

Evolution of Time in Neural Networks: From the Present to the Past, and Forward to the Future

Ji Ryang Chung, Jaerock Kwon, Timothy A. Mann, and Yoonsuck Choe

Abstract What is time? Since the function of the brain is closely tied in with that of time, investigating the origin of time in the brain can help shed light on this question. In this paper, we propose to use simulated evolution of artificial neural networks to investigate the relationship between time and brain function, and the evolution of time in the brain. A large number of neural network models are based on a feed-forward topology (perceptrons, backpropagation networks, radial basis functions, support vector machines, etc.), thus lacking dynamics. In such networks, the order of input presentation is meaningless (i.e., it does not affect the behavior) since the behavior is largely reactive. That is, such neural networks can only operate in the present, having no access to the past or the future. However, biological neural networks are mostly constructed with a recurrent topology, and recurrent (artificial) neural network models are able to exhibit rich temporal dynamics, thus time becomes an essential factor in their operation. In this paper, we will investigate the emergence of recollection and prediction in evolving neural networks. First, we will show how reactive, feedforward networks can evolve a memory-like function (rec-

Ji Ryang Chung

Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU College Station, TX 77843-3112 e-mail: jchung@cse.tamu.edu

Jaerock Kwon

Department of Electrical and Computer Engineering, Kettering University, 1700 W. University Avenue, Flint, MI 48504 e-mail: jkwon@kettering.edu

Timothy A. Mann

Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU College Station, TX 77843-3112 e-mail: mann@cse.tamu.edu

Yoonsuck Choe

Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU College Station, TX 77843-3112 e-mail: choe@tamu.edu

This paper is largely based on “Evolution of Recollection and Prediction in Neural Networks”, by J. R. Chung, J. Kwon and Y. Choe, which appeared in the *Proceedings of the International Joint Conference on Neural Networks*, 571–577, IEEE Press, 2009.

ollection) through utilizing external markers dropped and detected in the environment. Second, we will investigate how recurrent networks with more predictable internal state trajectory can emerge as an eventual winner in evolutionary struggle when competing networks with less predictable trajectory show the same level of behavioral performance. We expect our results to help us better understand the evolutionary origin of recollection and prediction in neuronal networks, and better appreciate the role of time in brain function.

1 Introduction

What is time? Since the function of the brain is closely tied in with that of time investigating the origin of time in the brain can help shed light on this question. Fig. 1 illustrates the relationship between the past, present, and future on one hand, and brain function such as recollection and prediction on the other hand. Without recollection (or memory), the concept of past cannot exist, and likewise, without prediction, the concept of future cannot either. Furthermore, recollection seems to be a prerequisite for prediction. With this line of thought, we can reason about the possible evolutionary origin of time in the biological nervous systems. In this paper, we propose to use simulated evolution of artificial neural networks to investigate the relationship between time and brain function, and the evolution of time in the brain.

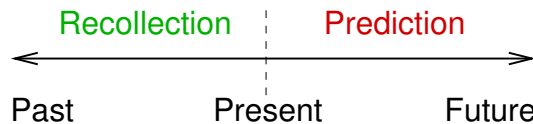


Fig. 1 Time, Recollection, and Prediction. The concept of past, present, and future and brain functions such as recollection (memory) and prediction are all intricately related.

Many neural network models are based on a feedforward topology (perceptrons, backpropagation networks, radial basis functions, support vector machines, etc.), thus lacking dynamics (see [Bishop, 1995], and selective chapters in [Haykin, 1999]). In such networks, the order of input presentation is meaningless (i.e., it does not affect the behavior) since the behavior is largely reactive. That is, such neural networks can only operate in the present, having no access to the past or the future. However, biological neural networks are mostly constructed with a recurrent topology (e.g., the visual areas in the brain are not strictly hierarchical [Felleman and Essen, 1991]). Furthermore, recurrent (artificial) neural network models are able to exhibit rich temporal dynamics [Elman, 1990, Elman, 1991, Beer, 2000]. Thus, time becomes an essential factor in neural network operation, whether it is natural or artificial (also see [von der Malsburg and Buhmann, 1992, Choe and Miikkulainen, 2004, Miikkulainen et al., 2005, Peck et al., 2008]).

Our main approach is to investigate the emergence of recollection and prediction in evolving neural networks. Recollection allows organism to connect with its past, and prediction with its future. If time was not relevant to the organism, it would always live in the eternal present.

First, we will investigate the evolution of recollection. We will see how reactive, feedforward networks can evolve a memory-like function (recollection), through utilizing external markers dropped and detected in the environment. In this part, we trained a feedforward network using neuroevolution, where the network is allowed to drop and detect markers in the external environment. Our hypothesis is that this kind of agents could have been an evolutionary bridge between purely reactive agents and fully memory-capable agents. The network is tested in a falling-ball catching task inspired by [Beer, 2000, Ward and Ward, 2006], where an agent with a set of range sensors is supposed to catch multiple falling balls. The trick is that while trying to catch one ball, the other ball can go out of view of the range sensors, thus requiring some sort of memory to be successful. Our results show that even feedforward networks can exhibit memory-like behavior if they are allowed to conduct some form of material interaction, thus closing the loop through the environment (cf. [Rocha, 1996]). This experiment will allow us to understand how recollection (memory) could have evolved.

Second, we will examine the evolution of prediction. Once the recurrent topology is established, how can predictive function evolve, based on the recurrent network's recollective (memory-like) property? For this, we trained a recurrent neural network in a 2D pole-balancing task [Anderson, 1989], again using neuroevolution (cf. [Gomez and Miikkulainen, 1998, Lim and Choe, 2006c, Lim and Choe, 2006a]). The agent is supposed to balance an upright pole while moving in an enclosed arena. This task, due to its more dynamic nature, requires more predictive power to be successful than the simple ball-catching task. Our main question here was whether individuals with a more predictable internal state trajectory have a competitive edge over those with less predictable trajectory. We partitioned high-performing individuals into two groups (i.e., they have the same behavioral performance), those with high internal state predictability and those with low internal state predictability. It turns out that individuals with highly predictable internal state have a competitive edge over their counterpart when the environment poses a tougher problem [Kwon and Choe, 2008].

