Neural Noise Induces the Evolution of Robust Behaviour by Avoiding Non-functional Bifurcations

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

Centre for Computational Neuroscience and Robotics, University of Sussex, UK
{jf76,ezequiel}@sussex.ac.uk

Abstract. Continuous-time recurrent neural networks affected by random additive noise are evolved to produce phototactic behaviour in simulated mobile agents. The resulting neurocontrollers are evaluated after evolution against perturbations and for different levels of neural noise. Controllers evolved with neural noise are more robust and may still function in the absence of noise. Evidence from behavioural tests indicates that robust controllers do not undergo noise-induced bifurcations or if they do, the transient dynamics remain functional. A general hypothesis is proposed according to which evolution implicitly selects neural systems that operate in noise-resistant landscapes which are hard to bifurcate and/or bifurcate while retaining functionality.

Keywords: Evolutionary robotics; Systemic robustness; Continuous-time neural networks; Neural noise; Bifurcations.

1 Introduction

The role of noise in systems with sensorimotor control has generated a growing interest in bio-inspired robotics – in particular, in its relation to systemic aspects of robust behaviour (e.g., [3, 4]). Neural noise is important to be studied in behavioural systems as it may result in movement inaccuracy (e.g., constant errors) and imprecision (e.g., variable errors and uncertainty) [3]. In the context of adaptive behaviour during goal-oriented tasks, Bays *et al.* (2007) propose that the strategy of the central nervous system for dealing with neural noise, i.e., the spontaneous neural background activity present in most brain tissues, is to 'optimally combine sensorimotor signals'. Despite this broad hypothesis, we have very little idea about how the algorithms underlying the management of the effects of neural noise are realized at the neuronal level because the majority of work in this area neither explains how these mechanisms emerge from sensorimotor interactions, nor analyses how such strategies may have originated during evolution. In the context of artificial evolution, evidence that noise also has some useful properties has been presented several times (e.g., [4, 6]) and this leads us to a second question: whether in natural systems noise should always be considered detrimental. Combining these two ideas, the question of *what sort of control-strategy emerges if neural noise is induced during the evolution of neurocontrollers* becomes one of conceptual and practical interest not only for evolutionary and autonomous robotics but potentially for neuroscience as well.

M. Asada et al. (Eds.): SAB 2008, LNAI 5040, pp. 32–41, 2008.

The use of noise is a widespread practice in evolutionary robotics (ER). Using his *minimal simulations* paradigm, Jakobi (1997) has investigated the uses of noise and parametrical uncertainty in the evolution of neurocontrollers and found that they have a significant rate of success when transferred from simulated solutions into real robots (where direct evolution is impractical or prohibitive). Minimal simulations work by avoiding the accurate but costly replication of the physical complexities of a real-world robot-environment system and instead 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 subject to large amounts of environmental noise and variability between evaluations.

Some of the lessons of the minimal simulations methodology may illuminate questions about natural robustness. Biological systems exhibit phenomena, such as sensorimotor robustness to noise [3] or robustness in functional terms [6], which may relate to the presence of neural noise and therefore warrant investigation. However, in an ER context, it is necessary to address the question of what mechanisms enable robustness of behaviour sustaining functionality in the presence of neural noise. Studying these mechanisms can inform our understanding of what to look for in natural systems and how to build better artificial ones. ER provides a useful, relatively assumption-free paradigm in which to do this (e.g., agent's dynamics maintaining functionally the same during behaviours).

This paper describes an attempt to understand how neural systems can maintain their function while dealing with neural noise. This is an exploratory piece of work aimed largely at generating hypotheses, and the motivations are conceptual as well as practical. We present results from ER simulations exploring the effects of neural noise on neurocontroller dynamics in order to investigate systemic robustness at the behavioural level. In this context, robustness refers to the ability to maintain performance in the face of perturbations (internal or external) and uncertainty [1]. In order to facilitate understanding of the results and comparative analysis, a simple phototaxis task is chosen. In the next section, the methods and experiments are introduced, and in the final section we examine the consequences of the results and discuss questions that remain open.

