocampal slices from mature rats (6-9 weeks). Because GABAA receptor currents are mediated primarily by chloride, we used a combination of gramicidin perforated patch recording to preserve intracellular chloride, as well as conventional whole-cell recording to manipulate intracellular chloride. We initially determined the reversal potential for GABAAR currents (~ 22 mV negative to action potential threshold) and the intracellular somatic chloride concentration (15 mM). We then investigated whether a spike-timing protocol could induce changes in GABAAR reversal potential. Surprisingly, pairing of presynaptic stimulation with postsynaptic spiking led to bidirectional changes in the reversal potential, with the direction of change being dependent on the interval between pre- and post stimulation. When the postsynaptic neuron was made to fire action potentials 5 ms after presynaptic stimulation (correlated), at 5 Hz for 90 seconds, a depolarization of the reversal potential was seen. However, when the interval was lengthened to 100 ms (uncorrelated), a hyperpolarization of the reversal was seen. Due to the interplay between excitatory and inhibitory transmission, we suggest that this form of GABAergic plasticity may contribute to the enhancement of excitatory transmission under certain conditions.

179. Neural activities of resolution of state uncertainty in a partially-observable maze task

Wako Yoshida, Shin Ishii

Nara Institute of Science and Technology

We are often required to make decisions in situations where the available information is limited. True states in such uncertain environments are unobservable; the best we can do is to estimate them and then act on the basis of our `belief'. Recent theoretical studies have shown that information processing in the brain is achieved by supplementing uncertain observation with such belief, and we suggest here that this is also the case for higher cognitive systems and could involve the prefrontal cortex (PFC). According to recent imaging studies, the anterior PFC (APF) is activated when subjects needed to maintain multiple rule candidates and evaluate them using feedback for the used rule; namely, estimating unobservable (hidden) variables (in this case, hidden rules) using observed information. Although these studies led to the hypothesis that APF activation is involved in hidden variable estimation, testing this is not straightforward because of the existence of an inverse problem: a subject's internal state - what she or he believes - is itself hidden from the observer and cannot be determined unequivocally from behavioral data. We attack this issue by combining an fMRI experiment using a newly developed partially-observable maze task with a computational technique based on the probabilistic model inference.

In scanning experiments, 13 normal subjects performed goal-search task in a maze whose topography had been learned in the training task on the day prior to scanning. At the beginning of the task, a goal position was presented on a maze map. Then, the subjects explored the maze from an unknown starting position based on the observed wire-frame scene representing the wall status of six grids surrounding their hidden current position, which constitutes a partially observable environment. To achieve the maze goal efficiently, it is essential to estimate the current position constituting grid position and body orientation, and we assumed a subject's processing model based on a hidden Markov model (HMM), in which the subject maintains position candidates whose scenes are consistent with the observed one, chooses one of them as a position estimate, and then, selects an action based on the estimate. After given the next scene, the subject examines whether the previous estimation is correct or not by matching the predicted scene and the actual one. If the prediction is correct, the position estimate is confirmed and the candidates are temporally renewed such that they are consistent with the new scene (updating: UD). If there is a discrepancy between the scenes, however, the position candidates are re-enumerated, based on the history of scenes and actions of past trials, and a new estimate is chosen from those candidates (backtracking: BT).

Based on the above model and behavioral data, our HMM reproduced all possible sequences of the subject's position estimate and the elemental processes, UD and BT, and calculated the posterior probabilities of the sequences using the Bayes theorem. However, since the most probable sequence could not be unequivocally determined, we calculated the expected subjects' brain loads by averaging over all possible sequences with respect to their posterior probabilities. This technique called marginalization enables us to access two important factors for the position estimation in a stable manner. One was the position entropy which evaluates an integrative computational load for the estimation as the amount of information for position candidates. The other was the BT probability which evaluates a load for position re-enumeration. We then analyzed behavioral and imaging data from subjects, in relation to those two factors.

The subjects' reaction times (RTs) in the task exhibited a gradual decrease with increasing time and a number of transient peaks occurring primarily in the early phase. Interestingly, the RT peaks synchronized with the peaks of BT probabilities estimated by our HMM In addition, the smoothed RT, in which the peaks responsible for BT were removed, correlated significantly with the position entropy. The tight correspondences between these model-based factors with the RTs indicated that the processes of position estimation and action selection by the subjects were accurately estimated by our HMM. The precision of our simulation was also confirmed by the high accordance (84.4%) between the subjects' actions and those predicted by our HMM.

We subsequently conducted a regression analysis using two factors as regression functions. The medial PFC (mPFC; BA9/6) on the medial frontal gyrus was mainly activated with the BT probability. The mPFC is known to be co-activated with self-referential stimuli, and have a monitoring function of one's individual performance based on reward prediction. The processes performed during BT trials, such as error detection and consequent action selection, corresponded to the mPFC function, and our results suggest that the mPFC is responsible for evaluation of individual cognition as well as action. For the position entropy, we only observed activations of the bilateral APF (BA10) and left APF (BA9) and this result supported our hypothesis that the APF is activated during the resolution of partial observability. Recent imaging studies using subtraction analyses demonstrated APF activation during cognitive tasks with multiple rules, and those tasks can be regarded as optimal decision problems in partially-observable environments; as the correct rule was not explicitly provided by the environment, the subjects were required to examine the possibility of rule candidates. Our regression analysis has clarified that APF activities vary according to the position entropy, which is a measure of the conflict among candidates of the hidden state, adding a dynamic aspect to previous results.

In addition, our results provide evidence that activities in different regions of the PFC reflect critical computational components involved in decision-making in uncertain environments. This fits well with the proposed role of these regions in decision-making, which is likely to be crucial in complex real-world environments. We also illustrate the utility of statistical model-based inference and regression in delineating key task parameters that may be represented in spatially distinct brain regions.

180. Dissociation of accuracy and reaction time in a two alternative odor mixture discrimination task

Hatim A Zariwala, Naoshige Uchida, Adam Kepecs, Zachary F Mainen

Cold Spring Harbor Laboratory, Cold Spring Harbor, NY

We are interested in the neural mechanisms underlying olfactory-guided decisions. Reaction time experiments provide basic constraints for models of decision-making as well as for the timing of neural activity relevant for the decision process.

We use a two-alternative choice discrimination task in which rats associate different odors with one of two response ports. Because rats often discriminate even closely related pure odors almost perfectly, we use binary mixtures to increase task difficulty. Eight different odor mixture ratios were randomly interleaved from trial-to-trial and rats were rewarded according to the dominant component of the mixtures. In a reaction time (RT) version of this task rats performed such two odor discriminations in ~300 ms or 1-2 sniffs. Surprisingly, in this paradigm rats failed to increase their RT when faced with harder problems, prompting Uchida & Mainen (2003) to propose that there may be some limits on the ability of rats to integrate information in this problem.

Abraham et al (2004) performed a similar study (but in a "go/no-go" task) and reached somewhat different conclusions. While the mice used in their study also performed discriminations rapidly, they appeared to slow down for more difficult problems and performed somewhat more accurately on these problems than the rats in the Uchida & Mainen study. Khan & Sobel (2004) proposed an explanation for the discrepancy between these two results based on the concept of speed-accuracy tradeoff. The rats in Uchida and Mainen (2003) might have been motivated to work more quickly at a lower accuracy while the mice in Abraham et al. (2004) were working at a slower but more accurate level.

In the present study, we set out to test whether speed-accuracy tradeoff could account for the differences between these studies. Therefore we manipulated several parameters of the task in order to 'instruct' the rats to perform more accurately and/or more slowly.

In the first series of experiments, we varied several task parameters that might indirectly influence the motivation to perform quickly. To discourage fast guesses, negative reinforcement was introduced by delivering aversive air puffs for incorrect choices. This manipulation produced an over all increase in RT of ~100 ms and therefore was clearly causing rats to slow down. However, this slowing did not improve accuracy and was independent of problem difficulty. Similarly, introducing a fixed delay to the reward, fixing the overall reward rate in a session and punishing errors with long time outs all decreased RT without a significant effect on accuracy.

Because motivational manipulations failed to induce a dependence of accuracy on speed, we next developed a task variant in which response times were more directly instructed using an auditory "go signal" (Rinberg et al. 2005). Using this approach, reaction times could be varied systematically from between 300 to over 1000 ms. We randomly interleaved go-signals of different latencies within a single session but tested only a single odor mixture problem in each session. Go signal latencies were drawn from an exponential distribution in order to produce a flat hazard rate (uniform temporal expectation of go-signal). The results of this experiment showed that rats performed at similar accuracy over the entire range of odor sampling durations. Moreover, when the same rats were subsequently tested on the hardest problem using a RT task (mean RT 275 ms), the performance accuracy was as high as for the longest go-signals (88% +/- 2% on RT task and 90% +/- 2% for t_{go}>700ms, n = 4 rats)

When comparing the results of the go signal experiments and the original RT experiments, we noted that the overall performance was much higher on the most difficult problems (88% vs. ~60%). A pertinent difference is the fact that in the original RT experiments problems of different difficulty (different mixture ratios) were interleaved in a session, while in the latter only a single difficulty (single mixture ratio) was presented. This "stimulus context" effect can potentially explain the higher accuracy observed by Abraham et al (2004), who also used a non-interleaved task.

In summary, rats failed to show an effect of problem difficulty on RT, even when given strong incentives to do so. Moreover, performance failed to increase with longer RTs achieved by either motivation or explicit instruction. Taken together, these results are inconsistent with the idea that speed-accuracy tradeoff accounts for the different conclusions reached by previous studies. Rather, changes in stimulus context (single vs. interleaved problems) appear to be a critical but overlooked variable. Furthermore, the dissociation of discrimination accuracy from odor sampling time indicates that integration is not a fundamental component of the neural mechanisms involved in high accuracy odor mixture discriminations by rats.

Abraham et al: Neuron, 2004 Dec 2;44(5):865-76
Khan & Sobel: Neuron, 2004 Dec 2;44(5):744-7
Uchida & Mainen: Nat Neurosci., 2003 Nov; 6(11):1224-9
Rinberg Soc. Neurosci. Abstr., 2005

181. Conserving mean activity through adaptive inhibition leads to temporal sharpening when combined with Hebbian enhancement of excitatory connections

Samat B Moldakarimov[1], James L McClelland[2], Bard G Ermentrout[1]

[1]University of Pittsburgh, [2]Carnegie Mellon Unibersity

Experience in a task requiring precise timing of the interval between transient stimuli leads to coherent and short-lived neural responses, providing finer timing information than in inexperienced animals. Such temporal sharpening of responses may play a role in many contexts where high temporal resolution is required, including speech perception. We have built a model designed to explore the possible neural processing and connection adjustment mechanisms that might underlie these effects. Our model is based on a recurrent network of excitatory and inhibitory units. We connect units via modifiable synapses. We examine the conditions under which repeated presentations of a transient signal to an excitatory unit can “teach” the network to respond to the signal earlier, with higher amplitude and shorter duration. Within the constraints of the architecture considered, strengthening both excitatory->excitatory and inhibitory->excitatory connections is required to obtain temporal sharpening along with an increase in amplitude.

182. Discrimination Training and Neural Coding of Speech Sounds in Rat Primary Auditory Cortex

Crystal T Novitski, YeTing H Chen, Amanda C Puckett, Vikram Jakkamsetti, Claudia A Perez, Matthew S Perry, Ryan S Carraway, Michael P Kilgard

The University of Texas at Dallas

To document perceptual learning ability on speech sounds, rats were trained on a variety of speech discrimination tasks. Fifteen rats were trained using an operant Go/No Go procedure to discriminate between multiple natural monosyllabic word target (CS+) and distracter (CS-) sounds. Initial results indicate that large spectral differences were easy to discriminate (dad vs. sad), while subtle spectral differences were harder (rad vs. lad). In addition, animals trained to discriminate these speech sounds were able to accurately generalize when presented with exemplars that varied in duration, pitch, and speaker identity. The results of these studies indicate that rats can learn to reliably discriminate many complex speech sounds. Following training, responses to trained speech sounds were recorded from primary auditory cortex (A1) and compared to responses in naïve rats. Preliminary analysis indicates that the response properties of A1 neurons are substantially altered by long-term training.

Supported by NIH R21 #1R15DC006624-01

183. Laminar model for cortical development with an emphasis on the macaque visual system

Andrew M Oster, Paul C Bressloff

University of Utah

Cortex displays a characteristic six-layer structure with each lamina possessing its own unique horizontal circuitry and receiving its own set of inputs from thalamus, other cortical areas, or inter-laminar projections. Because of this variability in cortical architecture, there may be variations in response properties within a cortical column. For instance in the macaque monkey primary visual cortex (V1), the property of direction selectivity is known to be sequestered in layers 4B/upper 4Ca and layer 6 within the depth of the column (Hawken et al., 1988). Furthermore, in the macaque monkey only weak orientation selectivity is found in upper 4C; not until the upper laminae does strong orientation selectivity emerge, illustrating the importance of the intricate laminar structure. Recent studies by Feldman et al. on rodent somatosensory cortex have shown that the inter-laminar vertical projections from layer 4 to layer 2/3 are neither innate nor concrete and found Hebbian-like plasticity in the inter-laminar vertical projections. Additionally, work done by Lund et al. has highlighted the diversity of the horizontal connections in the different laminae. Inspired by Feldman and Lund's findings, we introduce an activity-dependent laminar model for cortical development whereby functional differences between the laminae allow the cortex to produce intricate preference maps that more accurately resemble the biological circuitry than representing cortex as a single sheet. To demonstrate the virtues of a laminar cortical developmental model, we model the joint development of cytochrome oxidase (CO) blobs and ocular dominance (OD) columns in the macaque monkey. The macaque primary visual cortex displays a stripe-like OD pattern. The principal thalamic input to V1 arrives at layer 4C and is then projected vertically to the upper laminae, which also receive direct thalamic input from a different visual pathway. In the upper laminae, V1 stains a periodic pattern of CO blobs that are aligned with the centers of the OD stripes. In fact, the direct thalamic input to the upper laminae defines the CO blobs. We assume that the thalamic input to layer 4C segregates early, before the upper laminae exhibit OD. Our model has the property that the OD preferences of layer 4C are inherited by the upper laminae, followed by a segregation of the vertical projections from layer 4C and the thalamic input to the upper laminae, which ultimately results in the formation of the CO blobs.

184. Arbitrary Functions Learning with Neural Network Model Based on Spike-Timing Dependent Plasticity

Yefei Peng, Paul W. Munro

University of Pittsburgh

The computational properties of a network of interconnected neurons can be investigated at either a neuronal level or at a system/network level. The response properties of a neuron are generally characterized in terms of some set of stimuli. Whether the stimuli in the set are patterns of afferent activity or are further removed from the neuron, such as visual patterns driving a cell in V1, the response of a neuron is a function of the stimuli. The computational nature of these functions remains an open question.

STDP is a phenomenon in which repetitive spike pairing on the order of milliseconds leads to LTP and LTD. We build a feed forward neural network model based on STDP learning rule, and train this network to compute arbitrary functions. A post-synaptic output neuron R will be trained to respond to stimuli generated by a small number of extracellular stimulating electrodes, each driving multiple afferents to the target neuron, including excitatory and inhibitory inter neurons. In addition, the target neuron will be driven by a single action potential timed to fire either before (to train the cell not to respond) or following (to train the cell to respond) the stimulus. Activity and plasticity of these inter neuron 'hidden units' is expected to play a major role. In effect, the target neuron R computes a classification function; the stimulus is a member of the class if the neuron responds to it.

Based on this model, we carry out simulations investigating the role of neural plasticity. The goal of the work is to study the rules that specify the activity-dependence of synaptic change, and how the computational properties of neurons depend on the statistics of the afferent activity patterns. Three tasks were performed. First, AND and XOR. Second, three-bit parity. Third, continuous inputs. As results, we show that all the tasks can be learned by a network without inhibitory plasticity, as long as the parameters of the network are in appropriate range. Inhibitory plasticity could improve learning performance of the network.

185. Beyond Pair-Based STDP: a Phenomenological Rule for Spike Triplet and Frequency Effects

Jean-Pascal Pfister, Wulfram Gerstner

Brain-Mind and I&C, EPFL

While Classical experiments on spike-timing dependent plasticity analyzed synaptic changes as a function of the timing of a pair of pre- and postsynaptic spikes (Bi and Poo, 1998), recent experiments also point to the effect of spike triplets. Here we develop a mathematical framework that allows us to characterize timing based learning rules. Those learning rules are characterized as a function of their type of interaction (all spikes interaction versus only `neighboring' spikes interaction) and as a function of the order of interaction (pairs for second order, triplets for third order,...).

Moreover,we identify a candidate learning rule that captures a variety of experimental data, including the dependence of potentiation and depression upon pre- and postsynaptic firing frequencies observed in L5 pyramidal neurons (Sjöström et al., 2001). With different parameters, our model is also able to reproduce most of the recent triplet experiments performed in culture of hippocampal neurons (Wang et al., 2005). The explicit and simple formulation of the learning rule makes a lot of calculations feasible.For example, our phenomenological learning rule can be strictly mapped to the Bienenstock-Cooper-Munro rule.

186. Online Learning in a Model Neural Integrator

Srinivas C Turaga[1], Haim Sompolinsky[2], H. Sebastian Seung[1]

[1]MIT, [2]The Hebrew University

The oculomotor neural integrator is a brainstem nucleus that stores short term memory of eye position. The memory of eye position is used to drive eye muscles tonically, leading to gaze holding. This brain region receives input in the form of an eye velocity signal, which must be transformed into an eye position signal by integration. One class of model neural integrators performs the velocity to position transformation by means of positive feedback in a recurrent neural network (Cannon, Robinson & Shamma 1983). However, recurrent feedback models of integration are susceptible to a fine-tuning problem where changes in feedback of less than 10% can de-tune the system (Seung 1996).

The problem of de-tuning may be overcome by online re-tuning using an error signal. A mistuned integrator induces gaze drift during fixations known as retinal slip. Recent experimental evidence suggests that a visual error signal in the form of retinal slip might be involved in tuning the integrator in goldfish (Major et al. 2004). We propose and analyze a simple learning rule for online tuning of weights in the integrator circuit using the retinal slip error signal. Our learning rule is similar in spirit to that proposed by Arnold & Robinson (1997), but in contrast to their learning rule, we provide theoretical justification for our proposal.

Our learning rule is proved to be globally convergent. The rule requires only local spatial and temporal information, so has minimal computational overhead and can be implemented easily in a biological setting. We demonstrate that this simple learning rule can be used more generally in online learning for motor control problems where a continuous error signal is available.

187. A Hebbian reinforcement learning algorithm reproducing monkey performances in visuo-motor learning task.

Eleni Vasilaki[1], Stefano Fusi[2], Xiao-Jing Wang[3], Walter Senn[1]

[1]Institute of Physiology, University of Bern, Switzerland, [2]Institute for Neuroinformatics (INI), ETH, Zurich, Switzerland, [3]Volen Center for Complex Systems, Brandeis University, MA

Experimental work of Asaad et al (1998) has revelead mechanisms which might be involved in a visuo-motor learning task performed by monkeys (Fusi et al, 2005). In this task, the monkeys are trained to associate visual stimuli (pictures) with delayed saccadic movements, left or right, with associations being reversed from time to time. To successfully reproduce these data with an artificial neural network, we propose a Hebbian-based reinforcement learning algorithm which is Hebbian in case of reward and anti-Hebbian in its absence.

Reinforcement learning algorithms are based on a random exploration of network states. This exploration is usually believed to be driven by intrinsic noise which might be present either in the neuronal firing or the synaptic release. Here we suggest that the exploration is mediated by the induction of synaptic modifications which is inherently stochastic. While the direction of a putative synaptic modification depends on the pre- and postsynaptic activity and the reward signal, the actual induction of the modification is only performed with a certain probability (Amit and Fusi, 1994). The putative modification is proportional to the difference of the postsynaptic activity and some LTP/LTD modification threshold. This threshold reflects the averaged postsynaptic activity across all pattern presentations. For binary neurons and synapses, we have shown that the algorithm moves along the gradient of the expected total reward. However, the algorithm can be also applied to analogue synapses in a mean field approach.

