Noise and the Pursuit of Complexity: A Study in Evolutionary Robotics

Anil K Seth

Centre for Computational Neuroscience and Robotics
and School of Cognitive and Computing Sciences,
University of Sussex, Brighton, BN1 9SB, UK,
anils@cogs.susx.ac.uk

Abstract. This paper describes a new approach for promoting the evolution of relatively complex behaviours in evolutionary robotics, based on the use of noise in simulation. A 'homing navigation' behaviour is evolved (in simulation) for the Khepera mobile robot, and it is shown that high noise levels in the simulation promote the evolution of relatively complex behavioural and neural dynamics. It is also demonstrated that simulation noise can actually *accelerate* artificial evolution.

1 Introduction

The objectives of this paper are to illustrate two new ways in which the careful use of *noise* can be of benefit to evolutionary robotics; *a)* by promoting the evolution of relatively complex behaviours, and *b)* by accelerating the artificial evolution process. These objectives require that a particular approach to evolutionary robotics be adopted; namely the evolution of robot controllers in *simulation,* with successful subsequent transference of evolved controllers to the real world. The use of simulation permits the incorporation of specifiable amounts of noise[1], and the condition of transference ensures that evolutionary robotics remains faithful to real robots in the real world.

This approach, already shown to be viable (for example see Jakobi [4], Nolfi [9], Miglino et al. [8]), stands in contrast to two major alternatives. The first, evolution in real time on real robots (e.g. Floreano and Mondada [2]), does not easily allow for the explicit, quantitative specification of noise levels[2] and is also formidably time intensive. The second, evolution in simulation with significant abstraction from reality (e.g. Sims [11]), removes evolutionary robotics from the real world and therefore will not be pursued here.

In what follows, a simulation is used to evolve a 'homing navigation' behaviour; a behaviour originally evolved and investigated in the context

[1] Noise here consists of various aspects of a simulation incorporating degrees of randomness; these aspects are detailed in section 4. It does *not* refer to stochastic aspects of the search algorithm.

[2] Qualitative noise levels *can* be manipulated in the real world, for example by introducing flashing lights into the vicinity of the robot (see e.g. Jakobi et al. [6]).

of real-world evolution by Floreano and Mondada [2]. It is demonstrated that large amounts of simulation noise promote the evolution of robots with relatively complex behaviours and neural dynamics compared to those evolved in simulations with low noise levels. It is also demonstrated that high noise levels can accelerate the evolutionary process.

The rest of this paper is organised as follows: section 2 discusses the role of noise in simulation, and introduces some new perspectives. Sections 3 and 4 describe the experiments undertaken; briefly introducing the original real-world study by Floreano and Mondada, and then describing the simulation replication in detail. Section 5 presents some results, which are then discussed in section 6.

2 Noisy simulations

In contrast to the potential roles of noise outlined in section 1, evolutionary robotics has, to date, concentrated on the discovery that noise in simulation can help bridge the 'reality gap'. Controllers evolved in simulations with appropriate noise levels transfer to real robots and real situations (see e.g. Jakobi [5], Miglino et al. [8]).

Jakobi ([5], [4]) has formalised the use of noise for facilitating transference by distinguishing between *base set* features (those aspects of an agent-environment system that may come to play a part in the eventual behaviour) and *implementation* features (those aspects which are either simulation artefacts, or not relevant to the behaviour, or just real-world aspects that are difficult to model). The idea is then that the base-set features should be modelled noisily, so they that they will transfer to the real world, (this being base set 'robustness'), and that the implementation aspects should be made *very* unreliable, so that evolution cannot come to incorporate them in *any* viable controller, (this being base set 'exclusivity').

The use of this approach, with its integral role for noise, thus also permits the assessment of the effects of noise on the complexity of the evolved behaviour, and on the speed of evolution. Both these ideas are new to evolutionary robotics, but the former builds on previous work (Seth, [10]), in which noise is shown to promote the evolution of strategies of increasing complexity in co-evolving Iterated Prisoner's Dilemma ecologies. In these studies, the 'evolved complexity' is reflected in the strategies deployed by the agents, which in turn is reflected in the lengths of variable length genotypes. This paper is concerned with exploring the same effect in an evolutionary robotics context.

3 A 'homing navigation' experiment

The context for this study is drawn from an experiment by Floreano and Mondada 1996 [2]. Their work is briefly described below.

