Neural Uncertainty and Sensorimotor Robustness

Jose A. Fernandez-Leon and Ezequiel A. Di Paolo

Centre for Computational Neuroscience and Robotics (CCNR),
Department of Informatics, University of Sussex, Brighton, BN1 9QG, UK
{jf76, ezequiel}@sussex.ac.uk

Abstract. Real organisms live in a world full of uncertain situations and have evolved cognitive mechanisms to cope with problems based on actions and perceptions which are not always reliable. One aspect could be related with the following questions: could neural uncertainty be beneficial from an evolutionary robotics perspective? Is uncertainty a possible mechanism for obtaining more robust artificial systems? Using the minimal cognition approach, we show that moderate levels of uncertainty in the dynamics of continuous-time recurrent networks correlates positively with behavioral robustness of the system. This correlation is possible through internal neural changes depending on the uncertainty level. We also find that controllers evolved with moderate neural uncertainty remain robust to disruptions even when uncertainty is removed during tests, suggesting that uncertainty helps evolution find regions of higher robustness in parameter space.

Key words: neural uncertainty; minimal cognition; continuous-time recurrent networks; Evolutionary Robotics

Introduction

Real organisms live in a world full of uncertain situations and have evolved cognitive mechanisms to cope with problems based on actions and perceptions which are not always reliable. Both uncertainty and noise are intrinsic realities that all cognitive systems have to deal with in terms of the situation faced and the mechanisms deployed to cope with that situation [1]. What is not so often realised is that uncertainty and noise may also provide positive mechanisms for producing robust and adaptive behaviour. In particular, in Evolutionary Robotics the use of variability between evaluations and noise during a same evaluation are well-known factors relevant for the quality of the neurocontrollers that evolve.

There are two ways in which the careful use of noise and variability can be of benefit to evolutionary robotics; (a) by promoting the evolution of relatively complex behaviours, and (b) by accelerating artificial evolution. Large amounts of simulation noise may promote the evolution of robotic agents with relatively complex behaviours and neural dynamics compared to those evolved in simulations with low noise levels, accelerating the evolutionary process [2] and possibly enhancing robustness during evolution. In this work, robustness will refer to stability against external perturbations and internal variability [3].

Jakobi [4] has applied these ideas successfully to minimal simulations which allows neurocontrollers evolved in simulation to have a significant rate of success in

F. Almeida e Costa et al. (Eds.): ECAL 2007, LNAI 4648, pp. 786 – 795, 2007.

transferring to real robots (where direct evolution is prohibitive). Minimal simulations work by avoiding accurate replication of the physical complexities of a real-world robot-environment system and proceed to abstract a base set of factors upon which evolution must rely in order to produce the desired behaviour. All other factors in the robot-environment system are crudely modelled and subjected to large amount of noise and variability between evaluations with the result that the only suitable strategy for the evolutionary search process is to ignore them.

If minimal simulations work under the assumption that evolved solutions will ignore unreliable factors, no matter how complex, and hence be robust to their variability in real-world situations, we may ask what would be the consequence of making the neurocontrollers themselves also unreliable in some significant manner. Our hypothesis is that neural uncertainty will contribute to several aspects of robustness of performance in order to significantly and substantially improve robot performance for a given environment. Neural uncertainty could relate to the amount by which a value, typically the neuron output activity, may differ from those obtained when no significant variations in such value are produced in spite of changes in the neuron inputs.

This paper will explore previous hypothesis using a minimally cognitive model. It should be viewed as a preliminary exploration easily extendable to more complex systems and tasks. We test the role of neural uncertainty by biasing the dynamics of every single node in the agent's neurocontroller. The agent must still be able to perform the task in spite of the randomly distorted internal dynamics. We surmise that agents capable of good performance will also show several forms of robustness.

The following section describes the method proposed to induce uncertain internal activity in neurons. This method is applied on the evolution of Khepera-like simulated agents in a simple phototaxis scenario. This is followed by experiments testing the effects of uncertainty under morphological and environmental disruptions. Finally, the last section examines the consequences of the results and discusses some of the questions that remain open.