In sum, our results suggest how recollection and prediction may have evolved, i.e., how "time" evolved in the biological nervous system. We expect our results to help better understand the evolutionary origin of recollection and prediction in neuronal networks, and better appreciate the role of time in neural network models. The rest of the paper is organized as follows. Sec. 2 presents the method and results from the recollection experiment, and Sec. 3, those from the prediction experiment. We will discuss interesting points arising from this research (Sec. 4), and conclude our paper in Sec. 5.

2 Part I: Evolution of Recollection

In this section, we will investigate how memory-like behavior can evolve in a reactive, feedforward network. Below, we will describe the ball catching task, and explain in detail our neuroevolution methods for the learning component. Next, we will present the details of our experiments and the outcomes.

2.1 Task: *Catching Falling Balls*

The main task for this part was the falling ball catching task, inspired by [Beer, 2000, Ward and Ward, 2006]. The task is illustrated in Fig. 2. See the figure caption for details. The task is simple enough, yet includes interesting dynamic components and temporal dependency. The horizontal locations of the balls are on the two different sides (left or right) of the agent's initial position. Between the left and right balls, one is randomly chosen to have faster falling speed (2 times faster than the other). The exact locations are randomly set with the constraint that they must be separated far enough to guarantee that the slower one must go out of the sensor range as the agent moves to catch the faster one. For example, as shown in Fig. 2C, when there are multiple balls to catch and when the balls are falling at different speeds, catching one ball (usually the faster one) results in the other ball (the slower one) going out of view of the range sensors. Note that both the left-left or right-right ball settings cannot preserve the memory requirement of the task. The vertical location, ball speed, and agent speed are experimentally chosen to guarantee that the trained agent can successfully catch both balls. In order to tackle this kind of situation, the controller agent needs some kind of memory.

The learning of connection weights of the agents is achieved by genetic search where the fitness for an agent is set inversely proportional to the sum of horizontal separations between itself and each ball when the ball hits the ground. 10 percent of the best-performing agents in a population are selected for 1-point crossover with probability 0.9 and a mutation with the rate 0.04.

2.2 Methods

In order to control the ball catcher agents, we used feedforward networks equipped with external marker droppers and detectors (Fig. 3, we will call this the “dropper network”). The agent had five range sensors that signal the distance to the ball when the ball comes into contact within the direct line-of-sight of the sensors. We used standard feedforward networks with sigmoidal activation units as a controller (see e.g., [Haykin, 1999]):

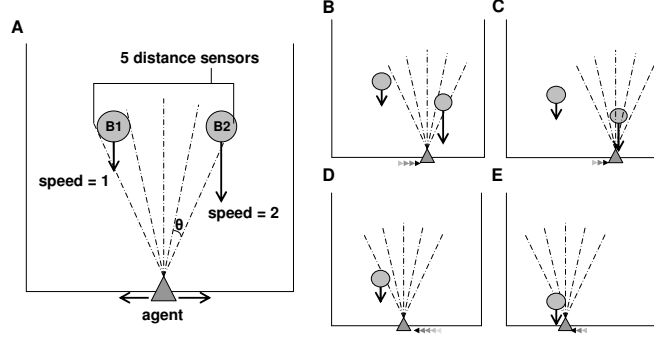


Fig. 2 Ball Catching Task. An illustration of the ball catching task is shown. The agent, equipped with a fixed number of range sensors (radiating lines), is allowed to move left or right at the bottom of the screen while trying to catch balls falling from the top. The goal is to catch both balls. The balls fall at different speeds, so a good strategy is to catch the fast-falling ball first (B and C) and then the go back and catch the slow one (D and E). Note that in C the ball on the left is outside of the range sensors' view. Thus, a memory-less agent would stop at this point and fail to catch the second ball. Adapted from [Chung et al., 2009].

$$\begin{aligned}
 H_j &= \sigma \left(\sum_{i=1}^{N_{in}} v_{ji} I_i \right) & j = 1, \dots, N_{hid} \\
 O_k &= \sigma \left(\sum_{j=1}^{N_{hid}} w_{kj} H_j \right) & k = 1, \dots, N_{out}
 \end{aligned} \tag{1}$$

where I_i , H_j and O_k are the activations of the i -th input, j -th hidden, and k -th output neurons; v_{ji} the input-to-hidden weights and w_{kj} the hidden-to-output weights; $\sigma(\cdot)$ the sigmoid activation function; and N_{in} , N_{hid} , and N_{out} are the number of input, hidden, and output neurons whose values are 7, 3, and 3 respectively.

The network parameters were tuned using genetic algorithms, thus the training did not involve any gradient-based adaptation. Two of the output units were used to determine the movement of the agent. If the agent was moved one step to the left when $O_1 > O_2$, one step to the right when $O_1 < O_2$, and remained in the current spot when $O_1 = O_2$.

If these were the only constructs in the controller, the controller will fail to catch multiple balls as in the case depicted in Fig. 2C. In order to solve this kind of problem, a fully recurrent network is needed, but from an evolutionary point of view, going from a feedforward neural circuit to a recurrent neural circuit could be non-trivial, thus our question was what could have been an easier route to memory-like behavior, without incurring much evolutionary overhead.

Our answer to this question is illustrated in Fig. 3. The architecture is inspired by primitive reactive animals that utilize self-generated chemical droppings (excretions, pheromones, etc.) and chemical sensors [Wood, 1982, Tillman et al., 1999, Conover, 2007]. The idea is to maintain the reactive, feedforward network archi-

texture, while adding a simple external mechanism that would incur only a small overhead in terms of implementation. As shown in Fig. 3, the feedforward network has two additional inputs for the detection of the external markers dropped in the environment, to the left or to the right (they work in a similar manner as the range sensors, signaling the distance to the markers). The network also has one additional output for making a decision whether to drop an external marker or not.

As a comparison, we also implemented a fully recurrent network, with multiple levels of delayed feedback into the hidden layer. (See [Elman, 1990, Elman, 1991] for details.) This network was used to see how well our dropper network does in comparison to a fully memory-equipped network.

