

Optimal noise in spiking neural networks for the detection of chemicals by simulated agents

Nicolas Oros, Volker Steuber, Neil Davey, Lola Cañamero, Rod Adams

Science and Technology Research Institute
University of Hertfordshire
AL10 9AB
United Kingdom

{N.Oros, V.Steuber, N.Davey, L.Canamero, R.G.Adams}@herts.ac.uk

Abstract

We created a spiking neural controller for an agent that could use two different types of information encoding strategies depending on the level of chemical concentration present in the environment. The first goal of this research was to create a simulated agent that could react and stay within a region where there were two different overlapping chemicals having uniform concentrations. The agent was controlled by a spiking neural network that encoded sensory information using temporal coincidence of incoming spikes when the level of chemical concentration was low, and as firing rates at high level of concentration. With this architecture, we could study synchronization of firing in a simple manner and see its effect on the agent's behaviour. The next experiment we did was to use a more realistic model by having an environment composed of concentration gradients and by adding input current noise to all neurons. We used a realistic model of diffusive noise and showed that it could improve the agent's behaviour if used within a certain range. Therefore, an agent with neuronal noise was better able to stay within the chemical concentration than an agent without.

Introduction

Animals are able to detect and react to chemicals (odours, pheromones...) present in the environment. The key sense to detect these chemical cues is smell rather than taste (Wyatt, 2003). Almost all animals have a similar olfactory system including olfactory sensory neurons (OSN) that are exposed to the outside world and linked directly to the brain. Pheromones and other odour molecules present in the environment are converted into signals in the brain by first binding to the olfactory receptor protein situated in the cell membrane of the OSN. Spikes are then sent down the axon of the OSN (Kandel et al., 2000). A chemical blend is composed of many molecules that can be detected with tuned odour receptors and therefore, activates a large range of olfactory sensory neurons. Odours are coded by which neurons emit spikes and also by the firing patterns of those neurons sending spikes to others during and after the stimulus. In many vertebrates and insects, oscillations of the neural activity have been recorded in the olfactory systems (Wyatt, 2003). Therefore, the synchronization of firing between different sensory neurons

seems to be very important for odour perception and interpretation. The firing rate and the number of sensory neurons are also important in odour recognition when stronger stimuli increase the frequency of firing of individual sensory neurons but also stimulate a larger number of them.

Different studies have been done on the perception of simulated chemicals using artificial neural networks where neural synchronization occurs (Brody & Hopfield, 2003; Hopfield, 1999; Hoshino et al., 1998) and also using robots (Kanzaki et al., 2005; Kuwana & Shimoyama, 1998; Payton et al., 2001; Pyk et al., 2006; Webb, 1998). We were interested in studying the perception and the behaviour of an agent in response to changes of its environment. The primary research question is how two encoding strategies can be used to integrate sensory information in order to control a simulated agent. To the best of our knowledge, no neural architecture, controlling a simulated agent, has been created that encodes the sensory information onto both the firing rate and the synchronization of firing (temporal coincidence of incoming spikes) depending on the environment. As the interaction between the two encoding strategies is complex, we decided to create a simple architecture using a spiking neural network. This model could encode the sensory information onto both the firing rate and the synchronization of firing depending on the environment. The neural network controlled the agent by encoding the sensory information onto temporal coincidences in a low concentration environment, and firing rates at high concentration.

It is well known that real neuronal systems contain noise (Kandel et al., 2000) which may improve the brain's ability to process information, a phenomenon also called stochastic resonance (Hänggi, 2002; Mori & Kai, 2002; Moss et al., 2004; Wiesenfeld & Moss, 1995). Researchers in robotics and artificial life have already implemented simple models of neural noise (Di Paolo, 2003; Florian, 2006; Jacobi et al., 1995). Here we study the effect of a more realistic noise model based on a diffusive OU (Ornstein-Uhlenbeck) process (Uhlenbeck & Ornstein, 1930). We added this noise in the neural network and studied its effect on the behaviour of the agent. Our results suggest a potential function for noise in real

biological systems, and highlight that features of biological systems can be used to construct better agents.

Environment

We created a simulation of a continuous world including an agent and a maximum of two chemicals. We decided to use a simple model of chemicals that are not diffused and evaporated but with concentrations that can be calculated directly at any given point. Our agent was equipped with two antennae and a differential steering system using two wheels. The two antennae were separated widely enough to detect the presence of the chemical concentration (Fig. 1). The left and right wheels were situated on the sides of the agents. To control the agent, we had to decide which neurons' model to use in order to study firing synchronization of the sensors.

