Frequency and Motivational State: Evolutionary Simulations Suggest an Adaptive Function for Network Oscillations

Bram T. Heerebout (B.T.Heerebout@uva.nl)
Psychonomics Unit, Roetersstraat 15
1018 WB Amsterdam, The Netherlands

R. Hans Phaf (R.H.Phaf@uva.nl)
Psychonomics Unit, Roetersstraat 15
1018 WB Amsterdam, The Netherlands

Abstract

Evolutionary simulations of foraging agents, controlled by artificial neural networks, unexpectedly yielded oscillating node activations in the networks. The agents had to navigate a virtual environment to collect food while avoiding predation. Between generations their neural networks were subjected to mutations and crossovers in the connection strengths. The oscillations drastically enhanced the agents' performance, which was due primarily to an increased switching efficacy from approach to avoidance behavior. In this paper we further analyzed networks from the last generation and found that oscillations modulated winner-take-all competition. On average the oscillations had a much higher frequency when an agent was foraging (i.e., in an appetitive state) than when it was trying to escape from a predator (i.e., in an aversive state). We suggest that also in real brains oscillation frequencies correspond with particular types of action preparations.

Keywords: Evolutionary Simulations; Competitive Networks; Oscillations; Affective Modulation

Emerging Oscillations

Insight into the selective pressures that shaped the brain's architecture is probably necessary for understanding its functional mechanisms. The brain, however, does not fossilize, and there is no fossil record of the brain's evolutionary development. Computational simulations may instead provide a suitable approach to investigate the evolutionary development of the brain's functional architectures. In the simulations discussed here, we extended earlier work (den Dulk, Heerebout & Phaf, 2003) on the evolutionary justification of LeDoux's dual-pathway model (1996). Agents were controlled by an artificial network and had to collect plants in a virtual environment while avoiding predation. We investigated the adaptive value of recurrent connections to a new layer in the indirect pathway of this dual-route model (see Figure 1). Surprisingly, the added layer not only led to a major increase in the agents' success, but simultaneously induced oscillations of activations in the network, which were not previously observed by den Dulk et al..

LeDoux (1996) suggested that affective stimuli are processed via two neural pathways running in parallel from the sensory thalamus to the amygdala. One pathway projects directly from the sensory thalamus, the other indirectly by running through the cortex before reaching the amygdala. The direct pathway is faster and is more coarsely grained than the indirect pathway. LeDoux gave an evolutionary justification of the dual-pathway architecture by arguing that in threatening situations the evolutionary cost of a miss exceeds the sum of the costs of the many false alarms produced by the short pathway. When there is more time to process stimuli extensively, the costs of false alarms may even be reduced through inhibition of the direct action tendencies via the long pathway. LeDoux's formulation in terms of evolutionary benefits and costs clearly facilitated the translation into evolutionary simulations.

After one thousand generations of random variation and selection of the fittest, the agents of den Dulk et al. (2003) developed networks with dual-processing dynamics resembling LeDoux's model. These agents were only found in simulations when the scents (i.e., the stimulus properties) of plant and predator were hard to distinguish and the fitness (i.e., "lifetime" multiplied by the total number of plants collected) reflected time pressures in escaping from the predator. Processing in the direct pathway induced fast avoidance of both predator and plant, whereas slower indirect processing led to plant approach and facilitated predator avoidance. The qualitatively different types of processing in the two pathways were further supported by lesion studies (of the separate pathways) in the artificial neural networks.

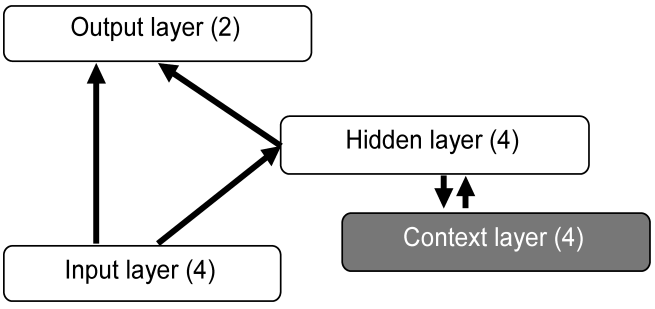

Figure 1: Schematic view of the network's architecture. After the initial simulations a 'context layer' (in gray) was added, allowing for recurrent processing.

