An Observation of a Fitness Landscape via a Cross-section

Akira Imada

Brest State Technical University
Moskowskaja 267, Brest, 224017 Republic of Belarus
ai@brpi.unibel.by

Abstract. When we search for a solution of a problem employing an evolutionary computation, information of a geometry of fitness landscape is sometimes very helpful. What does the fitness landscape look like? How many global peaks are there in the fitness landscape? Do Local peaks exist? How does the shape of these peaks look? To study these issues, we propose a simple technique of learning a shape of a peak in the fitness landscape by observing a cross-section of the peak by a hyperplane.

1 Introduction

In exploring a search space by an evolutionary computation, learning the geometry of its fitness landscape is very helpful. In this paper, we propose a simple method where a hyperplane cut the fitness landscape and observe its cross-sections. As a test bed, we use a fitness landscape of the fully connected neural network model of associative memory where, for example, the width of a peak suggests synaptic plasticity and steepness of the side wall implies the hardness of the search. We also observe some local peaks on the side wall of the global peak.

In addition to using the model as a test bed, we are interested in the model per se. Although the fully connected neural network model of associative memory is somehow an old topics, many interesting issues from physics point of view are still unknown. The number and distribution of attractors in synaptic weight space, or the fate of specific trajectories when some of the weight values are slightly modified are these examples. To attack these issues, we explore the fitness landscape defined on weight space of a fully connected neural network model of associative memory.

These studies are steps to learn a whole geometry of the fitness landscape.

1.1 Associative Memory

Associative memory is *a dynamical system which has a number of stable states with a domain of attraction around them* (Komlós and Paturi 1988). If the system starts at any state in the domain, it will converge to the stable state. Hopfield (1982) proposed a fully connected neural network model of associative memory in which information is stored by being distributed among neurons and is retrieved from dynamically relaxed neurons' states. In the model, some of the appropriate configurations of synaptic weights give the network a function of associative memory.

The Hopfield model consists of N neurons and N^2 synapses. Each neuron state is either active (+1) or quiescent (−1). When an arbitrary N-bit *bipolar pattern*, a sequence of +1 and −1, is given to the network as an initial state, the dynamical behaviors of states of neurons afterwards are characterized by the strengths of the N^2 synapses. The synaptic strengths are called *weights*, and the weight from neuron j to neuron i is denoted as w_{ij} in this paper. Provided the synaptic weights are determined appropriately, network can store some number of patterns as attractors. Hopfield employed the so-called Hebbian rule (Hebb, 1949) to prescribe the weights.[1] That is, to store p bipolar patterns ξ^ν:

$$\xi^\nu = (\xi_1^\nu, \cdots, \xi_N^\nu), \qquad \nu = 1, \cdots, p, \tag{1}$$

the weight values are determined as:

$$w_{ij} = \frac{1}{N} \sum_{\nu=1}^{p} \xi_i^\nu \xi_j^\nu \quad (i \neq j), \quad w_{ii} = 0. \tag{2}$$

An instantaneous state of a neuron is updated asynchronously (one neuron at a time) as:

$$s_i\,(t+1) = sgn\left(\sum_{j \neq i}^{N} w_{ij} s_j\,(t) \right), \tag{3}$$

where $s_i(t)$ is a state of the i-th neuron at time t. If an initial state converges to one of the stored patterns ξ^ν as an equilibrium state, then the pattern is said to be *recalled*. Furthermore, if an initial state chosen from the stored patterns remains unchanged from the beginning, then the pattern is said to be stored as a *fixed point*.