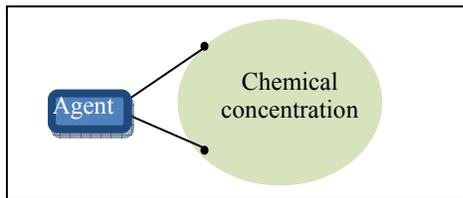


Figure 1. An agent equipped with two wheels and two antennae used to detect chemicals.

Neural Network

There are three main ways to encode the intensity of sensory information into spiking neurons based on biological evidences (Floreano & Mattiussi, 2001; Florian, 2003; Gerstner & Kistler, 2002; Izhikevich, 2003, 2004; Koch, 1999). The most commonly used method consists of mapping the stimulus intensity to the firing rate of the neuron (firing rate encoding). Another method encodes the intensity of the stimulation into the number of spikes sent by different neurons arriving at a pre-synaptic neuron at the same time (firing synchronization or temporal coincidence encoding). The last main encoding scheme maps the strength of the stimulation in the firing delay of the neuron (delay encoding). As we saw earlier, spatial configuration is an important feature in odour recognition of neurons as is the synchronization of firing between neurons (Kandel et al., 2000; Laurent et al., 1996; Wyatt, 2003). J. Hopfield and C. Brody (Brody & Hopfield, 2003; Hopfield, 1999) created simple neural networks using spiking neurons to simulate an olfactory process. In their system, the recognition of an odour was signalled by spike synchronization in artificial glomeruli. In our system, the neural network was supposed to detect the blend of two different chemicals and modify the agent's behaviour. We used a model of neural network that allowed us to study synchronization of firing in a simple manner. The neural network could control the agent by encoding the sensory information onto temporal coincidences in a low concentration environment, and firing rates at high concentration.

Models of Spiking Neurons

It is well known that compared to the complex and computationally slow Hodgkin and Huxley model, simple spiking models like integrate-and-fire neurons can run quickly enough and have a more realistic behaviour than firing rate ones (Floreano & Mattiussi, 2001; Florian, 2003; Gerstner & Kistler, 2002; Izhikevich, 2003, 2004; Koch, 1999). This is why more and more researchers are implementing spiking neurons in robots and simulated agents. Therefore, we decided to use a simple model of a spiking neuron. Our model is based on a leaky-integrator model which includes synaptic integration and conduction delays. The idea is that a spike sent by a neuron will take some time to arrive at another neuron. This time delay depends on the distance between the sender and the receiver. All the spikes arriving at a neuron are summed to calculate the neuron's input current density (in Amperes per Farad) and membrane potential (in Volts) after every time step ($\Delta t = 0.1ms$). Once the membrane potential reaches a certain threshold θ , the neuron will fire and then will be set to 0 for a certain time (refractory period). During this time, the neuron cannot fire another spike even if it is highly stimulated.

Many real neurons' membrane potential is around -70mV during resting state. When a neuron fires, its membrane potential will increase rapidly to about 30mV, so the height of a typical spike is approximately 100mV (Kandel et al., 2000). We set the resting potential to 0 and the potential of a spike to 100mV. It is reasonable to set the neuron's threshold at 20mV, the refractory period to 3ms and the membrane time constant τ_m to 50ms (Kandel et al., 2000). We also decided to set a synaptic time constant τ_s to 2ms: a spike that arrives at a synapse triggers a current given by:

$$I_j(t) = \left(\frac{t - (t_{spike} + delay)}{\tau_s} \right) \exp \left(1 - \frac{t - (t_{spike} + delay)}{\tau_s} \right) \quad (1)$$

where $I_j(t)$ is the synaptic input current, t_{spike} corresponds to the time a spike has been sent to the neuron, $delay$ is the time delay in seconds before the spike arrives to the neuron ($delay = coeff_delay * distance$) with $coeff_delay = 5 \cdot 10^{-5}$.

The change of membrane potential is given by:

$$\frac{dV}{dt} = - \left(\frac{V}{\tau_m} \right) + \sum (I_j W_j) \quad (2)$$

where V is the membrane potential, τ_m is the membrane time constant and W_j the synaptic weight.

Sensory Neurons

We created a model of a spiking sensory neuron in which the chemical concentration is processed so that a quasi-linear relationship between the concentration and the firing rate of the sensor is produced (Oros et al., 2008). Such relationships exist in biological systems. For example in humans, the relationship between the frequency of firing and pressure on the skin is linear (Kandel et al., 2000). We used a two step process where two biologically realistic non-linear mappings between sensory information and input current and between

input current and firing rate results in a linear relationship. Researchers in robotics and artificial life use a linear direct mapping between the sensory information and the firing rate (Di Paolo, 2002, 2003; Florian, 2006). The sensory neurons used in our model are able to encode the stimulus intensity, measured at the tip of the antenna, into sensory input current using a biologically plausible sigmoid function (Oros et al., 2008). This current is injected to the sensor's membrane potential that increases, making the sensor fire into appropriate firing rates. Therefore, the sensory neurons encode the concentration value onto the appropriate firing rate. The sensors were configured in order to distinguish a large range of concentrations between 1 and 300. Over 300, they were saturating.