We will first briefly describe the general simulation setup and its recent extensions (for a comprehensive specification, see Heerebout & Phaf, 2009). Next, the oscillatory effects of

the added recurrent connections will be presented, and analyzed in terms of competitive processes. Finally, we will argue that positive, appetitive stimuli lead to higher frequency oscillations than negative, aversive stimuli, and that these modulate attentional processing by facilitating disengagement or maintenance of focus, respectively.

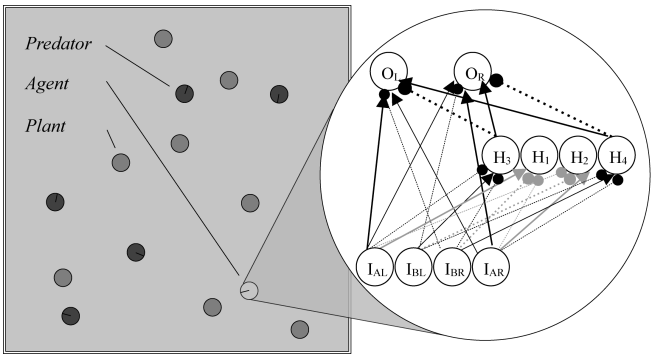

Figure 2: The artificial world with an agent, predators, and plants. The magnification shows an agent's network from the control simulation.

The Simulation Setup in Short

The setup largely followed the simulations of Beer (1990) in which he used Braitenberg vehicles (Braitenberg, 1984) to collect plants through chemotaxis. In our virtual environment the number of plants, ten, was kept constant. When eaten it would instantly be replaced by a new plant at a random location. Furthermore, there were six predators that the agents had to avoid (see Figure 2). Each agent was equipped with an artificial neural network that processed input from the agent's sensors, located at the left and right front side of the agent's round body. The sensors detected the scents emitted by all entities, with quadratically decreasing intensity with respect to the distance from the entity. The network produced output to the motor actuators on both sides of the agent, enabling a tank like propulsion. The strength of the network's connections, which were hereditary, and could only change over generations, determined the behavior of the agent. For instance, if the connections from each sensor relayed the sensors' activation to the motor actuators on the opposite side, the agent turned toward the source of the scent. The sensor closest to the source would detect the scent with greater intensity, resulting in a greater thrust by the motor actuator on the opposite side than the thrust on the sensor's side. In contrast, if the networks had parallel excitatory connections, the agent would turn away from the source. In the simulations of den Dulk et al. (2003) the network processed the input via feed-forward connections, directly to its two output nodes and, in parallel, indirectly via hidden nodes. Plants and predators both emitted mixtures of scent 'a' and scent 'b' and only differed with respect to the ratio of the smell intensities (0.5a and 1.0b for plants vs. 1.0a and 0.5b for predators). By mixing the scents the information processing performed by the agent was made more complex. This allowed for the evolutionary development of an indirect route that processed the sensors' information more elaborate, but slower than the direct route.

The evolutionary process was implemented by a genetic algorithm (Holland, 1975) which was applied to the weights of the agents' network connections. The networks of the first generation were initialized with small random numbers. To ensure that both plants and predators would be relevant to evolutionary development, an agent's fitness was measured as the number of collected plants multiplied by the number of time steps (i.e., the time needed for activation transfer by a node) the agent managed to evade predators. All agents were tested twelve times. A high fitness increased the odds of producing offspring and of remaining in the population for another generation. Agents selected to reproduce were copied and recombined into a pair of offspring. In addition, the offspring connection weights were mutated slightly. The mutations were drawn from a Poisson distribution (the average weight change, in absolute value, was 0.13, with connection weights ranging from -10.0 to 10.0) ensuring that small mutations were more likely to occur than large mutations. Some offspring would outperform others and so increase the probability to pass on their advantageous hereditary connection weights onto their offspring. This cycle was repeated for 10,000 generations, after which the final generation's networks were analyzed.