2 Methods and Experimental Setup

In order to avoid unnecessary complexity at this initial stage, a minimal approach is deliberately used [4, 6]. The aim is to evaluate the consequences of evolving networks under fixed and variable values of neural noise and to test the obtained solutions in terms of behavioural robustness with the purpose of uncovering the mechanisms at play. A population of simulated agents is evolved to perform phototaxis in normal body and environmental conditions while being disrupted by internal neural and external sensorimotor noise. In each test, one light source is presented every time step for an extended period. Limited random noise is applied locally to the dynamics of each neuron. The level of noise in each neuron ($y_0 \in [-A;A]$) is modelled either as a fixed or a constantly changing activation parameter selected every time step, where A is a fixed value for each experiment (A=0,1,2,3,4). The range (A) of y_0 is a control parameter in our studies.

Agents are modelled as solid circular bodies of radius 5 (arbitrary units) with two diametrically opposed motors that differentially steer the agent with their output (in range [0;1]) and two frontal light sensors positioned with a separation between sensors of 47.75°. The agents' motors can drive backwards and forwards in an unlimited 2-D arena. Agents have a very small mass, so motor output is directly indicates the tangential velocity at the point of the body where the motor is located. The sensors respond to the closeness of a point light source by linearly scaling the distance from the light to each sensor ((clutteredSensorMiss)*(1-(distanceToLightSource/diagonalArena))). Distance and time units are of an arbitrary scale. The model includes sensor shadowing when an agent body occludes light. When not otherwise specified, each evaluation consists of a serial presentation of 6 light sources for a relatively long fixed time (T_{ls}=50 time steps) during an agent's lifetime (T=300 time steps). An agent's task is to approach light sources as they appear. After T_{ls}, the light source is eliminated and another one appears at a random distance ([10;120]) and angle ([0;2π]). The intensity of each source is fixed and equal among them. Sensory inputs are on the range [0;1].

Agents are controlled by a continuous-time recurrent neural network (CTRNN). The dynamics of the network are governed by the following equations:

$$\tau_i \frac{dy_i}{dt} = -(y_i - y_0) + \sum_j^n w_{ji} z_j + I_i \qquad z_j = \sigma(y_j + \theta_j) \tag{1}$$

Using terms derived from an analogy with real neurons, y_i represents the cell potential of the i^{th} neuron (out of N) depending on a decaying time constant τ_i (scaled exponentially in range [1;2+e^2]), θ_i the bias (calculated by center-crossing), z_i the firing rate, w_{ji} the strength of synaptic connection from node j to node i (range [-10;10]), and I_i the degree of sensory perturbation as an incoming current, which is zero for non-input nodes. CTRNN are implemented using center-crossing (see [2, 5]). The center-crossing restriction helps to prevent the incidence of nearly saturated dynamics that would otherwise nullify the effects of neural noise. Time constants τ_i, sensory gain, and synaptic weight w_{ji} are genetically (real-valued) encoded and optimised using a genetic algorithm. The term y_0 represents the level of additive neural noise as described above.

The network topology consists of 2 motor neurons (#0 & #1), 2 input nodes (#2 & #3), and 2 internal neurons (#4 & #5). Full connectivity is used for connecting neurons, but only output neurons include self-connections. Left/right symmetry in synaptic weights is not enforced. We test the role of internal noise in neural systems by randomly biasing the dynamics of the neurons, considering it for each new evaluation in variable configurations of noise only. Even though the addition of the term y_0 could be simply considered as a perturbation on the current input (I_i), this parameter can be also interpreted as influencing the asymptotic behaviour of each neuron including those that receive no sensory input.

A population of 60 individuals is evolved using a steady state, rank-based (selection) genetic algorithm with elitism (50%). Each individual is run for a number of independent evaluations (10), and the fitness of each phenotype is calculated by averaging the fitness obtained in each evaluation. The mutation operator consists of the addition of a small vector displacement selected from a Gaussian distributed value in each gene (with mean 0.0 and standard deviation 1.0). When mutated genes are over or above their range, a non-reflective criterion is applied, generating a new

random value for affected genes. Crossover is not used. The network and other simulation variables are integrated with an Euler time step of 0.1. Fitness is calculated in the following manner: $F = 1\text{-}D_f/D_i$, (D_f: final distance to source; D_i: initial distance to source), and it is determined for each light source (goal) and then averaged for the whole evaluation. F is taken as 0 if $D_f > D_i$ and F is in the range [0;1]. The search algorithm is run for a fixed number of generations (1000), for a fixed number of iterations per generation (200), generally taking a few hundred generations to achieve a high level of average fitness.