In analyzing the Hopfield model, there have been basically two different approaches: one is to explore *pattern space* searching for attractors under a specific weight configuration, and the other is to explore *weight space* searching for an appropriate weight configuration that stores a given set of patterns. To be more specific, the former is an analysis of the Hamiltonian energy as a function of all the possible configurations of bipolar pattern to be given to the network, where synaptic weights are pre-specified using a learning algorithm, usually the Hebb's rule, so that the network stores a set of p given patterns. In this context, the model for $p = 1$ corresponds to the Mattis model of spin-glass (Mattis, 1976), in which the Hamiltonian energy has two minima, while the model for infinitely large p corresponds to the Sherrington-Kirkpatrick model (Sherrington and Kirkpatrick, 1975), in which the synaptic weights become Gaussian random variables. Analyses of the former type have been made in between these two extreme cases (Amit, 1989). The latter analysis, on the other hand, was addressed

[1] The rule was at first advocated by Hebb (1949) and later formalized by Cooper (1973).

by Gardner (1988). She discussed the optimal weight configurations for a *fixed* number of given patterns in terms of the volume of the solutions in weight space, suggesting that the volume shrinks to vanish when p approaches to $2N$. In short, the former approach searches for the optimal pattern configurations which minimize the Hamiltonian energy *in pattern space* with the weights being fixed, while the latter searches for the weight configurations *in weight space* that optimally store a set of given fixed patterns.

In this paper, our concern is on weight space where some points give a network a capability to store a pre-specified one pattern ($p = 1$), and we call these points solutions. We visualize these solutions on *fitness landscape.*

1.2 Fitness Landscape

The concept of the fitness landscape was first introduced by Wright (1932) to study biological evolutionary processes. Since then, this concept has been used not only in evolutionary biology but also in chemistry, physics, computer science and so on.

In chemistry, for example, a molecule can be represented as a string of N letters with each letter being chosen from an alphabet of size k (Macken et al. 1991) Twenty amino acids ($k = 20$) for proteins or four nucleotides ($k = 4$) for nucleic acids can be considered as examples of the alphabet. The k^N possible combinations of the letters construct a configuration space of the string. Then, for example, the free energy of RNA folding into secondary structures (Fontana and Schuster 1987), or the ability of peptides to bind to a particular substrate to catalyze a specific reaction (Maynard Smith 1970) is assigned as a fitness value to each configuration. In physics, the Hamiltonian energy of Ising spins defines a fitness landscape on the configuration space of N spins. Each spin takes the value either 1 or -1 ($k = 2$). Bray and Moore (1980) argued about the number and distribution of meta-stable states (local optima) of the Hamiltonian energies.

To explore these fitness landscapes, we need a rule by which a point in the space moves to one of its neighbors. Then, consecutive movements of a point to the neighbors form a *walk* on the landscape. Macken et al. (1991) used *random point mutation* that changes a single letter in the string to specify neighbors of the string. Then, by sampling points along an *evolutionary walk* in which point moves to the *firstly found fitter* neighbor, they studied the statistical properties of the landscape of the chemical affinity of antibody for antigen in immune response. Weinberger (1990) used two different walks: *"gradient walk"* in which the walker steps to the *best* of its neighbors and *"random adaptive walk"* in which the next step is chosen *at random* from the set of better neighbors, to investigate the Kauffman's NK landscape which is a model formulated in more general form (Kauffman and Levin 1987).

In our context, we have two possible fitness landscapes: the one defined on bipolar pattern space, and the other defined on synaptic weight space (see above). In both cases, top region is flat with the former top corresponding to *basin of attraction of a pattern* and the latter top to *plasticity of synaptic weights.* In this paper, we simply view fitness landscape as a surface determined by fitness value on all the possible weight configuration.