Motor Neurons

We decided that, in order to move, the agent should be driven by two wheels each controlled by two motor neurons: one to go forward, one to go backward. We created sensors able to detect a chemical gradient. But an agent equipped with such sensors will not move without any stimulus. So we decided for simplicity that an agent should always move forward in the absence of any external input. We performed this by adding a small baseline input current (0.5 A/F) in the motor neurons responsible to go forward. The final velocity of the wheels was calculated by subtracting the firing rate of the motor neurons, responsible for moving the agent forward and backward, running over a certain period of time. The agent was moved by calculating the velocity every 10ms.

Temporal Coincidence

We used the agent and world described above. The environment contained either one or two chemicals denoted by A or B. In this experiment, each chemical source had a circular shape and the same fixed value all over its surface. One agent, placed in the world, was controlled by a simple spiking neural network implementing the neurons described in the previous section.

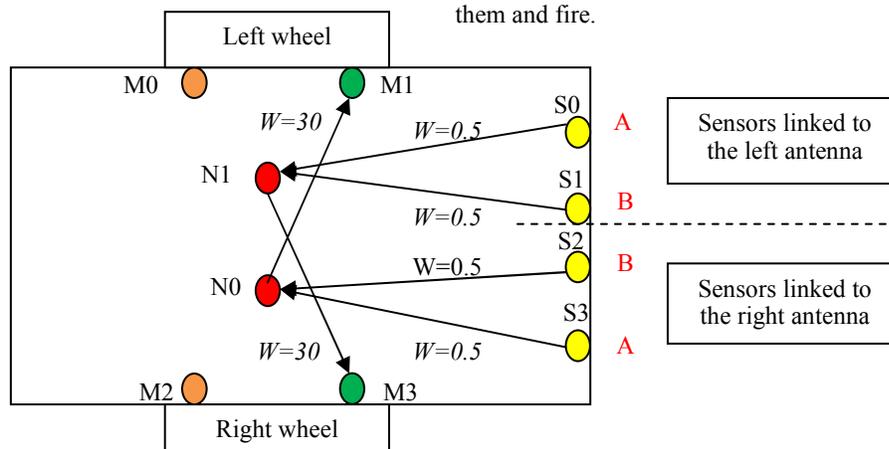


Figure 2. Agent's neural controller. The sensors S0 and S3 detect the chemical A and the sensors S1 and S2 detect the chemical B. The sensory axons' lengths are all similar (delays = 2.5ms). The motor neurons M1 and M3 are responsible to move the agent forward. The threshold of the neurons (N0 and N1) was set to 4.6 mV. W is the synaptic weight.

The neural controller was based on a Braitenberg vehicle (anger behaviour) (Braitenberg, 1984) where an agent moves faster toward a stimulus when it detects it (Fig. 2).

Our hypothesis was that by using this architecture, the sensory neurons needed to encode the sensory information onto the firing rates, and also onto temporal coincidences between spikes sent by sensors. To verify this hypothesis, we performed three series of tests to study the effect of the starting positions, the sensory delays and the value of the concentrations on the agent's behaviour.

Experiment I

The first test was to study the effect of the agent's starting position on its behaviour. Both concentration values for the chemicals A and B were set to be low. In all the experiments described in this paper, the concentration range was from 1 to 300. In this instance, A and B concentrations were set to 1 or 2. We tried ten different starting positions and five different settings for the environment: with one chemical A, one chemical B, and finally one concentration of the chemical A overlapping with one concentration of the chemical B. Each run lasted 600 seconds and the neural network was updated every 0.1ms (so the run lasted 6,000,000 time steps). Every 10ms, the agent was moved and the sensory inputs updated.

In these experiments, the agent could detect double concentrations of one chemical (A or B) but did not react to it. However, the agent was able to react only to the blend of both chemicals A and B, where it stayed inside the overlapping concentrations. We recorded the current density and membrane potential of the neuron N0 during a small interval of time when the agent was inside the blend of chemicals A and B (Fig. 3, top). The input current of the neuron N0 was increasing when spikes coming from both S2 and S3 arrived at the same time. Then, the membrane potential also increased and reached the threshold θ (0.0046 Volts) making the neuron N0 fire. The potential was then set to 0 during the refractory period. As the sensors were synchronized and the delay between them and the neurons were the same, the spikes arrived at the same time to the neuron allowing it to detect them and fire.