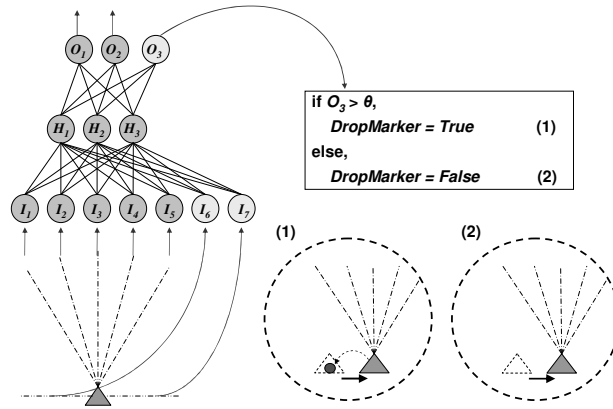


Fig. 3 Feedforward Network with Dropper/Detector (“Dropper Network”). A feedforward network with a slight modification (dropper and detector) is shown. The basic internal architecture of the network is identical to any other feedforward network, with five range sensor (I_1 to I_5), and two output units that determine the movement (O_1 and O_2). The two added input units (I_6 and I_7) signal the presence of a dropped marker on the bottom plane, and the one additional output unit (O_3) makes the decision of whether to drop a marker at the current location or not. Note that there are no recurrent connections in the controller network itself. Adapted from [Chung et al., 2009].

2.3 Experiments and results

The network was trained using genetic algorithms (neuroevolution), where the connection weights and the dropper threshold θ were encoded in the chromosome. The fitness was inversely proportional to the sum of the distance between the agent and the ball(s) when the ball(s) contact the ground. Each individual was tested 12 times with different initial ball position (which was varied randomly) and speed (1 or 2

steps/time unit), and mixed scenarios with fast left ball vs. fast right ball. We used one-point crossover with probability 0.9, with a mutation rate of 0.04.

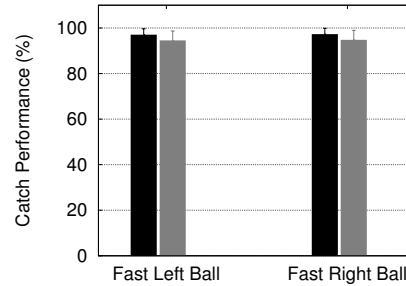


Fig. 4 Ball Catching Performance. The average ball catching performance of the dropper network is presented (gray bar), along with that of the recurrent network (black bar). The error bars indicate the standard deviation. The results are reported in two separate categories: fast left ball and fast right ball. This was to show that the network does not have any bias in performance. Both networks perform at the same high level (above 90% of all balls caught). This is quite remarkable for a feedforward network, although it had the added dropper/detector mechanism. We also tested a purely feedforward networks, but they were only able to catch 50% of the balls (catch one, miss one). Adapted from [Chung et al., 2009].

It is quite remarkable that feedforward networks can show an equal level of performance as that of the recurrent network, although the feedforward networks were equipped with the dropper/detector. For example, compared to the recurrent networks, the number of tunable parameters are meager for the dropper network since they do not have layers of fully connected feedback. Six additional weights for input-to-hidden, and three for hidden-to-output, plus a single threshold parameter (10 in all) is all that is needed.

One question arises from the results above. What kind of strategy is the dropper network using to achieve such a memory-like performance? We analyzed the trajectory and the dropping pattern, and found an interesting strategy that evolved. Fig. 5 shows some example trajectories. Here, we can see a curious *overshooting* behavior.

Fig. 6 shows how this overshooting behavior is relevant to the task, when combined with the dropping events. The strategy can be summarized as below: (1) The right ball fall fast, which is detected first. (2&3) The agent moves toward the right ball, eventually catching it (4). At this point, the left ball is outside of the range sensors' view, it overshoots the right ball, drops a marker there, and immediately returns back, seemingly repelled by the marker that has just been dropped. (5) The agents keeps on dropping the marker which pushing back to the left, until the left ball comes within the view of the range sensor. (6) The agent successfully catches the second ball. This kind of aversive behavior is quite the opposite of what we expected, but for this given task it seem to make pretty good sense, since in some way the agent is "remembering" which direction to avoid, rather than remembering where the slow ball was (compare to the "avoiding the past" strategy proposed in

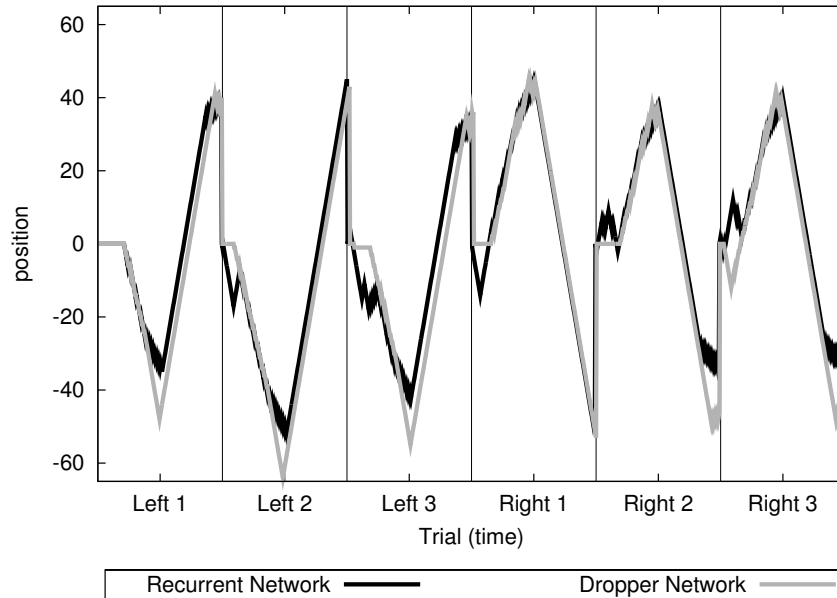


Fig. 5 Agent Trajectory. Agent trajectory during six ball catching trials are shown (gray: dropper network; black: recurrent network). The x axis represents time, and the y axis the agent position (0 marks the initial location of the agent). Within each trial, 200 time steps are shown. As the left and the right ball positions were randomized, the peak of the trajectories differ in their y values. The first three trials were the “fast left ball” condition and the last three were the “fast right ball” condition. Both networks are successful at catching both balls within each trial, but the dropper network shows a curious *overshooting* behavior (for example, near the half way point in each trial). See Fig. 6 for details. Adapted from [Chung et al., 2009].

[Balch, 1993]). Finally, we have also been able to extend our results reported here to a more complex task domain (food foraging in 2D). See [Chung and Choe, 2009] for details.