1.3 Fitness Evaluation of the Hopfield Model

Assume that p random bipolar patterns represented by the Equation (1) are stored as an associative memory in the network. When one of these memorized patterns ξ^ν is given to the network as an initial state, states of neurons vary from time to time afterwards (unless ξ^ν is stored as a fixed point). In order for the network to function as an associative memory, the instantaneous network state must be similar to the input pattern ξ^ν. The similarity is defined by,

$$m^\nu(t) = \frac{1}{N}\sum_{i=1}^{N} \xi_i^\nu s_i^\nu(t), \tag{4}$$

where $s_i^\nu(t)$ is the state of neuron i at time t. This $m^\nu(t)$ is referred to as *overlap*. The quality of retrieval of memorized patterns are represented by $\langle m^\mu \rangle$, a temporal average of $m^\nu(t)$ over certain time interval t_0. We evaluate the fitness value of each network by further averaging these $\langle m^\mu \rangle$ over all p memorized patterns. Namely, the fitness f is

$$f = \frac{1}{p \cdot t_0}\sum_{\nu=1}^{p}\sum_{t=1}^{t_0} m^\nu(t). \tag{5}$$

Then the fitness 1 implies all the p patterns are stored as fixed points, while all other cases have a fitness less than 1. In this paper, t_0 is set to $2N$, twice the number of neurons.

2 Experiment

In this paper, our simulations are with networks of 49 neurons and we focus on solutions that store just one pattern ($p = 1$). Solutions are normalized on the surface of the hypersphere of radius one whose center is at the origin, since two weight matrices that differ only by scaling factor are equivalent. In other words, a straight *ridge* emanating from the origin and points surrounding this ridge are mapped on the surface of the hypersphere and these points construct *peaks* in a usual sense.

2.1 Downhill Walk by EP

We employ EP to explore our fitness landscape.

EP is a technique for optimization, proposed by L. Fogel, Owens and Walsh (1966)[2] . In EP, a population is constructed of μ individuals. Each of these μ

[2] Although the population in EP is usually randomized at the beginning of a run, we start it with all identical points.

points is mutated once, which produces μ mutants. Besides its N^2 coordinates, each point has additional N^2 variables σ_i to be added to the i-th coordinate of the point. Since these σ_i are chosen from the Gaussian distribution with an appropriate standard deviation, this is called Gaussian mutation. After the mutation, these σ_i are also modified as:

$$\sigma_i^{\text{new}} = \sigma_i^{\text{old}} + 0.01 \cdot \sigma_i^{\text{old}} \cdot N_i(0,1). \tag{6}$$

All the σ_i are initialized to σ_0 at the beginning of a run. Now, we have the original μ points and their μ mutants. The fitness value of each of the 2μ points are compared to those of q points which are chosen randomly at every time of the comparison from the whole 2μ points. Then the 2μ points are ranked according to the number of wins, and the best μ points survive (q-tournament selection). The cycle of reconstructing the new population with better fitness points and restarting the search is repeated until one of the global optima is found or a set maximum number of generation has been reached. In this paper, both λ and μ are set to 200 and q is chosen to be 10.

In his Ph.D dissertation, Jones (1995) argued that *"a reverse hill-climbing algorithm allows the determination of details of the basin of attraction of points on a landscape."* The landscapes that Jones explored were defined on a discrete domain, that is, on vertices of a hyper cube. Here, we extended the idea to a landscape that is defined on a continuous domain. In order for a walker to move around in *discrete* domain, a point mutation is employed. Since our landscape is defined on real-valued variables, the Gaussian mutation is employed. Though many versions of such a walk can be thought as candidates, we use EP here expecting its path to be almost the steepest one, and call the walk "*reverse EP hill-climbing*". The reverse EP hill-climbing is essentially the same as the EP search described above, with the difference being the following two points: (i) the starting point is a point in a top region of a peak, and (ii) fitness function f is replaced with $(1 - f)$ so that the top of the peak corresponds to the lowest fitness value of zero.

3 A Visualization by Cross Section by a Hyperplane

In this section, we explore Hebbian peak which stores a specific one pattern as associative memory.

In order to learn what the landscape looks like, a single EP walker explores on the region starting from the original top until the walker found the point whose fitness value is less than one. This exploration is repeated 100 times, and each time when the walker reaches the point where fitness is less than one, the point is recorded as an edge point. In Figure 1, a series of fitness values on the hyper straight line in between and outside a pair of two points chosen out of these 100 points[3] . That is a cross section of the fitness landscape being cut by

[3] Although the two points whose distance is maximum are chosen here as an example, we observe a similar result with other set of two points including the minimum distance pair of points.

a hyper plane.

