A Minimal Neural Network for Target Searching and Danger Avoidance by Smell*

Wee Kheng Leow and Hsueh Chyi Yee
Department of Information Systems and Computer Science
National University of Singapore
Lower Kent Ridge Road, Singapore 119260, Singapore
leowwk@iscs.nus.sg

Abstract

Research on sensory-motor coordination has attracted an increasing amount of interests recently. However, these works have been focused almost exclusively on visuomotor coordination. Research on olfactory-motor coordination, on the other hand, has been lacking. This paper contributes to the study of olfactory-motor coordination by devising a minimal neural network that allows a simulated creature to search for food and to avoid danger by smell.

1. Introduction

To many animals, being able to detect distant food by smell is vital for their survival. The sense of smell can be more important than vision, especially for nocturnal animals and creatures that live underground. An understanding of how these animals search for food and avoid danger by smell can be useful for industrial applications such as tracking the source of air pollution by a mobile robot. Over the years, research on sensory-motor coordination has attracted an increasing amount of interests (e.g., [3, 5, 1, 8]). However, these works have been focused almost exclusively on visuomotor coordination. Studies of olfactory-motor coordination, on the other hand, have been lacking.

This paper describes the olfactory-motor network of a simulated creature that can search for food and avoid danger by smell. We are interested in not just *any* network that can perform the tasks, but in a *minimal* network. As pointed out by Meyer and Guillot [7], in order that we can understand what network architecture is important for solving what kind of behavioural problem, it is necessary to establish the minimal solution for a problem and to explore the limits of the solution.

2. Sensing odorants

Studies in neurophysiology [4, 6] showed that olfactory sensors' activities are related to both the type and the strength of odorants. The sensors are broadly tuned to respond to a variety of odorants. Some sensors respond more strongly to, say, camphor while others respond more strongly to musk. Therefore, the sensors will show different activity patterns in response to different types of odorants. The strength of the activity pattern depends on the intensity of the odorants, i.e., the amount of air-borned chemicals. The more the chemicals, the stronger are the sensors' activities.

In the presence of a variety of odorants, olfactory sensors' response turns out to be very complex. The response to a combination of odorants is not necessarily a linear combination of the responses to individual odorants [6]. Nevertheless, it is assumed in this paper that such a linear combination is approximately correct, so that essential network mechanisms required for target searching and danger avoidance can be easily identified. This assumption is justifiable since our work focuses on olfactory-motor coordination instead of odor discrimination.

3. Finding target by smell

Finding target by smell is inherently more difficult than searching by sight. Visual information is directional because light travels in a straight path. A creature that detects the light emitted or reflected from a target can easily compute its direction to the target,

*Supported by NUS Academic Research Grant No. RP950656.

Appears in *Proceedings of International Conference on Neural Networks*, 1997, vol. I, 580–585.

and then approach it in that direction. Olfactory information is, on the other hand, local and non-directional. An odorant source emits chemicals which are dispersed in all directions. The amount of chemicals at a particular location decreases with increasing distance from the source. A creature's olfactory sensors can only detect the type and the amount of chemicals at its present location, and this information is insufficient for it to locate the source.

There are many ways to obtain enough information to deduce where is the odorant source. The following is a list of some methods, in increasing amount of information or neural hardware requirement:

1. Temporal Difference. The creature can compare the odorant intensities sensed at the previous and current time steps. If the creature has moved a short distance since the previous time step, then it can decide its next step of movement based on the temporal difference. For example, it can decide to continue moving forward if it senses an increase in odorant intensity.

2. Spatial Difference. For a creature that has two sets of olfactory sensors located at a significant distance apart, it can compare the odorant intensities sensed by the two sensors and bias its movement towards the side which has a greater intensity.

3. Second-Order Difference. The creature can compute the local spatial differences of odorant intensity at a variety of directions and move along the direction of greatest increase. This method computes a second-order spatial difference of the intensity. Other types of second-order difference, such as spatial-temporal difference, are possible.