3 Part II: Evolution of Prediction

In this second part, we will now examine how predictive capabilities could have emerged through evolution. Here, we use a recurrent neural network controller in a 2D pole-balancing task. Usually recurrent neural networks are associated with some kind of memory, i.e., an instrument to look back into the past. However, here we argue that it can also be seen as holding a predictive capacity, i.e., looking into the future. Below, we first describe the 2D pole-balancing task, and explain our methods, followed by experiments and results. The methods and results reported in this part are largely based on our earlier work [Kwon and Choe, 2008].

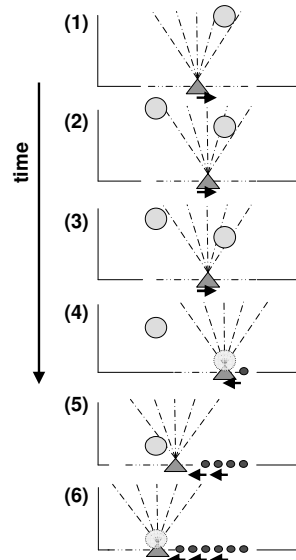


Fig. 6 Dropper Network Strategy A strategy that evolved in the dropper network is shown. (1) Fast ball enters the view. (2&3) Agent moves toward the fast ball. (4) Agent catches fast ball, lose view of the slow ball, overshoots, and start dropping markers (black dots). (5) Seemingly repelled by the markers, the agent moves back to the slow ball, continuously dropping the markers, and (6) eventually catches it. Adapted from [Chung et al., 2009].

3.1 Task: 2D Pole Balancing

Fig. 7 illustrates the standard 2D pole balancing task. The cart with a pole on top of it is supposed to be moved around while the pole is balanced upright. The whole event occurs within a limited 2D bound. A successful controller for the cart can balance the pole without making it fall, and without going out of the fixed bound. Thus, the pole angle, cart position, and their respective velocities become an important information in determining the cart's motion in the immediate next time step.

3.2 Methods

For this part, we evolved recurrent neural network controllers, as shown in Fig. 8A. The activation equation is the same as Eq. 1, and again, we used the same neuroevolution approach to tune the weights and other parameters in the model. One difference in this model was the inclusion of a facilitating dynamics in the neuronal activation level of the hidden units. Instead of using the H_j value directly, we used the facilitated value

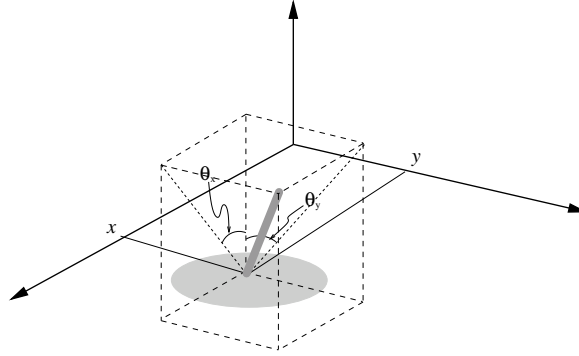


Fig. 7 2D Pole-Balancing Task. The 2D pole-balancing task is illustrated. The cart (gray disk) with an upright pole attached to it must move around on a 2D plane while keeping the pole balanced upright. The cart controller receives the location (x, y) of the cart, the pole angle (θ_x, θ_y) , and their respective velocities as the input, and generates the force in the x and the y direction. Adapted from [Chung et al., 2009].

$$A_j(t) = H_j(t) + r(H_j(t) - A_j(t-1)), \quad (2)$$

where $H_j(t)$ is the hidden unit j 's activation value at time t , $A_j(t)$ the facilitated hidden unit j 's activation value, and r an evolvable facilitation rate parameter (see [Kwon and Choe, 2007] for details). This formulation turned out to have a smoother characteristic, compared to our earlier facilitation dynamics in [Lim and Choe, 2005, Lim and Choe, 2006b, Lim and Choe, 2008].

One key step in this part is to *measure* the predictability in the internal state dynamics. That is, given m past values of a hidden unit H_j (i.e., $\langle H_j(t-1), H_j(t-2), \dots, H_j(t-m) \rangle$), how well can we predict $H_j(t)$. The reason for measuring this is to categorize individuals (evolved controller networks) that have a predictive potential and those that do not, and observe how they evolve. Our expectation is that individuals with more predictable internal state trajectory will have an evolutionary edge, thus opening the road for predictive functions to emerge. In order to have an objective measure, we trained a standard backpropagation network, with the past input vector $\langle H_j(t-1), H_j(t-2), \dots, H_j(t-m) \rangle$ as the input and the current activation value $H_j(t)$ as the target value. Fig. 9 shows a sketch of this approach. With this, internal state trajectories that are smoother and easier to predict (Fig. 10A) will be easier to train, i.e., faster and more accurate, than those that are harder to predict (Fig. 10B). Note that the measured predictability is *not used as a fitness measure*. Predictability is only used as a post-hoc analysis. Again, the reason for measuring the predictability is to see how predictive capability can spontaneously emerge throughout evolution.

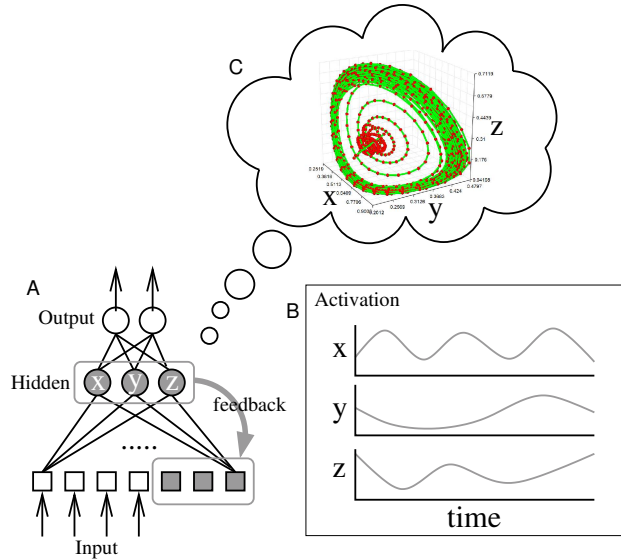


Fig. 8 Cart Controller and Its Internal State. A sketch of the cart controller network is shown. A. The network had 3 hidden units, which was fed back as the context input with a 1-step delay, to implement a recurrent architecture. The network had 8 inputs, each corresponding to the measures listed in Fig. 7. The two output units represents the force to be applied in the x and the y direction, respectively. B. The activity level of the hidden units can be seen as the agent’s internal state, which in this case, can be plotted as a trajectory in 3D (see C). Adapted from [Chung et al., 2009].