They demonstrate the real-world evolution of a 'homing-navigation' behaviour, using a Khepera robot [7]. This robot is equipped with an extra 'floor sensor' in addition to the usual array of 8 infra-red proximity (and ambient light level) sensors. The environment for the experiment was a 40cm by 45cm walled arena, situated in a dark room, but with a small light tower placed in one corner. This corner (denoted the 'charging area') also had black paint on the floor out to a radius of 8cm (the floor being otherwise white).

The robot had a simulated battery of 50 'actions'[3], after which it would 'die'. However, if it happened to pass over the charging area during its life, the battery would be instantaneously recharged and the robot could carry on for another 50 actions (up to an arbitrary maximum of 150).

The sensor arrangement of the robot and the (fixed) architecture of the neural net controller is shown in fig 1; with 8 input units corresponding to the 8 IR sensors, 2 input units for the ambient light sensors 2 and 6 (on the front and rear of the robot body), and one input unit each for the floor sensor and battery level. These inputs were fed through to a 5 unit internally recurrent hidden layer, which was in turn connected to a two unit motor output layer, which then set the wheel speeds. Sigmoid activation functions were employed at all layers except the input layer, which linearly scaled the sensory inputs to range from -0.5 to $+0.5$.

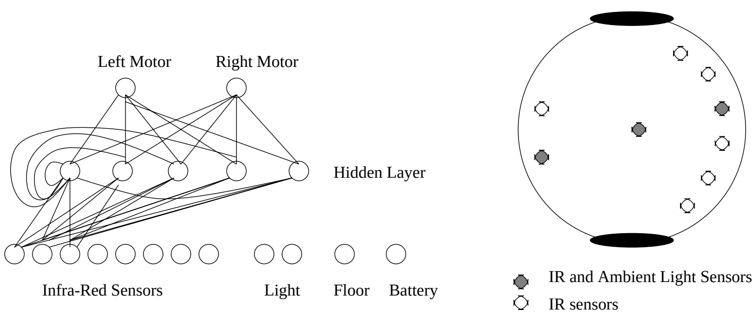

Fig. 1. *Network architecture and sensor layout. For clarity, not all network connections are shown. All input units are connected to all hidden units, and all hidden units are connected to themselves and to every other hidden unit, as well as to both motors. The central floor sensor is located beneath the robot base.*

The fitness function used by Floreano and Mondada was very simple, calculated incrementally at every step (except when the robot was directly over the charging area, where no score was awarded), and maximised by high speed and low IR input (V is the scaled average wheel speed (taken as a vector), and i the scaled activation value of the IR sensor with the highest value):

$$\theta = V(1-i), \qquad 0 \leq V \leq 1, \quad 0 \leq i \leq 1$$

[3] Each action corresponded to one update of the controlling neural network, with updates taking place about every 300ms.

There is nothing in this fitness function that explicitly specifies periodic return to the charging area, but robots that come to adopt this strategy will tend to live longer and thus accrue higher fitness than those that do not.

Floreano and Mondada performed artificial evolution in the real world, downloading candidate controllers onto real Khepera robots for each evaluation, and using a simple tournament GA to evolve the weights and thresholds of the network (which were then fixed for the duration of each individual). Over the course of ten days (200 generations) fit individuals evolved. These individuals typically explored their environment at high speed, returning to the charging area at periodic intervals, usually just in time to avoid running out of battery power. This work of Floreano and Mondada is therefore an impressive demonstration of real-world evolutionary robotics, and also provides convincing support for their primary thesis that simple fitness functions can be used to evolve relatively complex behaviours.

The thrust of the present work is different. A simulation of the Khepera robot, controller, and environment, supports the evolution of a similar 'homing-navigation' behaviour to that described above. The present hypothesis, however, is that the amount of *noise* present in the simulation strongly influences the complexity of the behaviours (and underlying neural dynamics) that evolve. It is also illustrated that noise can accelerate the artificial evolution process. A subsidiary motivation is to provide a further example (in addition to e.g. Jakobi [5], and Miglino et al. [8]) that long periods of real world evolution can be reduced to very short periods of simulated evolution, without the sacrifice of real-world efficacy.

4 The simulation

In order for the evolved behaviours to work in the real world, and to permit the incorporation of quantitatively specifiable amounts of noise, Jakobi's 'minimal simulation' methodology was followed [4]. Three important 'base set' aspects of the experiment were identified:

- the way in which the IR sensors respond, which depends on the orientation of the robot and the distance of the robot from a wall at a given angle or a corner of a given shape.
- the way in which the ambient light sensors respond to a distant light source, which depends on the orientation of the robot, and the angle of the robot to the light.
- the way in which the robot's position and orientation changes, which depend on the wheel speeds.