This paper focuses on the use of temporal difference for target searching and danger avoidance since it requires the least amount of sensor hardware and computation. Moreover, this paper will focus only on the case of a static environment in which the spread of chemicals from the source has reached an equilibrium state and there is no wind disturbance in the environment. Studies for the case of a dynamic environment and comparisons of the various difference measures for target searching will be described in separate papers.

4. Network architecture and dynamics

The sensory-motor coordination network of the simulated creature in divided into three main modules: the *Sensor Module*, the *Detection Module*, and the *Motor Control Module* (Fig. 1). The Sensor Module consists

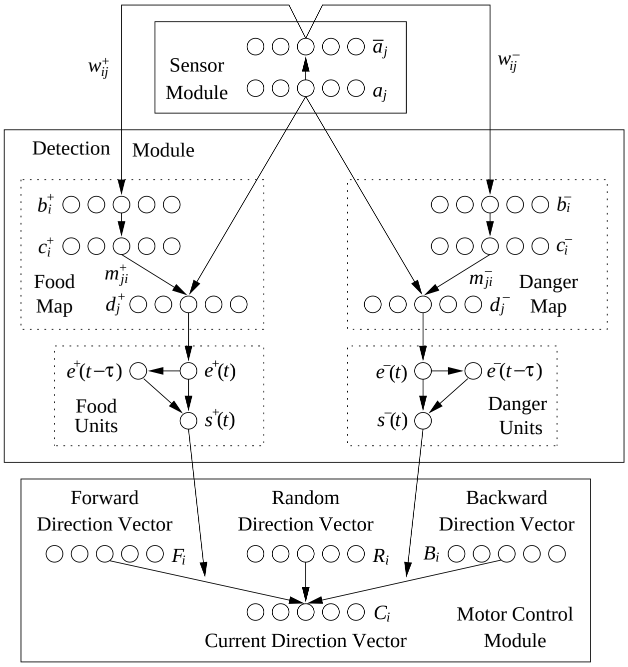

Figure 1. The neural network for sensory-motor control is divided into three main modules: (a) the Sensor Module which consists of smell sensors, (b) the Detection Module which constsis of the Food Map, Food Units, Danger Maps and Danger Units, and (c) the Motor Control Module which consists of various Direction Vectors.

of two sets of neural units denoted by a_j and $\bar{a}_j$. Each of the smell sensors a_j is broadly tuned to a variety of odorants (as discussed in Section 2). The sensors' activities are proportional to the amount of chemical present, which is assumed to be inversely proportional to the square of the distance from the source. The sensors' response to a combination of stimuli is a linear combination of the responses to the individual stimuli. The units $\bar{a}_j$ in the Sensor Module computes the normalized activity of a_j, i.e., $\bar{a}_j = a_j / \sum_k a_k$. The normalized activity is used in the Detection Module to compute the amount of appetitive stimuli (food) and aversive stimuli (danger) present.

The Detection Module consists of two identical subnetworks, one for detecting food and the other for detecting danger (Fig. 1). The only difference between them is that their connections from units $\bar{a}_j$ in the Sensor Module carry different weights such that they detect the presence of either food or danger but not both. The following discussion will just focus on the detection of food.

The Detection Module consists of the *Food Maps* and the *Food Units*. The Food Maps are used to extract the activity pattern corresponding to the dominant appetitive stimuli among the combination of patterns in the smell sensors. Each Food Map unit b_i^+ detects the presence of one type of known stimulus by computing the weighted-sum of the inputs from units $\bar{a}_i$:

$$b_i^+ = \sum_j w_{ij}^+ \bar{a}_j \ . \tag{1}$$

The weights w_{ij}^+ of the connections to unit b_i^+ encode the sensors' activity pattern in response to stimulus i. In other words, the weights encode the stimulus that unit b_i^+ detects. The weights are normalized such that $0 \leq w_{ij}^+ \leq 1$ for each i, j and $\sum_j w_{ij}^+ = 1$. Unit b_i^+ will be strongly activated if the stimulus that it detects is present. The Food Map units c_i^+ identify the most strongly activated unit b_i^+, which in turn, identifies the dominant stimulus:

$$c_i^+ = \begin{cases} 1 & \text{if } b_i^+ = \max_j b_j^+ \\ 0 & \text{otherwise} \end{cases} \ . \tag{2}$$

