

# Good vibrations switch attention: An affective function for network oscillations in evolutionary simulations

BRAM T. HEEREBOUT AND R. HANS PHAF

*University of Amsterdam, Amsterdam, The Netherlands*

In the present study, a new hypothesis on the neural mechanisms linking affect to attention was brought forward by evolutionary simulations on agents navigating a virtual environment while collecting food and avoiding predation. The connection strengths between nodes in the networks controlling the agents were subjected to random variation, and the fittest agents were selected for reproduction. Unexpectedly, oscillations of node activations emerged, which drastically enhanced the agent's fitness. We analyzed the mechanisms involved in the modulation of attention and found that oscillations acted on competitive networks. Response selection depended on the connection structure, but the speed and efficacy of switching between selections was modulated by oscillation frequency. The main focus of the present study was the differential emergence of stimulus-specific oscillation frequencies. Oscillations had a higher frequency in an appetitive motivational state than in an aversive state. We suggest that oscillations in biological networks also mediate the affective modulation of attention.

The method of computationally simulating evolutionary processes provides a unique opportunity for the automated development of models and hypotheses on cognitive and affective processes and their underlying neural mechanisms. The role of the modeler is limited to setting up the evolutionary selection procedure (e.g., a genetic algorithm; Holland, 1975), the initial state, and the environmental conditions. This has several advantages, of which the opportunity for innovation is not the least. Novel models with mechanisms and functions that have not been previously considered may emerge from these undirected optimization procedures.

In previous simulation work (Heerebout & Phaf, 2010), we serendipitously discovered "artificial neural oscillations" in agents that inhabited a virtual environment (see Figure 1). The oscillations were an emergent property, not intentionally built in, and not even thought of beforehand, which proved to be highly adaptive. These simulations extended our investigation of LeDoux's (1996) evolutionary justification of his dual-pathway model for the processing of emotional stimuli (den Dulk, Heerebout, & Phaf, 2003). The agents, which were controlled by artificial neural networks, increased their chances of reproduction and survival by collecting food while avoiding predators. Analyses of the oscillating agents showed that the evolutionary advantage was conferred by an enhanced capacity for attentional switching when in an oscillatory mode.

To test the switching efficacy of an oscillating agent, it was compared with a nonoscillating agent from a control simulation with a simpler network. Both agents responded to a plant that was placed in front of them at an angle

of 45° to its left, which would, as the agent approached the plant, suddenly be replaced by a predator. The agents' phenotypic behavior was described in terms of speed of movement and rotational speed, and revealed distinct approach and avoidance behaviors. The oscillating agent would first "cautiously" approach the plant following a swerving, "zig-zag" trajectory, and then, as it detected the predator, would make a sharp turn (at 0.042° per time step) and accelerate strongly (the speed increased 87% in the first 10 time steps). The nonoscillatory agent kept a more constant speed. When it detected a predator, it accelerated only slightly (its speed increased only 13% in the first 10 time steps) while turning away (with 0.034° per time step). In addition, this appeared to be a robust finding, because fast switching oscillating agents turned up in five out of the seven replications.