An Adaptive Function for Oscillations

In the new simulations of Heerebout and Phaf (2009) a four node 'context' layer was added (see Figure 1). Originally, we aimed at the development of a working memory capacity for previously processed stimuli. The activations stemming from the hidden nodes were sent back in the following time step, so that previous stimuli, which were kept active in the context layer, could influence current processing.

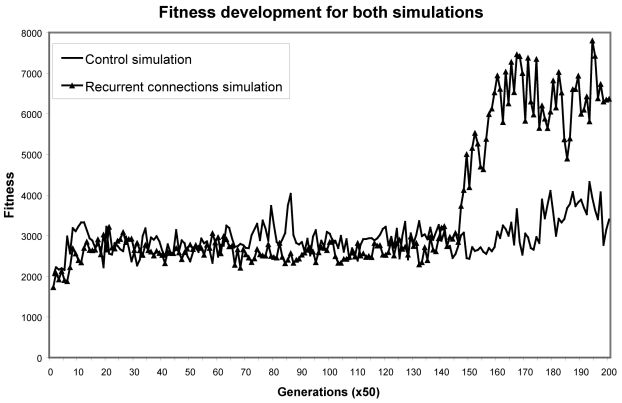

Figure 3: The fitness (in energy units x time steps spent in the environment) development as a function of generation (averaged over clusters of 50).

The fitness development for agents with the recurrent architecture initially resembled that of our control simulation, which replicated the results of den Dulk et al. (2003), but showed a sudden jump in fitness around generation 7500 (see Figure 3). The doubling in fitness (i.e.,

from about 3000 to 6000) indicated a qualitative change in the agents' weight configuration that had a major impact on their performance. Contrary to our initial expectations, the jump did not result from the development of some working memory capacity, but coincided with the emergence of the oscillations in the agents' networks (see Figure 4).

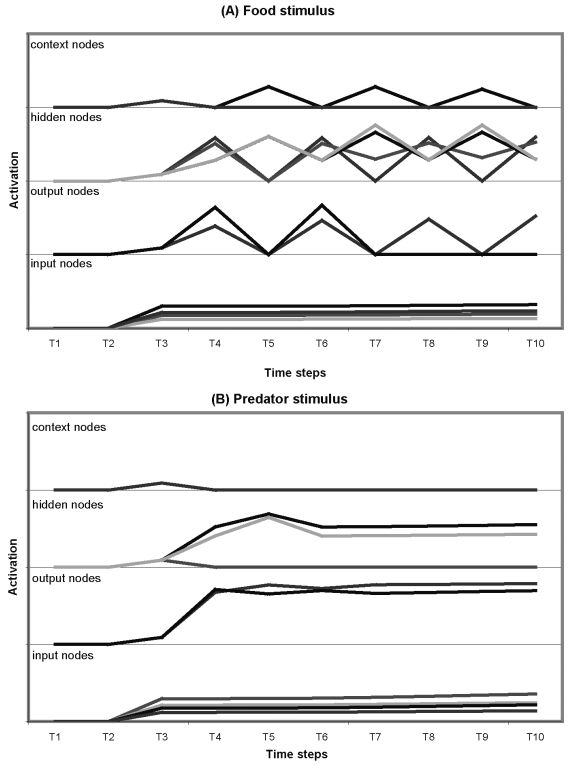

Figure 4: The recorded node activations sorted per layer when the agent was confronted with a plant (A) and a predator (B). The plant stimulus caused oscillating activations. With the predator stimulus the nodes in the context layer were not activated and the output nodes reached steady activation levels.