The monkeys learn the visuo-motor task within 30 trials on average. Interestingly, after making a mistake, the probability for a correct choice in the next presentation resets to chance level, independently of the previous success history (Fusi et al, 2005). We modeled the decision making process through competition of two neuronal populations acting as winner-take-all (WTA) mechanism. The monkey data can be reproduced only if the learning step in case of punishment is much larger than in case of reward. With this imbalance in reward, it is also possible to successfully learn the XOR-problem and a reduced MNIST data set (a standard digit classification task).

The small synaptic changes in case of reward versus punishment are consistent with the principle of "learning from mistakes" (Bak and Chialvo, 2001). If reward is given, there is little synaptic change; on the contrary a mistake may severely affect the system. This represents a stop-learning condition which limits the synaptic changes to those synapses which really need to be modified (Senn and Fusi 2005). The stop-learning condition in combination with a slowly moving threshold allows for imposing soft synaptic bounds.

References

1. Fusi S, Asaad W.F., Miller E.K. and Wang X.-J. (2005) A microcircuit model of arbitrary sensori-motor mapping: learning and forgetting on multiple timescales. Soc Neurosci Abstr 813:10.

2. Senn, W., & Fusi, S. (2005). Learning only when necessary: better memories of correlated patterns in networks with bounded synapses. Neural Computation 17, 2106-2138.

3. Wang, X-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. Neuron 36, 955-968.

4. Bak, P., & Chialvo D. R. (2001). Adaptive learning by extreme dynamics and negative feedback. Physical Revew E 63(3).

5. Asaad, W.F., Rainer, G. & Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. Neuron 21, 1399-1407.

6. Amit, D. & Fusi, S (1994). Learning in neural networks with material synapses. Neural Computation, 6, 957-982.

188. Eyelid conditioning, timing and the cerebellum

Horatiu Voicu, Tatsuya Ohyama, Michael D Mauk

UTH Health Science Center

Timing of the conditioned response in eyelid conditioning is the result of the competition between LTD and LTP at the granule to Purkinje synapses. While the synapses that undergo net LTP dominate the early part of the conditioned response, the ones that undergo net LTD help in producing an eyelid closure. The transition point between the two competing processes is revealed by the latency to onset of the conditioned response which depends on the ISI. The longer the ISI the later the transition point occurs. Our detailed model of the cerebellum suggests that the dependency between ISI and learning could happen if we assume that the cerebellum receives signals that become active at CS onset and turn off at different times. By subtracting signals of different lengths the granule-Golgi system produces parallel fibers that are active at different times within the trial. This enhances the temporal coding at the granule-to-Purkinje synapses, which in turn enhances learning and timing of cerebellar output. The activity at the cerebellar input that turns off at different times could be caused by the mossy fiber-to-UBC synapses that trap glutamate for hundreds of milliseconds.

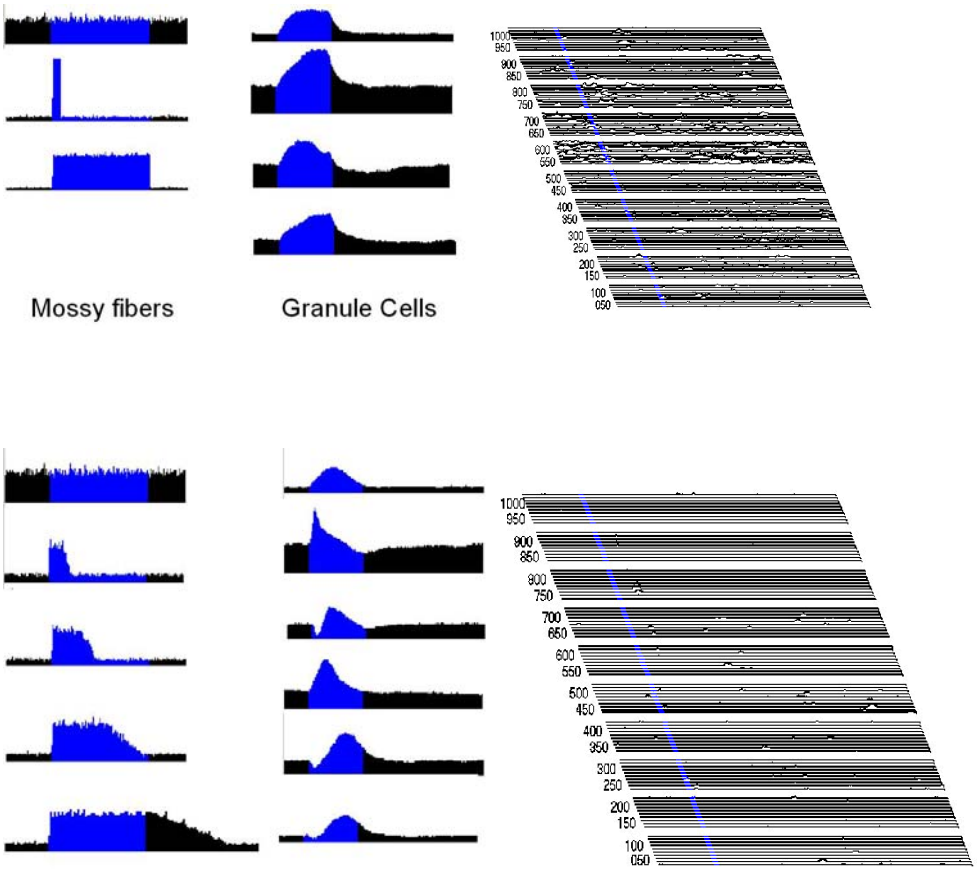

Figure 1. The first two columns represent the peristimulus histograms for two conditions: mossy fibers show phasic and tonic firing for the duration of the CS (above) and mossy fibers show phasic and tonic firing with different durations starting at the CS onset (below). The animated graphs show the acquisition of the conditioned response as a function of ISI. The width of the blue stripe encodes the length of the ISI. Each line represents the average of 10 trials.

189. Experience-Induced Arc Sharpens Orientation Ensembles in Visual Cortex

Kuan Hong Wang, Ania Majewska, Mriganka Sur, Susumu Tonegawa

M.I.T.

To investigate the molecular and cellular mechanisms underlying experience-dependent modifications of sensory representations in the cerebral cortex in vivo, we developed a new genetic system by replacing the neural activity-regulated gene Arc with a GFP reporter, and chronically imaged at a single cell resolution using two photon microscopy the activation patterns of endogenous Arc promoter-driven GFP (Arc-GFP) signals in primary visual cortex following daily visual stimulation in alert mice. Arc-GFP expression was induced in orientation-specific neuronal ensembles in the heterozygous mice. Repeated presentation of the same orientation led to the reactivation of a progressively smaller but more reliable population, illustrating a mechanism for adult cortical plasticity at the cellular ensemble level. The lack of Arc protein in the homozygous mice resulted in the activation of larger ensembles with reduced orientation specificity and broadened neuronal spike tuning curves, suggesting that Arc protein is a visual experience-induced molecule with a role in suppressing the activation of orientation-non-specific neurons in primary visual cortex.

190. Memory recall based on rebound conductances

Daniel Z Wetmore, Eran A Mukamel, Mark J Schnitzer

Stanford University

The neurobiology of memory retrieval remains poorly understood compared to memory formation. Retrieval can occur rapidly after presentation of a cue and judging from this rapidity likely involves mainly electrical dynamics. Moreover, stored memories are often retrieved unreliably. We present a theory based on neuronal rebound conductances that accounts for behavioral data on timing and reliability aspects of memory recall. We focus on two cerebellum-dependent forms of learning and memory, classical eyeblink conditioning and gain adaptation of the vestibulo-ocular reflex (VOR), for which retrieval occurs as a well-timed motor response. Cerebellar Purkinje cells involved in these behaviors project to neurons in the deep cerebellar nucleus (DCN) and medial vestibular nucleus (MVN) that exhibit rebound depolarizations mediated primarily by T- and h-type channels, respectively.

In eyeblink conditioning, for acquisition of a conditioned reflex to occur the onset of the conditioned stimulus (CS) must precede the onset of the unconditioned stimulus (US), even if by only ~50-100 ms. After a regiment of such 'forwards' training, a properly timed conditioned response slightly anticipates US arrival. By comparison, 'backwards' training in which CS onset occurs after or concurrently with US presentation does not lead to conditioned motor reflexes. A short delay of <100 ms between the CS and US onsets leads to unreliable conditioned responses that occur only a fraction of the time following a CS test presentation. By analogy, studies of VOR gain adaptation have shown temporal asymmetries in learned behavior emerge after training paradigms that resemble 'forwards' and 'backwards' conditioning (Raymond & Lisberger, 1996). Current theories of cerebellum-dependent learning and memory do not account for these striking asymmetries in learned behavior.

Through analytical and computational modeling, we tested the hypothesis that rebound excitation in cerebellar output nuclei may underlie the temporal and unreliability aspects of cerebellum-dependent memory recall. To begin, we created a one-compartment electrical model of a DCN neuron. The model cell contained both rebound and leak conductances and received excitatory input from CS-activated mossy fibers as well as GABAergic input from Purkinje cells. Bi-directional plasticity at parallel fiber to Purkinje cell synapses during training leads to biphasic modulation of Purkinje cell spike rates in response to a CS test presentation (Fig 1). After forwards training, presentation of the CS causes Purkinje output first to rise, which hyperpolarizes the DCN neuron and deinactivates its T-conductances. Purkinje activity then drops, which allows T-current activation and rebound production (Fig 2a). By comparison, after backwards training a CS presentation leads to a drop and then a rise in Purkinje activity, which modulates the DCN cell voltage in a biphasic manner poorly suited to induce rebounds. Forwards training with only a brief interval between the CS and US leads to rebounds of smaller amplitude, because T-channels lack sufficient time to deinactivate. This model contrasts with the prevailing view of long-term depression (LTD) and potentiation (LTP) as opposing processes, one for memory induction and the other for extinction. Here, the two forms of plasticity are complementary and are both invoked at memory storage as a requirement for reliable future retrieval.

Simulations using a two-compartment model of a DCN cell showed that the graded amplitude of rebounds generated in a dendritic compartment sets the probability of inducing an all-or-none Ca^{2+}-spike and an accompanying burst of somatic Na^{+}-spikes. This burst can drive conditioned responses. Thus, rebounds that occur selectively following forwards training appear capable of explaining the temporal asymmetries observed in eyeblink conditioning. Moreover, we find that in compartmental simulations of MVN neurons, h-mediated rebounds may also suffice to account for the temporal asymmetries observed in specific forms of VOR adaptation.

To explore further the issues of memory reliability and readout, we examined rebound generation using a two-dimensional phase-plane analysis that focuses on the slow variables setting the timescale for rebound generation. These variables are the membrane voltage, V, and the T-type channel inactivation variable, I, which has a voltage-dependent time constant of ~50-100 ms close to the minimum CS-US interval needed for reliable retrieval (Fig 2c). A family of phase plane trajectories undergo rebound depolarization prior to the arrival of the expected US. Rebound amplitude is a strict and graded function of the (V, I) values attained by the hyperpolarization that occurs during the CS-US interstimulus interval (ISI) (Fig 2b). A two-dimensional color map of rebound amplitude as a function of (V, I) reveals a simple basis for the sharp dependence of reliability on the ISI value and the levels of LTP and LTD attained during training (Fig 2c). Large amplitude rebounds initiate in a neighborhood of the (V, I) plane that may be viewed as a 'reliability zone' in that once a trajectory enters this zone a large rebound will occur without fail.

The present work suggests retrieval of cerebellum-dependent memories may be governed by a 'lock-and-key' mechanism. In this view, the lock resides in the DCN (or MVN) and requires a specific pattern of input activity to initiate rebound spiking and a subsequent motor response. Synaptic plasticity induced by training leads to temporal modulation of spike output during post-training testing trials (Fig 1). Such spike patterns represent keys, but only those produced by forwards training match the temporal filter imposed by the DCN. Thus, only the pattern of biphasic Purkinje activity produced by forwards training can unlock a behavioral response. Backwards training also triggers plasticity, but the key does not match the DCN temporal filter and conditioned responses do not occur.

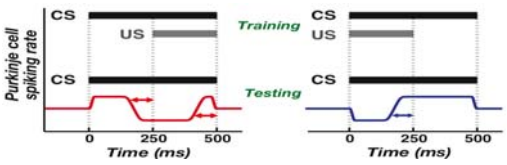

Figure 1. Pkj firing rates decreased during portions of a 500 ms CS that coincided with a 250 ms US delivered, reflecting PF-Pkj LTD, but Pkj output increased during portions of the CS that did not coincide with the US, reflecting PF-Pkj LTP.

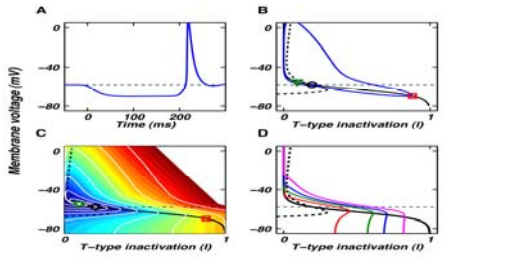

Figure 2. Phase-plane analysis of rebound depolarization in DCN neuron model. **(A)** Membrane voltage of a single-compartment model DCN neuron following release from a hyperpolarized voltage (-75 mV) shows a rebound with a latency of ~50 ms. **(B)** Trajectory in phase-space defined by membrane voltage (V) and T-type inactivation variable (I), together with the nullclines corresponding to constant I (solid) or V (dashed). **(C)** The color plot represents the magnitude of rebound for trajectories corresponding to different initial conditions (white lines). **(D)** The threshold for generation of an all-or-none high-voltage-activated (HVA) calcium spike is shown for several values of the density of HVA current.

191. Dendritic Morphology and Storage Capacity in Hippocampal Pyramidal Neurons

Xundong Wu, Bartlett W Mel

University of Southern California

The hippocampus is involved in the rapid storage of online, or episodic memories. The need to burn in new memories "at the speed of life" places a huge constraint on a neuromorphic learning system. In particular, the need to encode input patterns in "one shot" calls for high synaptic learning rates, but this leads to a high forgetting rate as well which severely compromises overall storage capacity. Assuming that CA1 pyramidal neurons must encode a continuous stream of sparse input patterns arriving from CA3 via the Schaffer collaterals, and that the goal is to store these patterns as faithfully as possible for as long as possible, we have studied how the storage capacity of this pathway depends on (i) the CA1 dendritic morphology, that is, the number and size of the integrative subunits within the dendritic arbor, (ii) the local voltage thresholds needed to trigger synaptic learning and/or to trigger a local dendritic spike, (iii) the synaptic learning step size, and (iv) the number of long-term stable weight levels available to each synapse. We obtained the following results: (1) Capacity depends strongly on the choice of learning and response thresholds; high thresholds leading to extremely sparse subunit activations are strongly preferred; (2) The optimal response threshold is always higher than the optimal learning threshold; this suggests there may be two separate nonlinear thresholding mechanisms within a dendritic branch (one pertaining to learning, the other to response); and (3) Small subunits (e.g. 100 syns or fewer) would be optimal in a noise-free environment, but even very small amounts of noise on active input lines push the optimal subunit size to the order of hundreds of synapses per branch (which corresonds to the number of synapses on a typical basal or oblique branch of a CA1 pyramidal cell).

192. Retinal lesion-induced receptive field reorganization in primary visual cortex is spike timing dependent

Joshua M Young[1], Bogdan Dreher[2], Klaus Obermayer[1]

[1]Neural Information Processing Group, Department of Computer Science, Berlin University of Technology, Germany / Bernstein Center for Computational Neuroscience Berlin, Germany, [2]Institute for Biomedical Research, The University of Sydney, Australia

Two alternative learning rules, sensitive to either the temporal correlation between pre- and post-synaptic spikes (correlative Hebbian plasticity), or the temporal sequence of pre- and post-synaptic spikes (spike timing-dependent plasticity), have been proposed to underlie receptive field reorganization in primary visual cortex (1-3). However, all aspects of receptive field change explored by computational models to date are compatible with both rules, leaving their relative importance for receptive field plasticity unclear. It has been observed that cortical neurons within a zone of feedforward input loss created by a retinal lesion undergo a rapid (within minutes to hours) phase of receptive field expansion. Evidence suggests that this expansion phase is followed by a longer (from weeks to months) competitive phase of receptive field refinement during which the efficacies of a subgroup of the horizontal inputs to a single neuron are preferentially enhanced, resulting in an apparent 'shift' of the neuron's receptive field. A leaky integrate-and-fire network simulation of layers 2/3 in primary visual cortex was used to model this second phase of reorganization and the associated receptive field shifts by combining heterosynaptic gain modulation with competitive synaptic plasticity at horizontal connections. The model consisted of a one dimensional chain of topographic positions, or 'columns', spaced at hypercolumn widths, with each column being represented by a reciprocally connected excitatory and inhibitory neuron. Both cell types received a set of horizontal inputs from excitatory neurons in neighboring columns, the weight distributions of which were based on in vivo intracellular measurements. The effect of a lesion was simulated by preventing a subset of columns within the center of the network from receiving feedforward input during training. Prior to training the effective gain of excitatory neurons within the deprived columns was increased via a mechanism based on the empirical phenomena of synaptic scaling, where low firing rates cause excitatory inputs to excitatory and inhibitory neurons to be potentiated and depressed respectively. We then compared the receptive field effects of competitive learning at the horizontal inputs to excitatory cells using either a spike timing-dependent or a correlation-based Hebbian rule. We found that the global patterns of receptive field shifts observed in vivo are robustly produced using the spike timing-dependent rule, but not using the correlative rule. We then compared the effects of these two learning rules in a larger two-dimensional rate model based on the same principles as the integrate-and-fire network and found that the same fundamental differences between the two learning rules were still apparent. We conclude that spike timing-dependent plasticity, combined with heterosynaptic gain modulation, is a critical determinant of long-term receptive field reorganization.

Support contributed by: BMBF 10025304 (Germany), ARC (Australia).

1. Song, S. & Abbott, L. F. Neuron 32, 339-50 (2001).

2. Sirosh, J., Miikkulainen, R. & Bednar, J. A. in Lateral Interactions in the Cortex: Structure and Function. (ed. Choe, Y.) (UTCS Neural Networks Research Group, Austin, Texas, 1996).

3. Andrade, M. A., Muro, E. M. & Moran, F. Biol Cybern 84, 445-51 (2001).

193. Learning by message-passing in networks of discrete synapses

Riccardo Zecchina[1], Alfredo Braunstein[2]

[1]International Centre for Theoretical Physics (ICTP), [2]Institute for Scientific Interchange (ISI)

Learning and memory are implemented in neural systems mostly through distributed changes of synaptic efficacy. The learning problem in neural networks (NN) asks whether one can find values for the synaptic efficacies such that a set of p patterns are stored simultaneously. Depending on the structure of the network - feed-forward or recurrent - the storage problem is either seen as a classification problem (input patterns are classified according to the output of the network) or as an attractor dynamics problem (patterns are the external stimuli which drive the dynamics of the network to the closest attractor). In any case, understanding the mechanisms underlying synaptic changes constitutes a crucial step for modeling real neural circuits (e.g. the Purkinje cells in the cerebellum [1]).

On the purely theoretical side many basic results have been derived, ranging from information theoretic bounds and statistical physics analysis of learning capabilities in model NN [2] to concrete algorithms, like artificial pattern recognition systems . Still there exist many open conceptual problems that are related to the need of satisfying realistic constraints.

Modeling material synapses is possibly one of the most basic ones [3], the discrete case (and specifically the switch-like binary one) being of particular experimental [4,5] and technological interest: recent experiments -- at the single synapse resolution level -- have shown that some synapses undergo potentiation or depression between a restricted number of discrete stable states through switch-like unitary events [4,5].

It is has been known since many years that the discreteness of synaptic efficacies makes the learning problem extraordinarily difficult: even the task of finding binary synaptic weights for a single layer network (the binary perceptron) which classifies in two classes a given set of patterns is both NP-complete and computationally hard on average (as observed in classical numerical experiments). In spite of the fact that binary networks can in principle classify correctly an extensive number p= alpha n of random patterns with n binary synapses, practically there exists no known algorithm which is able to store exactly more than just a logarithmic number as soon as a sub-exponential cut is put on their running time.