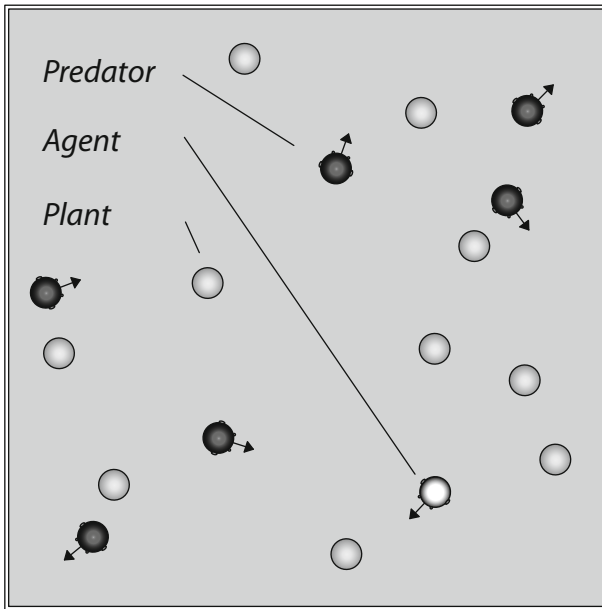
Although the behavioral consequences of the oscillations were evident, the underlying mechanism demanded further analysis. Moreover, the networks from the replicated simulations exhibited oscillations at different frequencies. If the oscillations are adaptive because they enhance the ability to quickly switch between competing behaviors, then the evolutionary demands imposed by the two types of stimuli might induce different oscillation frequencies. After all, increased flexibility may be advantageous when searching for food, but when fleeing from a predator, foraging constitutes a distraction with potentially fatal consequences.

The present study aimed to address these two issues. First, we will briefly reiterate the method of evolutionary simulations. This will be followed by the discussion of a

---

B. T. Heerebout, [b.t.heerebout@uva.nl](mailto:b.t.heerebout@uva.nl)

---



**Figure 1.** Graphic depiction of the virtual environment with one agent, six predators, and 10 plants. The arrows pointing away from the agent and the predators indicate the direction of movement. The environment is “torus” shaped (i.e., if a predator or agent leaves the environment on one side of the square, it immediately reappears on the opposite side).

simplified connection scheme, newly abstracted from previously evolved agents with oscillations in their networks. This simple connection scheme operated according to a winner-take-all competitive process, of which the switching speed was modulated by oscillation frequency. Similar competitive processes are often assumed to be involved with attentional selection in biological neural networks. More specifically, in our evolved networks, the competitive mechanism seems to be responsible for selection for action (cf. Allport, 1989). We further investigated how the oscillations are able to modulate selection speed without influencing the specific outcome of this selection process. In addition, the simulation was replicated another 25 times to test whether positive stimuli indeed yielded faster oscillations than negative stimuli. Finally, the results from the evolutionary simulations were combined with empirical data on attentional selection and affect into the hypothesis that also in biological neural networks, positive affect is associated with higher neural oscillation frequencies and with more efficient attentional switching than is negative affect, which conversely fosters the maintenance of attentional focus.

### Evolutionary Simulations

Computer simulations allow researchers to perform experiments that would otherwise be impossible because of temporary, spatial, monetary, or ethical constraints. Our evolutionary history is not directly observable; therefore, empirically investigating evolutionary explanations is a difficult task. To address this problem with computational means, we applied a genetic algorithm (GA; Holland, 1975) to agents navigating a virtual environment while

collecting food and avoiding predators. GAs are traditionally used as an optimization method, inspired by natural selection. Starting from an initial population that consists of a random combination of gene values, the search for optimal solutions proceeds through selection and reproduction. Fitness is defined as some measure of how well the parameter set solves the problem. Thus, better solutions get a higher chance of survival and reproduction. When a specific solution gets to be reproduced, its parameters are subject to mutations and crossovers. Mutation may improve the performance of the population by occasionally suggesting a new partial solution. With crossover, two individual solutions selected for reproduction are recombined. The new solutions replace the less fit solutions in a population. From generation to generation, this leads to a higher overall fitness.

The evolutionary-inspired optimization technique has been fruitfully applied to real-world problems. For example, the Aircraft Aerodynamics and Design Group at Stanford employed GAs for the optimization of aircraft design (Gage, 1995). A novel wing design was sought that minimized drag. The initial population resembled simple solutions, encoding the wing as a collection of tapered and twisted elements. Next, the solutions were slightly mutated, adding or removing a few random elements. The drag for each solution was then calculated using a vortex lattice analysis. This fitness function determined which designs were selected for reproduction (i.e., which were copied with slight alterations). Multiple repetitions of the process quickly led to the discovering of winglets—small vertical additions at the wingtips (a standard CO<sub>2</sub> emission-reducing measure on passenger jets). Later on, the algorithm found even more efficient “C”-shaped wingtips. This design concept was patented, and its application to future aircraft models is under investigation. Although the fitness function guided the evolutionary process, the random variations allowed for the emergence of a fundamental, innovative design concept.