In order to simulate these factors, three look-up tables were employed (adapted from [4]). One held the values that IR sensors would (roughly) hold if the robot were to be positioned at 10cm from an infinitely long flat wall. This table consisted of values for all 8 IR sensors, for each of 10 robot orientations from 0 to $\pi/2$. To then calculate the actual values of the IR sensors, given walls at different distances and orientations, and for robot orientations greater than 90 degrees, the values in this table were appropriately scaled (linear scaling between 0 and 1023). If the robot happened to be in a corner, two sets of IR sensor readings were calculated, one for each of the walls involved. The maximum value for each sensor was then taken to form a composite reading. The second look-up table simply held the angles of the sensors with respect to the centre line through the robot body. The angle of each sensor to the light source could then be calculated as a function of an angle from this table, the orientation of the robot, and the angle of the robot to the light. It was then a simple matter to calculate whether or or not this angle (of sensor to light) fell within the angle of acceptance of the sensor.

The third look up table simply held values for the changes in the x and y coordinates of the robot, if it were travelling with a given speed in each of 36 different orientations. This could then be linearly scaled according to the actual speed of the robot.

The other important aspect of the simulation was, of course, that a lot of noise was employed, both during each trial and between trials (each individual was evaluated over twelve separate trials in the GA). *Intra-trial* noise was applied to the IR, light, and floor sensor values, also the robot position, orientation, and rate of orientation change during turning, and finally the position of the robot after impact with a wall. In the last case, these collisions were modelled by just randomly repositioning the robot within about 2-3cm of the wall, with a large orientation and speed change (also randomly determined). *Inter-trial* noise was applied to the angle of acceptance of the light sensors, the arena dimensions (including charging zone radius), and also the actual *levels* of IR, background IR, and ambient light noise.

In terms of Jakobi's 'minimal simulation' methodology discussed in Section 2, the intra-trial noise was designed to deliver base set robustness, and the inter-trial noise was for ensuring base set exclusivity. Thus the simulation would be expected to deliver transferable controllers. For the purpose of investigating the complexity of the evolved behaviours and neural dynamics, and the speed of evolution, two conditions were investigated; one with high levels of both inter- and intra-trial noise, and another with zero inter-trial and very low intra-trial noise. The actual levels used are given in Appendix 1.

The experiment proceeded by using a distributed GA, with a population of 100, to evolve the weights and thresholds for the controlling network as shown in fig 1. These weights and thresholds were specified as unbounded floating point numbers on a 102 allele genotype, with mutation and crossover being the only genetic operators employed[4].

[4] Crossover probability was set at 0.95, with a 0.03 probability of point mutation per allele; a point mutation altered the value of the allele by a random value within the range ± 0.5 (alleles were initialised within the range ± 1).

5 Simulation results

5.1 Basic performance and transference

Many evolutionary runs were performed, in both high noise and low noise conditions. With high noise levels in the simulation, evolutionary runs of about 100 generations always produced very fit individuals. These runs took about 1 hour on a single user Sun SparcUltra (143MHz) workstation, orders of magnitude faster than the real world evolution reported in [2]. Equally fit individuals evolved with low noise levels, although not nearly as rapidly as with high noise levels (this result is discussed in section 5.4). An example graph charting the progress of a 'high-noise' simulation is shown in fig 2.

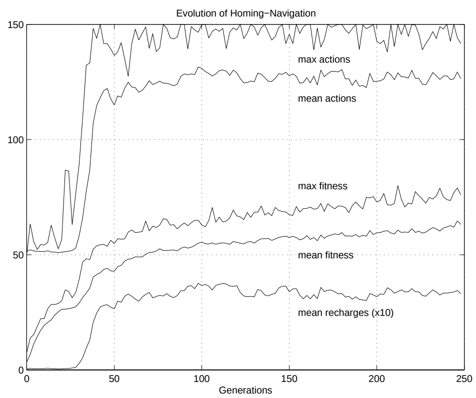

Fig. 2. *An example evolutionary run of a noisy simulation, demonstrating the evolution of fit robots in a very short time.*