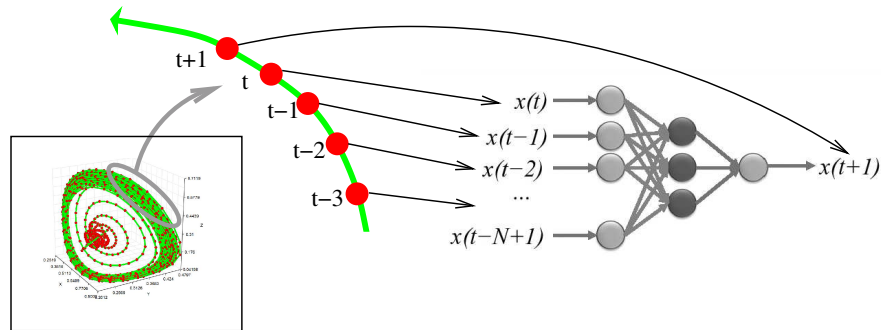


Fig. 9 Measuring Predictability in the Internal State Trajectory. A simple backpropagation network was used to measure the predictability of the internal state trajectory. A sliding window on the trajectory generated a series of input vectors (N past data points) and the target values (the current data point) to construct the training set. Those with a smoother trajectory would be easier to train, with higher accuracy. Adapted from [Chung et al., 2009].

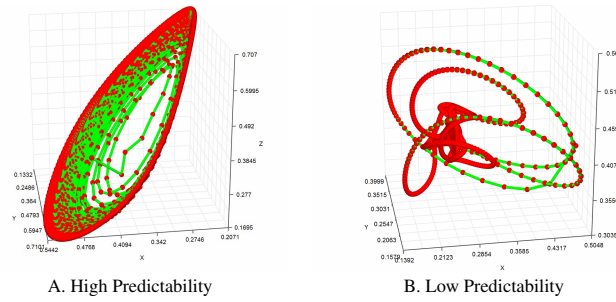


Fig. 10 Internal State Trajectories. Typical internal state trajectories from the hidden units of the controller networks are shown for A. the high predictability group and B. the low predictability group.

3.3 Experiments and results

Fig. 11 shows an overview of our experiment.

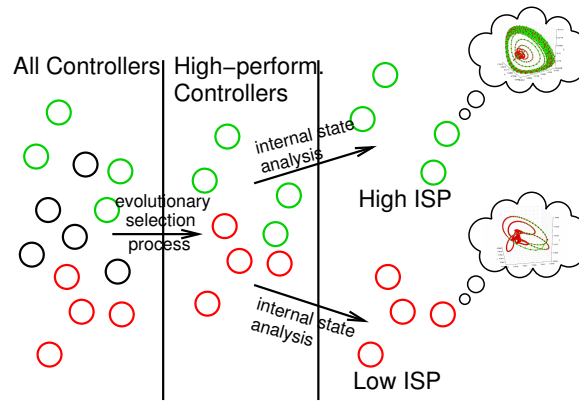


Fig. 11 Overview of the Experiment. An overview of the experiment is shown. First, high-performing individuals (capable of balancing the pole for over 5,000 steps) are collected throughout the generations. Next, the internal state predictability of the selected ones are measured to separate the group into high internal state predictability (High ISP) and low ISP groups. The High and Low ISP groups are subsequently tested in a tougher task. Adapted from [Chung et al., 2009].

The pole balancing problem was set up within a $3 \text{ m} \times 3 \text{ m}$ arena, and the output of the controller exerted force ranging from -10 N to 10 N . The pole was 0.5 m long, and the initial tilt of the pole was set randomly within 0.57° . We used neuroevolution (cf. [Gomez and Miikkulainen, 1998]). Fitness was determined by the number of time steps the controller was able to balance the pole within $\pm 15^\circ$ from the vertical. Crossover was done with probability 0.7 and mutation added perturbation with

a rate of ± 0.3 . The force was applied at a 10 ms interval. The agent was deemed successful if it was able to balance the pole for 5,000 steps.

For the backpropagation predictors, we took internal state trajectories from successful controllers, and generated a training set for supervised learning, using 3,000 data points in the trajectory data. We generated an additional 1,000 inputs for validation. Standard backpropagation was used, with a learning rate of 0.2. For each data point, if the error was within 10% of the actual value, we counted that as correct, and otherwise incorrect. With this, for each trajectory we were able to calculate the predictive accuracy.

We evolved a total of 130 successful individuals, and measured their internal state predictability. Fig. 12 shows the predictability in the 130 top individuals, which exhibits a smooth gradient. Among these, we selected the top 10 and the bottom 10, and further compared their performance. Note that since all 130 had excellent performance, the 20 that are selected in this way by definition have the same level of performance. The trick here is to put those 20 controllers in a harsher environment, by making the pole balancing task harder. We increased the initial pole angle slightly to achieve this. The results are shown in Fig. 13. The results show that the high internal state predictability (high ISP) group outperforms the low internal state predictability (low ISP) group by a large margin. This is a surprising outcome, considering that the two types of networks (high ISP vs. low ISP) had the same level of performance in the task they were initially evolved in. This suggests that certain internal properties, although only internally scrutinizable at one time, can come out as an advantage as the environment changes. One interesting observation we made in our earlier paper [Kwon and Choe, 2008] is that the high performance in the high ISP group is not due to the simpler, smoother internal state trajectory linearly carrying over into simpler, smoother behavior, thus giving it an edge in pole balancing. On the contrary, we found that in many cases, high ISP individuals had complex behavioral trajectories and vice versa (see [Kwon and Choe, 2008] for details). In sum, these results show how *predictive* capabilities could have evolved in evolving neural networks.