Because of their innovative power, GAs stand in stark contrast with more conventional cognitive modeling approaches. Connectionist models, but also models constructed with a symbolic formalism, such as ACT-R or SOAR, are usually analytical models that are constructed to explain empirical data, formalizing some cognitive theory. The evolutionary simulations, consisting of a GA applied to the agents’ neural networks, on the other hand, automatically synthesize models that produce the goal behavior, as specified in the fitness function. It should be emphasized that the modeler does not construct a model, but merely sets up conditions that are conducive to the evolution of the model that performs a particular function. When considering different models that explain behavioral data equally well, the model that has been shown to be more likely to develop in evolutionary simulations should be preferred above the others. Moreover, because of the innovative power of the GA, even new models and hypotheses that have not been previously considered by neuroscientists may emerge.

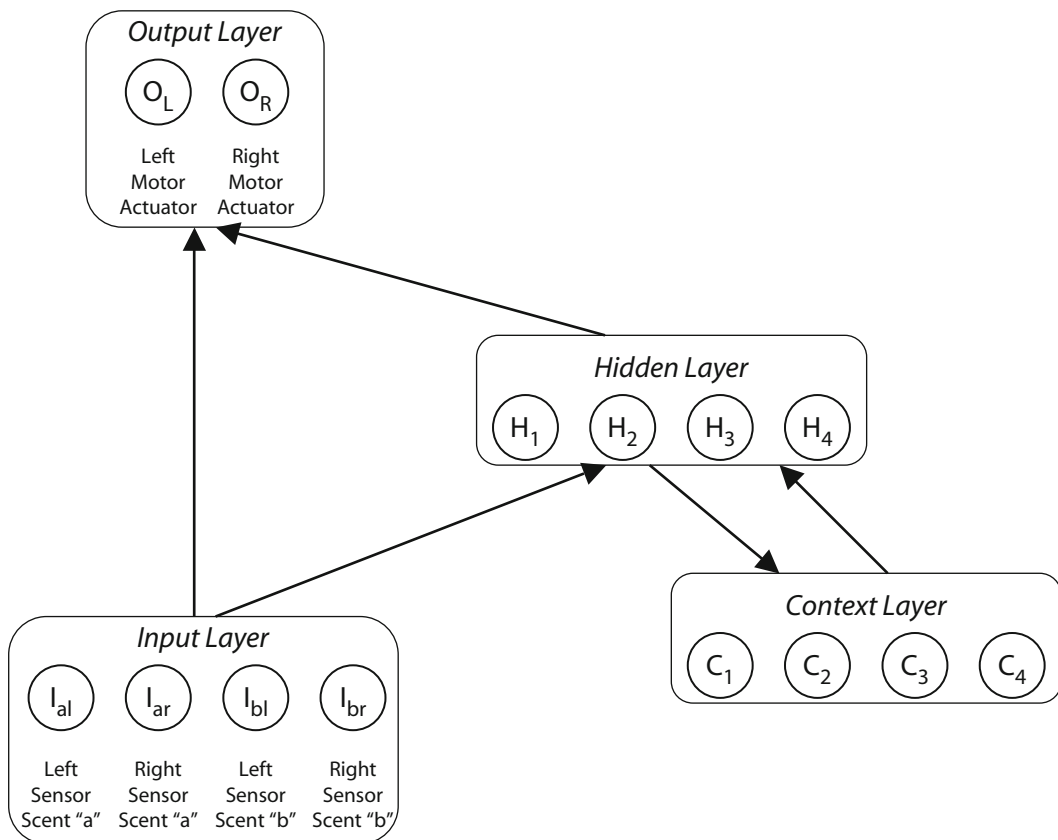
The conventional neural network modeling of affective processes (Armony, Servan-Schreiber, Cohen, & LeDoux, 1995; den Dulk, Rokers, & Phaf, 1997; Phaf, den Dulk,