In this contribution, we present a distributed message-passing algorithm of statistical physics origin which is able to store efficiently an extensive number (p=alpha n with alpha>0) of random patterns in binary NN characterized by a wide range of different topologies [5].

We consider single and multi-layer networks with local connectivities of the neurons ranging from finite to extensive. The typical computational complexity of the algorithm will be shown to scale roughly as O(n^2 log n), that is almost linearly on the size of the input for an extensive number of patterns. This fact together with the parallel nature of the algorithm allows to easily find optimal synaptic weights for systems as large as n=10^6 with alpha relatively close the critical value alpha_c above which perfect learning is no longer possible.

From the algorithmic viewpoint, our solution to the binary learning problem should be seen as an example of solution of constraint satisfaction problems over dense factor graphs (a graphical representation of combinatorial constraints used in information theory). As such, our result show how the recent progress in combinatorial optimization by statistical physics and message passing techniques which have allowed to solve efficiently famous combinatorial problems like random K-satisfiability [6] or random Q-coloring, can be extended to other classes of problems in which constraints involve an extensive number of variables.

Finally, we show how to greatly simplify the algorithm while retaining most of its computational performance. This fact represents a proof-of-concept of how highly non-trivial learning can take place by message-passing between simple devices disposed over the network itself. Hopefully these results could turn out to be of interest in the study of the biological treatment of information in neural systems.

References:

1) N. Brunel, V. Hakim, P. Isope, J.-P. Nadal, B. Barbour, Neuron 43, 745-757 (2004)

2) A. Engel, C. van den Broeck Statistical mechanics of learning (Cambridge University Press, 2001).

3) D.J. Amit, S. Fusi, Neural Computation 6, 957-982 (1994)

4) D.H. O'Connor, G.M. Wittenberg, S. S.-H. Wang, Proc. Natl. Acad. Sci. USA 102, 9679-9684 (2005);

5) C.C.H. Petersen, R.C. Malenka, R.A. Nicoll, J.J. Hopfield, Proc. Natl. Acad. Sci. USA 95, 4732-4737 (1998)

6) A. Braunstein, R. Zecchina, Learning by message passing in networks of discrete synapses, Physical Review Letters in press, http://arxiv.org/abs/cond-mat/0511159

7) M. Mezard, G. Parisi, R. Zecchina, Analytic and Algorithmic Solution of Random Satisfiability Problems , Science 297, 812 (2002)

194. Supervised STDP as a Means to Initiate and/or Coordinate Neuronal Maps

Leo van Hemmen

Physik Department, Technical University of Munich

Sensory modalities `observe' the outside world by giving rise to neuronal input, spikes. A `map' may be defined as a neuronal representation of the outside sensory world. Different sensory modalities may therefore give rise to different maps. If so, an animal has a problem if these modalities observe the same universe surrounding it, e.g., three-dimensional space. In vertebrates, for example, hearing and seeing often induce localization of the same object, a sound source, in space and time through spatio-temporal neuronal input.

We can now ask at least two questions. First, how does an animal's brain bring about its neuronal maps that provide a (quite often) detailed representation of space and time? Second, is there a means to coordinate these maps if they describe the same space-time?

Spike-timing-dependent synaptic plasticity (STDP) is a universal means to ``learn'' spatio-temporal neuronal activity patterns [1,2]. Quite often, however, STDP is not completely self-organizing but needs an external ``teacher'' to initiate and/or coordinate maps arising from, say, two (or more) sensory modalities. To solve the first problem STDP has been proposed [1,2], and proven, to be the way-out. Until recently, however, the answer to the second question was still open. We have now shown [3] that supervised STDP, also called SSTDP, may be an appropriate way to initiate and coordinate neuronal maps. Here we explain SSTDP through a concrete example, the clawed frog Xenopus.

During night 180 lateral-line organs allow Xenopus to localize prey by detecting water waves generated by insects floundering on the water surface. Not only can the frog localize prey but it can also to determine its character. This suggests waveform reconstruction and a key question is how to get the appropriate neuronal wetware. Catching time differences arising from the input on the skin is important and STDP seems to be the natural tool. We show how supervised STDP allows a frog to learn what is where in the dark. In the present case supervision may arise from the visual system of the young frog.

[1] W. Gerstner, R. Kempter, J.L. van Hemmen, and H. Wagner, A neuronal learning rule for sub-millisecond temporal coding, Nature 383 (1996) 76-78.

[2] C. Leibold and J.L. van Hemmen, Spiking neurons learning phase delays: How mammals may develop auditory time-difference sensitivity, Phys. Rev. Lett. 94 (2005) 168102-1/4

[3] J.-M.P. Franosch, M. Lingenheil, and J.L. van Hemmen, How a frog can learn what is where in the dark, Phys. Rev. Lett. 95 (2005) 078106-1/4

195. High frequency stimulation of the subthalamic nucleus restores thalamic relay reliability in a computational model

Jonathan E Rubin[1], Yixin Guo[2], Cameron McIntyre[3], Kresimir Josic[4], David Terman[5]

[1]University of Pittsburgh, [2]Drexel University, [3]Cleveland Clinic, [4]University of Houston, [5]The Ohio State University

Deep brain stimulation (DBS) of the subthalamic nucleus (STN) can alleviate motor symptoms of Parkinson's disease. In a computational model, we have demonstrated that under parkinsonian conditions, rhythmic inhibitory output from the internal segment of the globus pallidus (GPi) compromises the ability of thalamocortical relay (TC) cells to respond to depolarizing inputs. High frequency stimulation of STN regularizes GPi output, and this restores TC responsiveness, despite the increased frequency and amplitude of inhibition that result. These findings can be understood via phase plane and bifurcation analysis of a TC model. Further, under certain simplifying assumptions, a Markov chain approach allows for an analytical calculation of how TC firing probability to inputs, expected number of response failures between firings, and related firing statistics can be obtained from the statistics of the inhibitory input train and the signal to be relayed, supporting the generality of these findings. Finally, we consider a model TC cell response to GPi spike trains recorded experimentally from monkeys in parkinsonian and DBS conditions. We show that model responses to data generated under these conditions agree with those obtained from simulated data. Further, summation of responses over a heterogeneous TC population leads to significant fluctuations in reliablity in the parkinsonian case, as opposed to an overall reliable response in DBS.

In summary, while rhythmic inhibition compromises TC relay in parkinsonian conditions, regularization of this inhibition, seen for example in DBS data, restores TC relay reliability.

196. A biophysical approach to behavioral neuroscience in C. elegans

Aravinthan Samuel

Harvard University

Navigational behavior in C. elegans is the outcome of computations carried out by neural circuits which transform sensory information into movement decisions. The advantage of studying simple model organisms like C. elegans is that careful quantitative analysis of navigational behavior allows us to infer the algorithmic structure of these computations. Thermotactic navigation in C. elegans is a particularly attractive paradigm for studying computation and behavior. First, it is straightforward to provide thermal stimuli in a reliable and quantitative manner to behaving worms as well as quantify worm motile response using video microscopy. Second, thermotaxis is a complex, experience-dependent, and multifaceted behavior. Worms remember the temperature at which they were cultivated. When navigating spatial thermal gradients at temperatures above the set-point of thermotactic memory, worms exhibit cryophilic movements down spatial gradients. When navigating at temperatures near the set-point, they tend to track isotherms. When navigating below the set-point, they do not bias their movements up or down temperature gradients. Therefore, worms are capable of triggering different thermotactic modes based on comparison between current temperature and a remembered set-point. Distinct thermotactic behavioral modes imply separate strategies by which C. elegans transforms sensory perception of thermal gradients into movement decisions. I will discuss what biophysical analysis of thermotactic behavior is teaching us about how a multilayered navigational algorithm might be encoded in the neural circuit of C. elegans.

197. High-Fidelity Coding of Single Trials by Neurons in the Macaque Frontal Pursuit Area

David Schoppik, Katherine I Nagel, Stephen G Lisberger

HHMI/UCSF

In the face of repeated presentations of an identical sensory stimulus, both the motor response and its associated neural activity vary to some degree. The extent to which the variability in the neural activity accounts for variability in the motor response defines the fidelity of behavioral drive by individual neurons.

We had two monkeys repeat an oculomotor behavior: 200 ms of fixation, followed by 500 ms of smooth tracking of traditional step-ramp target motions. While the monkey performed the behavior, we made single unit recordings from 157 pursuit-sensitive neurons in the smooth eye movement region of the frontal eye fields. We modeled the neurons by performing optimal kernel analysis against the variability in eye position across the duration of the trial. Filters were based on the responses in 80% of the available trials and then were tested using the variability in the spike trains on the remaining 20%. Approximately 1/3 of the neurons yielded statistically significant filters, when compared to filters generated with randomly shuffled spikes. The R value between predicted and actual eye variability was 0.2 +/- 0.1, implying than the firing rate of an individual neuron is precise enough to capture ~5% of the trial-by-trial variability in behavior. Analyses are underway to describe precisely the features of the behavioral variability encoded by the neuron.

198. Dimensionality and Dynamics in the Motor Behavior of C. elegans

Greg J Stephens, William Bialek, William S Ryu

Princeton University

How do neural systems process current sensory information with past experience to produce desired behavior? While there has been substantial effort towards understanding the encoding of sensory information, investigations of motor output are comparatively underdeveloped, in part due to the complexity of natural behaviors and the lack of canonical descriptions of relevant behavioral states. Here we provide a quantitative and principled analysis of motor behavior in a domain which is both rich enough to be interesting yet simple enough so that movements can be explored exhaustively: the motions of *C. elegans* freely crawling on an agar plate. In this case, the worm movements are completely characterized by the dynamics of its two-dimensional shape. Using tracking video-microscopy with 4hz sampling we capture a worm's image and extract the skeleton of the shape as a head-to-tail ordered collection of tangent angles sampled along the curve. While previous work has described worm movements with ad hoc collections of descriptive statistics (e.g. head-to-tail distance, center of mass, orientation...) we use the tangent angle representation to explore the full space of shapes, making no a priori assumptions about what aspects of shape are important. Using principal components analysis we show that the space of shapes is remarkably low dimensional, with four dimensions accounting for ~95% of the shape variance. With this low dimensional description we analyze the dynamics of movement as time series for a small number of variables. Two dimensions exhibit a limit cycle oscillation corresponding to sinuous crawling movements. The frequency of these oscillations fluctuates with components that have many time scales, generating a ~1/f spectrum. The displacements along the other two dimensions correspond to stereotyped movements such as bends, reversals and Ω turns. We use our low dimensional description to provide an exhaustive characterization of *C. elegans* response to a short laser pulse which at low intensities evokes a thermal response and high intensities evokes an escape response. Our quantitative description of *C. elegans* movement should prove useful in a wide variety of contexts, from the linking of motor output with neural circuitry to the genetic basis of adaptive behavior.

199. A local tuning model predicts adaptation rate and degree of generalization of visuomotor rotation learning

Hirokazu Tanaka[1], Terrence J Sejnowski[1], John W Krakauer[2]

[1]Computational Neurobiology Laboratory, Salk Institute, [2]Motor Performance Laboratory, Columbia University

Visuomotor adaptation has been studied extensively, beginning with the seminal work of Helmholtz, yet the physiological mechanisms for adaptation remain largely unknown and no comprehensive computational model has been proposed. Here we propose a physiologically-plausible computational model that successfully explains results from a well-known visuomotor adaptation paradigm, visuomotor rotation, in which movement of a screen cursor is rotated with respect to the direction of hand movement. The model assumes units that are tuned to preferred target and hand-movement directions, and the output hand-movement vector is computed by summing their activity multiplied by their preferred hand-movement vectors (population vector). The model adapts to the rotation by modifying the units' preferred hand-movement directions in order to minimize the squared error between target and cursor directions, using the simple delta rule. The model explains two psychophysical findings for rotation adaptation--learning curves and generalization patterns--found in Krakauer et al. (2000), in which subjects were trained with single, two, four, or eight training targets. First, the model reproduces the observed exponentially-decaying learning curves and the decrease in rate of adaptation as more target directions are trained. Second, again reproducing experimental results, when the model is trained with a single-target direction, there is only local learning, which falls off rapidly as untrained directions get beyond 90° degrees from the trained direction. Complete generalization to all directions requires training with eight target directions, consistent with the experimental data. Furthermore, we show that the shape of tuning functions influences these two adaptation processes and find that units must be narrowly tuned to their preferred direction in order to reproduce the experimentally observed generalization pattern . This finding implies that the computation for rotation adaptation occurs in a brain area (s) that are locally tuned to hand-movement direction, not globally tuned such as primary motor cortex, which has cosine tuning, or the cerebellum, which has sum-of-two-Gaussian tuning. Although a candidate area has not yet been identified in single-unit studies, we postulate that parietal cortex may be the candidate region based on our previous PET-imaging data.

200. Space-time separability in goal-oriented motion generation

Elizabeth B Torres, Richard Andersen

CALTECH

A typical problem in control entails finding the optimal path joining two points. For example, one would want to find this path for a car while consuming minimum fuel/energy. In brain research, the field of computational motor control has posed the problem of goal-oriented motion generation (e.g. movements in the reaching family) in a similar way. Given a starting position of the hand and a target location, explore all possible trajectories to find the one that is optimal with respect to a certain criterion. In other words, solve a minimization problem involving an integral over a pre-defined time interval. The minimizing path gives information about where in space and exactly when in time to move for each point along the trajectory. When the problem is posed in this way, not only it becomes intractable for realistic degrees of freedom, but also difficult to test experimentally in the context of natural motions. More importantly, time does not enter into the motor error for learning, contrary to what recent empirical evidence has shown.

In this work we present experimental data from both human and non-human primates that invite a new way of posing the problem of motion planning in brain science. The data suggests that the system has a major simplification of the problem that allows us to reformulate it geometrically. Instead of having to explore all possible trajectories and recall the minimizing curve every time one is faced with motion generation, all that the system seems to need is information about the starting and final points, and a scalar quantity in the temporal domain tied to the first impulse of the motion. Such a quantity is the time that it takes to reach the peak velocity of the motion. This major simplification is possible because the system separates the spatial from the temporal components of the motion. The spatial strategy for a given task is invariant to changes in the temporal aspects of the trajectories. Our data clearly shows this space-time separability feature together with the cognitive control of the time to first peak and supports the geometric treatment of the problem. For a given starting location and possible spatial targets, a space-time surface characterizes each given task according to goals and instructions. Thus the problem of finding the minimizing path can be solved with a simple differential equation that generates geodesic curves between the starting and final points, thus allowing for a temporal re-parameterization of the curve, a feature recently observed when learning to acquire a new motor skill.

In order to address this proposition, we discuss behavioral data from various tasks in human and non-human primates as well as neural data obtained from 130 cells in the posterior parietal cortex (areas 5 and the Parietal Reach Region, PRR) of 2 rhesus monkeys.

The tasks in humans involved orientation-matching motions at 3 explicitly instructed speeds, and under initial posture manipulation, where the subjects implicitly changed the speed. The first task for monkeys involved straight reaches interleaved with reaches around obstacles. The second task interleaved reaches from a normal initial posture with reaches from an abducted posture. Parenthetically, this particular experiment revealed that in PRR there is already sensory information related to posture (before any visual cue) useful to transform the goal vector at the hand into the postural vector necessary to orient the arm correctly. In both cases, humans and monkeys, the spatial strategy was conserved as the temporal aspects significantly changed. In the neural data space-time separability manifested in the tasks facilitated the trajectory prediction one second and a half before movement initiation.

We believe that the space-time separability uncovered by our work is at the heart of autonomous behavior. There is a difference between movement generation that is purely subject to the physical laws of motions and movement that in addition to that is also subject to cognitive control, as when we speed up or slow down at will along the same spatial path, something that a solid ball falling down a ramp can not do.

201. Non-Parametric Methods for the Modeling of Neural Point Processes

Wilson Truccolo, John P Donoghue

Brown University, Neuroscience Department

Statistical non-parametric modeling tools that enable the discovery and approximation of functional forms relating neural spiking activity to relevant covariates, rather than imposing strong a priori parametric forms, are desirable tools in neuroscience. We explore and compare different non-parametric approaches involving: (a) generalized additive models using penalized regression splines and adaptive smoothing; (b) Bayesian extensions of penalized splines to fit generalized additive models via Markov Chain Monte Carlo sampling methods; and (c) gradient boosting of ensemble of regression trees with interaction terms and stochastic subsampling. All these approaches are formulated in terms of approximating the conditional intensity function of a point process in discrete time. The standard likelihood function for general point processes is used in each of the approaches. We illustrate and compare the approaches by modeling primary motor single cell spiking activity as a function of spiking history and hand kinematics during a 2-D reaching task performed by monkeys. We use the cross-validated log-likelihood and neural decoding via particle filters to compare the different non-parametric approaches. Overall, stochastic gradient boosting and Bayesian generalized additive models performed better in terms of cross-validated log-likelihood. Differences among the approaches in terms of neural decoding performance were less significant.

202. A fundamental ambiguity in model-based adaptive control

Douglas B Tweed, Mohamed N Abdelghani

University of Toronto

In adaptive control, the plant model learns the relation between the controller's commands and plant behavior. But there is an ambiguity in this learning task which arises because, in sensorimotor control, the motor command isn't usually an independent variable, but is a function of other signals that are also fed to the model. For example, consider the vestibulo-ocular reflex, or VOR, which measures head velocity and counterrotates the eyes so as to minimize retinal image slip. Retinal slip depends on the motor command produced by the VOR, and also on head velocity and on the state of the eye muscles, so all these variables should be inputs to the plant model. But the command should itself be a function of the latter two variables: head velocity and plant state. So it is hard for the model to sort out the contributions. For a simpler example, imagine a plant that adds two inputs, x and y, and suppose that y is a function of x; for instance, y = x. Then the model can't tell whether the plant is computing x + y, or 2x, or 2y, or 3x – y etc. Similarly in the VOR, the plant model can't deduce how the motor commands affect performance, so it can't properly train the controller. We will illustrate the problem with simulations of the VOR, showing that the model can't properly estimate ¶e/¶u – the partial derivative of the performance error e (retinal slip) with respect to the motor command u – even though its estimate of retinal slip itself is accurate. We show that this happens even when all the components of the system are working: given random u, the model learns ¶e/¶u; and given ¶e/¶u, the controller learns to make e = 0. But model and controller can't learn together because then u is a function of head velocity h, and the model can't distinguish the effects of u and h on e. One possible solution is to disambiguate the signals with noise. That is, while the plant model is learning – while the error in its estimate of e is larger than some small threshold g – add noise to u. This way, u isn't fully determined by h, so the model can deduce ¶e/¶u and train the controller. Whether it is achieved with noise or some other way, disambiguation is needed in any control task where u is determined by the variables that affect error, as it is in most sensorimotor systems; e.g. in reaching, error depends on target location and arm state, and on u, which is a function of target location and arm state. Disambiguation wasn't needed in past theories because they considered only parts of the control system or relied on innate knowledge – e.g. that retinal slip is 0 when eye velocity equals –1 times head velocity – in a way that is not realistic for most sensorimotor tasks.

-

203. Circuit geometry and the representation of time in cerebellar networks

Eran A Mukamel, Mark J Schnitzer

Stanford University

Behaviors involving adaptive motor control require neural circuits that encode timing information about sensory stimuli. The cerebellum plays a critical role in such behaviors that require temporal processing over the scale of about a second or less. Timing signals are likely generated within the cerebellum itself and may be an emergent property of the large network of granule and Golgi cells in the cerebellar cortex. However, the neural mechanisms that generate timing information remain unclear. Proposals published to date generally invoke either a population of cells that exhibits coherent oscillations of activity or a population that represents time in a sparse manner with activity in distinct cells encoding distinct temporal intervals [1,2]. In both types of models, individual cells provide timing signals that can be interpreted at the single unit level, without need for decoding widespread patterns of concerted activity. However, it is unclear whether either proposal fits with the well-established circuit properties of the cerebellar cortex. We suggest an alternative class of timing models that is explicitly based on the geometry of neural connections in the cerebellar cortex. In these models temporal information is represented in a spatially distributed manner across the Golgi-granule cell network. Decoding requires access to the patterns of activity across populations of cells.