At most one of the units c_i^+ is active at any one time. The weights m_{ji}^+ of the connections from unit c_i^+ to unit d_j^+ carry exactly the same values as w_{ij}^+. The active unit c_i^+ thus projects the activity pattern corresponding to the dominant stimulus into the Food Map units d_j^+:

$$d_j^+ = a_j \sum_i m_{ji}^+ c_i^+ \ . \tag{3}$$

Units d_j^+ also receive inputs from sensor units a_j. The activities d_j^+ are summed up by the Food Unit, $e^+(t) = \sum_j d_j^+(t)$, to extract the intensity of the dominant stimulus (see following example for explanation). The Food Unit $e^+(t-\tau)$ receives the activity of unit $e^+(t)$ with a time delay of τ, thus keeping track of the intensity of the dominant stimulus detected at the previous time step. The Food Unit $s^+(t)$ measures the increment of the stimulus intensity over one time step:

$$s^+(t) = \max(e^+(t) - e^+(t-\tau), 0) \ . \tag{4}$$

To see how the Food Maps and the Food Units extract the intensity of the dominant stimulus, let us consider the following example. The sensors' activity in response to a single stimulus p is given by

$$a_j = I_p u_{pj} \tag{5}$$

where I_p is the intensity of the stimulus and u_{pj} is a normalized pattern corresponding to the stimulus type. The pattern u_{pj} is identical to w_{ij}^+ for a particular i. In the presence of several stimuli, the sensors' activities become

$$\begin{aligned} a_j &= \sum_p I_p u_{pj} \\ \bar{a}_j &= \sum_p I_p u_{pj} \Big/ \sum_j \sum_p I_p u_{pj} \ . \end{aligned} \tag{6}$$

The food map units b_i^+ compute weighted-sums of the sensors' activities:

$$b_i^+ = \sum_j \sum_p I_p u_{pj} w_{ij}^+ \Big/ \sum_j \sum_p I_p u_{pj} \ . \tag{7}$$

Since u_{pj} and w_{ij}^+ are normalized, $\sum_j u_{pj} w_{ij}^+$ attains the largest value when $w_{ij}^+ = u_{pj}$, for each j. Suppose that stimulus k is the dominant stimulus, i.e., $I_k > I_p$ for $p \neq k$. Then, the largest b_i^+ is the one with $i = k$ and $w_{ij}^+ = u_{kj}$, for each j. As a result, $c_k^+ = 1$ and $c_i^+ = 0$, for $i \neq k$, and

$$d_j^+ = a_j m_{jk}^+ = \sum_p I_p u_{pj} w_{kj}^+ \ . \tag{8}$$

Assuming that the sensors' activity patterns in response to the individual stimuli do not overlap significantly, then $d_j^+ \approx I_k {w_{kj}^+}^2$, and $e^+(t) = I_k \sum_j {w_{kj}^+}^2$. Since w_{kj}^+ are normalized weights, $\sum_j {w_{kj}^+}^2$ is a constant for each type of stimulus. Thus, $e^+(t)$ is a good measure of the intensity I_k of the dominant stimulus.

In the Motor Control Module, the Current Direction Vector C_i encodes the current egocentric direction of motion. Each unit C_i is broadly tuned to a variety of directions, but is most strongly activated for its preferred direction, denoted by the unit vector $\mathbf{v}_i$. The resultant direction $\mathbf{C}$ encoded by the whole direction vector is the vector sum of each C_i's preferred direction, i.e., $\mathbf{C} = \sum_i C_i \mathbf{v}_i$. This kind of direction encoding has been observed in the Accessory Optic System of cats and primates [2].