3 Results

3.1 Evolutionary Robustness

We evolve populations with different levels of neural noise using independent random seeds, with either a variable or fixed y_0 value during evolution (A=0,1,2,3,4). The genetic algorithm is run 1000 generations. Each data point in Fig. 1-*left* represents the average fitness of the best 10 neurocontrollers for 20 independent experiments (6 randomly placed light sources) for different A values. The evolutionary algorithm generates relatively better results for phenotypes using variable instead of fixed values of y_0 (around 20% higher with A=4 noise range). Error bars in Fig. 1-*left* indicate standard deviation. Performance in the presence of sensorimotor disruptions (e.g., sensor inversion, sensors removal, etc.) is also higher for neurocontrollers evolved with A=4 than with lower values of A. As described in our previous work [2], neurocontrollers evolved with A=0 obtain 84.8% of robustness against neural noise, while controllers evolved with A=1, A=2, A=3, and A=4 obtain 86.9%, 96.4%, 98.5%, and 99.7% respectively. These percentages indicate the average robustness based on the level of performance under disruptions over the level of fitness in the control case.

Figure 1-*right* shows the performance of specific neurocontrollers evolved with A=0 and A=4 and achieving high fitness. These neurocontrollers, named NC#9, NC#10, NC#3, and NC#7 were selected because they present the highest or the lowest level of fitness during tests after evolution with A=0 or A=4. NC#9 and NC#10 were evolved with A=4, while NC#7 and NC#3 with A=0. All of these neurocontrollers demonstrate robustness against neural noise except NC#3. General better performance despite neural noise during tests after evolution are obtained when using A=4 for neurocontrollers evolved with A=4 (means: 0.89 for NC#9 and 0.87 for NC#10) than with neurocontrollers evolved with A=0 (means: 0.84 for NC#7 and 0.2 for NC#3) (Fig. 1-*right*). For neurocontrollers evolved with A=4 and when neural noise is applied in tests after evolution, agents continue to be able to move coherently toward light sources. Neurocontrollers evolved with A=4 also maintain this high performance with lower values of A (e.g., A=0 during tests after evolution, Fig. 1-*right*).

In summary, neurocontrollers evolved with A=4 neural noise remain robust to disruptions (e.g., sensorimotor and structural disruptions) even when noise is removed during tests, suggesting that the variability of neural noise helps evolution find regions of higher general robustness in parameter space. Nevertheless, while most phenotypes maintain high levels of fitness in spite of induced variability of y_0 (e.g., NC #9 with A=4 during tests in Fig. 1-*right*), other phenotypes show low level of

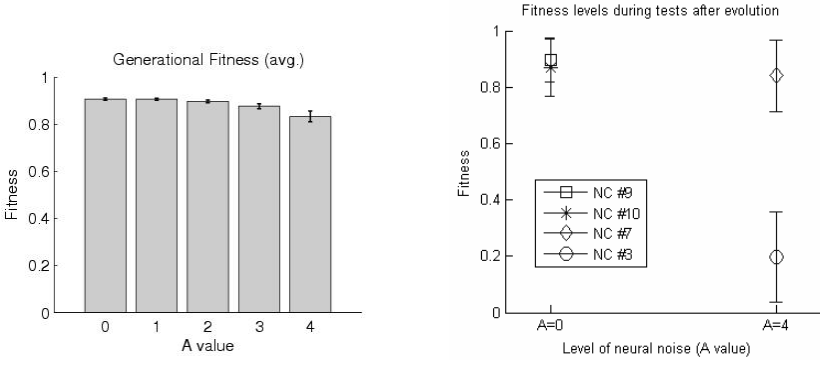

Fig. 1. Evolutionary and phenotypic performance in the presence of variable levels of neural noise (y_0), with randomly positioned 6 lights. *Left*: Fitness values obtained after evolution (avg. of 20 independent experiments), the *x*-axis is the level of *A*. *Right*: Comparison between phenotypes evolved with *A*=0 and *A*=4 (*x*-axis); *y*-axis is the mean fitness reached by neurocontrollers. Each data point represents the average fitness over 100 independent experiments.

fitness when neural noise is introduced (e.g., NC#3 with *A*=4 during experiments after evolution in Fig. 1-*right*). Understanding why some neurocontrollers perform differently than others under the influence of neural noise can provide clues toward the mechanisms that deal with neural noise.