To analyze these results, Heerebout and Phaf (2009) first averaged the networks of the last-generation agents, which all showed oscillations. The node activations of the resulting agent were recorded when confronted with a plant (see Figure 4a) or with a predator (see Figure 4b), both at 45 degrees left of the agent. With the plant the network showed steady oscillating activations with a period of two time steps. Because the context nodes were not activated by the predator stimulus, avoidance behavior was not modulated by oscillations. Hence, due to these 'food oscillations' the agent zigzagged toward plants, but, in contrast, followed a smooth curve when moving away from predators. Instead of temporarily maintaining activations, the recurrent connections between hidden and context layers acted as oscillation generators.

A core assumption in the evolutionary justification of LeDoux (1996) is that the dual-pathway architecture enables a fast switch in behavior as soon as stimuli that are crucial for survival are detected. Heerebout and Phaf (2009) also analyzed the agents' behavior when a food stimulus was suddenly switched to a predator stimulus. To investigate the agents' switching efficacy its translational and rotational speed were recorded right after a plant was replaced with a predator (switch condition), and these data were compared to when there was no stimulus switch (no-switch condition). The difference in translational speed between switch and no-switch condition was 5.57 times larger for the oscillating than for the non-oscillating agent from the last generation of the control simulation. The difference in rotational speed was 1.72 times larger for the oscillating agent than for the non-oscillating agent. In six out of seven replications oscillating agents emerged with similar increases in switching efficacy. A new type of dual-processing dynamics emerged in these simulations. The direct route was no longer biased towards avoidance, but had a general energizing function. Without direct activation, approach and avoidance were distinguished in the indirect route but not actually performed. These results show that oscillations have similar effects on switching speed as dual processing, as postulated by LeDoux, but that the former enhanced switching efficacy to a much larger degree. The enhanced ability to reorganize behavior after a threatening stimulus appears, thus, seems to account for the doubling of fitness.

Oscillations in Competitive Networks

In this study we derived an abstract connection structure (see Figure 5) from the last-generation agents of Heerebout and Phaf (2009) in order to investigate the underlying mechanisms. Negligible weights were set to zero, and nodes that received no activation due to this elimination of connections were omitted. The oscillatory mechanism was superimposed in this scheme on a competitive network with a winner-take-all function. The enhanced switching ability due to oscillations may solve one of the major problems of competitive networks. Such networks can, through a process of constraint satisfaction, settle in a steady state. Competitive networks do not, however, possess a supplementary mechanism to swiftly escape from the steady state in which they have settled. In order to start a new process of constraint satisfaction, all the activations in the network need to be reset to zero by the modeler, or must have decayed to zero after the stimulus has disappeared. If the winning state consisted of oscillations, it would, according to our hypothesis, be able to switch to another state in the periodically occurring troughs of low, and near-zero activation.

In the symmetrical connection scheme of Figure 5 (for reasons of computational parsimony all networks were symmetrical). Activations were mirrored with respect to input from node I_a and I_b. Activation of N_1, by input node I_a, for instance, triggered activation of N_2. Positive feedback from N_2 to N_1 helped to sustain the activation in both nodes. The excitatory connection to the inhibitory node N_3, however, implemented a "flip-flop" mechanism and caused these activations to oscillate. Interestingly, this type of

recurrent inhibition has also been identified as a neural oscillation generator by Ritz and Sejnowski (1997; see also Dupret, Pleydell-Bouverie & Csicsvari, 2008). In addition, the strong inhibitory crossed connections from N_2 to N_4 and from N_5 to N_1 caused N_1 and N_4 to compete for activation. When the context nodes N_3 and N_6 are removed from the connection scheme, a classical competitive network remains. The N_1 and N_4 nodes compete for activation when input activation is applied, and the most strongly activated node keeps a non-zero activation.