Experiment II

The second experiment was to test our hypothesis by modifying the sensory response delays to verify that our architecture necessarily needed to encode the sensory information onto temporal coincidence. We changed the delays by modifying the position of the sensors therefore modifying the length of their axons linked to the neurons. We only changed the delays of the sensors detecting the chemical B (S1 and S2).

We used one of the Experiment I's setups where the agent was staying in the chemical blend of the chemicals A and B having a concentration of 1 each. We tried different values of delays (from 1ms to 50ms) and we noticed that a small change (up to 7.5ms) did not modify the agent's behaviour. But a further change in the delays (from 7.5ms) made the agent unable to react to the blend of chemicals A and B so it could not stay inside the concentrations.

As in the Experiment I, we recorded the current density and membrane potential of the neuron N0 during 0.5s when the agent was inside the chemical blend.

In Figure 3 (bottom), we can see that the current of the neuron N0 increases when a spike coming from both S2 and S3 arrive but as the delay has been changed, the spikes do not arrive at the same time so the current is lower than in Experiment I. Therefore, the neuron's potential increases but never reaches the threshold so the neuron does not fire (Fig. 3, bottom).

Experiment III

In order to investigate the use of firing rate encoding, we used only one concentration of either A or B and increased it. When the concentration was augmented from 1 to above 50, the agent was then able to react to it. Therefore, the neural network showed much more sensitivity to two chemicals than to one. We also realized when using two overlapping chemicals A and B, as the concentration value increased, modifying the delays had a minor effect and the agent was still able to react to the chemicals. The firing rates were increasing too so the agent was moving faster. In these experiments, the temporal coincidence encoding was not necessary. The sensory information was encoded onto the firing rates of the sensors.