3.2 Analysis of Results

We analyze here neurocontrollers from Fig. 1-*right* in order to discover the properties which allow robust performance in the presence of neural noise. The analysis focuses both on applying neural noise to single neurons and on the general effects of noise in neurocontrollers. Particularly, neurocontrollers NC#9 and NC#3 are studied in detail because they present robustness (i.e., NC#9) or low performance (i.e., NC#3) during tests after evolution (Fig. 1-*left*). While they represent only particular instances, understanding the difference between these neurocontrollers may shed some light on how evolution works differently in the presence or absence of neural noise.

Most of evolved agents successfully acquired the capacity to perform phototaxis despite neural noise (Fig. 1). As expected, the approaching behaviour of agents is based on maintaining light sensory inputs regardless of neural noise effects, i.e. agents regulate their movements without losing signal from light source. For example, agents depicted in Fig. 2 tend to receive sensory stimulation mainly from one side, which is evidenced in agent's trajectories. The effects of inducing neural noise in neurons (i.e., variable levels noise in range [-4; 4] in neuron #5 and deactivating neuron #4's output), indicate that agents can approach to light or lose the light after turning in the 'wrong direction' and thus sensing it again (Fig. 2). Analysing the asymptotic response of neurocontrollers when inputs are forced to be constantly activated or deactivated for each sensor could show in more detail how different behavioural responses are generated in the presence of neural noise.

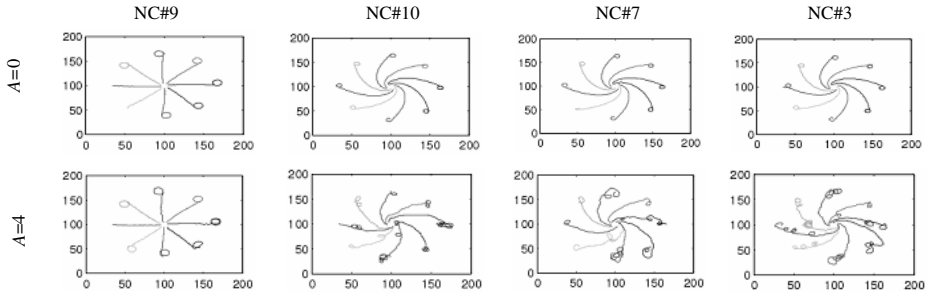

Fig. 2. Examples of behaviours affected by neural noise and neural disruptions. Columns correspond to each neurocontroller indicated at the top of the figure, and each row describes the level of neural noise (A) during tests. Agents start their trajectories to a light source (the small circle at the centre of each figure) from positions separated by 45° on each plot. Top row represents behaviours in normal operation but disabling neuron #4's activity; bottom row represents agent's behaviour after including noise in neuron #5 (neural noise y_0 in [-4; 4]) for each situation in the top row.

Because the noise term y_0 in Eq. 1 is an additive one, this means that nullclines in phase space will tend to be relatively displaced to each other for different values of y_0, but not warped or changed (as would be expected if noise were added to a weight term) [7]. Table 1 shows the asymptotic responses of NC#9 and NC#3 determined by the difference between left and right motor neuron activities (neuron #0's and neuron #1's outputs activations). This difference indicates the action that neurocontrollers generate after inducing different fixed values of y_0 (-4, 0, or 4) in neurons #4 and #5 (y_0 = 0 for the remaining neurons). To achieve these results, we forced inputs to be constantly activated (sensor input = 1) or deactivated (sensor input = 0). A positive difference means that agent turns left, while a negative difference that agent turns right in the asymptotic state.

The asymptotic response of neurocontrollers for each sensor and fixed y_0 configuration helps to differentiate the neurocontrollers' strategies for approaching light. Let us first consider NC#9 with no lights (sensor inputs = 0) and without noise in any neuron. In this situation, the agent moves slightly turning right (motor response -0.094) (indicated for (0;0) levels of noise in neurons #4 and #5 in Table 1). In the normal sensing situation, this motion will cause that right sensor to come into contact with the light. Then, the right sensor will start to receive more input, causing the agent to turn left slightly when approaching the light (motor response 0.199)(indicated for (0;1) levels of noise in Table 1). This will produce a decrease of the right sensing input up to a non sensing situation as in the starting condition, because sensor loses contact with the light. After approaching the light, the agent will generate a new movement to right (similar to that described before) generating a new increase of the right input sensing.