Tijsseling, & Lebert, 2001) has been strongly inspired by the neurobiological work of LeDoux (1986, 1996). According to LeDoux, the mammalian brain processes fear stimuli via two parallel pathways that run from the sensory thalamus to the amygdala. One pathway projects directly from the sensory thalamus to the lateral nucleus of the amygdala; the other projects indirectly, first passing through the cortex before reaching the amygdala. LeDoux suggested that these pathways might have evolved to functionally supplement each other. The direct pathway can produce a fast fear response consisting of autonomic, endocrine, and motor (e.g., avoidance) reactions to a potentially threatening stimulus. The slower, indirect pathway processes the stimulus more extensively and may strengthen the initial response, if appropriate, in the case of an actual threat, or it may inhibit the direct reaction in the case of a false alarm. LeDoux (1996) argued that the different functions of the two pathways are adaptive, because the evolutionary cost of a miss exceeds the total costs of the many false alarms produced by the direct pathway.

The straightforward formulation in terms of evolutionary benefits and costs by LeDoux (1996) facilitated the translation by den Dulk et al. (2003) into evolutionary simulations, which indeed resulted in the emergence of dual-processing networks. The setup largely followed the

simulations of Beer (1990), in which Braitenberg (1984) vehicles (i.e., the agents) were used to collect plants through chemotaxis. For the present study, the same software for the simulations as in the study of den Dulk et al. (2003) was used, but the networks of the agents were extended with an extra recurrent layer (see Figure 2) to investigate the adaptive value of such recurrent connections.

The present study completely followed the original design choices made by den Dulk et al. (2003). The oscillations thus resulted from a design process by the GA, but not from the part of the modelers. Relative to the methods of this previous study, no parameter adjustments were made, and the space of possible fitness functions was not explored in any manner to arrive at these results. Insofar as the range of parameter values was explored by den Dulk et al. (2003), it mostly proved of little influence when the parameters remained within reasonable bounds. For instance, whether the number of plants that were simultaneously present in the environment was four, six, or eight, the agents would always evolve to approach plants. The choice for the number of six plants, which was also used in our simulations, was therefore arbitrary. After initialization, the simulations followed a simple three-stage procedure, the gist of which will be described briefly in the present article, but more details can be found in Appendixes A, B, and C.



**Figure 2.** The four layers of the network architecture used in the simulations. The layers were potentially fully connected (i.e., all connections in the indicated direction could evolve to nonzero values in the course of the simulations). The arrows depict the direction of the connections. The context layer was added by Heerebout and Phaf (2010) to allow for recurrent processing.

During initialization, a population of agents was created with all of the agents' network connection weights set to zero. The potential network connections, which could evolve to nonzero values, specified a simple multilayered architecture with an input layer of four nodes projecting to two output nodes and to four hidden nodes. The hidden nodes projected to the output nodes and also to another layer of four hidden nodes, which recurrently projected back to the former hidden nodes. The connections between the two layers of hidden nodes could develop independently (i.e., they did not have symmetrical weight values). One time step was defined as the duration of the calculation of a node's activation. Therefore, the signals that a node received at time  $t$  determined its activation at time  $t$ , but this activation would be propagated over its sending connection at  $t+1$  (see also Appendix D).

In the first stage, the algorithm would select agents to create offspring until the population size had tripled. With the initial population, the selection is arbitrary, but with subsequent generations, the agent's fitness (measured in the second stage) determined the agent's probability of being selected (see Appendix A for details on the selection and reproduction process). In the second stage, the fitness of each agent would be determined individually in 12 different, random virtual environments, which always contained 10 plants and six predators. The movement of the predators was governed by networks similar to those of the agents. These networks, however, were preconfigured to let the predators approach the agents. This configuration would not change (i.e., there was no "arms race" between prey and predator).