In addition to the Current Direction Vector, there are three other direction vectors that govern that direction of motion. The Forward Direction Vector F_i and Backward Direction Vector B_i are constant vectors that encode the forward and backward egocentric directions. The Random Direction Vector R_i is used to generate random bias to the left or the right of the forward direction. The creature has touch sensors (not drawn in Fig. 1) all over its body to sense the presence of obstacles. When it detects an obstacle on the left side of its body, for example, the Random Direction Vector will generate a pattern encoding a direction towards the right. When no obstacle is detected, the Random Direction Vector will generate a random direction. The summed effect of these three vectors determine the current direction of motion:

$$D_i = \alpha^+ s^+ H F_i + \alpha^- s^- B_i + \beta R_i \qquad (9)$$

where H represents the creature's internal state of *hunger level*, and α^+, α^- and β are constant parameters. Equation 9 embodies the creature's strategy in food searching, danger avoidance, and obstacle negotiation:

- When the creature senses an increase in the amount of appetitive stimuli, it is *biased* towards moving forward. The higher the hunger level, the stronger is the bias.
- When the creature senses an increase in the amount of aversive stimuli, it is *biased* towards turning around and moving in the reverse direction.
- At all time, the direction of motion is *biased* by the Random Direction Vector, which is controlled by the sensing of obstacles.

To support the above strategy, the value of β should be comparable to the values of $\alpha^+ s^+ H$ and $\alpha^- s^-$ for very small values of s^+ and s^-.

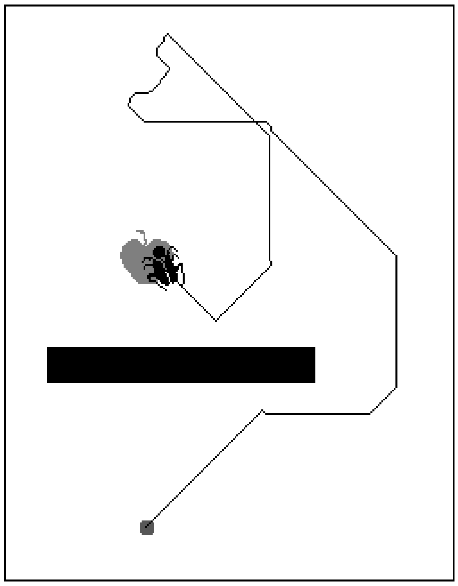

Figure 2. Experiment 1(a): Finding food and negotiating obstacle. The simulated creature (bug) was able to go around the obstable (bar) while searching for the food (apple). The dark spot denotes the starting point, and the curved line indicates the path taken by the creature.

5. Experimental results

Two sets of experiments were carried out to test various aspects of the simulated creature's behaviour.

5.1. Experiment 1: Finding food and negotiating obstacle

The first set of experiments test the creature's ability to negotiate obstacle while searching for food (Fig. 2). Initially, the creature took a random heading which turned out to be in the direction of increasing stimulus intensity. It continued moving in that direction until it encountered the obstacle. At this point, it took a right turn, and sensing no increase in stimulus intensity, immediately took a random turn which happened to be a left turn. It continued moving and turning until it reached the food. The creature's random turns were not always the correct turns. Nevertheless, it was able to go around the obstacle, without previously being taught to do so.

In the above experiment, the distribution of chemical stimuli obeys the inverse-square law of the distance to the food. In reality, however, local air perturbation can disturb chemical concentration thus invalidating the inverse-square law. To simulate this situation, random noise was introduced. Figure 3(a) shows the creature's behaviour in a noisy environment. The initial setup of the experiment was identical to the noise-free case except that random noise was added to the stimulus intensity. The peak amplitude of the random noise was equal to 1% of the stimulus intensity at the food, which was the peak stimulus intensity. Compared to the noise-free situation, the creature's path was a bit more erratic, but the general trend of the movement was similar, and the creature was able to find the food. Figure 3(b) shows the case in which the peak noise amplitude was equal to 10% of the peak stimulus intensity. This time the creature's path was very erratic. When the creature was some distance away from the food source, the random noise was sometimes large enough to mask the actual stimulus intensity, therefore misleading the creature. Neverthelss, when the creature was near enough, it had no problem going directly to the food. The experiments show that the sensory-motor coordination network is robust against the presence of noise in the stimulus intensity.