This right sensor strategy for approaching light in NC#9 we also observe during normal sensing when the fixed values of y_0 in the interneurons are (0;-4) and (0;4). From an asymptotic perspective, this strategy will cause the right sensor activity to be increased or decreased depending on robot's approach to light and in left and right

Table 1. Asymptotic responses of NC#9's and NC#3's turning behaviours. Table shows the difference between left and right motor activities (motor response), when sensory input nodes #2 and #3 are forced to be constantly activated (sensor input = 1) or deactivated (sensor input = 0) (indicated for each response data); y_0 #4 and y_0 #5 represent fixed levels of neural perturbation (y_0=-4, 0, or 4) in neurons #4 and #5, respectively.

#2	#3	y_0 #4	y_0 #5	NC#9	NC#3
				motor response	
0	0	-4	-4	0.103	0.200
0	0	-4	0	0.170	0.102
0	0	-4	4	0.211	-0.019
0	0	0	-4	-0.114	0.205
0	0	0	0	-0.094	0.108
0	0	0	4	-0.041	-0.012
0	0	4	-4	-0.116	0.210
0	0	4	0	-0.100	0.114
0	0	4	4	-0.075	-0.005
1	1	-4	-4	0.214	0.065
1	1	-4	0	0.233	-0.060
1	1	-4	4	0.242	-0.175
1	1	0	-4	0.214	0.071
1	1	0	0	0.233	-0.053
1	1	0	4	0.242	-0.168
1	1	4	-4	0.214	0.077
1	1	4	0	0.233	-0.046
1	1	4	4	0.242	-0.162

#2	#3	y_0 #4	y_0 #5	NC#9	NC#3
				motor response	
0	1	-4	-4	0.172	0.249
0	1	-4	0	0.202	0.171
0	1	-4	4	0.226	0.064
0	1	0	-4	0.160	0.253
0	1	0	0	0.199	0.176
0	1	0	4	0.225	0.070
0	1	4	-4	-0.099	0.258
0	1	4	0	-0.043	0.182
0	1	4	4	0.128	0.076
1	0	-4	-4	0.200	-0.018
1	0	-4	0	0.225	-0.139
1	0	-4	4	0.238	-0.236
1	0	0	-4	0.200	-0.011
1	0	0	0	0.225	-0.133
1	0	0	4	0.238	-0.230
1	0	4	-4	0.192	-0.004
1	0	4	0	0.223	-0.126
1	0	4	4	0.237	-0.225

responses that depends on right input sensing. However, values of y_0 are changed to (4;0) or (4;4) during normal sensing, the approaching strategy of NC#9 changes toward a left sensor configuration in order to develop phototactic behaviour. In this case, starting from the situation that the agent sees no light, eventually the right sensor will sense first more light producing that the agent moves to right (motor response -0.043) or moves left slightly (motor response 0.128) for (4;0) and (4;4) values for y_0, respectively (Table 1). When the levels of y_0 are (4;0), the agent will move producing a decrease in right sensor input due its approaching angle to the light. The left sensor input will increase generating a left movement (motor response 0.223) instead of a right one as explained before for the right sensor configuration. A similar strategy using the left sensor is observed with (4;4) levels of y_0, but it takes more time for generating such control because the robot turns left more slightly when the right sensor input increases during the approaching behaviour. Therefore, the strategy of NC#9 for maintaining phototaxis behaviour with left of right sensors is based on the combination between motor responses and the approaching behaviour that activate eventually right or left sensors as described before.

As described above, the agent falls into a left or right sensor configuration depending on the value of y_0 in the interneurons. The phototactic behaviour is not demonstrated in other asymptotic configuration of noise during normal sensing approach. For example, adding y_0=-4 to neuron #4, regardless of the level of noise in neuron #5, produces that the agent cannot perform phototaxis because it loses the capacity to turn right and also the ability to maintain right or left sensory inputs (see Table 1). Similarly, inducing y_0=-4 in neuron #5, and y_0=4 in neuron #4 also produces

a non-phototactic behaviour, because left movements are not reached correctly during normal sensing approach.