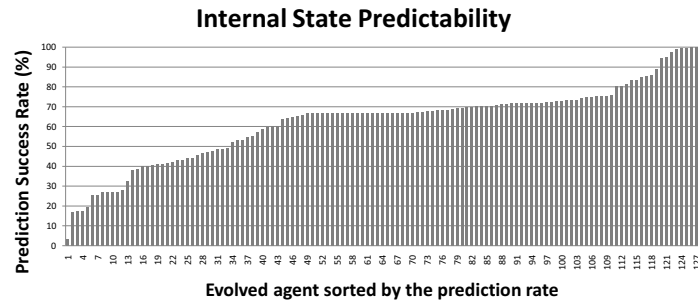


Fig. 12 Internal State Predictability. Internal state predictability of 130 successful controllers are shown, sorted in increasing order. Adapted from our earlier work [Kwon and Choe, 2008, Chung et al., 2009].

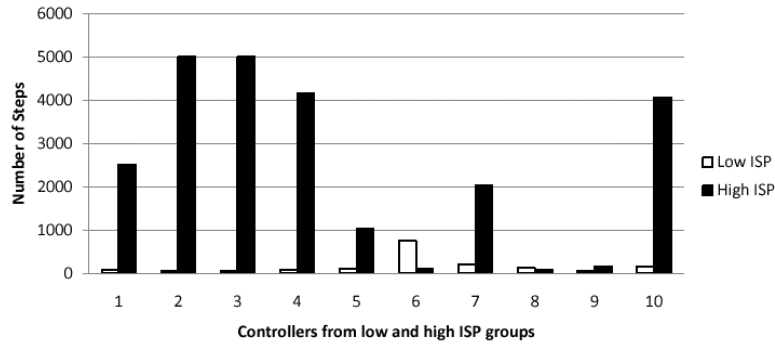


Fig. 13 Pole Balancing Performance. The performance (number of pole balancing steps) of the controller network is shown for the high ISP group (black bars) and the low ISP group (white bars). For this task the initial pole angle was increased to within $(\theta_x, \theta_y) = (0.14^\circ, 0.08^\circ)$. In all cases, the high ISP group does better, in many cases reaching the 5,000 performance mark, while those in the low ISP group show near zero performance. Note that these are new results, albeit being similar to our earlier results reported in [Kwon and Choe, 2008]. Adapted from [Chung et al., 2009].

4 Discussion

The main contribution of this paper is as follows. We showed how recollection and prediction can evolve in neural circuits, thus linking the organism to its past and its future.

Our results in Part I suggest an interesting linkage between external memory and internalized memory (cf. [Clark, 2008, Turvey and Shaw, 1979]). For example, humans and many other animals use external objects or certain substances excreted into the environment as a means for spatial memory (see [Rocha, 1996, Chandrasekharan and Stewart, 2004] for theoretical insights on the benefit of the use of inert matter for cognition). In this case, olfaction (or other forms of chemical sense) serves an important role as the “detector”. (Olfaction is one of the oldest sensory modalities, shared by most living organisms [Hildebrand, 1995, Vanderhaeghen et al., 1997, Mackie, 2003].) This form of spatial memory resides in the environment, thus it can be seen as external memory. On the other hand, in higher animals, spatial memory is also internalized, for example in the hippocampus. Interestingly there are several different clues that suggest an intimate relationship between the olfactory system and the hippocampus. They are located nearby in the brain, and genetically they seem to be closely related ([Machold et al., 2003, Palma et al., 2004] showed that the Sonic Hedgehog gene controls the development of both the hippocampus and the olfactory bulb). Furthermore, neurogenesis is most often observed in the hippocampus and in the olfactory bulb, alluding to a close functional demand [Frisén et al., 1998]. Finally, it is interesting to think of neuro-

modulators [Krichmar, 2008] as a form of internal marker dropping, in the fashion explored in this paper. Fig. 14 summarizes the discussion above. From the left to the right, progressive augmentation to the simplest feedforward neural network that enable memory of the past and facilities for prediction of future events is shown.

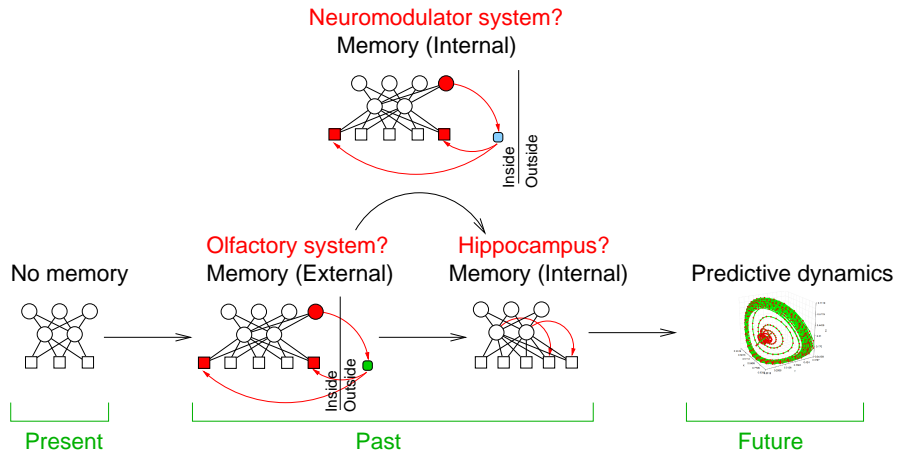


Fig. 14 From the Present to the Past, and Forward to the Future. Initially, only reactive behavior mediated by feedforward networks may have existed (left-most panel, bottom). By evolving external dropper/detector capability while maintaining the feedforward topology, simple memory function may have emerged (second panel, bottom), reminiscent of olfaction. Then, this kind of dropper/detector mechanism could have been internalized, resulting in something in between the dropper/detector (second panel, bottom) and the fully recurrent (third panel, bottom). Neuromodulators could be thought of as such an intermediate stage (top panel). Finally, a full-blown recurrent architecture may have resulted (third panel, bottom). The interesting thing is that the brain seem to have kept all the legacy memory systems (olfaction and neuromodulators), integrating them with the latest development (recurrent wiring).