(a) (b)

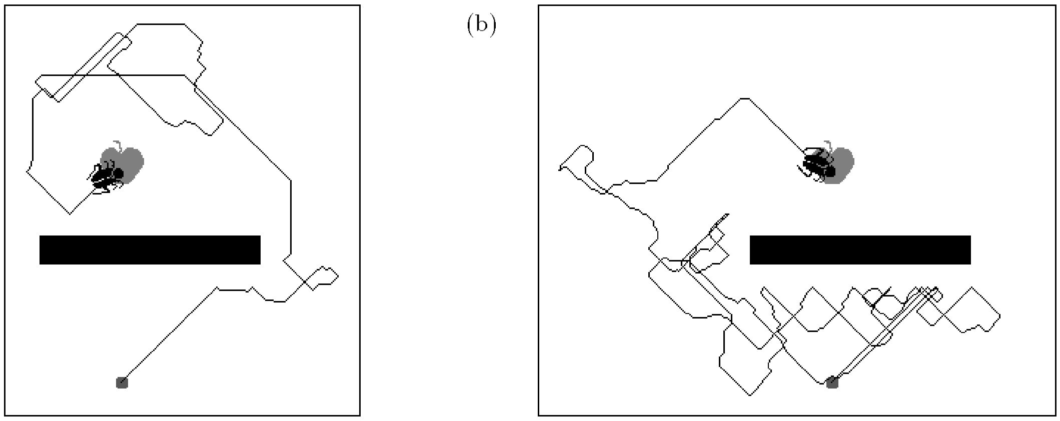

Figure 3. Experiment 1(b): Finding food and negotiating obstacle in noisy environment. (a) Peak noise amplitude = 1% of the peak stimulus intensity. (b) Peak noise amplitude = 10% of the peak stimulus intensity.

(a) (b) (c)

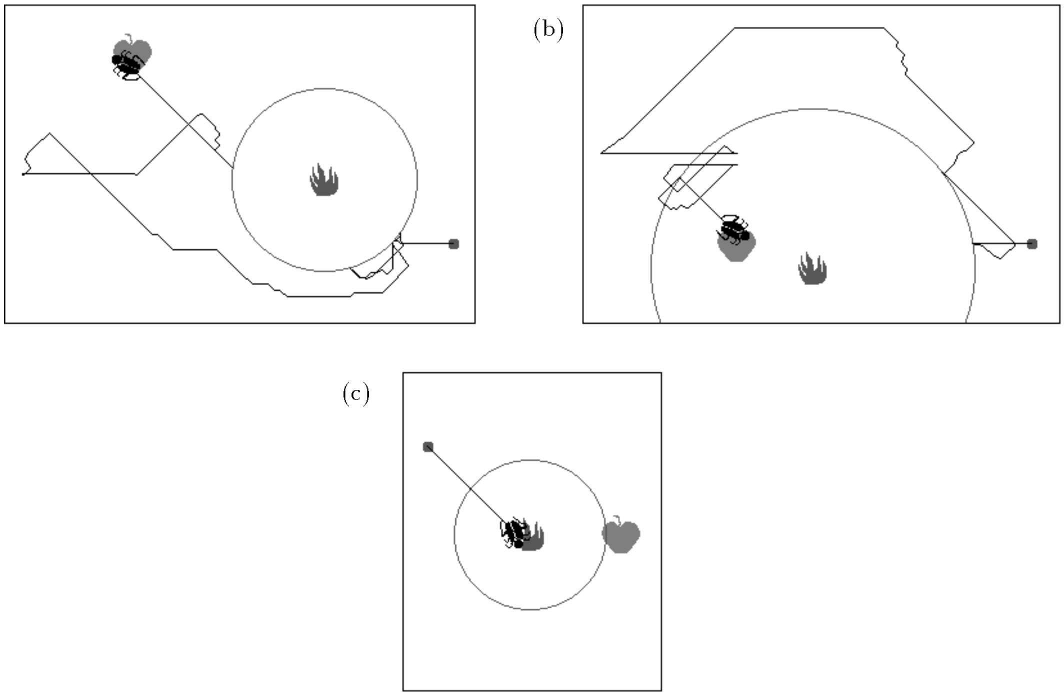

Figure 4. Experiment 2: Finding food and avoiding danger. The fire symbol denotes danger and the circle indicates the zone in which the aversive stimuli were strong enough to be detected by the creature. (a) When the creature was not hungry, it avoided the danger zone completely. (b) When the creature was hungry, it took risk by venturing into the danger zone. (c) When the creature was very hungry, it could become suicidal.