Let us now look at NC#3 defining its input sensing to 0 (meaning that agent sees no lights) and without neural noise. This agent moves turning left (motor response 0.108 indicated for (0;0) levels of noise in Table 1). Eventually, the left sensor will activate first abruptly in normal sensing situations, producing a right movement (motor response -0.133) (see Table 1) that will generate a decrease in left sensor input when agent approaches light. This use of the left sensor is also observed during normal sensing situations when the level of noise in internal neurons are (-4;0) and (4,0). In these cases, phototaxis is performed but not in the noisy configurations. This is the case mainly because the agent loses its capacity for turning left using its left sensor. For example, in (0;-4) y_0 configuration, the agent turns left (motor response 0.205) when it senses no lights, which eventually will produce an increase of left sensing and a slight right movement (motor response -0.011). However, this right movement is not enough to maintain left sensory input while the agent approaches light. This means that eventually the right sensor should become activated, producing a left movement (motor response 0.253). Thus, this also produces a new non-sensing situation with both sensors. According to the asymptotic response of NC#3 in Table 1, the agent sometimes turns in different directions with different levels of noise mainly when both sensors are activated or when no sensing is produced. Thus, the agent will in the end receive inputs from 'the wrong side' causing the agent turns to the wrong direction and cannot perform phototaxis as observed in successful situations.

The asymptotic analysis only provides an indication of how the attractor landscape is affected for different configurations of sensory input and y_0. Observations of the actual transient behaviour for the different configurations indicate that NC#9 performs well in 5 out of 9 cases ((0;0), (0;-4), (0;4), (4;0), and (4;4) values of y_0 in neurons #4 and #5) and NC#3 on 3 out of 9 ((0;0),(-4;0), and (4;0) values of y_0 in neurons #4 and #5). Moreover, NC#9 is able to deploy at least two different behavioural strategies while only one has been observed for NC#3. By taking the values of y_0 investigated as rough representatives of the whole space of variation for y_0, we can conclude that for most levels of noise (but not all) NC#9 will perform phototaxis in a combination of two strategies, but that this is not the case for NC#3. NC#9 is therefore sometimes undergoing bifurcations, but they are most of the time (roughly around two thirds of the time) functional allowing it still to perform phototaxis. Nevertheless, about one third of the time these bifurcations are non-functional as described before. For example, NC#9 is not significatively affected by noise when sensors are simultaneously activated, but it generates wrong long-term responses (e.g., turning left instead of turning right) more frequently than in NC#3 when input sensing are deactivated simultaneously depending on their levels of noise.

The proposed hypothesis is that *those controllers evolved with noise are not undergoing long-term dysfunctional bifurcations because of noise*. In this case, evolution finds networks that operate in regions of phase space for which moderate displacement of the nullclines does not significatively affect the functionality of the system. In fact, because NC#9 has two different strategies for approaching light in the presence of noise, this implies that noise can generate bifurcations but they happen to be also functional, meaning that perturbations in the noise range do not cause qualitative changes to system functionality. Evolution is therefore not only searching

for regions of space where bifurcations induced by noise are unlikely to happen, but also for regions where "neighbouring" bifurcations are also functional. As long as the balance of functionality over the possible bifurcations induced by neural noise remains positive, the controller will be able to cope with noise and will moreover be likely to cope with other perturbations (assuming these perturbations induce similar structural changes in the dynamical landscape).

We also observe that noise produces dysfunctional bifurcations in NC#3's dynamics more frequently than in NC#9. NC#9 demonstrates sharp changes in the behaviour of agents, however, indicating that there are still three possibilities to explain how fitness is maintained high despite neural noise. These are: (*I*) those particular bifurcations do not largely affect the transient dynamics of the network; (*II*) all (or most) bifurcations produce different forms of instantaneous phototaxis (they are mostly functional in themselves); (*III*) no bifurcations are produced. These two first possibilities also imply two explanations: (1) negative (non-functional) bifurcation may indeed happen in the range of the noise parameter, but they may be short lived while the agent performs phototaxis; (2) negative (non-functional) bifurcation may occur for significant amounts of time, creating bifurcations that lead asymptotically to non-phototaxis.