The cerebellum exhibits a highly stereotyped, near crystalline pattern of anatomical connections, and the idea that cerebellar network geometry may be linked to temporal coding is longstanding. Parallel fiber axons of the granule cells extend several millimeters along the cerebellar folia. This quasi-one dimensional geometric structure seems likely to determine the framework of neuronal dynamics which gives rise to computational properties. To explore such issues, we examined the dynamics of a linear firing rate network with a geometric structure mimicking that of Golgi-granule circuits. Using both analytic and computational approaches, including direct evaluation of the characteristic frequencies of large, sparse synaptic weight matrices, we studied the normal modes of these networks and simulated the dynamical behavior of networks containing >10,000 cells. Using values for network parameters that are physiologically reasonable, one finds it is plausible for a distributed processing to provide temporal representation over the 10-1000 ms scale. A main finding is that the number of Golgi cells limits the number of distinct modes of excitation available to the network. This in turn limits the precision of temporal representation. In mammalian cerebella there are about 10,000 granules per Golgi cell, which number about 400 per mm^2 of surface area within the cerebellar cortex [4]. This suggests behaviors or animal species that require particularly fine temporal coordination of motor behavior might have higher numbers of Golgi cells for such purposes, a prediction that might be testable. Although there is ample evidence indicating cerebellar neurons are not linear, forms of distributed temporal representation may persist in nonlinear networks. We are presently studying this issue and the role of circuit geometry in nonlinear cerebellar networks.

Towards isolating the influence of geometry on circuit performance, we also examined a linear network in which granule cell axons were extended in a random unparallel fashion. Detailed computational studies of the mode spectra showed that oscillatory frequencies of distributed modes depend on structural factors, including the orientation and length of parallel fibers. The relationship between geometric structure and dynamics is characteristic of the behavior of low-dimensional linear systems. Our results suggest that the organization of granule cell axons parallel to each other may be critical for creating a broad spectrum of oscillatory modes in the cerebellar input layer, and hence for representing a wide range of time intervals. We predict that manipulations that disrupt the geometrical structure of the cerebellar network could alter the range of stimulus frequencies represented by network dynamics. This could lead to behavioral deficits, such as mistimed responses following training in classical conditioning experiments.

1. Ivry, R. B. & Spencer, R. M. (2004). The neural representation of time. *Curr Opin Neurobiol,* 14, 22532.
2. Mauk, M. D. & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annu Rev Neurosci,* 27, 30740.
3. Braitenberg, V. (1967). Is the cerebellar cortex a biological clock in the millisecond range? *Prog Brain Res,* 25, 33446.
4. Palkovits, M., Magyar, P., & Szentagothai, J. (1971c). Quantitative histological analysis of the cerebellar cortex in the cat. ii. cell numbers and densities in the granular layer. *Brain Res,* 32, 1530.

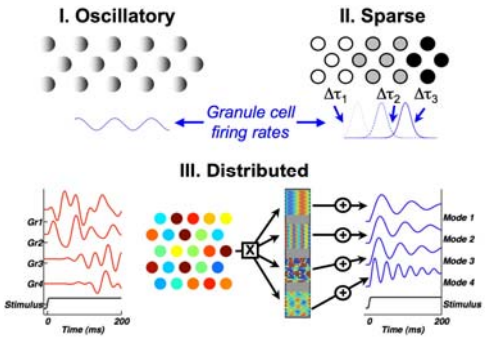

Fig. 1: Distributed representation for cerebellar temporal processing. In contrast to (I) oscillatory and (II) sparse coding schemes, which rely on coherent timing signals in individual granule cells, (III) distributed coding makes use of a population of granule cells. Sensory stimuli lead to irregular but reproducible patterns in the firing rates of granule cells (left). Plasticity at parallel fiber-Purkinje cell synapses leads to a spatial pattern of synaptic weights, which determines a regular dynamical mode at the Purkinje cell output.

204. Sensory encoding, relative spike-timing, and noise: lessons from the hermissenda eye

William H Nesse[1], Christopher R Butson[2], Gregory A Clark[2]

[1]University of Utah Department of Mathematics, [2]University of Utah Department of Bioengineering

Introduction
What makes neural networks more effective encoders than single neurons? Further, why do biological neural networks often outperform their human-engineered counterparts? One important emergent property exhibited by biological neural networks is the use of contextual spike-timing relationships-the timing of action potentials in one neuron, relative to the timing of action potentials in other neurons. Here we show that such contextual spike-timing relationships can improve neural information processing, and can help explain how biological nervous systems perform so splendidly in the presence of "noise", despite being inherently noisy themselves.

Methods
We used the eye of the marine nudibranch mollusk Hermissenda crassicornis as a simple model system for computational, physiological, and mathematical investigations. Simulations were implemented in Genesis and involved Hodgkin-Huxley level, multicompartment models of photoreceptors, using empirically derived biological parameters. Physiological investigations used conventional intracellular recording and stimulation techniques.

Results
The computational, physiological, and mathematical approaches support two key conclusions. First, in both the biological eye and the simulated network, feedback inhibition from type-A output photoreceptors onto type-B photoreceptors produced a striking absence of B-cell spikes shortly after A-cell spikes. Consequently, B cells fired later in the A-cell interspike interval, when the inhibitory B-to-A input is more potent in suppressing the next A-cell spike[1]. Hence, by altering contextual spike timing, feedback inhibition of type-B cells in turn yields greater inhibition of the type-A cell output. Thus, networks with feedback connections produce more effective modulation of type-A output.

Second, in contrast with the traditional view that noise degrades system performance, we found that random ionic noise, synaptic noise, and spike-timing noise improve, rather than impair, the ability of both simulated and real Hermissenda eyes to encode light intensity. In simulations, noise-induced improvements in light intensity encoding occurred across 8 light levels, were not confined to perithreshold stimulus intensities, and did not arise from simple stochastic resonance or DC-bias threshold effects[2,3]. Noise-free conditions produced "zones of stability"[4], characterized by phase locking, in which an increase in the frequency of inhibitory postsynaptic potentials (IPSPs) could paradoxically increase, rather than decrease, the firing rate of postsynaptic photoreceptors. In turn, this produced a non-monotonic relationship between light intensity and photoreceptor firing rate. In contrast, the addition of noise (random variations in ionic currents and IPSP amplitudes) disrupted the emergence of phase locking, yielding a more systematic relationship and an improvement in light-intensity encoding.

Similarly, in physiological investigations, noise-free conditions also produced phase locking within zones of stability. Consequently, increases in the frequency of artificial IPSPs (intracellular hyperpolarizing current injections) could paradoxically increase as well as decrease the postsynaptic photoreceptor firing rate evoked by light stimuli or by artificial light (depolarizing current) steps, thereby producing non-monotonic relationships that disrupted light-intensity encoding. Such anomalous effects were greatly reduced by the introduction of timing noise (IPSPs delivered at pseudo-random intervals, with the same mean frequency as in the noise-free condition). Thus, in noisy conditions, photoreceptor firing rate was more closely related to light intensity.

Mathematically, we can capture key features of the biological system by representing it as a network of phase-coupled oscillators, and we can derive a Fokker-Planck equation for the probability density of the phase state of the system under noisy conditions. This analysis supports the conclusions drawn above: Noise-free conditions display zones of stability in which phase-locking can disrupt efficient encoding. When noise is added to the system, this phase-locking is diminished. Nonetheless, statistical phase relationships persist that position the presynaptic neuron's spikes for effective inhibition of target neurons.

Discussion
These results indicate the importance of contextual spike timing relationships in the formation of emergent network properties, and help explain how biological neural networks, unlike most human-engineered devices, excel at accurate encoding of signals in noisy environments. They further suggest principles that may be advantageously incorporated into artificial intelligent systems, robotics, or neuroprostheses.

References
Fost, J.W. & Clark, G.A. (1996). J. Comput. Neurosci., 3: 155-172. 2. Butson, C.R. & Clark, G.A. (2002). Soc. Neurosci. Abstr., Program No. 848.6, online. 3. Clark, G.A. & Butson, C.R. (2004). Soc. Neurosci. Abstr., Program 870.21, online. 4. Perkel, D. et al. (1964). Science, 145: 61-63.

Supported by NIMH R01-MH068392 to GAC.

205. Smooth and Lurching Pulses in Two-Layer Thalamocortical-Reticular Integrate-and-Fire-or-Burst Networks

William Nesse, Paul C Bressloff

University of Utah

Ferret slices of lateral geniculate neucleus (LGN) and its companion reticular structure the perigeniculate nucleus (PGN) are known to generate "spindle" oscillations in the 7-10 Hz. range. These oscillations are due to synaptic connections between the glutametergic thalamocortical (TC) neurons in the LGN and GABAergic reticular (RE) neurons in the PGN, and intrinsic transient low-threshold calcium "T"-currents that when activated cause bursts of action potentials in both of these cell types[1]. Lateral connections in these structures serve to propagate this activity.

Studies of simulated integrodifferential equation models of propagating "spindle" waves in one-dimensional networks comprised of distinct thalamocortical (TC) and reticular (RE) cell layers[2,3] have observed both smooth and lurching wave pulses depending on the types of synaptic connectivity used. In the present study we use a reduced two-layer TC-RE network comprised of integrate-and-fire-or-burst (IFB) neuronal populations to explore the existence of smooth and lurching pulses analytically. First, we show under generic conditions that smooth pulses cannot exist with an "on-centered" synaptic connection topology consisting of a wide RE to TC synaptic footprint and narrow TC to RE synaptic footprint. Second, we solve for exact smooth and lurching pulse singular solutions in two-layer TC-RE IFB networks. To this end we use techniques similar to those employed by Coombes (2003) to study wave-pulses in single-layer IFB networks in combination with singular perturbation methods where synaptic variables tend quickly to their asymptotic values. Smooth pulses are solved for in the case of off-center wide TC to RE and narrow on-center RE to TC footprints. We solve for lurching pulses in the case of wide on-center RE to TC and narrow on-center TC to RE synaptic footprints. These solutions corroborate the numerical observations in Terman, Ermentrout, and Yew (2001). Furthermore, these exact solutions provide the wave speed of the pulse and how the speed depends on parameters of the synaptic footprint and physiological parameters. The solutions show that the speed of wave propagation in the smooth case depend on parameters in the RE layer whereas the propagation in the case of a lurching pulse depends primarily on parameters in the TC layer. Additionally these analytical results corroborate the findings from simulations of biophysically detailed models studied in Golomb, Wang, and Rinzel (1996), where the speed of the lurching front depends primarily on sum of the width of the synaptic footprints between the TC and RE layers and the T-current decay time constants.

References

1. Destexhe, A., Sejnowski,. Thalamocortical Assemblies, Oxford (2001).

2. Golomb, D., Wang, X., Rinzel, J.. Propagation of spindle waves in a thalamic slice model. J. Neurophysiol. 75: 750-768 (1996).

3. Terman, D.H., Ermentrout, G.B., \& Yew, A.C.. Propagating activity patterns in thalamic neuronal networks. Siam J. Appl. Math. Vol. 61, No. 5, pp, 1578-1604, (2001).

4. Coombes, S.. Dynamics of synaptically coupled integrate-and-fire-or-burst neurons. Physical Review E, 67, 041910 (2003).

5. Golomb, D., Ermentrout, B.. Continuous and lurching traveling pulses in neuronal networks with delay and spatially decaying connectivity. PNAS, 96 (23), 13480-13485, (1999).

6. Rinzel, J., Terman, D., Wang, X.J., Ermentrout, B.. Propagating activity patterns in large-scale inhibitory neuronal networks. Science, 279 (5355), 1351-1355 (1998).

206. Inferring causal subnetworks using point process models

Duane Q Nykamp

University of Minnesota

If one could measure simultaneously and individually the spiking activity of all neurons in a neural network, fitting a network model to the data might reveal causal connections among the neurons. A network model will contain terms corresponding to interactions among neurons, and the resulting values of these interactions terms may correspond to causal connections among the neurons. However, the reality is that one can measure the individual activity of only a small fraction of neurons. The presence of large numbers of unmeasured neurons could create the illusion of causal influence among the measured neurons when no such causal influence exists. Connections from unmeasured neurons could introduce dependencies among the activity of measured neurons. When fitting a network model to the activity of just the measured neurons, these dependencies could lead to large interaction terms of the model. In many cases, the interaction terms would not correspond to any causal influence from one measured neuron onto another measured neuron.

The simplest example of such a network configuration is a common input configuration where a single unmeasured neuron projects to two measured neurons. This common input configuration contains no causal influences between the two measured neurons. However, if one fit a network model to the spikes of just the two measured neurons, the resulting parameters would indicate an interaction between the two measured neurons. If the connection from the unmeasured neuron onto measured neuron one had a longer delay than the connection from the unmeasured neuron onto measured neuron two, the interaction term of the model could be erroneously interpreted as a causal influence from measured neuron two onto measured neuron one.

We have developed a framework through which one can begin to distinguish influences of unmeasured neurons in order to determine if observed interactions do indeed correspond to causal influence among the measured neurons. Our goal is to determine the causal subnetwork among the measured neurons, i.e., the network of causal influences among the measured neurons that is embedded in the larger network. The approach exploits predictions from a point process model of the relationship between neuron spikes and external variables such as a stimulus. Hence it is limited to applications where the one can find such a model that describes a neuron's response sufficiently well. Nonetheless, we have developed a modular approach where one can "plug in" many different models into the causal network analysis. The models can be of a relatively large class of point process models, and the models can be developed independent of the causal network analysis. Thus, the approach can be potentially applied to a large range of experiments where the spikes of multiple neurons are recorded simultaneously.

We have currently developed an analysis to distinguish common input configurations from causal connections between two measured neurons. The distinction can be made because the joint statistics among the two measured neurons' spike trains and the external variables should differ between the common input configuration and a causal connection configuration. Under certain circumstances, there is some ambiguity in the detection of which network configuration underlies the observed spikes and external variables. We argue that in many experimental contexts, the interpretation of the resulting causal connections is not affected by this ambiguity.

207. An Algebraic Approach to the Analysis of Network Functional Connectivity: Application on data from the basal ganglia

Naama Parush, Gali Heimer, Hagai Bergman, Naftali Tishby

Hebrew University

A major goal of many neurophysiological studies is to understand the collective activity of large populations of neurons. Mainly, however, only single cell activity and pairwise correlations can be reliably estimated from current experimental data.

The pairwise correlation of neurons in the globus pallidus nuclei has been the focus of several studies concerning the basal ganglia and Parkinson's disease [1,3]. These studies did not examine the relations between the correlations of different pairs of neurons, nor did they explore the dynamics of the correlations along time. Consequently, they disregarded the temporal and spatial behaviors of the correlations in the data.

In this work we propose a new approach to the analysis of network functional connectivity based on second order statistics. By simultaneously studying the connectivity of all pairs of neurons, we extract information on the connected components of the network and explore the dynamics of the network's connectivity. In addition we extend research by Heimer et al. [1] on globus pallidus neurons of Parkinsonian primates and expand it from pairwise correlations to the analysis of connectivity dynamics.

We first define a network connectivity representation based on all the distances between the activities of all pairs of neurons, and then provide a normalized measure for comparing two such representations. These tools enable us to study the dynamics of a series of momentary network functional connectivities.

The representation

The goal is to construct a representation that will capture the inter-organization of the neurons. Hence the representation is defined as a low dimensional embedding of the pairwise connectivity matrix that preserves maximal information on the connected components of the network.

The connected components of a network (graph) can be extracted from the spectral decomposition of the Laplacian of the similarity matrix (connectivity Laplacian) [2], where the similarity values between the neurons (S) are based on their correlation coefficient. Therefore, in a network of n neurons, we represent the network connectivity by an nxk matrix consisting of the k (k<n) orthonormal eigenvectors corresponding to the k lowest eigenvalues of the normalized connectivity Laplacian (we use the normalized Laplacian for convergence reasons). We regard k as the number of connected components in the network.

This representation preserves the inter-organization of the neurons. It can be shown that

$$\sum_{i,j} \| f(i) - f(j) \|^2 S(i,j) = trace(f^T L f)$$

, where L is the connectivity Laplacian. The minimal value of trace(f^TLf) is attained when f is the eigenvector corresponding to the lowest eigenvalue of L. Thus, the Euclidean distance between the embeddings of neurons with high similarity values will be small and vice versa. This embedding approach is similar to the approach taken in spectral clustering algorithms [2].

Normalized distance measure

We then define a normalized measure for comparing two such representations. The embedding is invariant to rotations (since the eigenvectors are arbitrarily chosen to be orthonormal); consequently it is not the specific embedding alone that represents the network but the whole subspace spanned by it. Therefore when comparing two networks we compare the two different subspaces spanned by their embedded representations. We define the distance between two networks ("network distance") as the square sine of the angle between the networks' embedding subspaces.

Using the new approach on recorded data from the basal ganglia

In a previous study, Heimer et al. [1] showed that dopamine replacement therapy (DRT) reverses abnormal synchronization of pallidal neurons in the (MPTP) 1-Methyl-4-Phenyl-1,2,3,6-tetrahydropyridine primate model of Parkinsonism. We divided the recorded spike train data into adjacent time windows, and calculated the correlations between each pair of neurons at all time windows. Examining the correlations along time revealed that in 6 out of 8 recording days during the clinical 'off' state of Parkinson's disease, a burst of synchrony that involved several neurons and lasted a number of seconds (~5s) repeated itself every 1 to 2 minutes. This finding provides temporal and spatial resolution to the pairwise synchrony in the globus pallidus of the MPTP treated Parkinsonian monkey. Such results could not have been predicted otherwise, and they call for further study of the unique mechanism (time constant > 1 minute) underlying this observation.

We conclude that synchronized multi neuron bursts are typical of MPTP treated primates, and that this neuronal behavior diminishes under the influence of DRT. The synchronized activity of several neurons forms a unique network connectivity state in which a cluster of cells is highly connected. Using our new method we are able to find the repetitions of this connectivity state during the entire recording time, and use this information to trace the neural-network correlates of the clinical transition between 'off' and 'on' periods.

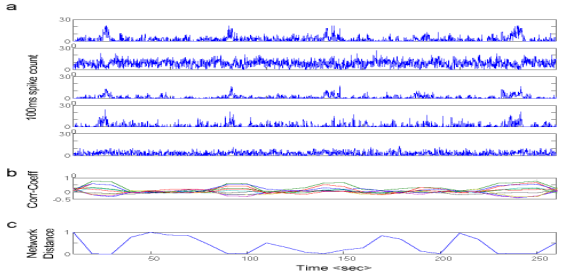

Fig. 1 - Simultaneous recordings of 5 neurons from the globus pallidus of a MPTP treated monkey, while the monkey was in a dyskinetic Parkinsonian 'off' state: (a) 100ms spike counts. (b) All pairwise correlation coefficients as a function of time (each of the 10 lines illustrates the correlation coefficient of a different pair). (c) "network distance" from the first synchronized activity state – the method found 5 repetitions of the synchronized state.

208. Brute-force computational exploration of calcium-based activity sensors in a model pattern-generating network

Astrid A Prinz

Emory University

Neural circuits, and especially pattern-generating circuits, perform their function reliably throughout an animal's life, in spite of ongoing molecular turn-over and developmental and environmental changes. At the cellular and synaptic level, this robustness has been shown to be at least in part due to activity-dependent regulation, i.e. the adjustment of cellular and synaptic properties in response to activity perturbations in a direction that acts to restore a target activity level.

Intracellular calcium dynamics are known to play an important role in activity-dependent homeostatic regulation, because calcium levels correlate well with a neuron's electrical activity, while at the same time playing a major role in intracellular signaling cascades.