Methods

We propose to study the effects of uncertainty in neural dynamics through a minimalist approach. Our aim is to test the consequences of evolving networks under uncertain dynamics and to test the resulting robustness of their performance. A population of simulated agents is evolved to perform phototaxis under normal bodily and environmental conditions on a series of light sources. In each test, one light source is presented at a extended-finite time. Limited, but random uncertainty is applied locally to the dynamics of each neuron. The uncertainty of each neuron (y_o) is modelled as a constantly changing activation parameter. The range of this random variable (A) is a control parameter in our studies.

In preliminary tests, we noticed that evolution avoids coping with the effects of uncertainty by selecting neurons working very high or very low in their dynamical range (in our case, near to -1 or 1 output value). The use of center-crossing techniques[5] makes sure that neurons avoid long-term saturation, where the random parameter y_0 has an effect on neural activities which has to be counteracted in order to perform the evolved task. Center-crossing condition is used to enforce neural dynamics away from saturation, enabling more interesting network dynamics throughout the evolutionary search.

Continuous-time recurrent neural networks (CTRNNs) using center-crossing are used to control the agents. The CTRNN topology consist of 2 input nodes, 2 inter-neurons and 2 motor neurons. Full connectivity is used between inter-neurons and output neurons including self connections for the inter-neurons.

In the proposed experiments, agents evolved for this normal phototaxis task are then tested for phototaxis under a series of disruptions of their sensory inputs (sensor gain and visual field inversion), and motor capabilities (motor gain), as well as internal dynamics (uncertainty variation). During evolution agents have not faced such disruptions.

Agent Model

Agents were modelled as solid circular bodies of radius five with two diametrically opposed motors and two frontal light sensors with 47.75° between them according to Khepera robot sensor layout. Motors can drive backwards and forwards in an unlimited 2-D arena. Agents have a very small mass, so that the motor output is the tangential velocity at the point of the body where the agent is located. Agents have a default velocity in each motor node. Light from point sources stimulates sensors with an intensity proportional to the inverse square of the distance from the sensor to the source. The model includes shadows on sensors produced when light is occluded by the agent body.

Neurocontroller Model

A fully connected 4-neuron CTRNN and 2 additional input nodes is used as the neurocontroller. The dynamics of each neuron is modelled as follows:

$$\tau_i \frac{dy_i}{dt} = -y_i + \sum_j^n w_{ji} z_j + I_i \qquad \textbf{(1)}$$

$$z_j = \sigma(y_j + \theta_j) \qquad \textbf{(2)}$$

where, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant (range [1;-1+e^2]), θ_i the bias (subject to a sigmoid function), z_i the firing rate, w_{ji} the strength of synaptic connection from node i to node j (range [-10;10]), and I_i the degree of sensory perturbation as an incoming current, which is zero for non-input nodes. Each sensor and motor has one sensory node, or effector neuron, respectively.

The transduction from effector neuron activity to motor output is directly applied from the neuron's output in the interval [-1;1], adding noise (a value uniformly distributed in [-0.25;0.25]). In addition, the incoming signals to sensory nodes are multiplying by a gain (equal to 2). Left/right symmetry in synaptic weights is not enforced. Some parameters are genetically encoded: $\tau_{i,}$, sensory gains, w_{ji}, and θ_i. The network is updated using the Euler integration method with a time step of 0.1. The center-crossing condition is θ^*_j =0.5 (see [6]). Due to the shape of z_j, unless a neuron's bias (θ^*_j) is properly turned to the range of inputs it receives, that neuron will simply saturate on or off. Thus, the richest dynamics should be found in the neighbourhood of the center-crossing networks in parameter space, and one would expect that an evolutionary algorithm would benefit from focusing its search there [5]. In our case, center-crossing increases the effect of neural uncertainty which is otherwise nullified when using nearly-saturated neurons for generating behaviour.