Prediction (or anticipation) is receiving much attention lately, being perceived as a primary function of the brain [Llinás, 2001, Hawkins and Blakeslee, 2004] (also see [Rosen, 1985] for an earlier discussion on anticipation). Part II of this paper raises interesting points of discussion regarding the origin and role of prediction in brain function. One interesting perspective we bring into this rich on-going discussion about prediction is the possible evolutionary origin of prediction. If there are agents that show the same level of behavioral performance but have different internal properties, why would evolution favor one over the other? That is, certain properties internal to the brain (like high ISP or low ISP) may not be visible to the external processes that drive evolution, and thus may not persist (cf. “philosophical zombies” [Chalmers, 1996]). However, our results show that certain properties can be latent, only to be discovered later on when the changing environment helps bring out the fitness value of those properties. Among these properties we found prediction. Our preliminary results also indicate an important link between delay in the nervous system and the emergence of predictive capabilities: Fig. 15 shows that neu-

ral networks controllers exhibit more predictable behavior as the delay in the input is increased in a 2D pole-balancing task. For more discussion on the relationship between delay, extrapolation, delay compensation, and prediction; and possible neural correlates, see [Lim and Choe, 2005, Lim and Choe, 2006b, Lim and Choe, 2006d, Lim and Choe, 2008, Kwon and Choe, 2009].

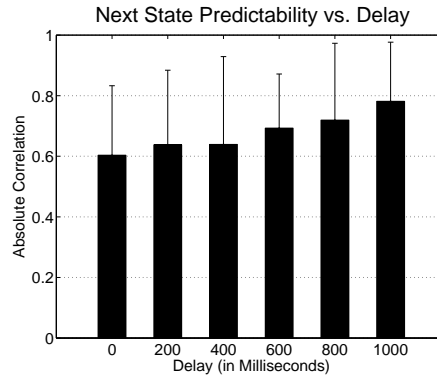


Fig. 15 Effect of Delay on State Predictability. The absolute correlations between hidden state activity at a given moment and true sensory state in the future time step are shown, for neural network controllers trained with different amount of delay in the sensory input lines (error bars indicate standard deviation). This correlation measures the degree of predictability in the internal dynamics. As the delay grows, the predictability also grows. Adapted from [Mann and Choe, 2010] (abstract).

There are several promising future directions. For Part I, recollection, it would be interesting to extend the task domain. One idea is to allow the agent to move in a 2D map, rather than on a straight line. We expect results comparable to those reported here, and also to those in [Balch, 1993]. Furthermore, actually modeling how the external memory became internalized (see Fig. 14, second to third panel) would be an intriguing topic (a hint from the neuromodulation research such as [Krichmar, 2008] could provide the necessary insights). Insights gained from evolving an arbitrary neural network topology may also be helpful [Stanley and Miikkulainen, 2002b, Stanley and Miikkulainen, 2002a]. As for Part II, prediction, it would be helpful if a separate subnetwork can actually be made to evolve to predict the internal state trajectory (as some kind of a monitoring process) and explicitly utilize the information.

5 Conclusion

In this paper we have shown how recollection and prediction could have evolved in neural network controllers embedded in a dynamic environment. Our main results are that recollection could have evolved when primitive feedforward nervous

systems were allowed to drop and detect external markers (such as chemicals), and that prediction could have evolved naturally as the environment changed and thus conferred a competitive edge to those better able to predict. We expect our results to provide unique insights into the emergence of time in neural networks and in the brain: recollection and prediction, past and future.

Acknowledgments

This chapter is a significantly expanded version of [Chung et al., 2009] (broader introduction, and new perspectives and pilot results shown in the discussion).

References

- [Anderson, 1989] Anderson, C. W. (1989). Learning to control an inverted pendulum using neural networks. *IEEE Control Systems Magazine*, 9:31–37.
- [Balch, 1993] Balch, T. (1993). Avoiding the past: a simple but effective strategy for reactive navigation. In *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, pages 678–685. IEEE.
- [Beer, 2000] Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4:91–99.
- [Bishop, 1995] Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Oxford University Press.
- [Chalmers, 1996] Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, New York and Oxford.
- [Chandrasekharan and Stewart, 2004] Chandrasekharan, S. and Stewart, T. (2004). Reactive agents learn to add epistemic structures to the world. In Forbus, K. D., Gentner, D., and Regier, T., editors, *CogSci2004*, Hillsdale, NJ. Lawrence Erlbaum.
- [Choe and Miikkulainen, 2004] Choe, Y. and Miikkulainen, R. (2004). Contour integration and segmentation in a self-organizing map of spiking neurons. *Biological Cybernetics*. In press.
- [Chung and Choe, 2009] Chung, J. R. and Choe, Y. (2009). Emergence of memory-like behavior in reactive agents using external markers. In *Proceedings of the 21st International Conference on Tools with Artificial Intelligence, 2009. ICTAI '09*, pages 404–408.
- [Chung et al., 2009] Chung, J. R., Kwon, J., and Choe, Y. (2009). Evolution of recollection and prediction in neural networks. In *Proceedings of the International Joint Conference on Neural Networks*, pages 571–577, Piscataway, NJ. IEEE Press.
- [Clark, 2008] Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognition*.
- [Conover, 2007] Conover, M. R. (2007). *Predator-Prey Dynamics: The Role of Olfaction*. CRC Press.
- [Elman, 1990] Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14:179–211.
- [Elman, 1991] Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7:195–225.
- [Felleman and Essen, 1991] Felleman, D. J. and Essen, D. C. V. (1991). Distributed hierarchical processing in primate cerebral cortex. *Cerebral Cortex*, 1:1–47.
- [Frisén et al., 1998] Frisé, J., Johansson, C. B., Lothian, C., and Lendahl, U. (1998). Central nervous system stem cells in the embryo and adult. *CMLS, Cellular and Molecular Life Science*, 54:935–945.