In the first case (1), bifurcation during a transient seems to be related with NC#9 because robustness against noise is functionally maintained despite increasing neural noise in most situations. The agent still performs phototaxis because it is held in a transient between attractors that are functional (see [8]). By contrast, NC#3 probably corresponds to case (2) with noise leading to the loss of performance when noise is increased. We have not ruled out transient effects for the situation described for NC#9, however. The pattern of sensor activation and of neural noise may induce bifurcations that are asymptotically non-functional (would not produce phototaxis in the long term) but their change keeps the neural and agent state in a functional transient when A=0 during test after evolution.

4 Conclusion

Experiments with neural noise have been presented here from an evolutionary and sensorimotor perspective. The simulation model in itself is minimal but results suggest that, at least in the experimental situations, evolution relies on mechanisms that maintain functional dynamics in transients, as shown for NC#9. Results also indicate that neural systems lose sensitivity to noise when systems are evolved with high levels of neural noise.

From an evolutionary perspective, the interesting lesson is that *neural noise in evolution seems to put pressure for selecting neural systems that are resistant to the effects of bifurcation*, and so their robustness lies in having a dynamic landscape that remains, in the overall balance, functionally the same during behaviour. This is evidenced by the noise robustness of NC#9 and the noise sensitivity of NC#3. The relationship between evolutionary mechanisms selected under noise processes has been minimally investigated in the simulation studies so far. In fact, mechanisms where noise is irrelevant could vary from the simple attractors' view where noise utility is removed because of convergence to stable system dynamics. Our results

suggest that the evolutionary process in the presence of neural noise – following the logic of Jakobi's minimal simulations – is finding robust neural dynamics. However, this robustness has a structure. It is a combination of locating the neurocontrollers in regions of parameter space where bifurcations are unlikely to occur and simultaneously where, if and when bifurcations occur, they remain in balance functional. This finding suggests that robustness to other sensorimotor perturbations may be a by-product of locating such regions of parameter space. If this is so, a prediction from this result is that a similar evolutionary process under parametrical uncertainty, but applied to non-additive parameters (such as weight values) may result, if successful, in even higher levels of robustness to sensorimotor perturbations.

In our results, and in accordance to the above explanation, adaptive performance was also observed when noise was removed, indicating that noise is not actively maintaining functionality in the analysed neurocontrollers. Nevertheless, we do not discard the idea that evolution may find solutions for which noise is advantageous, in which case our explanation will need to be appropriately modified. These questions will be further investigated in future work, including comparisons using embodied agents to perform different tasks.

Acknowledgments. Many thanks to Greg Studer and Matthew Egbert for comments of this paper, anonymous reviewers, 'The Peter Carpenter CCNR DPhil Award', and CONICET, Argentina, recognition (Type-I-Res.38-08/01/2004). This research was partially supported by the Programme Alβan, the European Union Programme of High Level Scholarships for Latin America (No. E05D059829AR).

References

1. Kitano, H.: Toward a theory of biological robustness. Mol. Syst. Biol. 3, 137 (2007)
2. Fernandez-Leon, J.A., Di Paolo, E.A.: Neural uncertainty and sensorimotor robustness. In: Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A. (eds.) ECAL 2007. LNCS (LNAI), vol. 4648, pp. 786–795. Springer, Heidelberg (2007)
3. Bays, P., Wolpert, D.: Computational principles of sensorimotor control that minimize uncertainty and variability. The Journal of Physiology 578(2), 387–396 (2007)
4. Jakobi, N.: Evolutionary robotics and the radical envelope of noise hypothesis. Journal of Adaptive Behaviour 6(2), 325–368 (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. Di Paolo, E.A., Harvey, I.: Decisions and noise: the scope of evolutionary synthesis and dynamical analysis. Adaptive Behavior 11(4), 284–288 (2004)
7. Beer, R.: On the dynamics of small continuous-time recurrent neural networks. Adaptive Behavior 3(4), 469–509 (1995)
8. Iizuka, H., Di Paolo, E.A.: Minimal agency detection of embodied agents. In: Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A. (eds.) ECAL 2007. LNCS (LNAI), vol. 4648, pp. 485–494. Springer, Heidelberg (2007)