5.2. Experiment 2: Finding food and avoiding danger

The second set of experiments test the creature's ability to avoid danger while searching for food (Fig. 4). The circles in the figures indicate the zones in which the aversive stimuli were strong enough to be detected by the creature. Figure 4(a) shows the case in which the creature was not hungry. Whenever the creature got to the boundary of the danger zone, it detected the presence of aversive stimulus and always moved away from the danger zone. In the case shown in Fig. 4(b), the creature initially avoided the danger zone. However, as it got more hungry, the detection of increase in appetitive stimulus had a stronger influence over the detection of increase in aversive stimulus. As a result, the creature would *take risk* by venturing deeper and deeper into the danger zone. In the case shown in Fig. 4(c), the creature was so hungry that the detection of increase in appetitive stimulus completely overwhelmed the detection of increase in aversive stimulus. The creature became misguided and *suicidal*—it plunged straight into the danger. This kind of risk taking and suicidal behaviour seems to be present in real animals.

6. Future work

This paper reports only the sensory-motor network's performance in a static environment with local perturbation of stimulus intensity. It will be interesting to see how well the network performs in a dynamic environment in which the food, the danger, and the obstacle can move around, and in which wind can disturb the stimulus distribution. The current network uses only temporal difference of sensed stimulus intensity to search for food and to avoid danger. It will be useful to compare the effectiveness of this method with those that employ spatial difference and second-order differences of sensed stimulus intensity.

7. Conclusion

This paper has described a neural network capable of target searching and danger avoidance by smell, and obstacle negotiation by touch, with emphasis on olfactory-motor coordination. The network uses only (1) the temporal difference of the stimulus intensity sensed by a single set of olfactory sensors and (2) a random direction generator contolled by touch. Items (1) and (2) are, respectively, the bare minimum required for olfactory-motor coordination and detecting obstacle by touch. Therefore, the network is indeed a minimal network capable of performing the above tasks in a noisy environment. The network will fail to operate successfully only when the internal bias towards the food (such as the hunger level) overwhelms the detection of increasing danger.

References

[1] M. A. Arbib and H. B. Lee. Anuran visuomotor coordination for detour behavior: From retina to motor schemas. In J.-A. Meyer, H. L. Roitblat, and S. W. Wilson, editors, *From Animals to Animats 2*, pages 42–51, Cambridge, Massachusetts, 1993. MIT Press.

[2] K. L. Grasse and M. Cynader. The accessory optic system in frontal-eyed animals. In A. G. Leventhal, editor, *The Neural Basis of Visual Function.* Macmillan Press, London, 1991.

[3] S. Grossberg and M. Kuperstein. *Neural Dynamics of Adaptive Sensory-Motor Control.* Pergamon Press, New York, 1989.

[4] E. R. Kandel, J. H. Schwartz, and T. M. Jessell. *Principles of Neural Science.* Appleton & Lange, Norwalk, Connecticut, 3rd edition, 1991.

[5] G. Manteuffel. A biological visuo-motor system: How dissimilar maps interact to produce behavior. In J.-A. Meyer and S. W. Wilson, editors, *From Animals to Animats 1*, pages 120–126, Cambridge, Massachusetts, 1991. MIT Press.

[6] C. Masson and C. Linster. Towards a cognitive understanding of odor discrimination: Combining experimental and theoretical approaches. *Behavioral Processes*, 1995.

[7] J.-A. Meyer and A. Guillot. Simulation of adaptive behavior in animals: Review and prospect. In J.-A. Meyer and S. W. Wilson, editors, *From Animals to Animats 1*, pages 2–14, Cambridge, Massachusetts, 1991. MIT Press.

[8] F. Mura and N. Franceschini. Visual control of altitude and speed in a flying agent. In D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson, editors, *From Animals to Animats 3*, pages 91–99, Cambridge, Massachusetts, 1994. MIT Press.