

Adaptive Olfactory Encoding in Agents Controlled by Spiking Neural Networks

Nicolas Oros, Volker Steuber, Neil Davey, Lola Cañamero, and 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 neural architecture that can use two different types of information encoding strategies depending on the environment. The goal of this research was to create a simulated agent that could react to two different overlapping chemicals having varying concentrations. The neural network controls the agent by encoding its sensory information as temporal coincidences in a low concentration environment, and as firing rates at high concentration. With such an architecture, we could study synchronization of firing in a simple manner and see its effect on the agent's behaviour.

Keywords: spiking neural network, neural encoding, firing rate, temporal coincidence.

1 Introduction

Animals are able to react to chemicals (odours, pheromones...) present in the environment. The key sense to detect these chemical cues is smell [18]. 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 [10]. 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 [18]. 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 [2, 6, 7] and also using robots [11, 13, 15-17]. We were interested in studying 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) 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.

2 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. In this experiment, each chemical source had a circular shape and the same fixed value all over its surface. 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 & 2). The left and right wheels were situated in the appropriate position (middle of each side) so the gravity centre was in the middle of the agent's body (excluding the antennae).

To control the agent, we had to decide which neurons' model to use in order to study firing synchronization of the sensors.

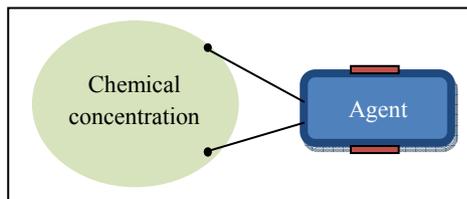


Fig. 1. An agent equipped with two wheels and two antennae used to detect chemicals

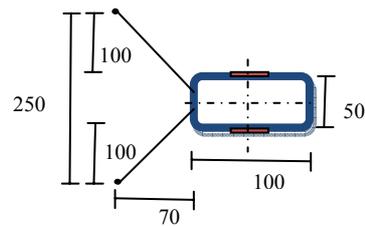


Fig. 2. Properties of an agent. Units are arbitrary.

3 Neural Network

There are three main ways to encode the intensity of sensory information into spiking neurons based on biological evidences [3-5, 8, 9, 12]. 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 [10, 14, 18]. J. Hopfield and C. Brody [2, 6] 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.

a) 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 [4, 5, 8, 9, 12]. 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 (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 [10]. When a neuron fires, its membrane potential will increase rapidly to about 30mV, so the height of a typical spike is approximately 100mV [10]. 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. 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.

b) 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. Such relationships exist in biological systems. For example in humans, the relationship between the frequency of firing and pressure on the skin is linear [10]. 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. 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. In this paper, we are using a very simple model chemical concentration that has only one value. Therefore, the sensory neurons encode this value onto the appropriate firing rate. The sensors were configured in order to distinguish a large range of concentrations between 1 and 300. Over this 300, they were saturating.

c) 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.

4 Experiments

We used the agent and world described in Section 2. The world contained either one or two chemicals denoted by A or B. One agent, placed in the world, was controlled by a simple spiking neural network implementing the neurons described in Section 3. The neural controller was based on a Braitenberg vehicle (anger behaviour) [1] where an agent moves faster toward a stimulus when it detects it (Fig. 3).

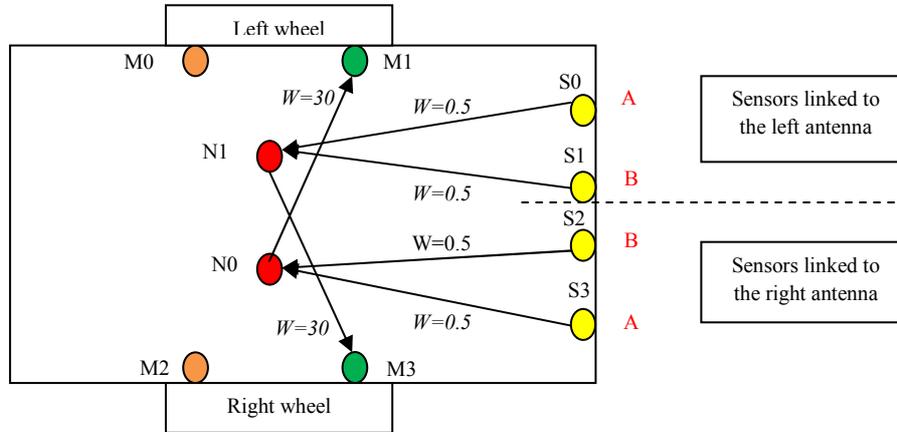


Fig. 3. Agent's neural controller. The sensors S0 and S3 detect the chemical A and the sensors S1 and S2 detect the chemical B. The sensory axons' lengths are all similar (delays = 2.5ms). The motor neurons M1 and M3 are responsible to move the agent forward. The threshold of the neurons (N0 and N1) was set to 4.6 mV. W is the synaptic weight.