The agents had olfactory sensors that were located on the front left and right of their round bodies, which picked up the smells emitted by the plants and predators. The agents could utilize the lateral difference in smell intensity to estimate the direction of the source. The sensors relayed the detected intensities to the input nodes of the agent's network. After processing, the output nodes' activation would drive the agents' left and right motor actuators, giving the agents tank-like propulsion. Appendix B describes the details of the smell distribution.

If an agent collided with a plant, the food was considered eaten, and one unit would be added to the agent's energy level. The plant would subsequently reappear at a randomly selected location. Movement by the agent led to a reduction of energy (the exact energy cost was calculated from the force that an agent exerted on the environment; see also Appendix B for details). An agent would be removed from the environment if it collided with a predator, if its energy was depleted, or if the maximum of 10,000 time steps had passed.

When the agent was removed, its fitness was calculated as its energy level multiplied by the total number of time steps spent in all environments, divided by the number of times it had been tested. In the third stage, agents were removed from the population, whereby the lowest scoring agents had the highest probability of being removed. Appendix C explains how the agents' fitness was used to calculate this probability.

After cycling through the aforementioned stages 10,000 times, den Dulk et al. (2003) obtained weight configura-

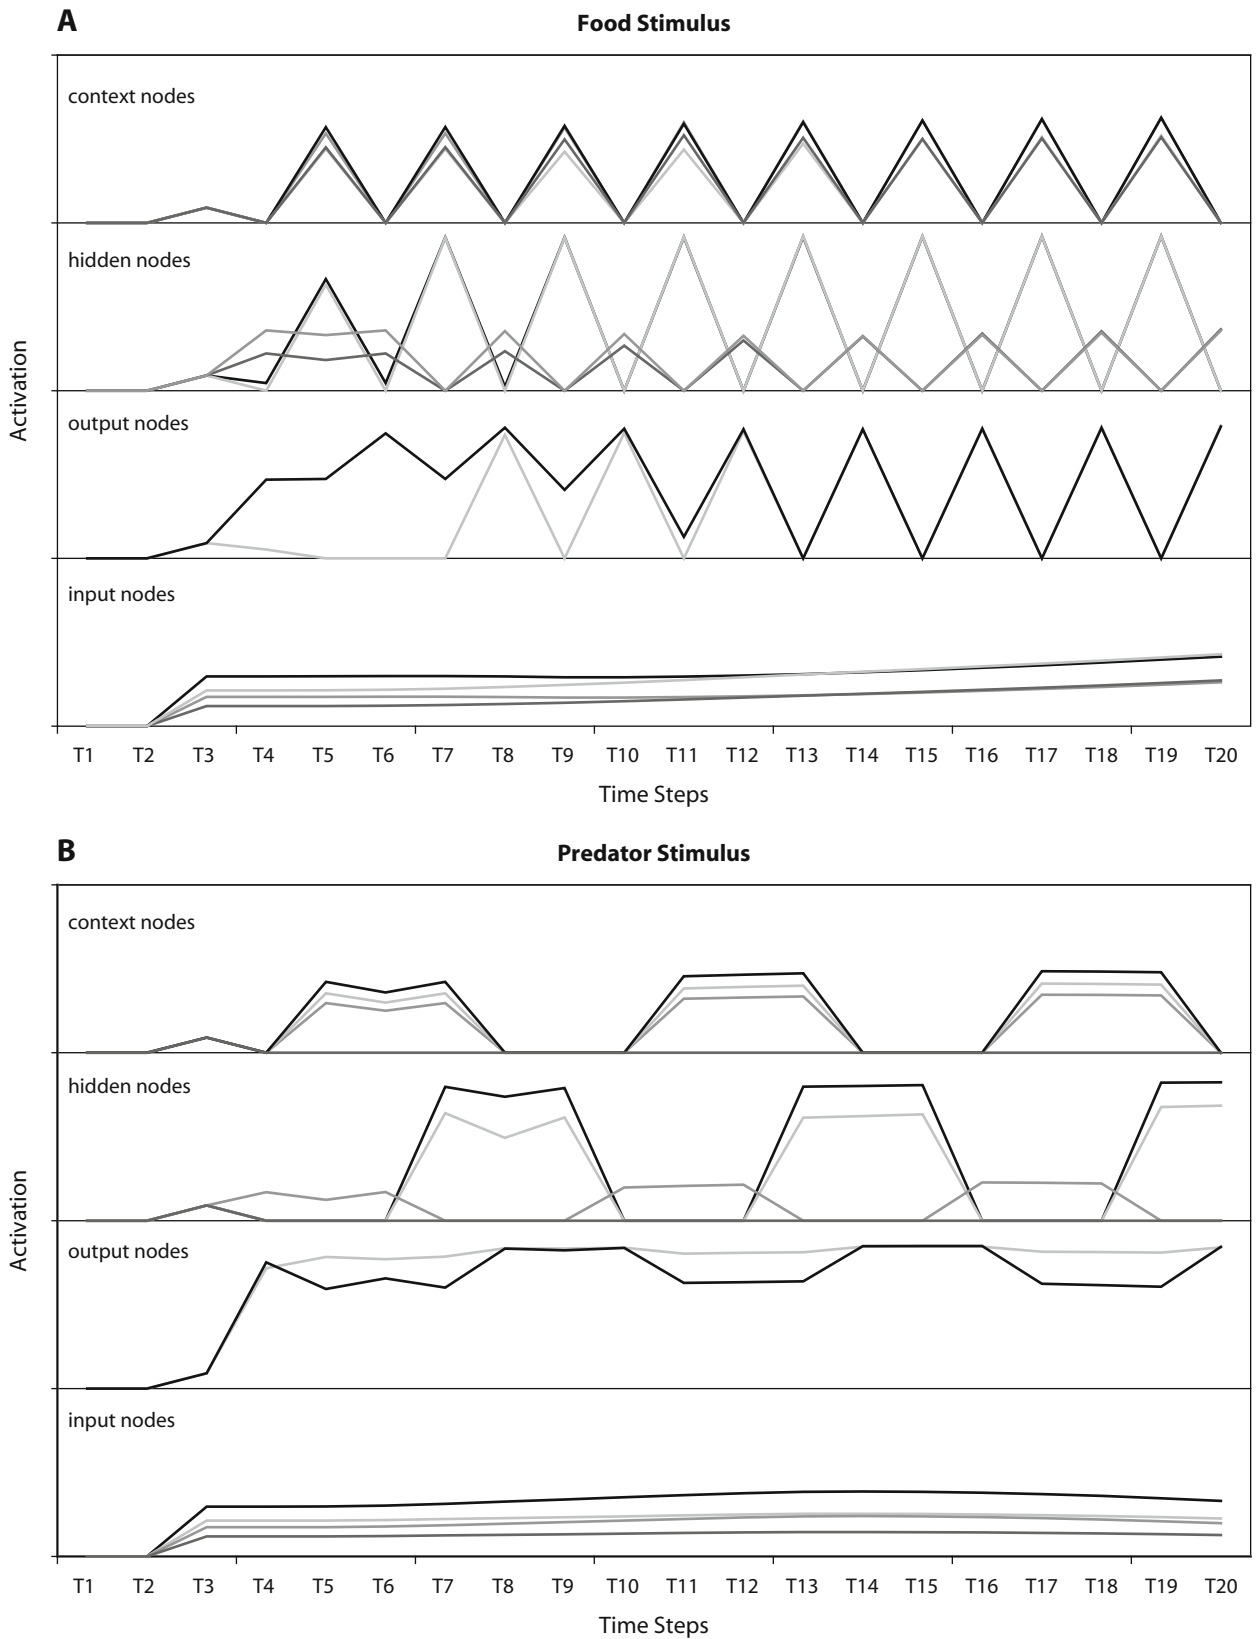
tions that had a functionality similar to LeDoux's (1996) dual-pathway model. This architecture was found only in simulations in which (1) the plants and predators were hard to distinguish, (2) processing via a hidden layer took more time than direct processing, and (3) both plant and predator were relevant to the agent's fitness. The qualitatively different types of processing in the two pathways were further supported by lesion studies of the separate pathways (i.e., setting all of the weights in one path to zero) in the artificial neural networks. These findings support LeDoux's (1996) hypothesis that the dual-pathway model is adaptive and that it can evolve over time through small mutations, when there is an exchange between time pressures and need for accuracy.