At the cellular level, models of calcium-based activity sensing and homeostatic regulation have been successful at stabilizing model neuron activity in the face of a variety of perturbations. However, generalizing this cell-based activity-dependent regulation to the network level has been challenging. Calcium-based activity sensors and regulatory pathways that work in models of individual neurons fail to stabilize network activity when implemented at the circuit level, which raises the question whether network-level activity-dependent homeostasis is possible solely on the basis of local, cell-based activity sensors and regulatory mechanisms.

To address this question, I simulated 354 different calcium sensors in each of the three cells of 20 million different models of the crustacean pyloric pattern-generating circuit, a network that generates a triphasic, rhythmic burst pattern. The 20 million networks differed in their cellular and synaptic properties and had previously been classified into 2.2% functional, "pyloric" networks and 97.8% non-functional networks based on whether they produced a rhythmic pattern that matched the biological pyloric rhythm for a variety of criteria. The 354 sensors differed in their sensitivity to calcium and in their temporal dynamics, covering a wide range of threshold concentrations and time constants.

I subsequently scanned the resulting database of sensor sets for sensor combinations that successfully distinguish functional from non-functional networks by providing non-overlapping sensor readings for the two categories. Such successful sensor sets form a potential basis for cell-based homeostatic regulation of network activity, and analysis of their common features will suggest constraints on the overall architecture of intracellular calcium-based regulatory cascades for homeostatic network regulation.

209. Retrieving the identity of a visual object while keeping information about its position in a model of IT cortex

Yasser Roudi, Alessandro Treves

Cognitive Neuroscience sector, SISSA

Since we can effortlessly recognise objects regardless of their position in the visual field, it may be thought that their neural representation in the final stages of the visual form processing stream in inferior temporal cortex (IT) should be independent of their position. In this view, the response of IT neurons, when retrieving object identity, should not contain any information about where the object is located, but only about identity.

Some studies have provided evidence for such position invariant representations in IT, whereas the results of some others have been discordant. The disagreement between studies has been associated to different experimental paradigms, relating mainly to the size of the stimuli to the behavioural task, to effects of learning and attention.

Here we show that a recurrent network, a model that has been considered in many previous studies of IT, is indeed capable of faithfully representing the spatial position of objects while retrieving information about their identity. In other words, the network can simply respond to a visual cue without caring about where the cue had been presented, as previously proposed, or it can retrieve object identity and distinguish it from that of other objects (what information) but while keeping at the same time information about the position of the cue (where information).

There are two essential ingredients that we add to a recurrent network model of IT to make it capable of representing both where and what information. These two ingredients are the existence of a minimum degree of metric organisation in the recurrent connections, and a spatially localised signal that changes the linear gain of neurons in a certain part of the network, corresponding to the position of the object in the world.

When considering solely metric recurrent connections, we have previously shown that a pattern of activity that had been stored without any spatial preference can be retrieved in the form of a localised bump of activity. However, such bumps of activity cannot be localised, it turns out, anywhere in the network: they rather favour a few distinct positions. We show here that a local change in the gain of neurons stabilises the retrieval bump at positions, that would have otherwise been unstable. In this way, the activity of the network can reflect where information without interfering with what information. Changes in the gain of neurons can in fact be produced by attentional mechanisms and/or by short time changes in the responsiveness of the neurons that receive the cue, through local inhibitory mechanisms or intrinsic membrane currents.

210. Distal gap junctions and active dendrites maximize stable, phase-locked network dynamics

Fernanda Saraga[1], Leo Ng[2], Frances K Skinner[3]

[1]Department of Zoology, Department of Physiology, Toronto Western Research Institute, University Health Network, University of Toronto, [2]Engineering Science Program, University of Toronto, [3]Toronto Western Research Institute, University Health Network, Department of Medicine (Neurology), Department of Physiology, Institute of Biomaterials and Biomedical Engineering, University of Toronto

Gap junctions (GJs) allow direct electrical communication between CNS neurons including those in hippocampus and neocortex. From theoretical and modeling studies it is well-known that although GJs can act to synchronize network output, they also allow many other network dynamic patterns, including antiphase and other phase-locked states, to occur. The particular network patterns that arise depend on cellular, intrinsic properties, which affect intrinsic firing frequencies, as well as the conductance strength and location of the GJs. Interneurons or GABAergic neurons in hippocampus are diverse in their cellular characteristics, and have been shown to have active dendrites. Furthermore, parvalbumin-positive GABAergic neurons in hippocampus, also known as basket cells (BCs), are known to contact one another via GJs on their distal dendrites. Using two-cell network models, we explore how distal GJ connections affect network output. We build compartmental models of hippocampal BCs using NEURON, and endow them with varying amounts of active dendrites. With homogeneous networks we find that active dendrites at specific levels work together with distal dendritic GJ connections to maximize the range of phase lags in phase-locked states for GJ conductance variations. These results hold when heterogeneity in the form of different levels of active dendrites is added to the two-cell network system. However, the range of possible phase lags is reduced, and the stability of phase-locked states is more restricted. In conclusion, our modeling work suggests that if distal, dendritic GJs are to give rise to stable, phase-locked, near-synchronous states, then active dendrites are required.

Support Contributed By: NSERC Canada

211. Periodic Bursting in Two Identical Coupled Cell Systems

LieJune Shiau[1], Marty Golubitsky[2], Kresimir Josic[2]

[1]University of Houston-Clear Lake, [2]University of Houston

Analysis in terms of fast-slow dynamical systems gives insights into mechanisms for generation of bursting behavior. Bursting behavior in neurons can be generated due to interactions between neurons in a network. For example, brain functions such as motor control, information processing and memory formation frequently involve bursting behavior of neurons. On the other hand, periodic bursting in fast-slow system can be viewed as closed paths through the unfolding parameters of degenerate singularities. Using this approach we show that bursting in coupled systems can have interesting behavior. Here we study and focus on the analysis of two identical cell systems, and use the Z_2 symmetry present in such systems to illustrate surprisingly interesting bursting phenomena. In particular, we show that Hopf bifurcation and Hopf bifurcation mode interactions can lead to bursting between in phase and out of phase periodic solutions, and symmetry-breaking Takens-Bogdanov singularities can lead to bursting that randomly chooses between two symmetrically related limit cycles.

212. Why Should Cortical Connectivity be Sparse and Predominantly Excitatory?

Armen Stepanyants, Darin R La Sota

Northeastern University

Two features of synaptic connectivity are common among many cortical circuits. First, cortical circuits are sparsely interconnected, where the probability for nearby neurons to be in synaptic contact is typically on the order of 10%. Second, cortical networks are predominantly excitatory with only about 20% of inhibitory synapses. In this study, we investigate the effect of sparseness and inhibition on the dynamics of artificial neural networks. We look at the expected numbers of attractors of different lengths and their corresponding basin sizes. Our results, based on the networks of 10-15 neurons, indicate that the values of sparseness and inhibition typically observed in the cortex maximize the expected numbers of attractor states and consequently result in richer dynamic properties of the networks.

213. Topological design of cortical networks that display power-law statistics of neuronal avalanches

Jun-nosuke Teramae, Tomoki Fukai

RIKEN Brain Science Institute

How cortical neurons process information crucially depends on how their local circuits are organized. Spontaneous synchronous neuronal activity propagating through cortical slices displays highly diverse, yet repeatable, activity patterns (“neuronal avalanches”). They obey power-law distributions of the event sizes and lifetimes, presumably reflecting the underlying structure of cortical networks. However, the explicit network structure and how such a network may be developed in a population of neurons remain unknown. Here, we present a network model that enables stable propagation of avalanche-like neuronal activity. We demonstrate a neuronal wiring rule that governs the cell-assembly formation during the development of this network model. The resultant network comprises predominantly feedforward chains of mutually overlapping cell assemblies with a wide variety of the size and length. In addition, our model exhibits the so-called small-world-like characteristics which are widely observed in real-world networks (large clustering coefficient and short characteristic path length). Our findings suggest that cortical circuits may have a more complex topological design than has previously been thought. We also discuss the possible circuit structures of inhibitory neurons and demonstrate that the balance between excitatory and inhibitory activities is an essential factor of stable propagations of activities and power-law statistics of them.

214. Initial neuronal group activity is precisely maintained during propagation within neuronal avalanches in vitro

Tara C Thiagarajan, Dietmar Plenz

NIMH/NIH

Cognitive processes, such as memory retrieval, have been shown to be fast and engage distributed groups of neurons on millisecond time scales. However, the engagement of neuronal groups and stable propagation of signals among them without decay or explosion has been difficult to achieve in network models. Recently, our lab described a previously unknown state in cortex, recorded in vitro from rat cortex, where spontaneously occurring synchronized groups of neurons, measured as local field potentials (LFP), rapidly triggered other synchronized groups to form irregular spatio-temporal patterns that were neither wave-like, nor rhythmic. We termed this propagated synchronized activity a 'neuronal avalanche'. One hallmark of the dynamics of this activity is that avalanches sizes (i.e. the number of electrodes engaged in the avalanche) are distributed according to a power law with an exponent of -1.5 (Beggs and Plenz, 2003). We now demonstrate an exceptional property of neocortical networks to maintain the precise size and pattern of synchronized activity during propagation within these neuronal avalanches. Avalanche activity was measured in organotypic cortex cultures (n = 6) and acute rat frontal cortex slices (n = 6; 3µM NMDA + 3µM SKF38393) using an 8x8 multi-electrode array. Naively, one would expect larger initiating LFPs to result in avalanches with greater spatial spread and longer lifetimes. To the contrary we found that small synchronized groups were equally likely to rapidly engage distant regions in the network as large groups; the distribution of avalanche sizes followed a similar power law with exponent -1.5 for all initial LFP areas. Increases in LFP area that occurred at any step in the process were restored in the subsequent step by an intrinsic mechanism that was independent of inhibition. In contrast, decreases were tolerated. Significantly, it was not only the area of the synchronized neuronal group or LFP that was preserved during propagation but also its precise waveform, suggesting the precise propagation of an underlying spatio-temporal pattern of activity. This particular feature was virtually destroyed by reducing GABAA inhibition. The preservation of the initial group properties imposes a limitation on the engagement of widely distributed neuronal groups. At the same time, the independence of the engagement process from group size introduces a degree of freedom in which the properties of a single synchronized group could be used to encode a stimulus without affecting the inherent propagation process.

215. Comparison of Neural Circuits that Estimate Temporal Derivatives

Bryan P Tripp, Chris Eliasmith

University of Waterloo

There has been great interest in how neural circuits perform basic mathematical operations, such as multiplication and integration over time. Mechanisms of temporal differentiation have not been analysed as extensively, despite the fact that differentiation appears to be an important neural process, for example in the vestibular system, and in cardiovascular and respiratory control. Here we characterize the performance of five neural circuits that perform differentiation, in terms of distortion errors and rejection of high-frequency noise. The circuits' frequency responses are of interest because an ideal differentiator scales periodic input by its frequency, so that a small amount of high-frequency input noise can obscure useful output. Circuits are analysed by treating each projection as a source of distortion error and noise, and determining how these errors propagate to the output. Example simulations are performed using ensembles of 500 neurons to represent input variables, and 2000 neurons distributed among the ensembles that perform differentiation. Five circuits are compared. Circuits 1-4 consist of spiking leaky-integrate-and-fire (LIF) neuron models, with exponentially decaying post-synaptic currents. In the first three, differentiation is achieved through parallel excitatory and inhibitory projections in which the inhibitory projection is lagged in some manner (i.e. by interneurons, conductance delay, and longer post-synaptic time constants, respectively). The fourth is a more sophisticated circuit that implements a 2nd order band-pass Butterworth filter. Circuit 5 consists of adapting and potentiating LIF neurons. In this circuit, synaptic weights are set so that the different dynamic responses of the neurons sum to approximate the same transfer function as that of circuit 4. This approximation holds for a family of inputs, but breaks down with sufficiently novel input. Different circuits have different advantages. Distortion errors are lowest in circuits 2 and 3, which resemble the differentiator in the vestibular system, but only if the strength of parallel excitatory and inhibitory weights can be modulated together. Circuit 4 is superior in the rejection of high-frequency noise. Circuit 5, which exploits short-term adaptation dynamics, is potentially very accurate for a limited family of inputs. Together, these results suggest that different mechanisms may be optimal within different neural systems.

216. The Role of Precise Thalamic Spike Timing in Generating Nonlinear Cortical Responses in the Rat Vibrissa System

Roxanna M Webber[1], Garrett B Stanley[2]

[1]Harvard - MIT Division of Health Sciences and Technology, [2]Division of Engineering and Applied Sciences, Harvard University

Rats and other rodents use their vibrissae to actively explore the external environment and can discriminate between similar textures using only their vibrissae. As a rat whisks against an object, patterns of vibrissa deflections are produced that reflect the interaction between the vibrissa movement and the textural properties of surfaces. The peripheral representations of the stimulus are transformed by the excitatory-inhibitory interactions of the thalamocortical circuitry, producing representations that eventually give rise to the sensory percept. We previously showed that the response of cortical neurons to temporal patterns of vibrissa deflection was nonlinearly dependent upon the amount of suppression induced by previous deflections. Specifically, many neurons exhibit a "lifting of suppression" phenomenon in which if the excitatory response to a given deflection is suppressed by prior activity, the suppression that would normally be induced is also decreased and allows for a larger response to a subsequent deflection. This phenomenon results in nonlinear neuronal responses to complex stimulus patterns and is critical in the prediction of responses to arbitrary stimuli. It is not clear, however, if the observed cortical responses are generated through intra-cortical connections or partially inherited from thalamic projections. Lifting of suppression was observed for a significantly smaller percentage of neurons recorded in the ventral posterior medial (VPM) nucleus of the thalamus as compared to cortex, implying that the lifting of suppression phenomenon may be generated through thalamocortical transformations. However, further analysis of the thalamic responses revealed a nonlinear modulation of response timing, which suggested that precise timing of thalamic spikes may contribute to the cortical phenomenon. Here, we explored these thalamocortical transformations through a simple biophysical network model consisting of reciprocally connected excitatory and inhibitory neurons. Responses of single units in the VPM to temporal patterns of vibrissa deflection were used as templates for the input to the cortical excitatory and inhibitory neurons. The model gives rise to many of the cortical phenomena observed experimentally in response to temporal deflection patterns. Therefore, we can use this model to test the hypothesis that precise timing of thalamic inputs can cause excitatory cortical neurons to fire before feed-forward cortical inhibition suppresses further spikes and investigate which aspects of the cortical response are due to changes at the thalamocortical synapse versus intra-cortical connections.

This work was supported by NIH R01NS48285-01A1 and an NSF graduate research fellowship to R. M. Webber.

217. Synaptic Mechanisms of Thalamocortical Auditory Processing

Mike Wehr

University of Oregon

Receptive fields depend critically on the interplay of excitation and inhibition. An explanation of how excitatory and inhibitory local circuits shape cellular response properties is fundamental to our understanding of sensory information processing in the brain. Yet virtually all studies of this subject have used indirect methods to assess inhibition. For example, the effects of inhibition are often inferred from suppression of firing, such as that caused by stimuli outside the center of the receptive field. This phenomenon is called *surround suppression*, and is often taken to be evidence of a particular configuration of inhibitory circuitry known as *lateral inhibition*. However, by measuring excitation and inhibition directly (using *in vivo* whole cell methods), we have previously shown that excitation and inhibition are balanced throughout the receptive field in auditory cortical neurons. Thus lateral inhibition does not occur in local circuits of the auditory cortex. This suggests that the surround suppression seen in cortical cells may be inherited from lower levels, such as the auditory thalamus. If so, we should observe lateral inhibition in thalamocortical neurons. We are currently recording intracellularly from thalamocortical neurons to test the hypothesis that lateral inhibition in local thalamic circuits generates the surround suppression seen in cortical neurons downstream.

218. Recurrent Network Models for Working Memory of Temporal Sequences

Olivia L White[1], Avigail Ben Or[2], Haim Sompolinsky[2]

[1]MIT, Department of Physics, [2]Hebrew University

While it is thought that firing patterns of neurons in frontal cortical areas subserve short-term memory, both how this is achieved and how to characterize these firing patterns are unknown. Although the capability of recurrent neuronal networks to integrate past stimuli has been the subject of considerable theoretical and and experimental work, the capability of recurrent circuits to reconstruct the history of past stimuli is poorly understood. In this work we investigate the short-term memory capacity of biologically plausible network models. We first consider rate-based models, extending our recent theoretical work (Phys. Rev. Lett. 92:148102 (2004)), we characterize the requirements on recurrent synaptic connections that are necessary for generating a sizable short-term memory capacity. Our theory predicts that short-term memory networks have recurrent synapses that are primarily inhibitory with a substantial component of anti-symmetric connectivity. Further, we develop a method for analyzing the robustness of memory capacity in recurrent networks to noise and structural perturbations. We compare the model behavior with known features of sequence recall and recognition in human memory tasks, such as recency, primacy, and sequence length effects. We then explore extensions of our theory to spike-based network models. We show that the memory capacity of a recurrent network of integrate-and-fire neurons does not grow with the size of the network. We conclude by proposing physiological measurements on groups of active neurons in frontal cortex that could distinguish transient responses from activity corresponding to a dynamical fixed point.

Acknowledgements: O.W. acknowledges support of the NSF and NIH. H.S. acknowledges support of the ISF and BSF.

219. Stability analysis of cooporative algorithms

Junmei Zhu

Computer Science Department, University of Memphis

Cooperative computation achieves a global order of a system through local interactions of its elements. Many cooperative algorithms have been developed for stereo correspondence, for example [1,3]. Different algorithms differ in the form of local interactions, i.e., the constraints represented by the excitatory and inhibitory regions. There has been considerable debate on which constraint is correct, and the criteria for comparison are rather subjective. A more objective criterion is to compare the final stable states of the systems.

General results on the convergence of cooperative algorithms are difficult to obtain. The convergence of small discrete systems can be obtained by combinatorial analysis [1]. Marr et al. developed a probabilistic analysis [2], but it applies only to inputs with a random structure. The standard approaches to dynamical system analysis fail because the Lyapunov function does not exist in general [2]. There are some analytical tools developed in the context of synergetics [4] and pattern formation. The idea is to represent the state as a superposition of modes that are solutions to the linearized dynamical equations, and study the dynamics of mode coefficients. However, the analysis is performed only at the initial homogenous state, and similar analytical results are difficult to obtain at desired final states.

In this work we present a method to analyze the stability of any fixed point in a dynamical system. This method is essentially a numerical version of the linear analysis in the synergetic approach. We consider a continuous formulation of the cooperative algorithm. Local interactions, both excitatory and inhibitory, are represented in kernel K. The change of a state W is modeled by the convolution of K with W passing through a nonlinear function f, commonly a sigmoid function, plus the initial value W0 of W:

$ \dot W = -W + K* f(W) +W0 $

At any state W*, we can compute the Jacobian J of this system evaluated at W*. W* is a stable state if all eigenvalues of J have negative real parts. In addition to determining stability, the eigenvalues also provide information of the local landscape around the fixed point, such as the speed of convergence.

This numerical method of stability analysis is very general. Here we apply it to compare different stereo cooperative algorithms. Especially we test whether true disparity maps are stable states in the systems.

[1] D. Marr and T. Poggio, Science, 194:283-287, 1976.

[2] D. Marr, G. Palm and T. Poggio, MIT AI Memo No. 446, 1977.

[3] M. Arbib, C. Boylls, and P. Dev, Cybernetics and bionics, 216-231, 1974.

[4] H. Haken, Synergetics, 1977.

220. Temporal Processing and Adaptation in the Zebra Finch Auditory Forebrain

Katherine Nagel, Tatyana Sharpee, Allison J Doupe

UCSF

Adaptation to stimulus statistics allows sensory systems with a limited number of neurons to encode a larger range of stimuli. To examine the effects of stimulus statistics on temporal processing in the songbird primary auditory area Field L, we probed neurons (n=35) with a stimulus that samples a range of amplitude modulations similar to those in song. Using reverse correlation we characterized the responses of single units in terms of a linear filter and a nonlinear gain function. The filter represents the stimulus feature that best drives the neuron, while the gain function describes the dependence of neural firing on filter output and captures the threshold and gain of the cell. We then examined how changes in the mean and variance of amplitude modulations affected these two parts of the neurons' responses.