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.

a) 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. 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 agent's neural activity during each run. Figure 4 shows an example where an agent starts from the position P2. In this case, the agent was able to stay in the overlapping area.

By looking at Figure 5, we can see that the agent begins by moving horizontally left to right until its right antenna detects the chemicals A and B (T1, Fig. 4 and 5). At this point, the sensors S2 and S3 fire and the temporal coincident arrival of their

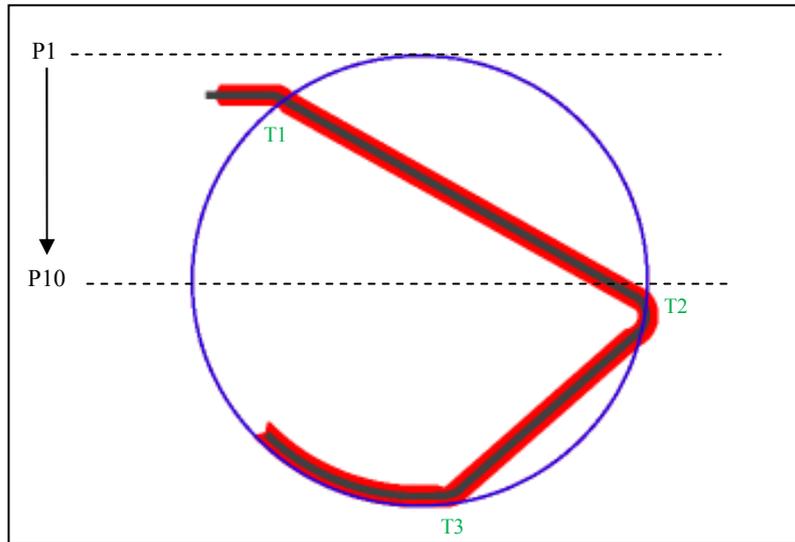


Fig. 4. Path of the agent starting from the left at P2. The circle in the centre represents the two overlapping concentrations of chemicals A and B.

spikes causes N0 to fire. M1 is then stimulated and increases its firing rate turning the left wheel faster than the right one. Soon after this, both antennae detect the chemicals causing also the neuron N1 to fire so the agent moves straight forward again. At T2, the left antenna of the agent goes outside the overlapping area so the sensors S0 and S1 stop to fire and therefore, do not stimulate the neuron N1. The motor neuron M3 then fires at a lower rate than M1 resulting in a left turn of the agent to stay inside the area. Finally, from T3, the interaction between the left antenna and the concentration causes the edge-following behaviour.

We also 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. 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 (Fig. 6).