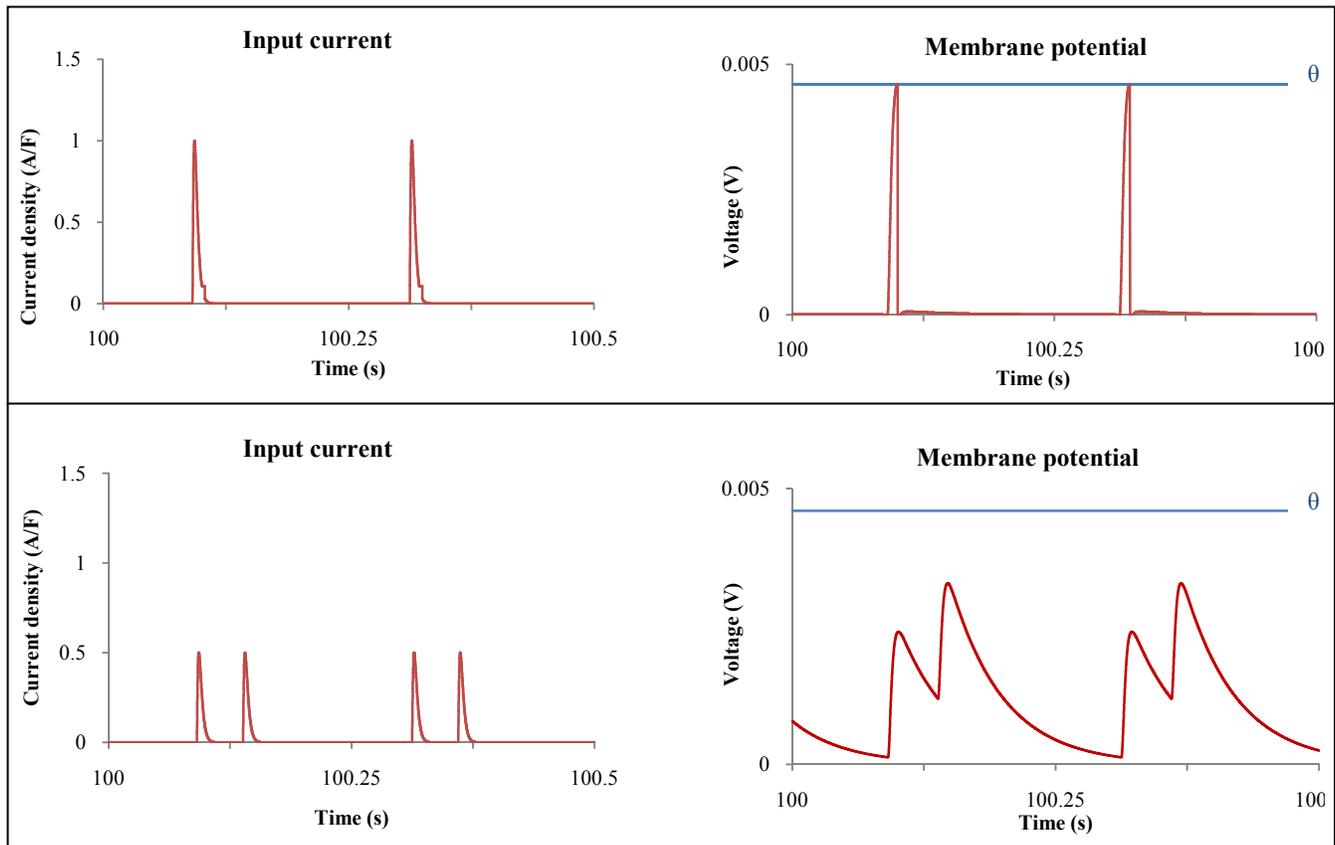


Figure 3. Current density (in Amperes per Farad) and membrane potential (in Volts) of the neuron N0 recorded between 100s and 100.5s. On the top panel (Experiment I), the spikes sent by the sensors arrived at the same time increasing the current density to 1 A/F. The membrane potential was then increased and reached the threshold making the neuron N0 fire. On the bottom panel (Experiment II), the spikes sent by the sensors were not coincident as the delays between the sensors (S1 and S2) and the neurons (N0 and N1) were changed (to 50ms in this case). Therefore the current was never above 0.5 A/F so the membrane potential could not reach the threshold to make the neuron N0 fire.

Diffusive noise

In the previous experiments, we presented a simple neural architecture where temporal coincidence and firing rate encoding strategies were both important mechanisms used in different environmental settings. In a low concentration setting, synchronization of spikes sent by the sensors was essential to allow the agent to detect the blend of two chemicals. We changed the sensory delays and noticed that the agent was then not able to react to the chemicals anymore. In a high concentration setting, the temporal coincidence between the firing of the sensors was not a necessary condition and the agent was able to stay inside the chemical concentration using just a firing rate encoding strategy. Interestingly, the model showed much more sensitivity to the presence of two chemicals than a single chemical. To this point, we have used uniform concentrations to simplify the study of the different encoding strategies. However, this model of chemical concentration was not realistic, so we decided to use an environment comprising two non uniform chemical concentration gradients. We tested our architecture in the new environment and noticed that the agent moved outside the concentration when its trajectory was along the direction of the gradient since both of its antennae were instantaneously outside the chemical concentrations. For this reason, we decided to add noise to the neural network.

We used a realistic model of noise in the form of a diffusive OU current (Uhlenbeck & Ornstein, 1930). This form of colored noise characterizes the subthreshold voltage fluctuations in real neuronal membranes (Rudolph & Destexhe, 2003). We added this noise to the total current calculated in Equation (2) in each neuron. The noise is described by:

$$\frac{dI(t)}{dt} = -\frac{1}{\tau_I}(I(t) - I_0) + \sqrt{\frac{2\sigma^2}{\tau_I}}\xi(t) \quad (3)$$

where τ_I denotes the current noise time constant (2ms in our case), I_0 is the mean synaptic current (0 in our case), σ is the noise diffusion coefficient and $\xi(t)$ is a white Gaussian noise (with mean = 0 and standard deviation = 1).

We performed different series of tests to find appropriate level of noise, by modifying σ , in order to have an agent that stays in the gradient chemical blend. We placed the agent at three different positions (Fig. 6) and tried eight different levels of noise (Fig. 4 and 5). For each level, we performed 100 runs per position. Each run lasted 300s and we recorded the fitness of an agent during the last 100s. The fitness function was very simple and consisted of the sum of the distance between the agent and the centre of the concentrations measured every time the agent moved. The maximum value of both concentrations was set to 25.