Inducing Neural Uncertainty

Neural uncertainty is represented by a random local parameter (y_0) in the dynamics of every single neuron in the agent's neurocontroller. This random parameter is selected every time step from the range [-A; A], where A is a fixed value in each experiment (A=0,1,2,3,4). Therefore, we test the role of neural uncertainty by randomly biasing, for each new evaluation, the dynamics of neurons. The random variable y_0 affects neurons as follows where i denotes the node index:

$$\tau_i \frac{dy_i}{dt} = -(y_i - y_0) + \sum_j^n w_{ji} z_j + I_i \qquad \textbf{(3)}$$

Even though the addition of the term y_0 could be interpreted simply as a perturbation on the current input (I_i), this interpretation is only valid for sensor nodes. However, in this model the perturbation is applied to all neuron in the network. This uncertainty parameter can be mathematically interpreted as influencing the asymptotic behavior of each neuron.

Genetic Algorithm

A population of 60 individuals (neurocontrollers) is evolved using a steady-state, rank-based genetic algorithm with elitism (50%). All real-value parameters are encoded in a real-valued vector, where each component encodes a single parameter in the neural network. These components belong to the ranges previously named for each parameter. A global mutation operator is used [6][7] adding a small random vector to the real-valued genotype when mutation occurs defined by a fixed mutation rate in each generation with $\sigma(0,1)$. The mutation vector direction was uniformly distributed, and its magnitude was a Gaussian random variable with mean 0 and variance 1. Each individual is run for a number of independent evaluations (typically 10). The fitness of each phenotype is calculated by averaging the fitness obtained in each evaluation. Each evaluation consists of a serial presentation of 6 light sources for a relatively fixed long time period (T_{ls}=50 time steps) during life time (T=300 time steps). The agent must approach in turn and remain close to light source. After T_{ls}, the light source is eliminated and another one appears at a random distance ([10;100]) and angle ([0;2π]). The intensity of each source is fixed. When light is sensed, the sensory input varies between 0 (lowest level) and 880 (highest level), with a noise level of 0.1.

Fitness is calculated in the following manner. $F = 1 - D_f/D_i$, (D_f: final distance to source; D_i: initial distance to source). This term is taken as 0 if $D_f > D_i$, and it is calculated for each source and then averaged for the whole evaluation. Therefore, F is in the [0;1] range.

The search algorithm is run for a fixed number of generations (typically 1000), for a fixed number of iterations (typically 200) per generation, taking usually a few hundreds of generations to achieve a considerably high level of average fitness, depending on the configuration used in each experimental test. The capability of the evolved agents to deal with internal uncertainty with and without disruptions is then tested individually.

Results

A series of experiments have been carried out for studying the effects of internal uncertainty in neurocontroller evolvability (the effects of changing uncertainty during evolution) and lifetime performance (robustness against neural uncertainty and sensorimotor disruptions). Statistical assessment of the various types of robustness is proposed based on statistical significance tests (*one-tailed t-test*) and *bootstrapping t-test* [9]. First, *one-tailed t-test* relates to compare whether mean of A=0 is lower to those under uncertainty cases. As confidence interval, 95% was taken for the true difference (usually averaging the means of 10 independent experiments). For this test, we assume that two samples come from normal distributions with unknown and possibly unequal variances (Behrens-Fisher problem using Satterthwaite's approximation for the effective degrees of freedom). Second, the *bootstrapping t-test* or *studentized bootstrap* does not assume that the sampling distributions are roughly normal and have the same variance. The number of bootstrap samples is 2000 with confidence level for the confidence intervals 95%. As summary, our analysis shows similar conclusions in both tests.

The Effects of Uncertainty During Evolution

Preliminary experiments testing evolvability were performed with non-center-crossing controllers, testing them on the phototaxis task. It was found that evolution effectively minimised the effects of neural uncertainty by saturating neural firing rates during most of the performance trial. Therefore, experiments without center-crossing were not suitable to test our hypothesis (i.e. that actively coping with internal uncertainty will affect behavioural robustness).