Neurons showed distinct types of changes in processing to increases in mean and in variance. An increase in mean led to a rapid change in firing rate, followed by a slow decay to a new mean rate. Under high mean conditions, filters became narrower (faster) and more differentiator-like--with stronger inhibition--, while several cells showed inversions of the filter shape, e.g. from onset to offset sensitivity. In contrast, an increase in variance produced fewer changes in firing rate and filter shape, but reduced the gain of the nonlinearity in most cells. By extracting filters and nonlinearities from different time points after the change in statistics, we found that changes in filter shape and nonlinear gain occur within 100msec of the statistical change, suggesting that they represent nonlinearities of processing, rather than time-dependent adaptations. Horizontal shifts in the nonlinearity occur more slowly, and may correspond to the slow changes observed in the mean firing rate. Both the increase in inhibition with increased mean, and the reductions in gain with increased variance, may permit neurons to signal effectively over a wider dynamic range, and are reminiscent of similar findings in other sensory systems. These results suggest that similar computational principles shape processing across sensory systems.

Supported by HHMI, Swartz Fdn, NIH

221. Choice probabilities in V2 reflect task strategy, as measured psychophysically

Hendrikje Nienborg, Bruce G Cumming

Laboratory of Sensorimotor Research/NEI/NIH

We recently reported that a substantial fraction (41%) of disparity selective neurons in V2 is correlated on a trial-by-trial basis with the monkeys' choices in a disparity discrimination task (quantified as choice-probabilities). Task difficulty was varied by changing the proportion of correlated dots at the signal disparity. Here we show that these choice probabilities occur predominantly in neurons tuned for near disparity, resulting in a negative correlation between the preferred disparity and choice-probability (since we define near disparities as negative; $r=-0.48$, $p<0.001$). Restricting the analysis to neurons whose maximum response was to a far disparity ('far-preferring') yields no significant choice-probability on average (mean choice-probability is 0.50, $n=25$), while the mean choice-probability for near-preferring neurons is 0.61 ($n=31$). One interpretation of this difference is that the monkeys' strategy in the task was to use mainly the responses from near-preferring neurons while ignoring the activity from far-preferring neurons. We tested this hypothesis psychophysically using reverse correlation. As in the recording-experiments, stimuli were dynamic random dot stereograms whose circular center consisted of 'signal' and 'noise' dots, surrounded by an annulus at zero disparity. The signal dots were presented either at a near or at a far disparity, which the monkey had to discriminate in a forced choice task. (For stimuli consisting entirely of noise, the monkeys were rewarded randomly.) Noise dots were drawn randomly from a discrete distribution of disparities. We sorted the stimuli according to the animals' choice, and calculated the average distribution of the noise disparities for 'near' and for 'far' choices. For both monkeys, these distributions show a pronounced peak at near disparities for near choices, and a trough at those same near disparity for far choices. Conversely, the distribution of far disparities is relatively flat for both types of choice. This indicates that the monkeys treat near and far disparities differently in this task. They appear to choose near in the presence of near disparity dots and far in the absence of near disparity dots, while ignoring far disparity dots for both choices. This asymmetrical feature of the monkeys' behavior mirrors the asymmetry of the choice probabilities. This suggests that the choice probabilities inform us not only about the usefulness of neuronal signals for a particular task, but also about the way in which those signals are actually used by the subject.

222. Visual working memory and attention in early visual cortex

Shani Offen, Denis Schluppeck, David J Heeger

NYU

Objective:

We investigated whether early visual cortex (V1-V3) plays a role in maintaining visual working memory and/or visual attention by measuring activity during a delay period. By varying delay duration we could distinguish delay period activity from sensory and motor responses.

Methods:

Subjects were scanned (3T fMRI, BOLD) while performing each of four tasks (two of which are illustrated in Figure 1) designed to probe visual working memory, attention, or both. Delayed comparison: A high-contrast grating (randomized orientation and spatial frequency) was presented within an annulus (1-3°) around fixation. It was followed, after a variable delay (1-16s), by a second high-contrast grating with near-threshold change in orientation and spatial frequency. Fixation point color then cued subjects to indicate either the orientation or spatial frequency change. Cued detection: Stimuli were identical, except the contrast of the final target grating was at detection threshold and subjects indicated its presence or absence. Detection: Same as cued detection except the probe and test orientations were chosen independently. Discrimination: Stimuli were identical to the detection task except the target had one of two orientations (tilted slightly right or left of vertical), and subjects indicated its orientation. Retinotopic maps (V1-V3) were determined using standard procedures and analyses were restricted to the cortical representation of the stimulus, separately for each visual area.

Results:

Despite the similarity between the stimuli and tasks, we found a striking difference in cortical activity (shown in Figures 2 and 3). Areas V1-V3 exhibited no sustained activity during the delay period for the delayed comparison task (which required working memory for easily visible stimuli), but robust sustained activity for the other three tasks (which all had barely visible targets). The sustained response was restricted to the retinotopic area responsive to the stimulus annulus, as shown in Figure 4.

Conclusions:

The results imply a dissociation between working memory and attention. We hypothesize that attentional selection, critical for scrutinizing stimuli that are barely visible, is implemented by sustained boosting of relevant neuronal signals in early visual cortex.

Support contributed by: NEI (R01-EY11794), NIMH (R01-MH69880), NSF GRF, and the Seaver Foundation.

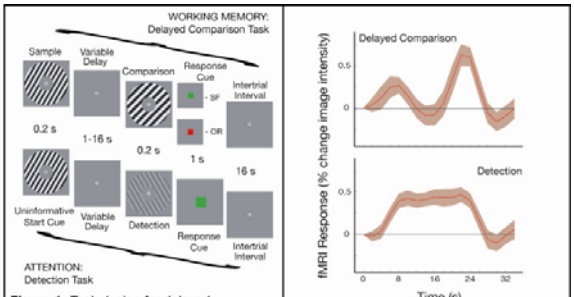

Figure 1: *Task design for delayed comparison and detection tasks*

Figure 2: *Example V1 responses. Activity returns to baseline for the delayed comparison task, but remains elevated for the detection task.*

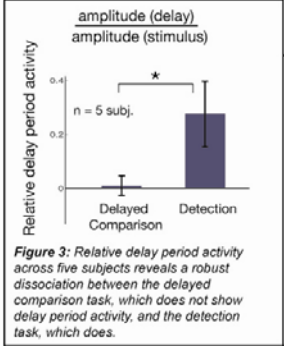

Figure 3: *Relative delay period activity across five subjects reveals a robust dissociation between the delayed comparison task, which does not show delay period activity, and the detection task, which does.*

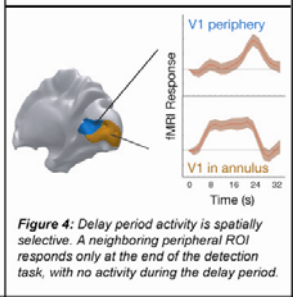

Figure 4: *Delay period activity is spatially selective. A neighboring peripheral ROI responds only at the end of the detection task, with no activity during the delay period.*

223. Early Visual Responses: More than "Low-Level Features"

Cheryl A Olman

University of Minnesota

It is not uncommon for a researcher studying image perception to want to match “low-level features” of images in a set, to eliminate effects of early visual responses in studies of perception. However, there is a necessary limit to this approach. As a trivial example: images matched for RMS contrast can have very different local contrast properties (Fig. 1A and D). Similarly, matching Fourier amplitude spectra equates spatial frequency content but gives no control over spatial distribution of contrast energy. A reasonable goal might be to modify the images so that responses in an array of filters tuned to location, spatial frequency, and orientation have similar distributions, thereby controlling for both global spatial frequency content and the amount of distributed of local contrast energy (Fig. 1B and E). However, two problems are obvious. First, this manipulation can have a striking and often undesirable effect on image appearance. Second, these images may still fail to evoke comparable early visual responses. Whether contour completion occurs in V1 or beyond, phenomena such as pop-out effects with images composed of Gabor patches demonstrate that corners, T-junctions and curves, modulate early visual responses even in the absence of scene perception or object recognition. Responses to junctions and extended contours should therefore be included in “low-level” responses when generating an image set that can not be distinguished based on pre-attentive processes. Equating images for features like the number of corners and the length of curving line segments would distort most images beyond recognition. Therefore, either an impoverished stimulus set should be used, which can be matched for features that are known to modulate early visual areas, or the different responses of early visual areas should be allowed in the experimental model. A working model of V1 (Fig. 1C and F) can account for the feed-forward differences in an image set and permit quantitative study of feedback modulation of V1 responses to individual image features during scene segmentation and object recognition. Such a model of V1 is of course an incomplete approximation, but enables the use of rich image sets in studying both V1 responses and scene perception.

224. Bayesian model learning in human visual perception

Gergo Orban[1], Jozsef Fiser[2], Richard N Aslin[3], Mate Lengyel[4]

[1]Collegium Budapest, [2]Volen Center for Complex Systems, Brandeis University, [3]Department of Brain and Cognitive Sciences, University of Rochester, [4]Gatsby Computational Neuroscience Unit, University College London

There is a growing number of studies supporting the classical view of perception as probabilistic inference [1,2]. These studies demonstrated that human observers parse sensory scenes by performing optimal estimation of the parameters of the objects involved [3-5]. Even single neurons in primary sensory cortices have receptive field properties that seem to support such a computation [6]. A core element of this Bayesian probabilistic framework is an internal model of the world, the generative model. The generative model serves as a basis for inference by specifying how the different sources of currently available sensory evidence are integrated with prior expectations about the external world. Thus, in order to understand the computational principles of perception, it is important to characterize the forms of generative models that are available for perceptual inference.

Most previous studies testing the Bayesian framework in human psychophysical experiments used fundamentally restricted generative models of perception. The generative models considered in these studies consisted of a few observed and hidden variables, and only a limited number of parameters that needed to be adjusted by experience [3-5,7]. More importantly, these generative models were tailor-made to the specific psychophysical tasks presented in the experiments. However, in principle, inference can be performed on several levels: the generative model can be used for inferring the values of hidden variables from observed information, but also the generative model itself may be inferred from previous experience [8]. Thus, it remains to be shown whether more flexible, `open-ended' generative models could be used and learned by humans during perception.

We used an unsupervised visual learning task to show that a general class of generative models (a Sigmoid Belief Network variant of Automatic Relevance Detection in a Product of Gaussians) quantitatively reproduced experimental data, including paradoxical aspects of human behavior, when not only the parameters of these models but also their structure (ie. the number and identity of hidden variables) was subject to learning. We validated our theory by quantitatively fitting behavioral data from the same individual subjects in a range of different experiments. Crucially, the applied Bayesian model learning embodied the Automatic Occam's Razor (AOR) effect that preferred the models that were `as simple as possible, but no simpler'. This process led to the extraction of independent causes that efficiently and sufficiently accounted for sensory experience, without a pre-specification of the number or complexity of potential causes.

All the presented experimental results were reproduced and had identical roots in our simulations: the model that was most probable based on the training data developed hidden variables corresponding to the real chunks that were originally used to generate the training scenes. These results demonstrate that humans can infer complex models from experience and implicate Bayesian model learning as a powerful computation underlying such basic cognitive phenomena as the decomposition of visual scenes into meaningful chunks.

This work was supported by IST-FET-1940 program (GO), NIH research grant HD-37082 (RNA, JF), and the Gatsby Charitable Foundation (ML).

[1] von Helmholtz HLF. Treatise on Physiological Optics. New York: Dover, 1962.

[2] Barlow HB. Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. Vision Res 30:1561-1571, 1990.

[3] Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-433, 2002.

[4] Koerding KP, Wolpert DM. Bayesian integration in sensorimotor learning. Nature 427:244-247, 2004.

[5] Kersten D, Mamassian P, Yuille A. Object perception as Bayesian inference. Annu Rev Psychol 55:271-304, 2004.

[6] Olshausen BA, Field DJ. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature 381:607-609, 1996.

[7] Weiss Y, Simoncelli EP, Adelson EH. Motion illusions as optimal percepts. Nat Neurosci 5:598-604, 2002.

[8] MacKay DJC. Bayesian interpolation. Neural Comput 4:415-447, 1992.

225. State-dependence differences in evoked responses in auditory cortex

Gonzalo H Otazu, Anthony M Zador

Cold Spring Harbor Laboratory

The traditional view of primary auditory cortex is that it is mainly involved in representing auditory information. However, it is known from other primary cortical areas that an animal's attentional state affects the neural representation of the presented stimulus. We want to know how the state of alertness can affect neural representations in auditory cortex.

We compared two pairs of states: 1) anesthetized (kethamine-medetomidine) versus awake, and 2) active attention during an auditory task versus awake but not performing any task. Tetrode wires were implanted to allow chronic recordings in rat auditory cortex. The auditory stimuli were delivered through earphones in the unrestrained animal, allowing us to measure cortical responses to identical acoustic stimuli while the animal was free to move.

We observe a clear difference in the magnitude of both the local field potential and the level of multiunit activity; evoked responses were larger in the anesthetized state vs the awake state. Similarly, they were larger in the non-attending condition vs the attending condition.

Responses in primary auditory cortex are affected by the animal attentional state.

226. Illusory Percepts from Auditory Adaptation

Lucas C. Parra[1], Barak A. Pearlmutter[2]

[1]City College New York, [2]Hamilton Institute, NUI Maynooth, Ireland

We find that phenomena resembling tinnitus and Zwicker phantom tone result from an auditory gain adaptation mechanism that attempts to make full use of a fixed-capacity channel. In the case of tinnitus, the gain adaptation enhances internal noise of a frequency band otherwise silent due to damage. This would generate a percept of a phantom sound as a consequence of hearing loss. In the case of Zwicker tone, a frequency band is temporarily silent during the presentation of a notched broad-band sound, resulting in a percept of a tone at the notched frequency when the stimulus is terminated. In our model, the same mechanism leads to a transient phantom percept following the notched sound. The model predicts that the Zwicker tone should be abolished by a short masking sound in the notched band. To test this prediction, we performed a psychoacoustic study with short masking sounds following notched noise. The subjective responses match this prediction for subjects with self-reported tinnitus, but not for normal hearing subjects. This experimental result support the hypothesis that the Zwicker tone percept and tinnitus are linked, and support the theory that the phantom percept is a consequence of an adaptation mechanism confronted with a sensory apparatus which has been degraded below the level of performance for which the adaptation mechanism was designed.

Although the adaptation model we use is quite simple (assuming Gaussian noise and Gaussian signals independent at each frequency channel except for an overall signal strength modulated in common, with adaptation on both a short and a long time scale and a fixed target mean bit rate) the results should be robust, in that more sophisticated models that incorporate deeper knowledge of the structure of sounds should display the same qualitative phenomena---assuming they adapt at multiple time scales in an attempt to fully utilize a fixed capacity channel.

227. Physics of active touch: What does the vibrissa system sense?

Jason Ritt, Christopher I Moore

McGovern Institute for Brain Research, MIT

Every sensory process starts with transduction, and neural activity reflects not only the external object sensed but the physical properties of the sensor itself. The notion of "natural scenes" - that the nervous system evolved to exploit sensory regularities - therefore should include sensor as well as stimulus properties. Moreover, in many cases active sensor control refines sensory processing even before the onset of neural activity (e.g., via contact velocity during touch).

We explore these principles in the rodent vibrissa (whisker) system. Rats can perform high-precision tactile tasks, using a combination of body, head and "whisking" motions to engage objects of interest with their vibrissae. We previously demonstrated elastic (resonant) oscillations and friction-elicited "stick-slip" vibrissa motions, with corresponding tuning in cortical units in anesthetized animals. However, it remains unclear what contribution mechanical properties make, if any, to sensory behavior in the awake rat, or if rats actively manipulate the expression of vibrissa mechanics by adapting vibrissa motions to sensory context.

Here we describe high-speed (>3000 Hz) videography of rats in a 2-alternative texture discrimination task. We show the first direct measurement of texture-driven vibrissa responses, including resonance and stick-slip motions, in awake behaving animals. Oscillation frequencies appear to be vibrissa dependent, while the occurrence of oscillations depends to some degree on behavioral choice. We observe traveling waves from the vibrissa tip (at the object) to the base (and mechanoreceptors), with travel time ~1ms. We conclude, in extension of previous work, that vibrissae are generally underdamped, and develop significant contact friction. We discuss implications of this physical characterization for the "natural scene" for vibrissa sensation.

228. Using Illusions to Model Top-Down Biasing in V1

Robert Rohrkemper, Hava Siegelmann

UMASS

We created a biologically-inspired algorithm which demonstrates subjective contour illusions. The algorithm attempts to classify the input image. Using the Kanizsa triangle as an example, it can detect either a single shape at the center, or multiple shapes in the periphery. With only bottom-up processing, our model provides several guesses based on the input image. However, with the addition of top-down attention, implemented through a template-matching scheme, the algorithm chooses the most likely categorization of the input image. The model can also be biased by changing the location of the illusion-inducing contours. The most likely interpretation can be chosen based on previous knowledge. If applicable previous knowledge is not available, it can choose based on simplicity.

Several stages of processing are involved. First, the model creates a saliency map using a neural network with local competition (Li, 1999). Each point in the input image is assigned a saliency score. Li's model draws boundaries between regions of an input image, which have different orientations.

This underlying V1 model is more biological than other proposed models because it assumes that all image attributes are simultaneously combined in order to create a single saliency map. Another method, implemented by Koch relies on separate saliency maps for each image attribute and later unifies the value for each point without a particular biological mechanism.

With an automatic thresholding method, developed by the authors, a subset of the salient points are chosen. Next, a top-down process attempts to match this subset to possible solutions from a template library. The most likely solution is chosen and is used to adjust the segmentation threshold. The method orders all points in terms of saliency and then organizes them in ranked clusters. First, whole clusters are added to the segmentation, and then it is fine-tuned by adding or removing individual elements through an interactive process. In our model, previously learned shapes are stored in a library. In reality, these shapes could be combined by various functions or even stored as more abstract, invariant representations.

The model shows that a top-down influence is a possible mechanism for increasing the accuracy of segmentation. With only partial information, it can accurately choose among shapes from a template library. It then uses correlations between the chosen template and the saliency map to create subjective lines.

Our work provides a possible explanation of electrophysiological experiments by Lee and Nguyen, 2001. Their work suggests that illusory contours are created by inter-cortical interaction and require top-down influence. Our model also agrees with results from Lee and Mumford, 2003. Borenstein and Ullman, 2002 have modeled this interaction, but their segmentation process was not biological. Our model uses a more realistic implementation of V1 to do the initial bottom-up segmentation, and refines the output based on top-down information.

Class-Specific, Top-Down Segmentation

E Borenstein, S Ullman, ECCV, 2002

Dynamics of subjective contour formation in the early visual cortex

TS Lee, M Nguyen, PNAS, 2001

Visual segmentation by contextual influences via intra-cortical interactions in the primary visual ...