b) 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

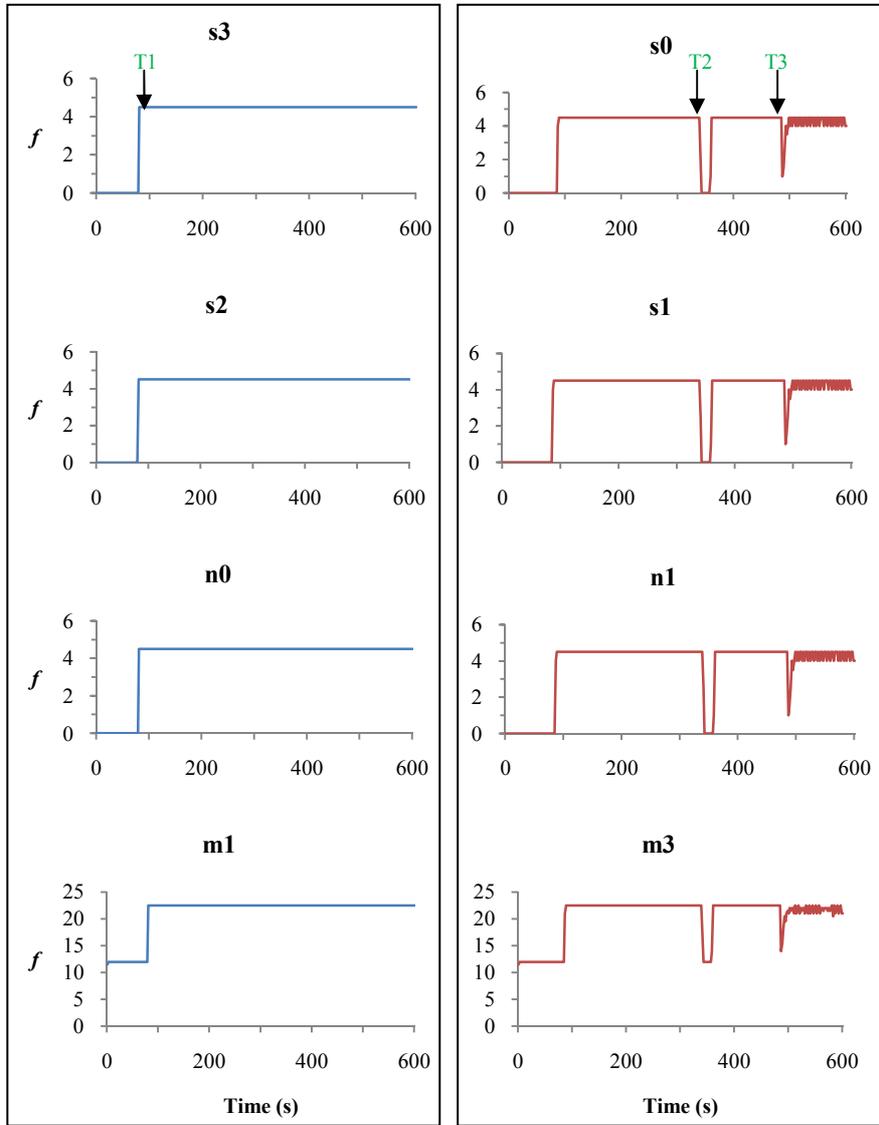


Fig. 5. Firing rates f (in spikes/s) of the neural network cells recorded every 2s during one run (Experiment shown in Fig. 4). The motor neurons M0 and M1 are not shown here as they do not fire. On the left panel, the sensors detecting the chemicals A (S3) and B (S2) from the right antenna activates the neuron N0 that stimulates the motor neuron M1 controlling the left wheel to move forward. On the right panel, the sensors detecting the chemicals A (S0) and B (S1) from the left antenna activates the neuron N1 that stimulates the motor neuron M3 controlling the right wheel to move forward.

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 the same set up as for the experiment shown in Figure 4.

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 6, 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. 6).

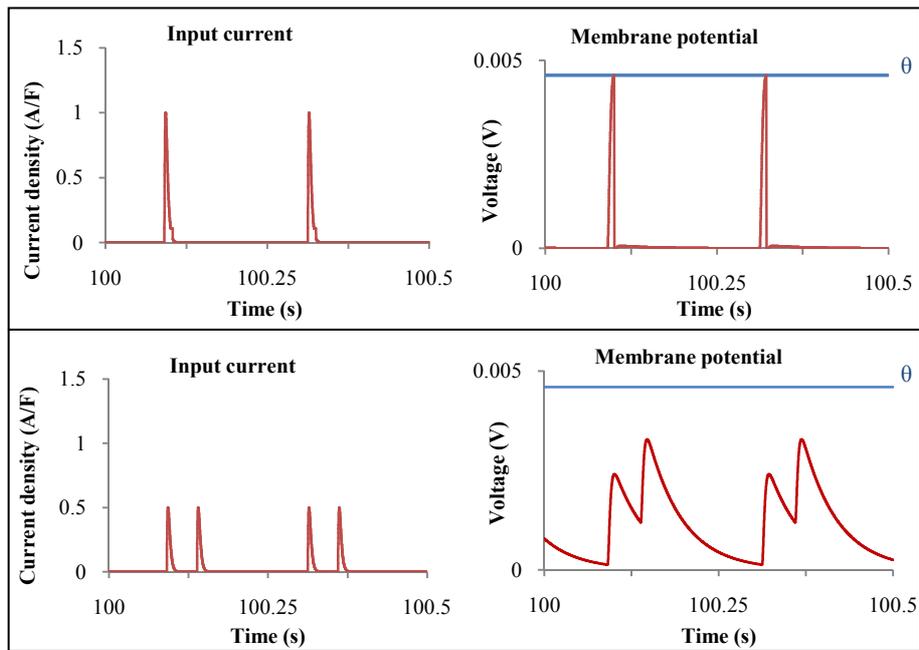


Fig. 6. Current density (in Amperes per Farad) and membrane potential (in Volts) of the neuron N0 recorded between 100s and 100.5s. On the top panel (Experiment I), the spikes sent by the sensors arrived at the same time increasing the current density to 1 A/F. The membrane potential was then increased 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 were changed to 50ms in this case, so the current was never above 0.5 A/F and therefore, the membrane potential could not reach the threshold to make the neuron N0 fire.

c) 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.