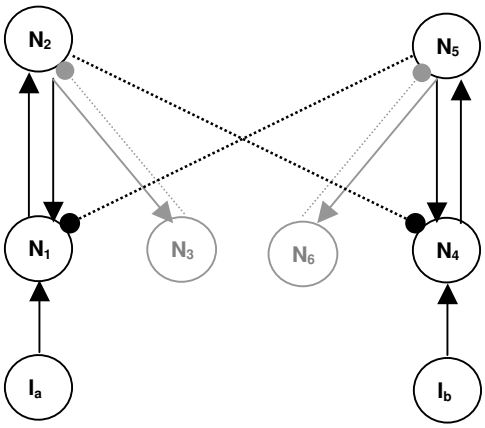

Figure 5: A schematic model of the network configuration that generates the oscillations. The 'I' nodes provide input of scent a and scent b. The solid lines with the arrowheads depict excitatory connections and the dashed lines with the round endings depict inhibiting connections.

The switching efficacy of the oscillatory network is shown by comparing its performance on a switching task with the performance of a non-oscillatory competitive network (see Figure 6). The artificial input consisted of a steadily increasing signal for the input node I_a, which was abruptly replaced by a strong signal for the I_b input node (see upper panel (A) of Figure 6). Panels B and C show the ensuing activations of the nodes N_1, N_2, N_3, and N_4 of the non-oscillatory and oscillatory competitive networks, respectively. The non-oscillatory network is unable to overcome the winner-take-all mechanism and remains stuck. The oscillatory network, however, smoothly switched in a few time steps due to the inhibitory pulses from node N_3 which caused the intermittent deactivation of node N_2 and provided an opportunity for N_4 to become activated.

Stimulus-specific Oscillation Frequency

According to LeDoux (1990) activation via the faster, more coarsely grained, direct pathway is biased to invoke an aversive response, because the costs of hesitation are higher when confronted with a threat than with food. We wondered whether a similar bias towards attending negative stimuli might be found in the oscillating networks. When a predator approaches there is little need to switch attention to other stimuli. A plant stimulus, on the other hand, should induce high frequencies so that the agent can switch its behavior rapidly if it should suddenly encounter a predator. Thus, we would expect to find higher oscillation frequencies for foraging agents than for fleeing agents.

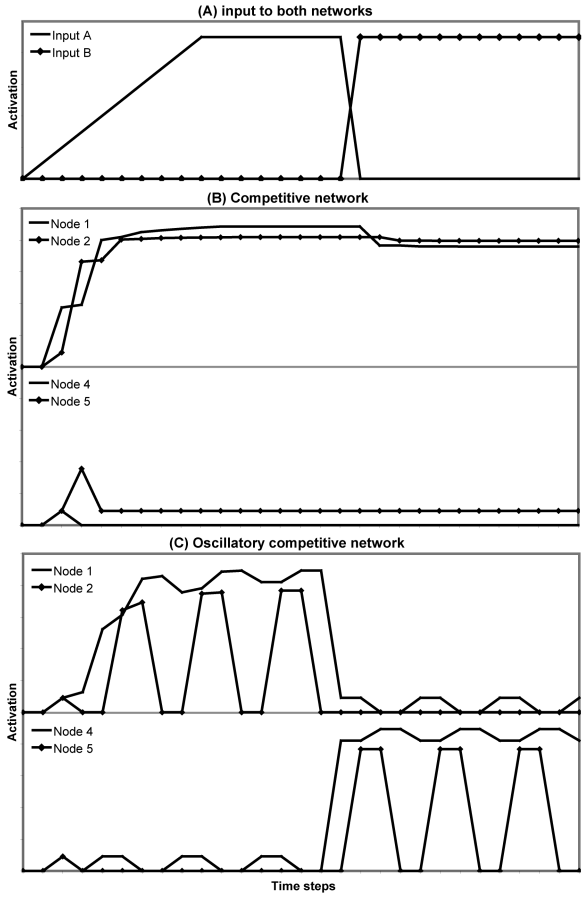

Figure 6: The input activations to both networks (A), activations in the non-oscillatory network (B), and in the oscillatory network (C). The scale of the x-axis (time steps) and the y-axis (activation) is the same for the three panels.