Z Li, Network: Comput. Neural Syst, 1999

Hierarchical Bayesian inference in the visual cortex

TS Lee, D Mumford - Journal of the Optical Society of America A, 2003

229. Inferring Neural Circuitry from Modulation Metrics: Lessons from a Computational Model of Primary Visual Cortex

Jim Wielarrd, Paul Sajda

Columbia University

The extracellular ratio of the first harmonic to the mean response, F1/F0, is traditionally used to classify cells in primary visual cortex as simple or complex. Recently, the existence of distinct circuitry leading to these cell classes has been brought into question based on theoretical work (Mechler and Ringach, Vision. Res. 2000) and intracellular recordings (Priebe et al. Nat. Neurosci. 2004). In contrast, experimental observation of a bimodal distribution for the intracellular subfield overlap (Martinez et al., Nat. Neurosci. 2005) seems to confirm distinct circuitry as the basis for simple and complex cells in V1. We use a large-scale neural network model of primary visual cortex (see Wielaard and Sajda, Cerebral Cortex 2006), and conduct a simulation study of extracellular and intracellular response modulations for drifting and contrast reversal grating stimuli, specifically considering the dependence of these modulations on the neural circuitry. We find that the intracellular (sub-threshold) F1/F0 distribution is highly insensitive to circuitry. For example it is unsuitable for distinguishing whether or not the dichotomy of simple and complex cells originate from a distinct (all-or-none) or a more egalitarian LGN axon connectivity in our model V1. For what concerns the interpretation of the unimodal nature of the intracellular F1/F0 distribution, we believe our simulation results clarify at least two misconceptions. Firstly, it is not correct to conclude (Priebe et al., Nat. Neurosci. 2004) that the unimodal shape of this distribution, nor of the other observed intracellular distributions (of mean, first harmonic and spike threshold potential) is evidence for the existence of a continuum of circuitry in V1. Our simulations show that two distinct classes of circuitry can just as well result in unimodal distributions for these quantities. Secondly, a unimodal intracellular distribution of F1/F0 does not necessarily provide much less information regarding circuitry than the corresponding bimodal extracellular distribution of F1/F0. When we introduce two classes of circuitry in the model, their identification can be made just as well in terms of the core and tails of the intracellular F1/F0 distribution as in terms of the bimodal extracellular F1/F0 distribution. In fact, for experimental data (Priebe et al., Nat. Neurosci. 2004) the same seems to hold true with regard to intracellular identification of simple and complex cells (when classified extracellularly). Finally, our simulations show that a much better indicator of V1 circuitry related to the simple/complex classification, can be obtained from intracellular responses to contrast reversal stimuli. We show that the intracellular F2/F1 distribution for contrast reversal stimuli, when averaged over spatial phases, is highly sensitive to whether simple and complex cells result from continuous circuitry or from distinctly binary circuitry. Our simulations predict that in the latter case the intracellular distribution of F2/F1 is profoundly bimodal, whereas when simple and complex cells result from continuous circuitry this distribution is profoundly unimodal. To our knowledge experimental data for this distribution is not yet available.

230. Adaptation in stimulus amplitude coding in rat barrel cortex

Jan WH Schnupp, Simon SM Ho, Jose A Garcia-Lazaro

Oxford University

During last year's CoSyNe conference, Dean, Harper and McAlpine reported that responses of midbrain neurons of the guinea pig rapidly adjust the dynamic range of their responses to the statistics of the ongoing sound level distributions. These adjustments alter the coding properties of the neural population, such that coding accuracy is increased near to the most commonly-occurring sound levels. To investigate the possibility that this type of rapid adaptation might also be found in other sensory modalities, we characterized the intensity/response functions of neurons in the primary somatosensory cortex (S1) of 3 anesthetized Long Evans rats to sinusoidal vibration of the whiskers with amplitudes that were changed every 40 ms. Amplitude values were selected randomly from a defined distribution with a high-probability region where levels were selected with 0.8 overall probability. The remaining levels were selected with an overall probability of 0.2. We show that somatosensory responses of rat S1 neurons undergo a process of rapid adaptation very similar to the sound level adaptation described for auditory guinea pig midbrain (IC) neurons. This process appears to improve coding efficiency by adapting the neuron's dynamic range to the distribution of stimulus amplitudes present within an ongoing dynamic stimulus.

231. Dissociation of Stimulus-Driven and Attention-Driven Activity in Macaque Primary Visual Cortex

Jitendra Sharma[1], Beau Cronin[1], Klaus Wimmer[2], Konrad Körding[1], James Schummers[1], Klaus Obermayer[2], Mriganka Sur[1]

[1]Dept. of Brain and Cognitive Sciences, Picower Center for Learning and Memory, MIT, Cambridge, MA, USA, [2]Dept. of Computer Science and Electrical Engineering, Bernstein Center for Computational Neuroscience, Berlin, Univ. of Technology, Berlin, Germany

Attention is a top-down process, informed by behavioral goals, which interacts with stimulus-driven and local, recurrent activity in the processing of visual information. While this general description is widely accepted, many important details concerning the relationship between these signals have been contested in the literature for some time. Here we describe the results of an experiment in the awake macaque which aims to quantify the relative contribution and time course of each type of processing.

The experiment involved a simple bar-release task in which the animal was required to maintain central fixation while focusing attention on a specific region in the visual field; attention was thus directed either towards or away from the receptive field of the neuron currently being recorded. During the course of each trial, oriented gratings were displayed in the receptive field, either alone or along with an isogonally- or orthogonally-oriented surround gradient. Spatial attention was assayed by the animal's ability to quickly respond after the cessation of a small stimulus placed very near to the receptive field. The difficulty of this task was tuned (through changes to the size and luminance of the attention spot) to require focused, behaviorally relevant attention.

To optimally distinguish between different theories of how stimulus, context and attention interact, we performed a fully Bayesian analysis of the data. The power of this analysis method lies in its ability to combine direct experimental observations with any additional prior knowledge. Within this framework we implemented a sampling based approach for the analysis of neural data. This method not only provides a fit to the data which combines prior knowledge with the data, but also results in precise estimates of the uncertainty concerning each model parameter. In fact, within the assumptions made about the system it provides the best estimate of the parameters and their uncertainty that is mathematically possible.

In particular, for this experiment we analyzed changes in the four parameters which characterize a Gaussian tuning curve (orientation and spatial frequency tuning, baseline and peak firing rates) as well as more complicated measures such as Fisher information. We find that even with the newly developed method, the most sensitive possible approach given the assumptions about the experiment, in the early, stimulus-driven phase of the response, attentional modulations are inconsistent on cell by cell basis. The presence of surround stimuli, on the other hand, changes both individual tuning curve parameters and alters the overall information content of the population response. The effect of attention becomes apparent later in the response, however, during the portion of the trial when the animal must perform a timely behavior in response to a visual signal in order to receive a reward. During this period, average activity in many cells gradually increases until the behavior is performed; this attentional modulation indicates that attention does affect visual processing, although in a manner which appears to be untuned with respect to stimulus orientation.

These results suggest that attention does not play a major role in the rapid processing of stimulus orientation in primary visual cortex. Rather, its population-wide effects are confined to the later part of the response, during the portion of the trial in which the visual input has direct behavioral ramifications. This stands in contrast with higher visual areas such as V4, in which it has been demonstrated that focal attention causes significant changes in the processing of individual stimulus features. Finally, this analysis represents one of the first applications of a Bayesian, sampling-based approach to the interpretation of neurophysiological recordings, an approach which provides a unifying framework for the investigation of tuning curve properties.

Support contributed by: Wellcome Trust 050080/Z/97/Z and NIH EY07023

232. Essential Features of Temporal Processing in the Songbird Auditory Forebrain

Tatyana O Sharpee, Katherine Nagel, Allison J Doupe

University of California, San Francisco

The modulation of sound amplitude in time is a prominent feature of many natural auditory signals including zebra finch song. Traditionally, neural processing of this aspect of sound has been characterized as a simple function of amplitude modulation frequency, using sinusoidally modulated tones or noise. Here we show that neural processing of amplitude modulation is not only strongly nonlinear with respect to a particular stimulus feature, but also inherently multidimensional, with at least two separate stimulus features influencing the spike probability.

We probed neurons in the zebra finch auditory Field L using stimuli designed to capture some of song's temporal properties. These stimuli were composed of a broadband noise carrier with Gaussian modulations of the log envelope. We searched for features of the envelope that maximized the mutual information between the stimulus and the neural response. We consistently found that the two most relevant stimulus features were the local mean amplitude and its rate of change. Together, these two features accounted, on average, for 67% of the mutual information between the log envelope and single spikes. In the majority of neurons (>75%, N=30), this two-dimensional description accounted for >50% of the information, reaching values close to 100% for some cells. Typically, one of the two features was dominant, accounting for ~52% of the overall information, with the second feature contributing another 7%; which particular feature was dominant varied between cells and with stimulus statistics. A complementary analysis based on spike-triggered covariance with respect to isolated spikes (preceded by 40 msec silence) confirmed that more than two significant dimensions are not common (<10 %).

While optimal stimulus features were qualitatively similar in shape between neurons, the relationship between these features and the neural response – the so-called response function – could vary. In some cases, the response functions were threshold-linear, so that the firing rate increased as the stimulus became more similar to the feature. In other cases, the functions were symmetric, so that the presence in the stimulus of a feature or of its opposite could both drive (or inhibit) the cell. When the neural responses were examined with respect to both features, the result seldom reflected a simple multiplication of the one-dimensional response functions. The 2D response functions were qualitatively consistent with types expected from optimality arguments. Thus, temporal processing in the auditory forebrain appears to be highly nonlinear, with contributions from both the amplitude and its rate of change, rather than simply showing selectivity for a particular modulation frequency.

Support Contributed By: NIH and Swartz Foundation.

233. A theoretical model of cochlear processing improves simulated cochlear implant hearing

Evan C Smith, Lori L Holt

Carnegie Mellon University

A theoretical model of cochlear processing improves simulated cochlear implant hearing

Cochlear implants are neuroprosthetic devices that use direct, electrical stimulation of auditory nerve fibers along the tonotopic axis of the cochlea to restore some degree of hearing to individuals with profound peripheral hearing loss (Zeng et al., 2004). Despite twenty years of research and wide clinical application, speech perception in cochlear implant users is highly variable and often quite degraded. Although a variety of reasons have been proposed for the poor performance of cochlear implant users, one clear factor is their limited spectral resolution. Normal human hearing has a rich spectral representation of the auditory world (30,000 spiral ganglion cells lie along the tonotopic axis of the cochlea), but cochlear implants carry very few frequency channels, often fewer than eight. The design of the filterbank, which determines the content of these channels, has a potent effect on performance. Here we present a theoretically motivated method for filterbanks design and a pair of experiments demonstrating a substantial increase in speech intelligibility when using these filters compared with standard filters. These finding represent some of the first behavioral evidence of efficient coding in human perception.

Recently, Smith and Lewicki (2004, 2005 & 2006) used an information theoretic model to show that mammalian hearing follows an efficient coding principle (Barlow, 1961; Atick, 1992; Simoncelli & Olshausen, 2001; Laughlin & Sejnowski, 2003). Neurons in the inner ear and the spikes with which they communicate form an efficient code for natural sounds in the environment, maximizing information while minimizing coding cost. Applying the same analysis to speech coding suggests that speech acoustics are optimally adapted to the mammalian auditory code (Smith and Lewicki, 2005 & 2006). We now seek to apply efficient coding theory to the problem of speech perception in individuals using cochlear implants.

Using the techniques described in Smith and Lewicki (2006), we can learn an efficient n-channel filterbank for any given class of sounds. Trained on speech with the constraint that there are only six channels (as might be the case for a cochlear implant), the learned filters come to represent the six most informative dimensions of speech acoustics.

In order to compare speech perception given an efficient 6-channel filterbank versus a standard filterbank, we employed a technique developed in Shannon et al. (1995) to simulate cochlear implant hearing in normal hearing adults. The time-varying amplitude envelope of each channel is used to modulate band-limited noise; then the channels are linearly combined to make the stimuli. The resulting sounds can initially be very difficult to understand but are clearly speech-like.

In our first experiment, the stimuli were whole sentences spoken in English from the TIMIT speech corpus (Garofolo et al., 1990). We created four classes of stimuli by using different filterbanks: two sets of standard cochlear implant filterbanks with either linear or logarithmic frequency spacing, the learned filters from our algorithm or no filter (undistorted speech). Sixteen participants listened to 42 examples of each class of stimuli and typed what they heard. We treated each word as independent and computed the probability of typing a correct word in each condition. The mean percent correct for each class across participants was 13%, 37%, 56% and 86% for the linear, log, learned and normal speech respectively (all results were highly significant). Performance improved by an average of 20% (56% versus 37%, respectively; p-value = $3.2\times10-5$) using the learned filters versus the standard set.

The second experiment used non-lexical vowel-consonant combinations spoken in isolation (Shannon et al., 1999). Three stimulus classes were used: standard logarithmic, efficient filters and undistorted. Thirteen participants listened to the stimuli and both verbally repeated and typed what they heard. Accuracy was not as high as in experiment one (26%, 39% and 76%, for the standard, efficient and unmodified conditions respectively), but the performance using the efficient filters was still significantly greater (13% increase; p-value = 0.03) using our learned filters compared to the standard set. Importantly, we found that the performance gains differed based on the acoustic properties of the speech sounds. Performance on speech transients (e.g., stop consonants) increased by 39% (p-value = $1.7\times10-5$). Steady state sounds, like vowels and glides, improved by only 5% (p-value = 0.08), though there clearly were ceiling effects with near perfect performance in most cases. In contrast, performance with noise-like sounds (e.g., fricatives) was not significantly different between the efficient and standard conditions.

234. Interactions between eye movements, receptive fields, and processing streams in primary visual cortex of alert monkeys

Max Snodderly[1], Igor Kagan[2], Moshe Gur[3]

[1]Medical College of Georgia, [2]Caltech, [3]Technion, Israel

Fixational eye movements are a ubiquitous component of visual tasks. They contribute retinal image motions that place stimuli at different locations within and outside of neuronal receptive fields and they impart motion to the images of otherwise stationary stimuli. The eye movements in turn are modified by attentional demands and by the visual environment. We have studied how these movements influence the activity of neurons in primary visual cortex of macaque monkeys. Many of the effects of fixational saccades can be understood by analyzing the ways that saccades move and locate stimuli relative to the receptive fields of the neurons. Although bursts of spikes often occur immediately after abrupt movements of the stimulus associated with saccades, most cells also exhibit sustained firing in the intersaccadic drift periods if the stimulus remains on the receptive field. In fact, many cells do not respond to saccades and have sustained responses only during drift periods. We have postulated that cells giving the sustained responses are important for perception of the details of a stimulus, whereas the cells giving only transient responses are primarily signaling an abrupt retinal image movement or the presence of an object.

Fixational saccades are accompanied by extraretinal effects that modulate the firing of both lateral geniculate (LGN) and V1 neurons. The saccadic modulation in the LGN consists of a weak suppression peaking around 20-30 ms after saccade onset, followed by a much stronger enhancement peaking around 80 ms after saccade onset. In V1, similar effects accompany both fixational and voluntary saccades, but the time course is quite different from LGN. In V1 the enhancement peaks at about 200 ms after the saccade. Since response latency for a flash is about 25 ms in the LGN and about 50 ms in V1 there is a time delay of only about 25 ms between responses in LGN and V1. Consequently the time difference of 120 ms in the extraretinal effect in LGN and in V1 suggests that the two structures are getting extraretinal inputs from different brain regions, both of which could enhance the late phase of V1 responses to saccadic movements of the retinal image. These extraretinal effects may combine to enhance visibility of the image immediately after saccades and in the intersaccadic intervals of fixation.

By precisely recording and compensating for fixational eye movements we have discovered a new pathway through primary visual cortex that is selective for the direction of stimulus motion. This pathway passes through layer 3, which is the main output from V1 to the cortical ventral stream for object recognition. The layer 3 neurons are direction-selective, silent, orientation selective, and end-stopped with small receptive fields. They are well suited to encode information about small-scale motions such as changes in facial expressions.

235. Heterogenous firing rate dependencies in simultaneously recorded neural populations in cat area 17

Martin A Spacek[1], Timothy J Blanche[2], Nicholas V Swindale[1]

[1]University of British Columbia, [2]Hanse-Wissenschaftskolleg, Brain Research Institute, University of Bremen

In the retina, it has been shown that cells act as independent encoders, where the spike trains are generally independent across cells. (Nirenberg et al, 2001). We wanted to see what kind of dependencies, if any, might exist in the firing rates of cells in primary visual cortex.

We used silicon polytrodes with closely spaced electrode sites (50-75 um) to simultaneously record from dozens of cells over multiple cortical layers in area 17 in anesthetized cat (Blanche et al, 2005). Electrode sites were closely spaced so as to isolate many adjacent cells within the recordable volume of roughly 2000 x 200 x 130 um. We stimulated with contrast normalized natural scene, pink noise (1/f amplitude spectrum), and white noise movies on a 200Hz monitor. Stimulation area was about 3 times the classical receptive field area.

We calculated instantaneous firing rate instead of PSTH rate because we wanted to track rate dependency over time, not over trials. Due to spontaneous variations in brain state and depth of anesthesia, the two are not necessarily the same. Instantaneous rates were calculated using partially overlapping bins of varying duration such that a fixed number of spikes n fell in each bin. We set n=4, but results were generally invariant for different values of n.

The joint firing rate distribution for each cell pair was constructed by stepping through every time point in the rates of both cells and building up a 2D histogram of joint rate probabilities. The x and y axes represented the rates of each cell in log-spaced bins, and probabilities were represented on the z axis. The factorial firing rate distribution was calculated by first finding the rate distribution for each cell separately, and then taking the outer product of the two distributions. Doing so assumes that the rates are independent. The difference between the joint and the factorial rate distributions was used as an indication of the level of dependency in the firing rates within the pair.

To initially quantify this difference, for each cell pair the joint entropy and the sum of the individual entropies were compared. Across all cell pairs, entropies were generally lower for the joint than the sum, and their difference (the Kullback-Leibler distance, which amounts to a measure of mutual information between the rates of the cell pair) was generally greater during natural scene stimulation than during white noise or pink noise stimulation. However, mutual information does not effectively measure the difference in shape between the joint and factorial distributions, yet the shape difference determines dependency.

To effectively compare distribution shapes, the distributions were smoothed and 95% probability contours were plotted. The shapes of these contours for joint and factorial distributions were then compared by measuring their percentage of overlap.

We found that dependencies exist in the firing rates of roughly one quarter of cell pairs in area 17, even though most of these pairs lacked a peak in their cross correlograms. The presence or absence of such dependencies may reveal more about neural architecture than cross correlation analysis alone.

We also found a much greater incidence and degree of dependency during natural scene stimulation than during white noise stimulation, and significantly greater dependency than during pink noise stimulation. We suggest that the higher order statistics in natural scenes, which presumably the network is specially wired to deal with, are responsible for the additional dependency in natural scenes versus pink noise.

References:

Nirenberg, Carcieri, Jacobs, Latham. Retinal ganglion cells act largely as independent encoders. Nature 411:698, 2001.

Blanche, Spacek, Hetke, Swindale. Polytrodes: High-density silicon electrode arrays for large-scale multiunit recording. J Neurophysiol 93:2987, 2005.

236. Adaptation within a Bayesian Framework for Perception

Alan A Stocker, Eero P Simoncelli

Howard Hughes Medical Institute and Center for Neural Science, New York University

A growing number of studies support the notion that humans apply an optimal or near-optimal strategy when performing a perceptual estimation task that combines the sensory observations with a priori knowledge as defined by Bayes' rule. A Bayesian framework provides a principled yet simple computational framework for perception that can account for a large number of known perceptual effects and illusions.

Adaptation is a fundamental phenomenon in sensory perception that seems to occur at all processing levels and modalities. A variety of computational principles have been suggested as explanations for adaptation. Many of these are based on the concept of maximizing the sensory information an observer can obtain about a stimulus despite limited sensory resources, which is similar to the concept of redundancy reduction and efficient representation. More mechanistically, adaptation can be interpreted as the attempt of the sensory system to adjusts its limited dynamic range such that it is maximally informative with respect to the statistics of the stimulus. Perceptually, adaptation seems to have two fundamental effects. First, subsequent stimuli are repelled by the adaptor stimulus, i.e. the perceived values of the stimulus variable that is subject to the perception task are more distant to the adaptor value after adaptation. Second, adaptation leeds ot an increase in the observer's discrimination ability around the adaptor value, whereas it decreases further away from the adaptor.

If a Bayesian framework is to provide a valid computational explanation of perceptual processes, then it needs to account for the behavior of a perceptual system, regardless of its adaptation state. So far, it has not been shown how adaptation could be in agreement with a Bayesian framework of perception.

We extend our previously developed Bayesian framework for perception [1] to account for adaptation. We first note that the perceptual effects of adaptation seems inconsistent with an adjustment of the internally represented prior distribution. Instead, we postulate that adaptation increases the signal-to-noise ratio of the measurements by re-allocating the limited resources of the measurement stage to the input range. We show that this re-allocation changes the likelihood function in such a way that the Bayesian estimator model accounts for reported perceptual behavior. In particular, we compare the model's predictions to human motion direction discrimination data and demonstrate that the model well predicts the characteristics of the observed adaption effects of repulsion and changes in discrimination threshold. The proposed model shows for the first time how adaptation can be incorporated into a Bayesian framework of perception. It represents a link between bottom-up processing guided by efficient coding and a top-down Bayesian estimation process.