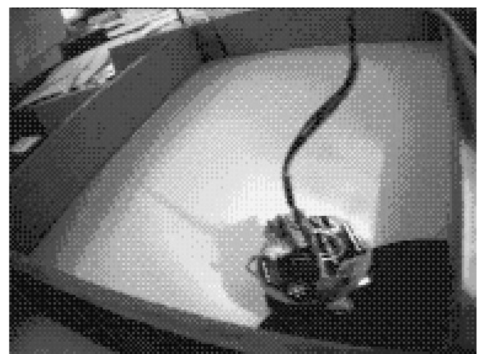

Fig. 3. *A real khepera about its business, controlled by a network evolved in a very noisy simulation.*

Successful transfer to reality was consistently observed when networks from the fittest robots, evolved in noisy conditions, were downloaded onto real Kheperas (see Jakobi [4] for further examples of successful transference using the same simulation methodology). This transference was effortless, with evolved networks effectively controlling real robots without any further 'tinkering' of the evolved network or the real environment (which consisted of a cobbled together arrangement of wooden walls and a torch balanced at one corner). Fig 3 illustrates a real Khepera (powered externally, but with all processing on-board) just leaving the recharging area halfway through a demonstration. The trajectories traced by real Kheperas were very similar to those seen in simulation, except that the real robot very occasionally impacted with the arena walls. These collisions did not normally prevent the robot from recovering and reaching the charging area when necessary. Robot controllers evolved in non-noisy conditions, however, did *not* transfer effectively to the real world. The relevance of this is discussed in section 6.

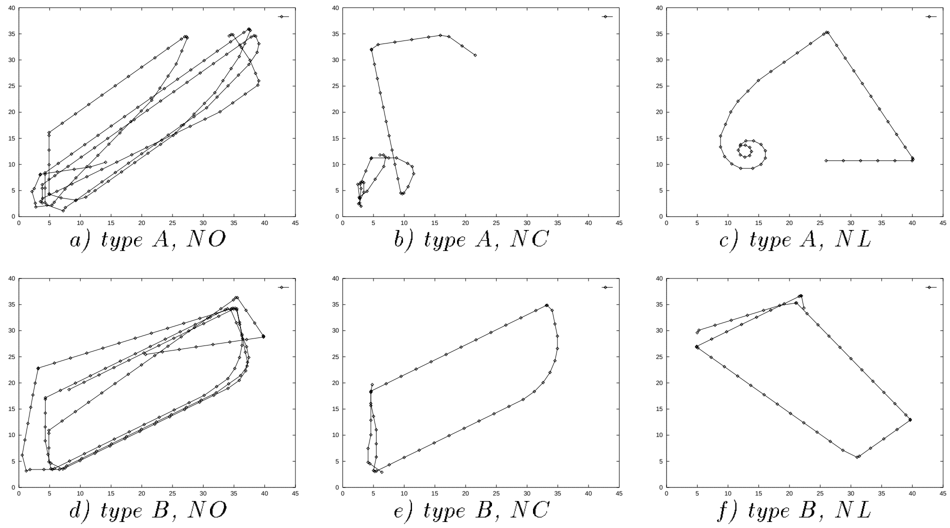

a) *type A, NO* b) *type A, NC* c) *type A, NL*

d) *type B, NO* e) *type B, NC* f) *type B, NL*

Fig. 4. *Trajectory plots (in simulation) of robots evolved under noisy (a-c) or non-noisy (d-f) conditions in either NO (a,d), NC (b,e), or NL (c,f) conditions. The charging area is located in the bottom left-hand corner of each plot, out to a radius of 8cm. The type B robots maintain a simple trajectory regardless of the environmental manipulations, but the type A robots clearly deploy more complex 'searching' and 'circling' behavioural strategies.*

The following analyses assess the contribution of high noise levels to the evolution of neural and behavioural complexity.

5.2 Behavioural analysis

A set of twelve evolved robots were analysed - six from noisy simulations (henceforth type A robots), and six from non-noisy simulations, (type B robots). All analysis took place in simulation. Initially, three environmental conditions were analysed for each robot[5]; a normal (NO) condition (with light source and charging area - this is the condition in which the robot controllers were originally evolved), a 'no charging area' (NC) condition, where the black paint is removed and the robot cannot recharge, and a 'no-light-source' (NL) condition where, although the charging area is present, the light source at the corner is removed. Low noise levels were employed in all these conditions.

Fig 4 (a-c) illustrates typical overhead trajectory plots for the robots evolved in noisy environments (type A) in the three conditions, and (d-f) illustrate the same for robots evolved in non-noisy environments (type B). In the NO condition, both A and B robots can repeatedly find the charging area (situated in

[5] These tests were also performed by Floreano and Mondada [2], who observed similar results to those of the 'noisy' robots in the present study. However, the conclusions drawn about the nature of the controlling network are different in the present work.

the lower left hand corner), and their trajectories are not obviously different. However in the NC and NL conditions, there are clear differences. The B robots maintain a behaviour pattern qualitatively similar to that displayed in the NO condition, but the A robots do nothing of the kind.