The effects of uncertainty on fitness are shown in Fig. 1. Each bar indicates the final fitness after 1000 generations for different values of A averaged over 20 independent evolutionary runs. The performances indicate that evolution can find efficient strategies when internal uncertainty is included, which are quantitatively comparable to the control situation of A=0. There is a tendency to decrease in performance when uncertainty values are higher than A=2. We notice that variance also has a tendency to increase. Both these observations confirm the expectation that with higher levels of uncertainty evolvability will be reduced. This situation is slightly significant in A=2 evolution (Fig. 1), where the performance of evolution shows a decay. We therefore confine uncertainty to less than $A = 2$ in subsequent experiments.

Based on results shown in Fig. 1, the t-test for A=0 situation and cases with uncertainty during evolution, were as follows. Results indicate that null hypothesis (equal means) could be rejected after comparing A=0 against all other situations under uncertainty except for the comparison of A=0 and A=1. The significance of the alternative hypothesis (the mean of A=0 is higher than those under uncertainty) is accepted in every case with a significance level of 60%, 100%, 100%, 100% for A=1, A=2, A=3,

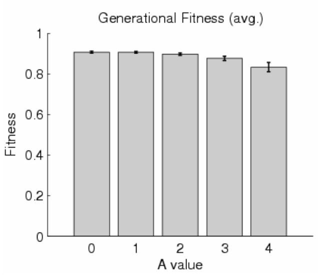

Fig. 1. Fitness values obtained averaging the fitness of 20 independent experiments after 1000 generations for different ranges of uncertainty. Horizontal axis represents the range [-*A; A*] of uncertainty

and A=4, respectively. These percentages mean that by chance is possible to see values more extreme than the one in tested experiments. In other words, percentages represent the significance level to obtain lower levels of performance when compared to A=0 situation. Similar conclusions were obtained after bootstrapping t-tests.

The relation between fitness and uncertainty during lifetime is shown in Fig. 2. strategies with A=1,2,3,4, and without uncertainty (A=0) during evolution were tested against different levels of uncertainty during lifetime when performing phototaxis using 10 independently evolved neurocontrollers in each case (Fig. 2). Experiments suggest that certain levels of uncertainty (A=1, A=2, and A=3) slightly improve or maintain agent performance in comparison with evolution without neural uncertainty (A=0). In other words, the A=0 neurocontroller obtained 84.8% of performance against uncertainty variability, meanwhile 86.9%, 96.4%, 98.5%, and 99.7% were obtained by A=1, A=2, A=3, and A=4, respectively. Such percentages indicate the robustness proportion (average) based on reached fitness under disruptions over the obtained one in the control case.

An important point is that in the absence of uncertainty, neurocontrollers evolved with uncertainty maintained their performance (Fig. 2), expressing that such neurocontrollers do not use internal noise functionally to perform the task.

For statistical significance, an analysis of the effects of uncertainty during lifetime was developed. Results indicate that the null hypothesis (equal means) could be acepted after comparing the control case against A=1 under every uncertainty situation in lifetime. Null hypothesis should not be accepted when comparing A=0 to A=2 for A=3 and A=4 levels of uncertainty during lifetime. Obtained significances for A=1 and A=2 represent the difficulty to observe values lower than the A=0 mean. Similar conclusions were obtained with bootstrapping t-tests.

Adaptation to Sensorimotor Disruptions in Lifetime

Fig. 3 shows the performance of neurocontrollers evolved with A=0, A=1, and A=2 in 20 independent evolutionary runs with different initial random seeds using the same evolutionary process previously explained; these controllers are then tested for different disruptions. The disruptions tests were mostly based on experiments proposed in [7] and [8]: (a) *neural lesions*: (a.1) sensor inversion; (a.2) node disruption (input node1 and 2); (a.3) neural uncertainty change (interneuron 1 and 2 with(out) uncertainty