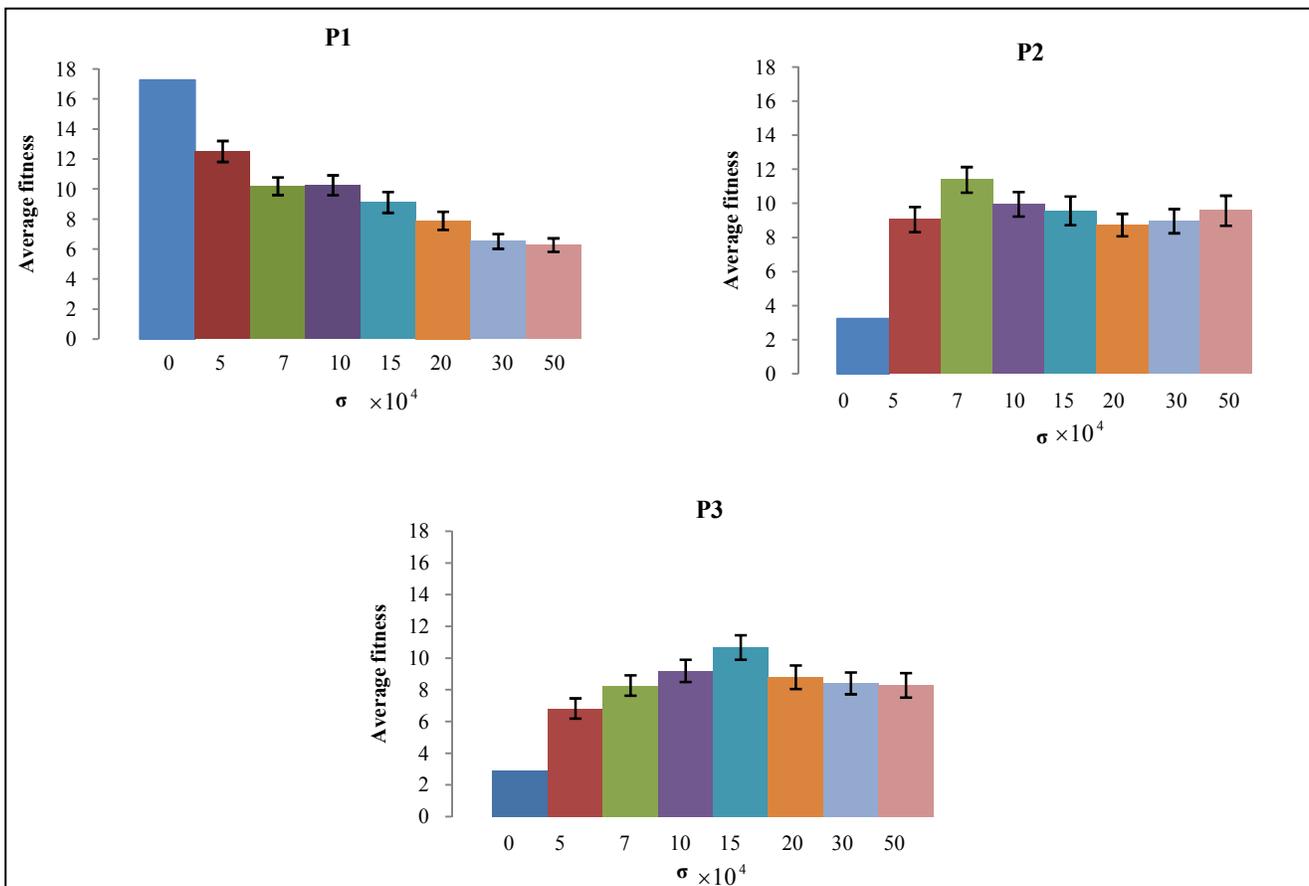


Figure 4. Mean fitness values recorded during 100s for an agent starting at the positions P1, P2 and P3 using different levels of noise ($\sigma \times 10^4$). The error bars represent standard errors.

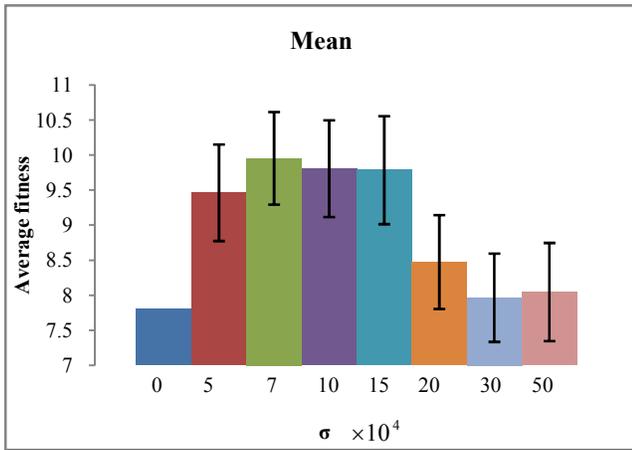


Figure 5. Mean of the fitness values displayed in Figure 4. ($\sigma \times 10^4$).

By looking at Figures 4 and 5, we can see that when the agent was starting from P2 or P3, an appropriate level of noise allowed it to stay within the concentration having a higher fitness than an agent without neural noise. We also note that the level of noise needed to be within a certain range as a low value did not improve the agent's behaviour and a high value disturbed it. We noticed as well that the agent was more sensitive to noise in low concentration areas than in high concentration areas.