The B robots seem only to have evolved to move in straight lines and to turn upon encountering walls; a strategy which does indeed periodically return the robot to the charging area in predictable environments. The A robots, by contrast, are clearly affected by the presence (or absence) of the black charging area and the light source. In the NC condition, these robots navigate towards the charging area and remain in the vicinity, performing what could be described as a *searching* behaviour. In the NL condition, the robots begin, as in normal conditions, with a fairly linear trajectory, but shortly begin to *circle,* giving the impression that the robot is trying to orient to a light source.

These searching and circling behaviours were also observed in real-world Khepera behaviour, when the environment in fig 3 was manipulated in the appropriate way. All six A robots displayed similar searching and circling behaviours, and all six B robots displayed the simple behaviour (as in fig 4 (d-f)).

In a further example of how the A robots deploy more complex behaviours than B robots, one of each kind were compared in a (low noise) condition in which the walls were removed (again in simulation), with the charging area then extending in a complete circle around the light source. Fig 6 illustrates that the B robot was completely impotent in such circumstances, hinting at its reliance on IR stimulation. Only one typical run is shown, but out of 40 tests the robot reached the charging area just once. By contrast, the A robot reached the charging area 10 times out of 40, and in 4 cases returned more than once. Fig 5 illustrates a particularly impressive A robot trajectory, and although in general the robot is undeniably adversely affected by the lack of walls, the considerably greater success rate enjoyed by the A robot suggests that a greater range of environmental stimuli (not just IR) is being assimilated in the determination of its behaviour. Indeed, the A robot is clearly able to turn in the absence of a wall, and the B robot is not.

5.3 Neural analysis

In this section, it is shown that the neural dynamics of the type A robots are also more complex than those of the type B robots, and in ways commensurate with their behavioural differences.

Initially, activation plots for all 19 neurons in all three (simulation) conditions (NO, NC, and NL) for all of the 12 robots (6 A, and 6 B) were collected. Figs 9 and 10 present example plots for one robot of each type in the NO condition. These plots illustrate, (and this is true for *all* the plots), that whereas the hidden units (HUs) of the B robots react almost solely to IR stimulation, the HUs of the A robots react much more strongly to light, battery, and floor sense data. For example, at time-steps 35, 60, 95, and 125 in fig 9, hidden unit and motor output can be seen when there is *no* IR input. This is never the case in fig 10. Indeed, in the 'no-wall' condition discussed in section 4.3 (and therefore in the

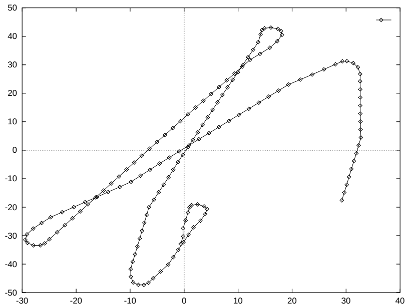

Fig. 5. *Type A robot in no-wall test; note that the edge of the graph does* not *represent a wall, and that the charging area is situated in a circle around the origin.*

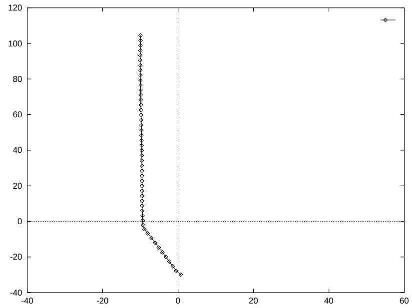

Fig. 6. *Type B robot in no-wall test; again the edge of the graph does* not *represent a wall, and the charging area is a full circle around the origin.*

absence of any IR input) only the *A* robots display any significant HU activity (figs 7 and 8).