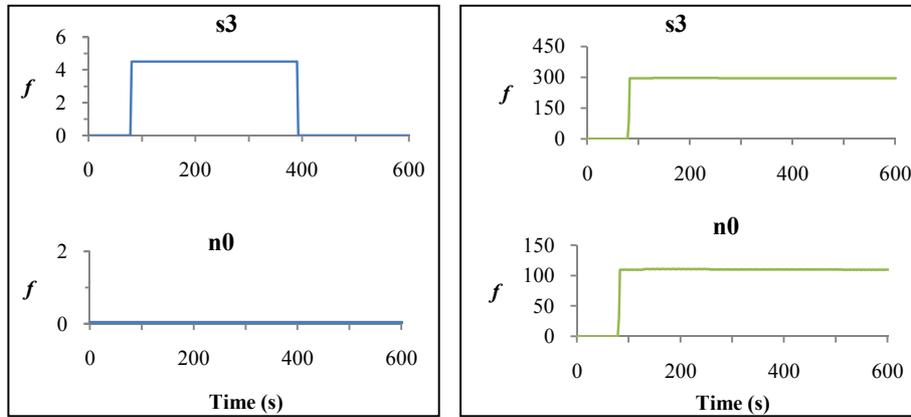


Fig. 7. Firing rates of the sensor S3 and neuron N0 recorded every 2s. On the left panel, as the environment contained a low concentration ($=1$) of chemical A only, the neuron could not detect it and therefore, the agent did not stay within the chemical source area. On the right panel, the concentration was high ($=300$) so the neuron could detect it and the agent stayed inside the area.

5 Conclusion

We 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. In principle, more than two chemicals can be detected and processed. The architecture presented here also works when the chemical concentration has a linear gradient.

Future work will investigate evolving such architectures using a developmental model (evolving the number of neurons and their connections, the synaptic weights, and delays of the neural network). Moreover, we will add noise to the neural network and use a more complex environment.

References

1. Braitenberg, V.: *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge (1984)
2. Brody, C.D., Hopfield, J.J.: Simple Networks for Spike-Timing-Based Computation, with Application to Olfactory Processing. *Neuron* 37, 843–852 (2003)
3. Floreano, D., Mattiussi, C.: Evolution of Spiking Neural Controllers for Autonomous Vision-Based Robots. In: Gomi, T. (ed.) *ER-EvoRob 2001*. LNCS, vol. 2217, pp. 38–61. Springer, Heidelberg (2001)
4. Florian, R.V.: Biologically inspired neural networks for the control of embodied agents. Technical report Coneural-03-03 Version 1.0 (2003)
5. Gerstner, W., Kistler, W.M.: *Spiking Neuron Models. Single Neurons, Populations, Plasticity*. Cambridge University Press, Cambridge (2002)
6. Hopfield, J.J.: Odor space and olfactory processing: Collective algorithms and neural implementation. *PNAS* 96, 12506–12511 (1999)
7. Hoshino, O., Kashimori, Y., Kambara, T.: An olfactory recognition model based on spatio-temporal encoding of odor quality in the olfactory bulb. *Biological Cybernetics* 79, 109–120 (1998)
8. Izhikevich, E.M.: Simple model of spiking neurons. *IEEE Transactions on Neural Networks* (2003)
9. Izhikevich, E.M.: Which model to use for cortical spiking neurons. *IEEE Transactions on Neural Networks* (2004)
10. Kandel, E.R., Schwartz, J.H., Jessell, T.M.: *Principles of Neural Science*. McGraw-Hill, New York (2000)
11. Kanzaki, R., Nagasawa, S., Shimoyama, I.: Neural Basis of Odor-source Searching Behavior in Insect Brain Systems Evaluated with a Mobile Robot. *Chemical Senses* 30, 285–286 (2005)
12. Koch, C.: *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, New York (1999)
13. Kuwana, Y., Shimoyama, I.: 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 (1998)
14. Laurent, G., Wehr, M., Davidowitz, H.: Temporal Representations of Odors in an Olfactory Network. *Journal of Neuroscience*. 16, 3837–3847 (1996)
15. Payton, D., Daily, M., Estowski, R., Howard, M., Lee, C.: Pheromone Robotics. *Auton. Robots*. 11, 319–324 (2001)

16. Pyk, P., Bermúdez i Badia, S., Bernardet, U., Knüsel, P., Carlsson, M., Gu, J., Chanie, E., Hansson, B., Pearce, T.J., Verschure, P.: An artificial moth: Chemical source localization using a robot based neuronal model of moth optomotor anemotactic search. *Autonomous Robots* 20, 197–213 (2006)
17. Webb, B.: Robots crickets and ants: models of neural control of chemotaxis and phonotaxis. *Neural Networks* 11, 1479–1496 (1998)
18. Wyatt, T.D.: *Pheromones and Animal Behaviour, Communication by Smell and Taste*. Cambridge University Press, Cambridge (2003)