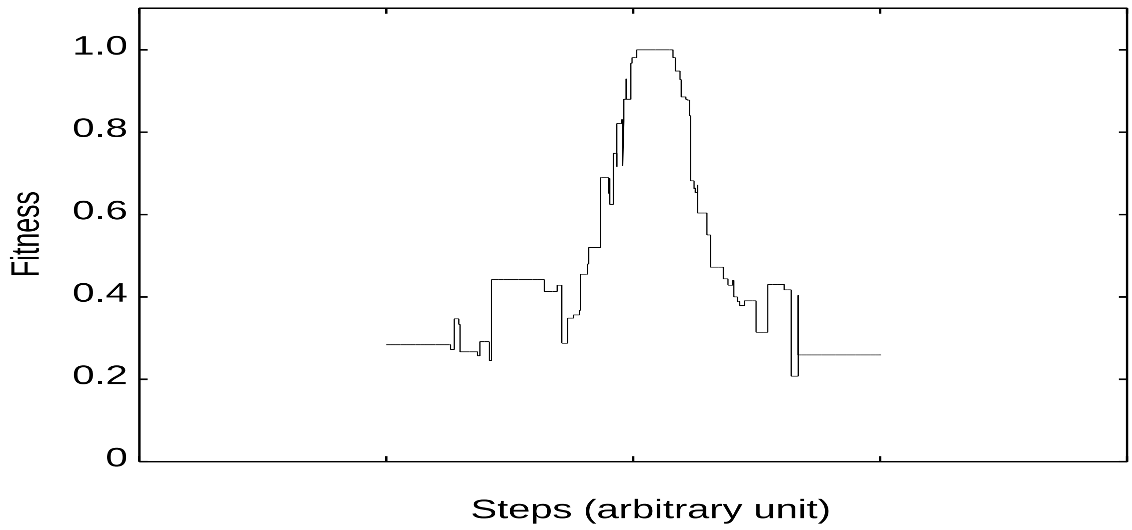

Fig. 1. A cross section of a Hebbian peak.

As shown in the figure, we can also observe some local peaks along a sidewall. The method is very simple and the cross section is not the representative one. However, when we want the rough image of a fitness landscape which is often of our case, we could obtain enough information by this method. For example, in our problem, top region of peaks is flat, which corresponds to *plasticity of synaptic weights.* As shown in Figure 2, which is a cross section around the peak of a Hebbian top, we can visualize the top region of a peak. This enable us to compare synaptic plasticity of two different solutions. We can also compare the hardness to be located by observing steepness of sidewall of two peaks. Or we can observe the number and distribution of local peaks around a global peak in this way (not shown here). By observing these two characteristics, for example, we can learn why one peak is difficult to be searched for than another peak (not shown here either).

4 Conclusion

We have proposed a simple method to learn a geometry of fitness landscapes. Though we focus here on our fitness landscape defined on weight space for the fully connected neural network model of associative memory, the method can be