To investigate stimulus-specific oscillations, we here further replicated the simulations with the recurrent connections 32 times. These were exact replications of the previous simulations by Heerebout and Phaf (2009), but new pseudo-random numbers were of course used in the stochastic decisions (e.g., the weight mutations). Although the initial chance of the evolutionary 'discovery' of oscillations may not be very high, we expected the trait would prosper once one individual had acquired it. Therefore, the replications were allowed to continue for 20.000 generations.

We found 438 agents with oscillating networks in the last generation of 27 out of 32 simulations. The oscillation frequencies, both with food and predator stimuli, ranged from zero to 0.5 time step^{-1} (frequency was measured as the inverse of the number of time steps in a complete activation cycle; periods ranged from infinite, 6, 5, 4, 3 to 2 time steps). Some (n = 347) networks oscillated both with plant

and predator, sometimes with different frequencies, whereas 62 networks showed only plant oscillations and 29 networks showed only predator oscillations. The presence of oscillations had a clear effect on the agents' fitness. The oscillating agents (n = 438; *average fitness* = 20,112; *SD* = 17,833 energy x time step) reached significantly higher levels of fitness ($F(1,546) = 7.67$; $p < 0.01$, $\eta_p^2 = 0.014$) than the non-oscillating agents (n = 110; *average fitness* = 15,220; *SD* = 12,493 energy x time step).

Over the total 548 agents the average frequency was significantly higher ($F(1,547) = 25.27$, $p < 0.0001$, $\eta_p^2 = 0.044$) when an agent detected a plant (f = 0.22, *SD* = 0.19 time step^{-1}) than when an agent detected a predator (f = 0.18; *SD* = 0.17 time step^{-1}). The frequency difference between stimulus types supports the conclusion that oscillations increase the agents' ability to switch behaviors, and that the difference arises from the differential environmental demands posed by food and predators.

Neural Oscillations

In contrast to Heerebout and Phaf (2009), Börgers, Epstein and Kopell (2005) reasoned on the basis of their (non-evolutionary) simulations that oscillations suppressed distracters and thus enhanced, inhibition and competition. We argued that the main function of these neural oscillations is to provide an organism with the ability to switch effectively from one type of behavior to another. Processing in an oscillating network is never completely fixed to one type and it does not take much change of input to tip over to another type. Our simulations show that this ability is highly adaptive in realistic environmental conditions when sudden switches are sometimes required. Competitive neural processing has been hypothesized to play a crucial role in attention (i.e., biased competition; Duncan, 1996; see also Phaf, van der Heijden, & Hudson, 1990). Competition and oscillation appear to have complementary functions with regard to attention. The main function of competition may be to direct attention towards a particular stimulus object, whereas oscillations come into play when attention has to be disengaged and redirected to a new stimulus object.

Due to the high level of abstraction of our model, the oscillations cannot be directly be associated with specific frequency bands found in biological neural networks. However, the proposed function of the emerging oscillations may provide important clues. Several roles have been proposed for the oscillations in different areas of the mammalian brain. These possible functions include the binding of cell assemblies (Gray, König, Engel & Singer, 1989), input selection (Hutcheon & Yarom, 2000), representation by phase information (Buzsáki & Draghun, 2004), selective amplification (Lengyel, Huhn & Erdi, 2005) and sequence learning (Ulanovsky & Moss, 2007). The enhanced capability to organize behaviour in response to specific stimuli of the agents in our simulations corresponds with the suggestion of Schaefer, Angelo, Spors, and Margrie (2006) who argued that oscillations serve to discriminate between stimuli. According to Fries, Nikolic, and Singer (2007) stimuli can be distinguished because the interaction between excitatory pyramidal neurons and inhibitory interneurons results in a time critical competition. Only the few pyramidal cells that are able to spike sufficiently early in the oscillation cycle, are able to spike at all. In this manner the weaker cells are suppressed (see also Brody & Hopfield, 2003, who showed that simple oscillating network models implemented sensory segmentation). Fries et al. (2007) tie this winner-take-all mechanism specifically to oscillations with frequencies in the gamma band (20-70 Hz). Because the mechanism closely resembles the flip-flop mechanism we encountered in the agents' networks both in structure and in function, gamma oscillations thus seem a likely candidate for the oscillations we obtained. Similar to the stimulus dependency of the frequencies in our simulations, the frequency of actual gamma oscillations also appears to depend on the state of the organism (Dupret et al., 2008).