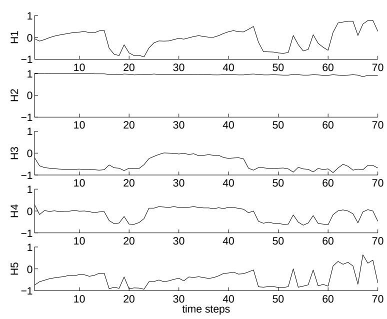

Fig. 7. *Type A robot HU activation in the no-wall test; significant levels of HU activation (only first 70 time-steps shown).*

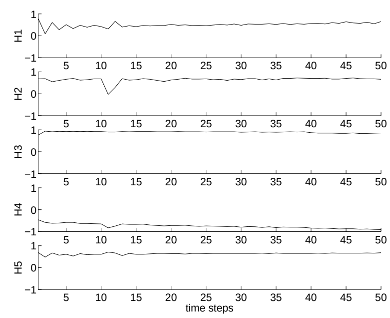

Fig. 8. *Type B robot HU activation in the no-wall test; very little activation in any units.*

To explore these results in a non-behavioural context, short periods of high activity (spikes) were injected into six combinations of input units, with the subsequent activations of the hidden and motor units being recorded. The first two conditions consisted of IR inputs only; with either all 8 inputs active, or all except the rear two[6]. The next two conditions tested combinations of ambient light input spikes in the absence of IR input. The fifth condition injected a negative floor sensor spike (as if the robot were over the charging area); again in

[6] 'Non-active' units, in all conditions, were set to 0, except for floor sensor and battery units, which were set to 1 (signifying a full battery, and a robot position away from the charging area).

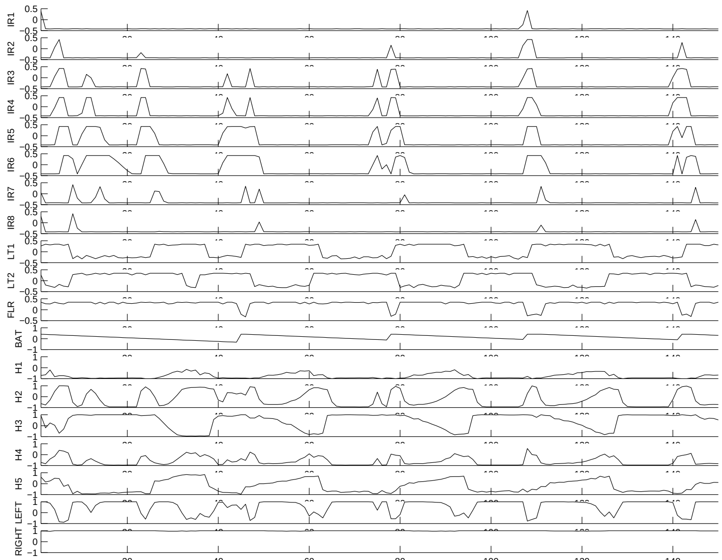

Fig. 9. *Neuron profile for type A robot (evolved in noisy conditions), tested in the NO condition. HUs (H1-H5) and motor units display activity* not *correlated with IR input (notably H1,H3,H5 - see time steps 35, 60, 95, and 125).*

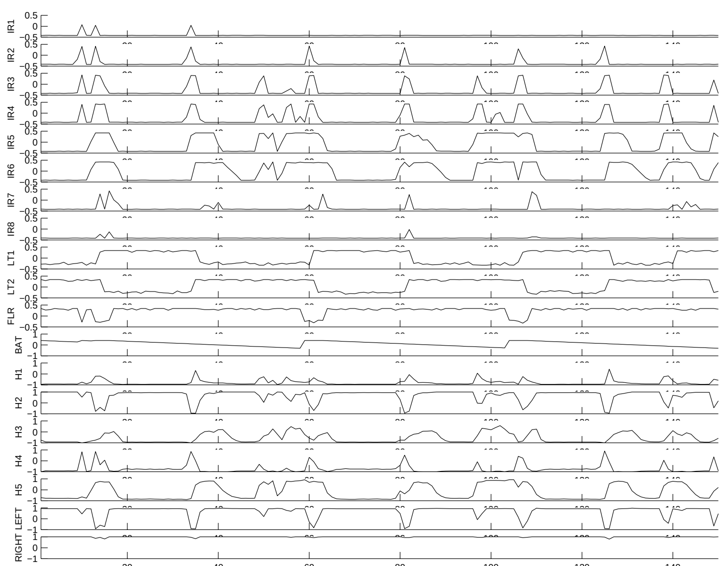

Fig. 10. *Neuron profile for type B robot (evolved in non-noisy conditions), tested in the NO condition. HUs (H1-H5) seem to be predominantly reacting to IR input.*

the absence of IR input, and the last condition injected a negative battery spike (signifying an empty battery), also in the absence of IR input.