Conclusion

We first presented in this paper a simple neural architecture where temporal coincidence and firing rate encoding strategies were both important mechanisms used in different environmental settings. In a low concentration setting, synchronization of spikes sent by the sensors was essential to allow the agent to detect the blend of two chemicals. We changed the sensory delays and noticed that the agent was then not able to react to the chemicals anymore. In a high concentration setting, the temporal coincidence between sensors firing was not a necessary condition and the agent was able to stay inside the chemical concentration using just the firing rate encoding strategy. Interestingly, the model showed much more sensitivity to the presence of two chemicals than a single chemical. Our results showed that a spiking neural network could be used to control an agent and could encode external stimuli in more than one way. The second study was on the effect of noise on the agent's behaviour using the same neural architecture. We used a more complex environment using chemical gradients and a realistic model of neural noise. We found that the overall fitness of the agent was better when a certain amount of noise was added in the neural network. Our results suggest that a realistic model of noise can improve an agent's behaviour. This is further evidence that adding biologically realistic features can be beneficial for certain engineering tasks, and suggests a potential function of noise in real biological systems. The effect of biologically realistic noise should be an interesting topic of research in other artificial life scenarios.

Our future work will be to see if we can evolve such architecture using a developmental model (evolving the number of neurons and their connections, the synaptic weights, and delays of the neural network).

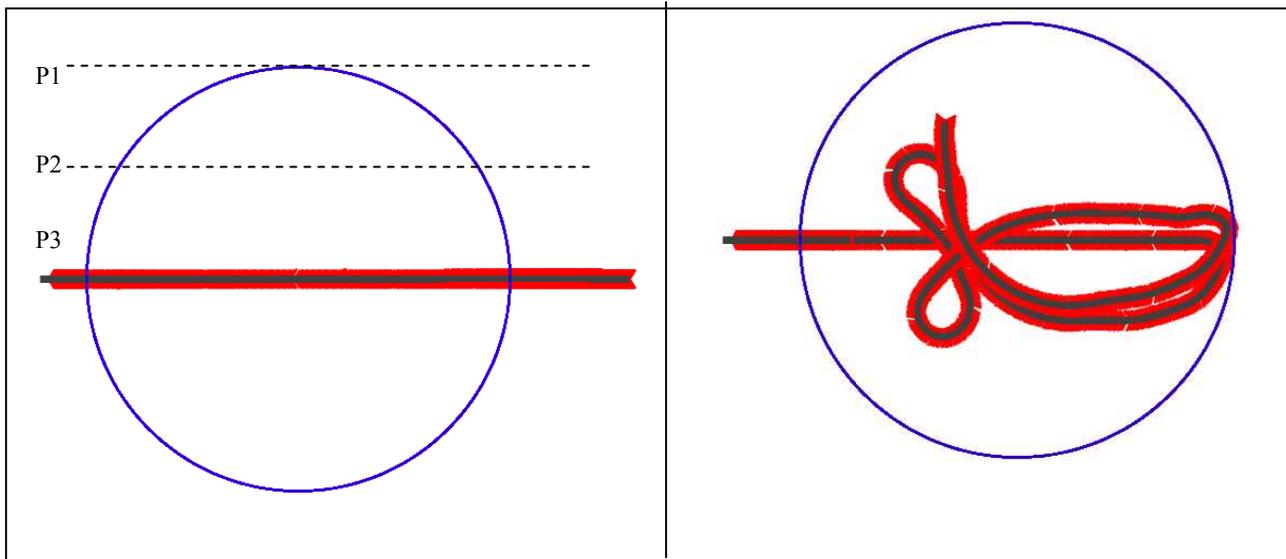


Figure 6. Left panel: path of an agent moving across the blend of chemicals A and B. The agent's neural controller doesn't have any noise so the agent goes straight as both of its antennae arrived at the same time outside the concentration. Right panel: path of an agent running over 300s. The agent's neural controller has noise so the agent does not go exactly in a straight line and therefore, can react to the absence of the chemical concentration to stay inside.