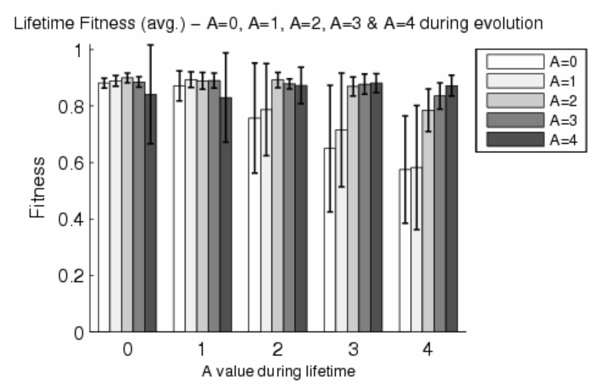

Fig. 2. Variable uncertainty levels during lifetime tests showing the obtained lifetime fitness for controllers evolved with A=0,1,2,3,4; each data point indicates the average fitness of 10 independent experiments.

A=1); (b) *sensorimotor disruptions*: (b.1) sensors gain (sensor left and right in [-75%, +100%] and in [-25%, +50%]); (b.2) motor gain (motor left and right ±50%). Experiments on the inversion of the visual field (*test2*) consist of exchanging the position of sensors. In *test3* and *test4*, the output of a single interneuron is set to zero. From *test5* to *test8* represent situations in which a random increase or decrease to sensors (right or left) is given according to [-75%, +100%] (high range) or [-25%, +50%] (low range). Finally, *test9* and *test10* depicts random right or left motor gains changes over a range of ±50%.

It is important to mention that for the neurocontrollers shown in Fig. 3, the uncertainty level during lifetime was maintained at the value at which it was evolved as indicated in the legend. Fig. 3 depicts the performance of neurocontrollers evolved with those values of uncertainty during tests A=0, A=1, and A=2 and for A=1 and A=2 with no uncertainty in lifetime. Following a short initial time period of normal phototaxis (before t=30), an agent is tested against a disruption. It is possible to estimate from this series of tests an idea of robustness corresponding to the different values of A. This can be done by averaging the ratio of performance between the disrupted and the normal conditions. For the case of A=0 this robustness indicator is 85.2%; for A=1, 85.9% with A=1 and 85.9% with A=0 in lifetime; for A=2 with A=2, 88.6%, and 90.3% with A=0 in lifetime (Fig. 3). Neurocontrollers evolved with A=0 do not show more robust performance than those evolved with A=1 or A=2. This result indicates that uncertainty has a beneficial effect on robustness against disruptions and lends some initial support to our original hypothesis. Agents with A=2 during evolution were more robust and their average performances are comparable to the cases with A=0 in control case. Results of sensor disruptions (*test5* to *test8*) are similar probably because strategies coping with the environment could implicate some degree of robustness against increases or decreases of sensing. In other words, sensory perturbations had not enough influence during tests. We can also observe that neurocontroller A=2 with uncertainty 2 in lifetime demonstrated a high level of robustness against morphological disruptions (77.6%), and 72.2% for A=1 (A=1 in lifetime) in comparison with the control case (71.5%), meanwhile A=2 without uncertainty in lifetime