Fig 11 presents summary data for all 12 robots over all these six conditions, in terms of the HU activity elicited. For example, the third and fourth conditions involved light input spikes, and 45 percent of the type A robot HUs responded strongly in these conditions, compared to 18 percent of the type B robot HUs. Similarly, the sixth condition tested the responses to battery sense data, and again more A HUs responded than B HUs. Thus, fig 11 makes it clear that the A robots take greater account than the B robots of the light and battery sense data. These conditions were statistically significant according to Mann-Whitney U tests ($(U = 57.0; df = 6,6; p < 0.01)$, $(U = 56.5; df = 6,6; p < 0.01)$ respectively). And although the statistical test is not significant, the B robots do appear to rely more heavily on IR input than the A robots.

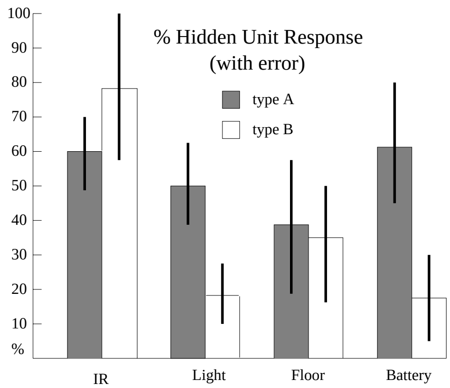

Fig. 11. *Hidden unit response patterns for both type A and type B robots. See text for details.*

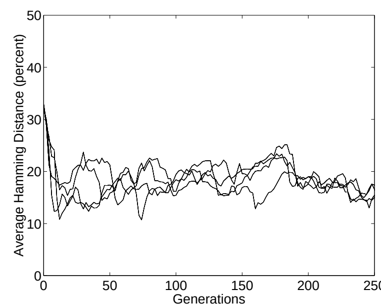

Fig. 12. *Convergence statistics for 4 runs of a noisy simulation (superimposed). Population converges rapidly to an average Hamming distance of about 20 percent.*

The results in this section are clearly suggestive of more complex evolved neural dynamics from noisy simulations than from non-noisy simulations (of course, the neural architecture itself is the same in both cases). This additional complexity is revealed primarily through type A HUs that respond to a wider range of environmental stimuli. The neural dynamics must therefore also cope with melding these multimodal inputs into a coherent motor output. These observations corroborate the behavioural data in that only the type A robots were strongly affected by manipulations of the light source and charging area. This suggests that these robots navigated to the charging zone by using the light, and could act on the basis of some internal state influenced by remaining battery level.

5.4 Noise accelerates artificial evolution

As suggested in section 5.1, it was indeed observed that high noise actually *accelerated* the evolutionary process, Figs 13 and 14 illustrate this effect, superimposing plots (of many evolutionary runs) of the number of actions taken (the lifetime) for the fittest robot of each generation, for many generations. This result is discussed in the following section.

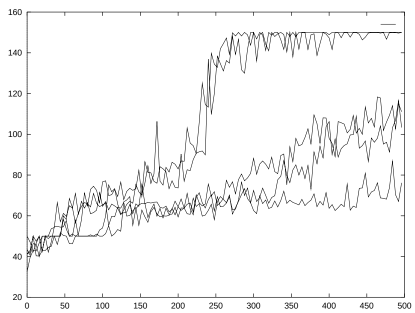

Fig. 13. *Evolution without noise. The fittest robot of each simulation reaches the maximum lifetime relatively slowly.*

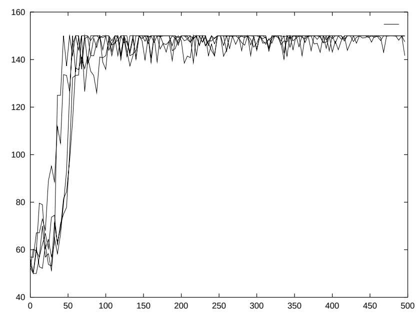

Fig. 14. *Evolution with noise. The fittest robot of each simulation reaches the maximum lifetime relatively quickly.*

6 Discussion

This paper has demonstrated how the appropriate use of noise in simulation can enhance the behavioural and neural complexity of the evolving robot controllers, and can even accelerate the evolutionary process itself. This section briefly discusses how this enhanced complexity may relate to issues of robustness, and (informally) how the effects of noise, in the present context, may be considered in terms of evolutionary dynamics.