References

- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*: MIT Press, Cambridge, MA.
- Brody, C. D., & Hopfield, J. J. (2003). Simple Networks for Spike-Timing-Based Computation, with Application to Olfactory Processing. *Neuron*, 37, 843-852.
- Di Paolo, E. A. (2002). *Evolving spike-timing dependent plasticity for robot control*. Paper presented at the EPSRC/BBSRC International Workshop: Biologically-inspired Robotics, The Legacy of W. Grey Walter, WGW'2002, Bristol.
- Di Paolo, E. A. (2003). Spike-Timing Dependent Plasticity for Evolved Robots. *Adaptive Behavior*, 10, 243-263.
- Floreano, D., & Mattiussi, C. (2001). Evolution of Spiking Neural Controllers for Autonomous Vision-Based Robots. *Proceedings of the international Symposium on Evolutionary Robotics From intelligent Robotics To Artificial Life*, 2217, T. Gomi, Ed. *Lecture Notes In Computer Science Springer-Verlag, London*, 38-61.
- Florian, R. V. (2003). Biologically inspired neural networks for the control of embodied agents. *Technical report Coneural-03-03 Version 1.0*.
- Florian, R. V. (2006). *Spiking neural controllers for pushing objects around*. Paper presented at the Proceedings of the Ninth International Conference on the Simulation of Adaptive Behavior (SAB'06), Rome.
- Gerstner, W., & Kistler, W. M. (2002). *Spiking Neuron Models. Single Neurons, Populations, Plasticity*: Cambridge University Press.
- Hänggi, P. (2002). Stochastic resonance in biology: how noise can enhance detection of weak signals and help improve biological information processing. *Chemphyschem* 3(3), 285-290.
- Hopfield, J. J. (1999). Odor space and olfactory processing: Collective algorithms and neural implementation. *PNAS*, 96(22), 12506-12511.
- Hoshino, O., Kashimori, Y., & Kambara, T. (1998). An olfactory recognition model based on spatio-temporal encoding of odor quality in the olfactory bulb. *Biological Cybernetics*, 79, 109-120.
- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*.
- Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons. *IEEE Transactions on Neural Networks*.
- Jacobi, N., Husbands, P., & Harvey, I. (1995). Noise and the Reality Gap: The Use of Simulation in Evolutionary Robotics, *Proceedings of the Third European Conference on Advances in Artificial Life*: Springer-Verlag.
- Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2000). *Principles of Neural Science* (4th ed.): McGraw-Hill.
- Kanzaki, R., Nagasawa, S., & Shimoyama, I. (2005). Neural Basis of Odor-source Searching Behavior in Insect Brain Systems Evaluated with a Mobile Robot. *Chemical Senses*, 30(1), 285-286.
- Koch, C. (1999). *Biophysics of Computation: Information Processing in Single Neurons*: Oxford University Press: New York, New York.
- Kuwana, Y., & Shimoyama, I. (1998). A Pheromone-Guided Mobile Robot that Behaves like a Silkworm Moth with Living Antennae as Pheromone Sensors. *The International Journal of Robotics Research*, 17, 924-933.
- Laurent, G., Wehr, M., & Davidowitz, H. (1996). Temporal Representations of Odors in an Olfactory Network *Journal of Neuroscience*, 16(12), 3837-3847.
- Mori, T., & Kai, S. (2002). Noise-Induced Entrainment and Stochastic Resonance in Human Brain Waves. *Physical Review Letters*, 88(21), 218101.
- Moss, F., Ward, L., & Sannita, W. (2004). Stochastic resonance and sensory information processing: a tutorial and review of application. *Clin Neurophys*, 27, 677- 682.
- Oros, N., Steuber, V., Davey, N., Cañamero, L., & Adams, R. G. (2008). *Optimal receptor response functions for the detection of pheromones by agents driven by spiking neural networks*. Paper presented at the 19th European Meetings on Cybernetics and Systems Research, Vienna.
- Payton, D., Daily, M., Estowski, R., Howard, M., & Lee, C. (2001). Pheromone Robotics. *Auton. Robots*, 11(3), 319-324.
- Pyk, P., Bermúdez i Badia, S., Bernardet, U., Knüsel, P., Carlsson, M., Gu, J., et al. (2006). An artificial moth: Chemical source localization using a robot based neuronal model of moth optomotor anemotactic search. *Autonomous Robots*, 20(3), 197-213.
- Rudolph, M., & Destexhe, A. (2003). Characterization of subthreshold voltage fluctuations in neuronal membranes. *Neural Computation*, 15(11), 2577-2618.
- Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of Brownian Motion. *Physical Review*, 36, 823-841.
- Webb, B. (1998). Robots crickets and ants: models of neural control of chemotaxis and phonotaxis. *Neural Networks*, 11, 1479-1496.
- Wiesenfeld, K., & Moss, F. (1995). Stochastic resonance and the benefits of noise: from ice ages to crayfish and SQUIDS. *Nature*, 373, 33-36.
- Wyatt, T. D. (2003). *Pheromones and Animal Behaviour, Communication by Smell and Taste*: Cambridge University Press.