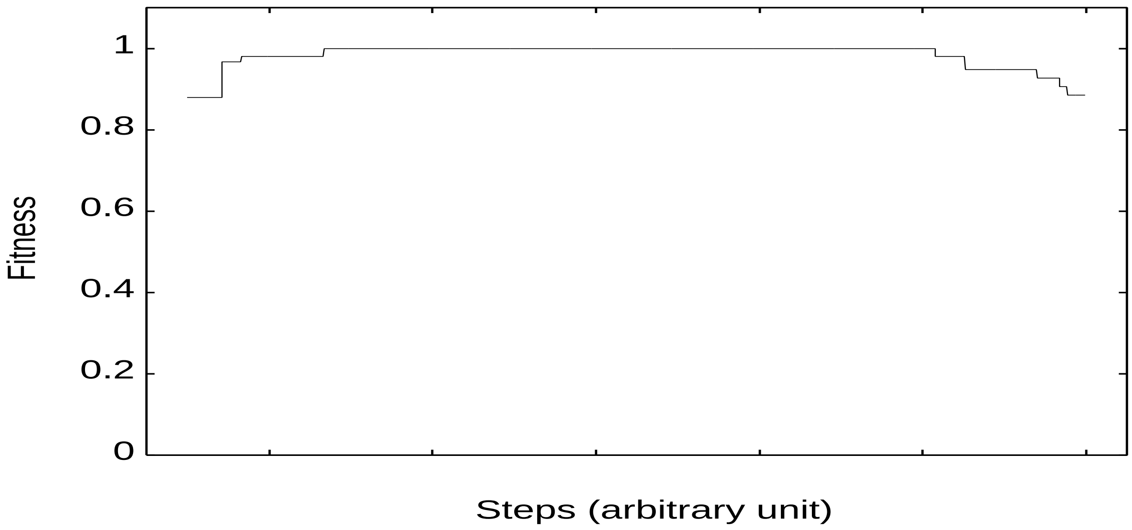

Fig. 2. Width of top of a Hebbian peak.

applied to fitness landscape of any problem and we can obtain information such as: What does a peak in the landscape look like? How large is the top area of the peak? How steep is the side wall? How does a peak surrounded by local peaks?

We now are analyzing our fitness landscape of associative memory from these view points.

Reference

Amit, D. J. (1989) *Modeling Brain Function: The World of Attractor Neural Networks.* Cambridge University Press.

Bray, A., and M. Moore (1980) *Metastable States in Spin Glasses.* Journal of Physics C: Solid State Physics 13, pp.L469–L476.

Cooper, L. N. (1973) *A Possible Organization of Animal Memory and Learning.* Proc. of the 24th Nobel Symposium on Collective Properties of Physical Systems, pp.252–264.

Fogel, L. J., A. J. Owens, and M. J. Walsh (1966) *Artificial Intelligence through Simulated Evolution.* Wiley.

Fontana, W., and P. Schuster (1987) *A Computer Model of Evolutionary Optimization.* Biophysical Chemistry 26, pp.123–147.

Gardner, E. (1988) *The Phase Space of Interactions in Neural Network Models.* Journal of Physics, 21A, pp.257-270.

Hebb, D. O. (1949) *The Organization of Behavior.* Wiley.

Hopfield, J. J. (1982) *Neural Networks and Physical Systems with Emergent Collective Computational Abilities.* Proceedings of the National Academy of Sciences, USA 79, pp.2554-2558.

Jones, T. (1995) *Landscapes and Search.* Ph.D dissertation, The University of New Mexico.

Kauffman, S. A., and S. Levin (1987) *Towards a General Theory of Adaptive Walks on Rugged Landscapes.* Journal of Theoretical Biology 128, pp.11–45.

Komlós, J., and R. Paturi (1988) *Convergence Results in an Associative Memory Model.* Neural Networks 1, pp.239–250.

Macken, C., P. Hagan, and A. Perelson (1991) *Evolutionary Walks on Rugged Landscapes.* SIAM Journal of Applied Mathematics 51, pp.799–827.

Mattis, D. C. (1976) *Solvable Spin Systems with Random Interactions.* Physics Letters, 56A, pp.421–422.

Maynard Smith, J. (1970) *Natural Selection and the Concept of a Protein Space.* Nature 225. pp.563–564.

Sherrington, D., and S. Kirkpatrick (1975) *Solvable Model of a Spin Glass.* Physical Review Letters 35, pp.1792–1796.

Weinberger, E. D. (1990) *Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference.* Biological Cybernetics 63, pp.325–336.

Wright, S. (1932) *The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution.* Proc. of the 6th International Congress of Genetics, pp.356–366.