- [Gomez and Miikkulainen, 1998] Gomez, F. and Miikkulainen, R. (1998). 2-D pole-balancing with recurrent evolutionary networks. In *Proceedings of the International Conference on Artificial Neural Networks*, pages 425–430, Berlin; New York. Springer-Verlag.
- [Hawkins and Blakeslee, 2004] Hawkins, J. and Blakeslee, S. (2004). *On Intelligence*. Henry Holt and Company, New York, 1st edition.
- [Haykin, 1999] Haykin, S. (1999). *Neural Networks: A Comprehensive Foundation*. Prentice-Hall, Upper Saddle River, NJ, 2nd edition.
- [Hildebrand, 1995] Hildebrand, J. G. (1995). Analysis of chemical signals by nervous systems. *Proceedings of National Academy of Sciences, USA*, 92:67–74.
- [Krichmar, 2008] Krichmar, J. L. (2008). The neuromodulatory system: A framework for survival and adaptive behavior in a challenging world. *Adaptive Behavior*, 16:385–399.
- [Kwon and Choe, 2007] Kwon, J. and Choe, Y. (2007). Enhanced facilitatory neuronal dynamics for delay compensation. In *Proceedings of the International Joint Conference on Neural Networks*, pages 2040–2045, Piscataway, NJ. IEEE Press.
- [Kwon and Choe, 2008] Kwon, J. and Choe, Y. (2008). Internal state predictability as an evolutionary precursor of self-awareness and agency. In *Proceedings of the Seventh International Conference on Development and Learning*, pages 109–114. IEEE.
- [Kwon and Choe, 2009] Kwon, J. and Choe, Y. (2009). Facilitating neural dynamics for delay compensation: A road to predictive neural dynamics? *Neural Networks*, 22:267–276.
- [Lim and Choe, 2005] Lim, H. and Choe, Y. (2005). Facilitatory neural activity compensating for neural delays as a potential cause of the flash-lag effect. In *Proceedings of the International Joint Conference on Neural Networks*, pages 268–273, Piscataway, NJ. IEEE Press.
- [Lim and Choe, 2006a] Lim, H. and Choe, Y. (2006a). Compensating for neural transmission delay using extrapolatory neural activation in evolutionary neural networks. *Neural Information Processing—Letters and Reviews*, 10:147–161.
- [Lim and Choe, 2006b] Lim, H. and Choe, Y. (2006b). Delay compensation through facilitating synapses and STDP: A neural basis for orientation flash-lag effect. In *Proceedings of the International Joint Conference on Neural Networks*, pages 8385–8392, Piscataway, NJ. IEEE Press.
- [Lim and Choe, 2006c] Lim, H. and Choe, Y. (2006c). Facilitating neural dynamics for delay compensation and prediction in evolutionary neural networks. In Keijzer, M., editor, *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation, GECCO-2006*, pages 167–174.
- [Lim and Choe, 2006d] Lim, H. and Choe, Y. (2006d). Facilitating neural dynamics for delay compensation and prediction in evolutionary neural networks. In Keijzer, M., editor, *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation, GECCO-2006*, pages 167–174. **Nominated for Best Paper Award.**
- [Lim and Choe, 2008] Lim, H. and Choe, Y. (2008). Delay compensation through facilitating synapses and its relation to the flash-lag effect. *IEEE Transactions on Neural Networks*, 19:1678–1688.
- [Llinás, 2001] Llinás, R. R. (2001). *I of the Vortex*. MIT Press, Cambridge, MA.
- [Machold et al., 2003] Machold, R., Hayashi, S., Rutlin, M., Muzumdar, M. D., Nery, S., Corbin, J. G., Gritli-Linde, A., Dellovade, T., Porter, J. A., Rubin, S. L., Dudek, H., McMahon, A. P., and Fishell, G. (2003). Sonic hedgehog is required for progenitor cell maintenance in telencephalic stem cell niches. *Neuron*, 39:937–950.
- [Mackie, 2003] Mackie, G. O. (2003). Central circuitry in the jellyfish *aglantha digitale* iv. pathways coordinating feeding behaviour. *The Journal of Experimental Biology*, 206:2487–2505.
- [Mann and Choe, 2010] Mann, T. A. and Choe, Y. (2010). Neural conduction delay forces the emergence of predictive function in an evolving simulation. *BMC Neuroscience*, 11(Suppl 1):P62. Nineteenth Annual Computational Neuroscience Meeting: CNS*2010.
- [Miikkulainen et al., 2005] Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (2005). *Computational Maps in the Visual Cortex*. Springer, Berlin. URL: <http://www.computationalmaps.org>.
- [Palma et al., 2004] Palma, V., Lim, D. A., Dahmane, N., Sánchez, P., Brionne, T. C., Herzberg, C. D., Gitton, Y., Carleton, A., Álvarez Buylla, A., and Altaba, A. R. (2004). Sonic hedgehog controls stem cell behavior in the postnatal and adult brain. *Development*, 132:335–344.

- [Peck et al., 2008] Peck, C., Kozloski, J., Cecchi, G., Hill, S., Schürmann, F., Markram, H., and Rao, R. (2008). Network-related challenges and insights from neuroscience. *Lecture Notes on Computer Science*, 5151:67–78.
- [Rocha, 1996] Rocha, L. M. (1996). Eigenbehavior and symbols. *Systems Research*, 13:371–384.
- [Rosen, 1985] Rosen, R. (1985). *Anticipatory Systems: Philosophical, Mathematical and Methodological Foundations*. Pergamon Press, New York.
- [Stanley and Miikkulainen, 2002a] Stanley, K. O. and Miikkulainen, R. (2002a). Efficient evolution of neural network topologies. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC'02)*, Piscataway, NJ. IEEE. In press.
- [Stanley and Miikkulainen, 2002b] Stanley, K. O. and Miikkulainen, R. (2002b). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127.
- [Tillman et al., 1999] Tillman, J. A., Seybold, S. J., Jurenka, R. A., and Blomquist, G. J. (1999). Insect pheromones - an overview of biosynthesis and endocrine regulation. *Insect Biochemistry and Molecular Biology*, 29:481–514.
- [Turvey and Shaw, 1979] Turvey, M. T. and Shaw, R. (1979). The primacy of perceiving: An ecological reformulation of perception for understanding memory. In Nilsson, L.-G., editor, *Perspectives on Memory Research: Essays in Honor of Uppsala University's 500th Anniversary*, chapter 9, pages 167–222. Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ.
- [Vanderhaeghen et al., 1997] Vanderhaeghen, P., Schurmans, S., Vassart, G., and Parmentier, M. (1997). Specific repertoire of olfactory receptor genes in the male germ cells of several mammalian species. *Genomics*, 39:239–246.
- [von der Malsburg and Buhmann, 1992] von der Malsburg, C. and Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67:233–242.
- [Ward and Ward, 2006] Ward, R. and Ward, R. (2006). 2006 special issue: Cognitive conflict without explicit conflict monitoring in a dynamical agent. *Neural Networks*, 19(9):1430–1436.
- [Wood, 1982] Wood, D. L. (1982). The role of pheromones, kairomones, and allomones in the host selection and colonization behavior of bark beetles. *Annual Review of Entomology*, 27:411–446.