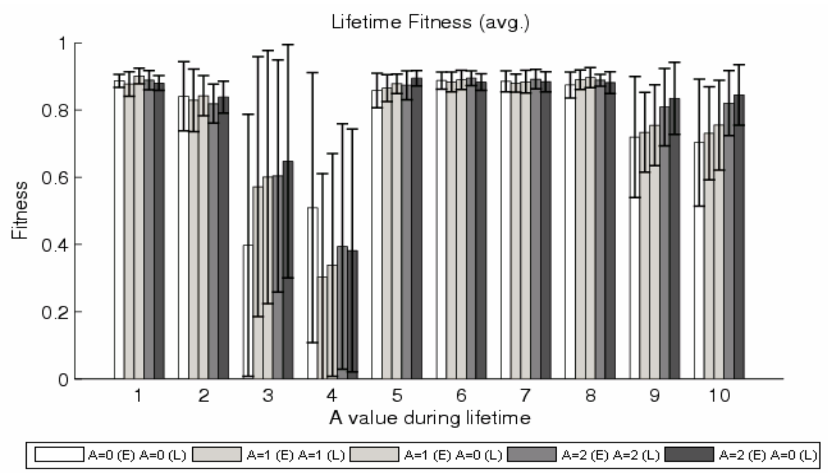

Fig. 3. Experiments with center-crossing neurocontrollers evolved (E) with A=0, A=1, and A=2; tests against lesions and sensorimotor disruptions during lifetime (L) with uncertainty values according to the legend; 20 independent evolutionary runs per plotted bar. *Test1*: no disruption; *Test2*: inversion; *Test3*: damage neuron 1; *Test4*: damage neuron 2; *Test5*: sensor left high; *Test6*: sensor left low; *Test7*: sensor right high; *Test8*: sensor right low; *Test9*: gain left; *Test10*: gain right.

obtained 80.6%. Respect of environmental disruptions indexes, they were 98.8%, 99.6%, 98.4%, 99.5% and 99.9% for A=0, A=1 with uncertainty A=1 and A=0, and A=2 with uncertainty A=2 and A=0, respectively. According to data shown in Fig. 3, results of t-test and bootstrapping support these ideas.

A full behavioural and dynamical analysis of how agents perform these tasks under different conditions remains to be carried out. Modifications to other body parameters and internal lesions (removal of connections or modifications to the behaviour of single neurons) have not been explored either.

In Fig. 3, the performance of neurocontrollers evolved with A=1, but tested with no uncertainty during the disruption (A=1 (E) A=0 (L)), results in very similar levels of robustness to the case where the same controller was evaluated using uncertainty A=1 during the test phase. This indicates that the result of including moderate levels of neural uncertainty in the evolutionary process is not only capable of producing higher robustness when the uncertainty is present, but also when the uncertainty is absent during the test phase. Similar results could be observed with A=2 neurocontroller. The presence of uncertainty seems therefore to act in ways that enhance the exploration of regions of the fitness landscape where controllers are more robust per se. This result gives us an insight into how moderate levels of uncertainty affect evolutionary search.

Conclusions

Experiments testing the capacity of CTRNN's center-crossing to deal with neural uncertainty during evolution were developed with special interest in exploring the robustness of the resulting behaviours. To this end, experimental conditions implying lesions and sensorimotor disruptions for which the neurocontrollers have not been evolved were tested and analysed.

Neurocontrollers with moderate uncertainty levels (A=1, A=2, and A=3) are more robust than those evolved with no uncertainty (A=0) or higher levels (more than A=3) in evolution and they remain robust even when uncertainty is reduced to zero during tests (Fig. 2). It means that the effect of uncertainty during evolution could help to find more suitable solutions to the task evolved. In other words, the evolutionary process seems to selects neuron dynamics that can deal with a wide range of internal and external changes. Evolution purposes combinations of neurons that can adapt themselves, and as a group. However, when uncertainty is higher than A=3 in evolution, such benefits are affected by a slightly variability.

A contrast between results could be shown. On the one hand, based on previous robustness indexes, we can tentatively conclude that a little uncertainty positively affects the robustness of neurocontrollers, but when it is considerably increased, some visible negative effects appeared in agent's performance (e.g., low performance levels, and erratic behaviours). Therefore, returning to our initial hypothesis about whether uncertainty is a prospective mechanism for obtaining more robust artificial systems, we can say that it is convenient (in robustness terms) to have a little "neural noise" (uncertainty) during evolutionary and lifetime processes, but just a few noise.

On the other hand, an important observation is that tests implying A=1 and A=2 with and without uncertainty during lifetime presented similar levels of robustness against uncertainty in Fig. 3. Some tentative suggestion about why both configurations performed well indicates that evolution generated feasible strategies against internal and external perturbations using uncertainty. Nevertheless, when uncertainty is high, it is probable that the trade-off between neural noise and center-crossing properties could be lost, meaning that evolution could not be able to generate better individuals in each generation due to the high variability caused by uncertainty. In spite of obtained results, further studies in this context must be carried out.