The sensorimotor role of gamma oscillations is, moreover supported by a study of Rougeul-Buser and Buser (1997). They observed 40 Hz oscillations in the motor, parietal, and visual cortices when a cat was waiting in front of a hole in the wall from which at times a mouse could pop out and then quickly disappear. When the cat was simply watching the mouse in a perspex box, however, lower frequencies of 10-15 Hz showed up. We would argue that the cat in the former situation has prepared itself to quickly switch from immobility to vigorous attack, whereas in the latter situation, where the cat cannot reach the mouse, both the need for this preparation and the corresponding oscillation frequency may be lower.

Frequency Reflects Affective and Motivational State

Specific stimuli have acquired their positive value, because they generally raised fitness levels throughout evolutionary history (Johnston, 2003). Inversely, negative stimuli are characterized by their recurrent potential to lower fitness. Also appetitive, approach, and aversive, avoidance, tendencies have likely developed in evolution to maximize fitness benefits and to minimize fitness costs, respectively. Such action tendencies at least emerged from the undifferentiated initial networks in our simulations as a response to fitness-relevant stimuli. Fitness costs are minimized by a focusing of attention on the negative stimulus and low distractibility. Fitness benefits are maximized by an approach to all positive stimuli and a high level of distractibility by negative stimuli. The level of distractibility in our simulations corresponded to oscillation frequency. On average, oscillations had a higher frequency when the agent was foraging than when it was trying to escape from a predator. We suggest that a specific oscillation frequency sets the network in a suitable attentional mode to deal with affectively valenced stimuli.

The neural architectures resulting from these simulations, and the mechanisms they implement (i.e., competition and

oscillations), are also likely to have evolved in biological neural networks. The genetic algorithm produced them under environmental conditions that also applied to much of our evolutionary history. Empirical research, moreover, supports the role of competition (e.g., Duncan, 1996) and of neural oscillations (e.g., Rougeul-Buser & Buser, 1997) in attention. Recent advances in human experimental psychology have also revealed similar modulatory influences of motivational (e.g., Förster, Friedman, Özelsel & Denzler, 2006) and affective state (e.g., Gasper & Clore, 2002; Rowe, Hirsh & Anderson, 2007) on the breadth of attention. The present simulation results integrate these different viewpoints by suggesting that high oscillation frequencies are elicited by positive, appetitive, stimuli and lower frequencies by negative, aversive, stimuli and that the latter help to sustain attentional focus, whereas the former facilitate attentional disengagement.

Conclusion

This study illustrates the productive capacity of evolutionary simulations. The model was not built by the modeler but by the simulated evolutionary process. The solutions produced by the genetic algorithms are unpredictable, but generally have much higher fitness levels than the initial configurations. Novel models with unexpected mechanisms and functions may emerge from these undirected optimization procedures. The simulations thus contradict the often heard prejudice that with neural modeling you get what you put into the model. Evolutionary computation adds automatic model production to the tools of the model builder, and ensures the emergence of the most biologically plausible models within the available search space.

References

Beer, R.D. (1990). Intelligence as adaptive behavior: an experiment in computational neuroethology. *Artificial Intelligence, 72,* 173-215.