[1] A.A. Stocker and E.P. Simoncelli. "Constraining a Bayesian Model of Human Visual Speed Perception". In: Advances in Neural Information Processing and Systems NIPS 2004, vol.17, p.1361-1368

237. Single unit and local field potential characterization of contrast dependent responses in area V4 of the macaque

Kristy A Sundberg, Jude F Mitchell, John H Reynolds

The Salk Institute

Luminance contrast is one of the most basic visual stimulus feature dimensions and the contrast dependent responses of neurons in early visual areas (such as LGN and V1) have been well characterized. In V4 a number of studies have sought to characterize responses to basic features (such as color and orientation), as well as more complex feature dimensions (such as degree of curvature). Additionally, a number of studies have examined how attention modulates V4 responses to visual stimuli. The contrast dependence of visual responses in area V4, however, has yet to be characterized. We presented gratings at 7 levels of luminance contrast to two awake monkeys, and examined how the response strength and response latency of single neurons and stimulus-triggered average local field potentials (LFP's) changed as a function of contrast. Similar to responses in area V1, V4 single unit responses generally increased with luminance contrast, and saturated at high contrast. A substantial fraction of neurons recorded, however, (approximately 20%) exhibited a peak in their contrast response function, beyond which additional increases in contrast resulted in a decrease in firing rate. Response latencies of single units decreased as stimulus contrast was increased, even among neurons with peaked contrast response functions. Thus, this reduction in latency depends on contrast, not firing rate. The magnitude of the stimulus-triggered average LFP typically increased with contrast, and showed a contrast-dependent reduction in latency that mirrored the latency reduction observed in individual neurons. We also examined the contrast dependence of the lfp power spectrum. Consistent with a recent report in area V1 (Henrie and Shapley, 2005), we find contrast-dependent increases in LFP power in the range of frequencies between 25 and 80 hz. Although this contrast-dependent increase in LFP power is similar to the increases we observed in firing rate of individual neurons, LFP power continued to increase strongly at the highest contrast values, where the population average response had saturated.

238. Wireless Multi-Unit Recording from Unconstrained Animals

Tobi Szuts[1], Edward Soucy[2], Alan Litke[3], Athanassios Siapas[4], Markus Meister[2]

[1]Program in Biophysics, Harvard University, [2]Department of Molecular and Cellular Biology, Harvard University, [3]Santa Cruz Institute for Particle Physics, University of California Santa Cruz, [4]Division of Biology, California Institute of Technology

Studying the neural populations of awake animals responding to natural stimuli is crucial to understand the brain. Multi-electrode techniques require the animal to be attached to the recording apparatus through a wire, if not also simultaneously head-fixed and/or anesthetized. In practice, this prevents recording in any environment more natural than the laboratory and severely limits the animal's range of motion. Wireless recording techniques to date have offered at most 4 channels, hardly enough to sample the state of ensemble activity. To combine the strengths of both methods, we present preliminary results from a 64-channel wireless recording system designed for, but not limited to, the rat.

The system consists of five previously developed components: a 28-tetrode microdrive, a 64-channel pre-amplifier/filter/multiplexer (the Neurochip, developed by Alan Litke), a miniature video FM transmitter, an FM receiver, and a computer with a high-speed data acquisition card for online analysis and data storage. The circuitry for the amplifier/multiplexer will be mounted directly on the microdrive, while the transmitter and battery are carried on a backpack to avoid overburdening the animal's head.

Initial tests have verified the fidelity of transmission; signals input to the neurochip can be recovered from the transmitted data, and a single channel of (previously recorded) spike data was recovered from a field of white noise without any loss of quality. The signal-to-noise ratio for transmission is around 100, considerably better than required for extracellular recordings. No deterioration in signal-to-noise is observed for distances up to 10 m. Present work involves integrating the disparate components: standardizing the connections between microdrive and neurochip, designing the power supply, dealing with radio frequency interference, and writing controlling software.

239. Modulation of auditory responses by modality-specific attention in rat primary auditory cortex

Lung-Hao Tai, Anthony M Zador

Cold Spring Harbor Laboratory

How does attention modulate sensory representations? In order to probe the underlying neural mechanisms, we established a simple rodent model of modality-specific attention. Here we describe preliminary results of experiments in freely moving rats in which we have used tetrodes to record neural responses in primary auditory cortex (area A1) while subjects performed this behavior.

Subjects were first trained to perform distinct auditory and olfactory two alternative forced-choice (2AFC) tasks. Training and testing were conducted in a custom three-poke computer-controlled behavioral apparatus. Subjects initiated trials with a center-poke, which triggered presentation of a tone (either 5 or 15 Hz), an odor (either R(-)-2-Octanol or S(+)-2-Octanol), or both. Subjects responded moving to the left or right poke. Correct responses were rewarded with water. Initially, training for the two tasks (auditory or olfactory) was performed on separate days. Once performance on each task in isolation had achieved criterion (1-2 weeks), the two stimuli (one auditory and one olfactory) could be presented simultaneously, with auditory and olfactory blocks (of 50 trials each) interleaved in a single session. In auditory blocks, only pure tones were presented, and subjects were cued to perform the task based on auditory stimuli. In olfactory blocks, both odors and pure tones were presented simultaneously, and subjects were cued to perform the task based on olfactory stimuli. (We are currently testing a more symmetric version of the task, in which both stimuli are present on all trials).

After subjects reached consistent performance on the interleaved blocks, tetrode drives were implanted in primary auditory cortex of the left hemisphere. Single unit responses to tones were heterogeneous, and included transient, sustained, and suppressed. Some responses showed modality-specific attentional modulation; in most cases, the responses to a particular auditory stimulus was enhanced in the auditory block (or, equivalently, suppressed in the olfactory block).

Our results, although preliminary, suggest that shifting attention from audition to olfaction and back can modulate the activity of single neurons in primary auditory cortex.

240. Multiple S-cone pathways in the macaque visual system

Chris Tailby[1], Samuel G Solomon[2], Peter Lennie[1]

[1]New York University, [2]University of Sydney

Given the symmetry in the organization of on-center and off-center cells in the Parvo- and Magnocellular pathways in primate, we might expect a similar complimentary organization in the receptive fields of neurons that carry S cone signals. The fact that cells that receive strong inhibitory input from S cones (S-) cells are less frequently encountered than those that receive strong excitatory S input (S+) cells suggests that this might not be true. We explored this question in Lateral Geniculate Nucleus (LGN) of macaque, where we recorded from cells with strong S-cone input. These were encountered in or near the interlaminar zones abutting the parvocellular layers. We show that on several important dimensions the properties of S- neurons are quite unlike those of S+ neurons.

The organization of chromatic inputs differs substantially in S+ and S- neurons. In S+ (n= 36) cells, signals from S cones were opposed by a combined signal from L cones and M cones (mean normalized L, M and S cone weights of -0.29, -0.02, and 0.56, respectively). The chromatic signatures of S- cells (n = 42) were more heterogeneous than those of S+ cells. In S- cells, signals from L cones were opposed to a combined signal from S and M cones (mean normalized L, M and S cone weights of 0.25, -0.27, and -0.29, respectively).

To pure S cone modulation S+ cells are twice as sensitive as S- cells (contrast gain of 0.91 ± 0.12 [impulses/sec]/percent contrast vs. 0.53 ± 0.06 [impulses/sec]/percent contrast), and unlike S- cells, their contrast response curves saturate at high contrast. Both S+ and S- cells become less responsive following prolonged exposure to S cone contrast. To achromatic modulation S- cells are much more sensitive than S+ cells (S-: 1.16 ± 0.39 [impulses/sec]/percent contrast; S+ 0.25 ± 0.01 [impulses/sec]/percent contrast). For both cell types the contrast response functions saturate at high contrast. This behavior is quite unlike that of parvocellular cells.

When driven by drifting achromatic gratings, responses often showed a pronounced direction bias that was spatiotemporal frequency dependent. Moreover, spatial frequency tuning curves sometimes had 2 peaks. Direction bias was not observed when the gratings modulated only the signals from S cones. Spatial frequency tuning curves for S cone isolating stimuli were low pass for both S+ and S- cells. These results are consistent with the chromatically opponent inputs arising in mechanisms that are not concentric, and are temporally offset.

The temporal properties of S+ and S- cells differed reliably, with the S- cells preferring higher temporal frequencies (mean preferred TF, S-: 19 Hz, S+: 13 Hz).

Along all dimensions on which we compared S+ and S- cells, the properties of the S+ cells were relatively tightly clustered, suggesting a homogenous class. The properties of the S- cells were much more widely distributed, and it is unclear whether they constitute a single group. Regardless of whether the S- cells constitute a single group, the difference between S+ and S- cells are large and consistent, and make it likely that they encode different aspects of the visual image.

241. Attentional modulation of stimulus competition in a large-scale model of the visual pathway

Calin I Buia, Paul H Tiesinga

University of North Carolina, Chapel Hill, Physics & Astronomy

Neurons in cortical area V4 are sensitive to shape and have large receptive fields. In a typical visual scene there are multiple objects in the V4 cell's receptive field, only a few of which may be behaviorally relevant. The visual system is capable of selecting relevant objects by increasing the neural response to them and reducing the response to non relevant objects. Neuronal synchrony may play an important role in this process. Using a large-scale network model of the visual pathway, we study the emergence of shape selectivity in V4, the competition between different objects for control of the firing rate of individual V4 neurons, the attentional modulation of this stimulus competition and the role of synchrony.

In previous V1 models for the emergence of contrast-invariant orientation tuning either only a small part of the visual field was represented or the model was studied using stimuli that covered the entire visual field. Hence, it could not be tested whether both stimulus orientation as well as stimulus position were correctly represented. We constructed a detailed biophysical model that responded correctly to small bars of different orientations at all positions in the represented visual field (4 by 4 degrees). The model consisted of 1024 cortical columns, each comprised of 84 excitatory and 21 inhibitory Hodgkin-Huxley-type neurons. The position and orientation tuning was obtained using a cortical connection pattern consisting of medium-range inhibition and short-range recurrent excitation. For simplicity, we assumed that the V2 response was similar to that of V1. We constructed a model for cortical area V4 that was driven by the output of the V1 model. The V4 cells were tuned to respond to specific conjunction of bars (such as "L" and "+" shaped stimuli). Because of the inhibition, there already was some stimulus competition between objects at the level of V1. The competition was further modulated in V4 by selective activation of groups of interneurons.

242. Rapid Adaptation from the Non-Classical Receptive Field in area MT of the Macaque

Pascal Wallisch, Gopathy Purushothaman, David C Bradley

University of Chicago

We are interested in the spatiotemporal dynamics of the antagonistic surround of neurons in area MT of the macaque, an area which plays a crucial role in the perception of visual motion (Born & Bradley, 2005). In a previous study, we found no such temporal dynamics under any of the conditions that we tested (Wallisch et al., 2003).

However, due to the large number of experimental conditions used, we were unable to obtain a large number of samples for any particular condition. Moreover, we compared firing rates taken from the entire duration of the stimulus. This might have obscured subtle effects in the first 150 ms or so of the response which have been shown to contain the most neural information (Osborne et al., 2004).

In this study, we densely sampled a small number of key conditions in an actively fixating macaque to study the time course of the response at a higher resolution. We report data from 150 single units that show a number of small but significant effects. First, neural response transients are significantly suppressed if a test stimulus is preceded by a moving stimulus in the center or surround of the receptive field of the neuron. Second, the magnitude of the effect is larger for cases in which the stimulus was preceded by a stimulus in the center of the receptive field while the effect size resulting from a preceding stimulus in the surround is comparable to the suppression resulting from simultaneous presentation of the stimuli. Finally, we report that these effects depend on the direction of the priming stimulus and can't be explained in terms of spatial attention.

We take these findings to suggest that rapid adaptation in MT might share neural substrates with surround suppression, which could facilitate the detection of novel vs. continuously moving stimuli. Moreover, these effects can be interpreted in terms of the source of adaptation effects in area MT.

243. Observers' decisions in a simple visual task are consistent with Bayesian processing of early perceptual uncertainty

Louise Whiteley, Maneesh Sahani

Gatsby Computational Neuroscience Unit, UCL

Behavioural evidence from domains such as motor planning, depth perception, and crossmodal cue combination (reviewed by Knill & Pouget, 2004) suggests that human observers act and respond close to optimally under conditions of motor and sensory uncertainty. Unfortunately, the neurophysiological basis for this Bayesian behaviour is not yet well understood. A decade or so of theoretical work has developed competing suggestions for how neurons might represent probability distributions (reviewed by Pouget, Dayan, & Zemel, 2003), but relating these to specific behavioural tasks and corresponding neurophysiological data has proven difficult.

In part, this difficulty has arisen from a mismatch between the domains studied behaviourally, and those which have been well characterised neurophysiologically. Perhaps our best experimental understanding of neural population codes comes from the early visual system; investigations of behaviour under uncertainty, on the other hand, have largely focused on more complicated inferential perception (e.g. depth), on other modalities (e.g. proprioception) or on uncertainty in motor output. One exception is a recent experiment by Stocker and Simoncelli (2005), in which a unimodal visual task was used to constrain a Bayesian model of speed perception.

Here we develop another purely visual task, but one involving a minimal judgement that relates to quantities thought to be represented in the earliest parts of visual cortex. This enables us to ask whether uncertainty in the simplest kind of perceptual judgements is also processed in a Bayes optimal fashion. It also provides a very tightly constrained domain for testing competing coding schemes, and allows contact with the wealth of neurophysiological data on early visual perception.

The task we used was a simple positional 'Vernier' judgement (Spillman & Werner, 1990), in which subjects were asked whether one Gabor patch was offset to the left or right of another patch. The difficulty of the task was controlled in order to produce a complete psychometric curve, with a significant range of offsets over which subjects made errors (see Green & Swets, 1989). We then introduced a penalty structure into the environment and asked if participants could maximise reward, thereby demonstrating at least implicit knowledge of their perceptual uncertainty (see Trommershauser et al., 2003a; 2003b; 2005). Participants were asked to maximise the points gained, where a correct answer yielded a fixed number of points, and an incorrect answer incurred a varying penalty. For example, the penalty for answering 'right' incorrectly might be twice as large as the penalty for answering 'left' incorrectly. Five different penalty conditions were used, and participants were informed and regularly reminded of the penalty values, but, as discussed further below, feedback regarding accumulated points was provided only occasionally.

To determine whether subjects were behaving optimally, we estimated the degree of perceptual uncertainty that would accurately reflect their observed error rates. We then predicted from this uncertainty how the mean of each psychometric curve should shift in order to maximize reward under each penalty condition. We found a close match between prediction and performance in all four subjects tested.

To help exclude strategies that rely on error-driven learning rather than direct representations of uncertainty we used periodic cumulative feedback. If trial-by-trial feedback had been used, it would have been possible for subjects to optimise their performance simply by shifting a perceptual threshold in response to that feedback. With periodic feedback, subjects do not know which trials led to penalties, and so cannot use this strategy efficiently, instead being forced to rely on their internal estimates of confidence. On the other hand, periodic feedback is sufficient to keep subjects motivated, which is crucial for this kind of paradigm.

Two further aspects of the results are of interest. First, all subjects reported a simple cognitive strategy in which they gave the answer with the lowest penalty when they were 'uncertain'. None reported adjusting this strategy according to the relative values of the penalties in the five different conditions, which is what the behavioural results demonstrate. This suggests that subjects were unconsciously computing optimal performance based on their perceptual uncertainty, rather than using a conscious cognitive strategy to produce the observed pattern of results. Second, the subjects' perceptual errors were observed to decrease as they become more practised. We therefore analysed the data in two halves, and found that the match between prediction and performance was uniformly good, despite the changing uncertainties.

Our results therefore support the generality of Bayesian behaviour, even in a very simple unimodal task. It appears that the brain must represent and use knowledge about perceptual uncertainty to produce the observed performance, although future work will address whether, or how much of, the entire probability distribution is represented.

Green, D.M., & Swets, J.A. (1989). Signal Detection Theory and Psychophysics. Peninsula Publishing.
Knill, D.C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. Trends in Neuroscience, 27, 712-719.
Pouget, A., Dayan, P., & Zemel, R.S. (2003). Inference and computation with population codes. Annual Reviews of Neuroscience, 26, 381-410.
Spillman, L., & Werner, J.S. (1990). Visual Perception: the Neurophysiological Foundations. San Diego: Academic Press.
Stocker, A.A., & Simocelli, E.P. (2005). Constraining a Bayesian model of human visual speed perception. In Lawrence K. Saul, Yair Weiss, and Leon Bottou (Eds.) Advances in neural information processing systems, vol. 17.
Trommershauser, J., Gepshthein, S., Maloney, L.T., Landy, M.S., & Banks, M.S. (2005). Optimal compensation for changes in task-related movement variability. The Journal of Neuroscience, 25, 7169-7178.
Trommershäuser, J., Maloney, L.T., & Landy, M.S. (2003a). Statistical decision theory and trade-offs in the control of motor response. Spatial vision, 16, 255-275.
Trommershäuser, J., Maloney, L.T., & Landy,M.S. (2003b). Statistical decision theory and the selection of rapid, goal- directed movements. Journal of the Optical Society of America A: Optics, image science, and vision, 20, 1419-1433.

244. A Bayesian View of Sensory Conflicts in Decision-Making

Angela J Yu[1], Peter Dayan[2], Jonathan D Cohen[1]

[1]Princeton University, [2]Gatsby Computational Neuroscience Unit, UCL

In this work, we present a formal analysis of the computations underlying decision-making tasks in which the sensory inputs are in conflict. These tasks have been used to study a wide variety of related cognitive functions, such as decision making, error detection, conflict monitoring, automaticity, selective attention, etc. In particular, we concentrate on the Eriksen task, a classical paradigm for studying the integration of conflicting sensory information, and the role of selective attention in controlling this integration. We propose two distinct Bayesian models that explain existent behavioral data, one focusing on stimulus compatibility and the other on spatial uncertainty. We demonstrate that under certain extensions to the classical Eriksen Task, the two models make distinct predictions that can be experimentally verified. We also extend the model to more general experimental manipulations, and examine the formal relationship between the Eriksen task and the Stroop task from a Bayesian viewpoint.

245. Neuronal sensitivity and choice probability in macaque VIP during a heading discrimination task

Tao Zhang, Ken H Britten

Center for Neuroscience and Section of NPB, UC Davis

Heading direction can be estimated very accurately based on optic flow information. The ventral intraparietal area (VIP) in macaque posterior parietal cortex is one of those candidate areas assumed to be involved in such heading direction judgments. Previously we have documented that VIP neurons are well tuned to stimuli simulating heading, and that microstimulation in VIP can bias monkey choices in a heading discrimination task. Those results indicated that VIP neurons might provide the signals supporting behavioral choice in heading discrimination tasks. To extend those observations, we now quantitatively related neuronal responses in VIP to monkey's perceptual performance. We trained our monkey on a two-alternative-forced-choice (2AFC) heading discrimination task, and recorded neuronal activities in VIP during task performance. This allowed us to obtain both psychophysical and physiological data in the same animal, on the same sets of trials and using the same visual stimuli. Thus, we could estimate neuronal sensitivity to the small heading differences the animal was reporting, and uncover correlations between neuronal discharge and the choices the monkey made on a trial-by-trial basis, which we measured using choice probability analysis.

VIP neurons are on average substantially less sensitive to small heading changes than is the monkey, although a small number of neurons in our sample were approximately equal in sensitivity to the monkey. Neuronal sensitivity was greatest in the middle of sigmoidal tuning functions, and on the flanks of band-pass-tuned cells.

Significant choice probabilities were common and stable across stimulus conditions. Although both positive and negative correlations were observed, positive choice probabilities (> 0.5) were more frequent in our sample.

Our results suggest that VIP activity is involved in heading perception. The relatively low neuronal sensitivity, combined with the observed choice probabilities, are consistent with the idea that the signals from many neurons are pooled together to form behavioral judgments about heading.