6.1 How do these results relate to robustness?

In section 5.1 it was observed that networks evolved in non-noisy environments did *not* transfer effectively to the real world (this, however, was to be expected from Jakobi et. al [6]). One effect of noise is therefore to endow the evolving networks with real-world robustness. Indeed, one way to understand the results presented here is that robustness is being achieved through evolution exploiting the properties of a qualitatively different, more complex, behaviour. This evolutionary strategy is possible because the fitness function employed is sufficiently simple and non-specific, so to allow different behaviours to achieve similar fitness scores[7] (as noted in section 5). Future work will further address this issue of robustness, by investigating behavioural differences engendered by two different levels of noise, with *both* levels being sufficiently high to support transference to the real world.

6.2 Some informal speculations in evolutionary dynamics

Hinton and Nowlan [3] show how *learning* can influence the course of evolution, through the Baldwin effect [1]. Specifically, they show that sharp peaks in a fitness landscape can be considerably smoothed if lifetime learning is allowed. As an informal speculation, it may be that simulation noise is playing a similar role in the present experiments, in ‘smoothing out’ any steep valleys that may lie between maxima representing type B and type A behaviours[8]. This assumes that the population travels across a landscape as a relatively converged mass (fig 12 illustrates that the populations in the present experiments do converge quickly).

To continue this informal discussion, noise selectively applied to only *parts* of a simulated agent-environment system, could smooth some parts of a fitness landscape more than others, and perhaps entirely eliminate certain local maxima. In the present example, sufficiently high noise levels ensure that the simple type B behaviour is not viable, and the corresponding maxima would no longer be present in the landscape. Evolution could then proceed directly towards the maximum representing the type A behaviour. Finally, and informally once again, smoothing a fitness landscape with noise may reduce the total number of maxima present in the landscape, thus allowing evolution to proceed with greater speed to its ultimate maximum. This may help to explain the results presented in section 5.4.

Appendix 1

The noise levels used in the simulations are given overleaf:

[7] It is, of course, possible that a more specific fitness function would allow a type A behaviour to evolve in a non-noisy environment; however this paper does *not* claim that noise can permit the evolution of behaviours that would be impossible to evolve without noise.

[8] The learning algorithm employed by Hinton and Nowlan is simply random search, which could be informally construed as ‘fitness evaluation noise’.

	noisy levels	*non-noisy levels*
IR (and background IR)	$\pm 50 (\pm 10)$	$\pm 10(0)$
light (and floor sensor)	± 50	± 5
robot position	$\pm 0.1 cm$	0
robot orientation	$\pm 0.02 rad$	0
turning noise	$\pm 0.2 rads$	0
friction	$\pm 3 cm$	0
arena size	$\pm 5 cm$	0
charge radius	$\pm 1 cm$	0
light angle of acceptance	$\pm 0.25 rad$	0

Acknowledgements

Thanks to Nick Jakobi and Adrian Thompson for help with the robots, and also to Matt Quinn, Andy Philippides, Tom Smith, Phil Husbands, Inman Harvey, and Hilary Buxton for good advice. Financial support was provided by the EPSRC award no. 96308700.

References

1. J.M. Baldwin. A new factor in evolution. *American Naturalist*, 30:441–451, 1896.
2. D. Floreano and F. Mondada. Evolution of homing navigation in a real mobile robot. *IEEE transactions on systems, man, and cybernetics: part B; cybernetics*, 26(3):396–407, 1996.
3. G.E. Hinton and S.J. Nowlan. How learning can guide evolution. *Complex Systems*, 1:495–502, 1987.
4. N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Journal of Adaptive Behaviour*, 6(2), 1997. (forthcoming).
5. N. Jakobi. Half-baked, ad-hoc and noisy: Minimal simulations for evolutionary robotics. In P. Husbands and I. Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life*, pages 348–357. MIT Press, 1997.
6. N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: the use of simulation in evolutionary robotics. In F. Moran, A. Moreno, J. Merelo, and P. Chacon, editors, *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life*. Springer-Verlag, 1995.
7. K-Team. Khepera: the user's manual. Technical report, LAMI-EPFL, 1993.
8. O. Miglino, H.H. Lund, and S. Nolfi. Evolving mobile robots in simulated and real environments. *Artificial Life*, 2(4):417–434, 1996.
9. S. Nolfi. Evolving non-trivial behaviours on real robots: a garbage collecting robot. Technical Report 96-04, Institute of psychology, NRC, Rome, Italy, 1996.
10. A.K. Seth. Interaction, uncertainty, and the evolution of complexity. In P. Husbands and I. Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life*, pages 521–530. MIT Press, 1997.
11. K. Sims. Evolving 3d morphology and behaviour by competition. In *Proceedings of Alife IV*, pages 28–39. MIT Press, Cambridge, Mass., 1994.