Börgers, Epstein and Kopell (2005). Background gamma rhythmicity and attention in cortical local circuits: A computational study. *Proceedings of the National Academy of Sciences, 102,* 7002-7007

Brody, C.D., & Hopfield, J.J. (2003). Simple Networks for Spike-Timing-Based Computation, with Application to Olfactory Processing. *Neuron, 37*, 843-852.

Braitenberg, V. (1984). Vehicles: Experiments in Synthetic Psychology. Cambridge, MA: MIT Press

Buzsáki, G., & Draghun, A. (2004). Neuronal Oscillations in Cortical Networks. *Science, 304,* 1926-1929.

den Dulk, P., Heerebout, B.T., & Phaf, R.H. (2003). A computational study into the evolution of dual-route dynamics for affective processing. *Journal of Cognitive Neuroscience, 15,* 194-208.

Duncan, J. (1996). Cooperating brain systems in selective perception and action. In T. Inui & J.L. McClelland (Eds.). *Attention and Performance, Vol. XVI. Information integration in perception and communication* (Pp. 549-578). Cambridge, MA: MIT Press.

Dupret, D., Pleydell-Bouverie, B., & Csicsvari, J. (2008). Inhibitory Interneurons and Network Oscillations. *Proceedings of the National Academy of Sciences, 105,* 18079-18080

Förster, J., Friedman, R.S., Özelsel, A., & Denzler, M. (2006). Enactment of approach and avoidance behavior influences the scope of perceptual and conceptual attention. *Journal of experimental Social Psychology, 42*, 133-146.

Fries, P., Nikolic, D., & Singer, W. (2007). The Gamma Cycle. *Trends in Neuroscience 30,* 309-316.

Gasper, K., & Clore, G. L. (2002). Attending to the big picture: Mood and global versus local processing of visual information. *Psychological Science*, *13*, 34-40.

Gray, C.M., König, P., Engel, A.K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature, 338*, 334-337.

Heerebout, B.T. & Phaf, R.H. (2009). Emergent Oscillations in Evolutionary Simulations: Oscillating Networks Increase Switching Efficacy. *Journal of Cognitive Neuroscience.*

Holland, J.H. (1975). *Adaptation in natural and artificial systems.* Ann Arbor: University of Michigan Press.

Hutcheon, B., & Yarom, Y. (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends in Neurosciences, 23*, 216-222.

Johnston, V.S. (2003). The origin and function of pleasure. *Cognition and Emotion, 17*, 167-179.

LeDoux, J.E. (1996). *The emotional brain*. New York: Simon & Schuster.

Lengyel, M., Huhn, Z., & Érdi, P. (2005). Computational theories on the function of theta oscillations. *Biological Cybernetics, 92,* 393-408

Phaf, R.H., Van der Heijden, A.H.C., & Hudson, P.T.W. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology, 22*, 273-341.

Ritz, R., & Sejnowski, T.J. (1997). Synchronous oscillatory activity in sensory systems; new vistas on mechanisms. *Current Opinion in Neurobiology, 7*, 536-546.

Rougeul-Buser, A., & Buser, P. (1997). Rhythms in the alpha band in cats and their behavioural correlates. *International Journal of Psychophysiology, 26*, 191-203.

Rowe, G., Hirsh, J.B., & Anderson, A. K. (2007). Positive affect increases the breadth of attentional selection. *Proceedings of the National Academy of Sciences, 104,* 383-388

Schaefer, A.T., Angelo, K., Spors, H., & Margrie, T.W. (2006). Neuronal Oscillations Enhance Stimulus Discrimination by Ensuring Action Potential Precision, *Public Library of Science, 4,* 1010-1024.

Ulanovsky, N., & Moss, C.F. (2007). Hippocampal cellular and network activity in freely-moving echolocating bats. *Nature Neuroscience, 10*, 224-233