The results suggest that the mechanisms at work in the production of robust controllers is an enhancement of the evolutionary search under moderate levels of uncertainty leading to the discovery of zones of the fitness landscape that correspond to neurocontrollers that are robust both *with* and *without* uncertainty. For higher levels of uncertainty (more than A=3) this effect seems to break down (Fig. 2) leading to controllers that can evolve high levels of fitness but are less reliable when uncertainty is removed.

The significance of this work serves to demonstrate that evolutionary neurocontrollers can evolve robustly under lower levels of uncertainty, in spite of internal noise in every neuron and sensorimotor disruptions. The consideration of neural noise in every neuron is the novel aspect of this paper, as opposed to previous works which only permitted noise in sensorimotor signals (see [10]).

Different issues are opened for discussion. Firstly, a proper analysis of neural dynamics and behaviour in the presence of different levels of uncertainty and for the cases of disruption must be planned in future works. Some preliminary observations have indicated that functional redundancy may be used to increase robustness based on analysis of neural dynamics (e.g. neurons acting in the same way). For example, this study will indicate how isolated neuron activities could affect other neurons under different levels of uncertainty, and how neurons will react to specific internal changes. Secondly, other forms of uncertainty must be considered. For example, uncertainty could relate to a constant value y_0 for each neuron, but uncertain from one evaluation to the next. This representation of uncertainty should facilitate the analysis of neural dynamics because a fixed level of uncertainty could generate clear patterns of neuron's dynamics (e.g. increasing or decreasing their activities). Therefore, uncertainty would bias the activity of neurons. Finally, experiments evolving different tasks should be proposed for generalizing how uncertainty affects evolution. Such generalization process should cover a wide range of tasks, mainly in the minimal cognitive context (e.g. learning and action-selection problems).

All of these issues deserve further attention, as well as clearer studies for their practical application. However, one of the goals of this paper is to open questions for discussing ideas about uncertainty in neurocontroller evolution.

Acknowledgments. Thanks to A. Furey, P. Fine, M. Rohde, and S. McGregor for their corrections of this paper. This work was partially supported by the Programme Alβan, the European Union Programme of High Level Scholarships for Latin America, scholarship No. E05D059829AR, and also recognized by CONICET, Argentina, Type-I-Res.38-08/01/2004.

References

1. Di Paolo, E.A., Harvey, I.: Decisions and noise: the scope of evolutionary synthesis and dynamical analysis. Adaptive Behavior 11(4), 284–288 (2004)
2. Seth, A.K.: Noise and the pursuit of complexity: A study in evolutionary robotics. In: Husbands, P. (ed.) Evolutionary Robotics. LNCS, vol. 1468, pp. 123–137. Springer, Heidelberg (1998)
3. Kitano, H.: Biological robustness. Nature Reviews 5, 826–837 (2004)
4. Jakobi, N.: Evolutionary robotics and the radical envelope of noise hypothesis. Journal of Adaptive Behaviour 6(2) (1997)
5. Mathayomchan, B., Beer, R.: Center-crossing recurrent neural networks for the evolution of rhythmic behavior Source. Neural Comp. 14(9), 2043–2051 (2002)
6. Beer, R.: On the dynamics of small continuous-time recurrent neural networks. Adaptive Behaviour 3, 469–509 (1995)
7. Di Paolo, E.: Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In: From Animals to Animats 6, SAB'2000, pp. 440–449 (2000)
8. Beer, R.: Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroscience. Academic Press, San Diego (1990)
9. Efron, B.: The Jackknife, the Bootstrap and Other Resampling Plans. CBMS-NSF Regional Conference Series in Applied Mathematics, Monograph 38. SIAM, Philadelphia (1982)
10. Nolfi, S., Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press, Cambridge (2000)