### Oscillations in Competitive Networks

Heerebout and Phaf (2010) subsequently extended the simulation setup to investigate the adaptive value of recurrent connections via a new hidden layer in the indirect pathway. The recurrent processing could hypothetically lead to a working memory capacity for previously processed stimuli by allowing activations from previous time steps to contribute to the present network state. Unexpectedly, the added layer caused the networks to oscillate (see, e.g., Figure 3). These oscillations clearly had an adaptive function, more than doubling the agents' performance relative to that of agents from the control simulation, without recurrent connections.

In comparison with those in LeDoux's (1996) dual-pathway model, an alternative type of dual-processing dynamics emerged in the present simulations. The direct route was no longer biased toward avoidance, but had a general energizing function. After lesioning the direct route, approach and avoidance tendencies could be distinguished in the indirect route, but the actions were not actually performed. For the action tendency to actually develop into an action, "arousal" from the direct route was needed. The nodes in the indirect route entered into a competitive process whenever input was delivered to them. The selection of specific approach or avoidance action tendencies was determined by the evolved weight configuration. The oscillations added speed to this selection to a similar extent as was added by preparation through the direct route in the classical dual-pathway model.


In classical competitive network models (see, e.g., McClelland & Rumelhart, 1981; Rumelhart & Zipser, 1985), input causes the nodes in a network to compete for activation, and a steady state can mostly be reached through a process of multiple constraint satisfaction, in which the connection weights constitute the constraints. Although networks of this kind have been applied to model visual attention (see, e.g., Phaf, Van der Heijden, & Hudson, 1990), the inability of competitive networks to swiftly leave a winning state remains a major problem. Switching in these networks usually requires a period of activation decay or even manual intervention by the modeler. The oscillations might suggest a solution to this switching problem by providing a supplementary mechanism to reset the competition. With oscillating competitive networks, the winning state could easily shift to another set of nodes



**Figure 3.** An example of oscillating activations in the nodes of the network of a last-generation network. The nodes' activations are shown over time, separated per layer. The activations shown in (A) were recorded when the agent was confronted with a plant, and those in (B) were recorded when the agent was confronted with a predator. In this example, the period of the oscillation was two time steps with the plant stimulus and six time steps with the predator stimulus.

in the troughs of the periodical activations. To examine how the oscillations influenced response switching, we distilled an abstract connection structure (see Figure 4) from the oscillating agents. Because the competitive behavior was preserved in the abstract network, even when the oscillation-generating connections were removed, the impact of the oscillations on the competitive process could be investigated.

The abstract network was constructed by identifying the connection sets that produced the oscillations in the networks. To this end, negligible and noneffective weights (without which the oscillatory properties of the remaining network did not change) were set to zero, and nodes that consequently did not receive any activation were omitted from the abstract network. With these steps, the idealized symmetrical connection scheme of Figure 4 (for reasons of computational parsimony, all networks were symmetrical with respect to the left and right nodes) was formed. 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  $N_2$ . Excitatory 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.

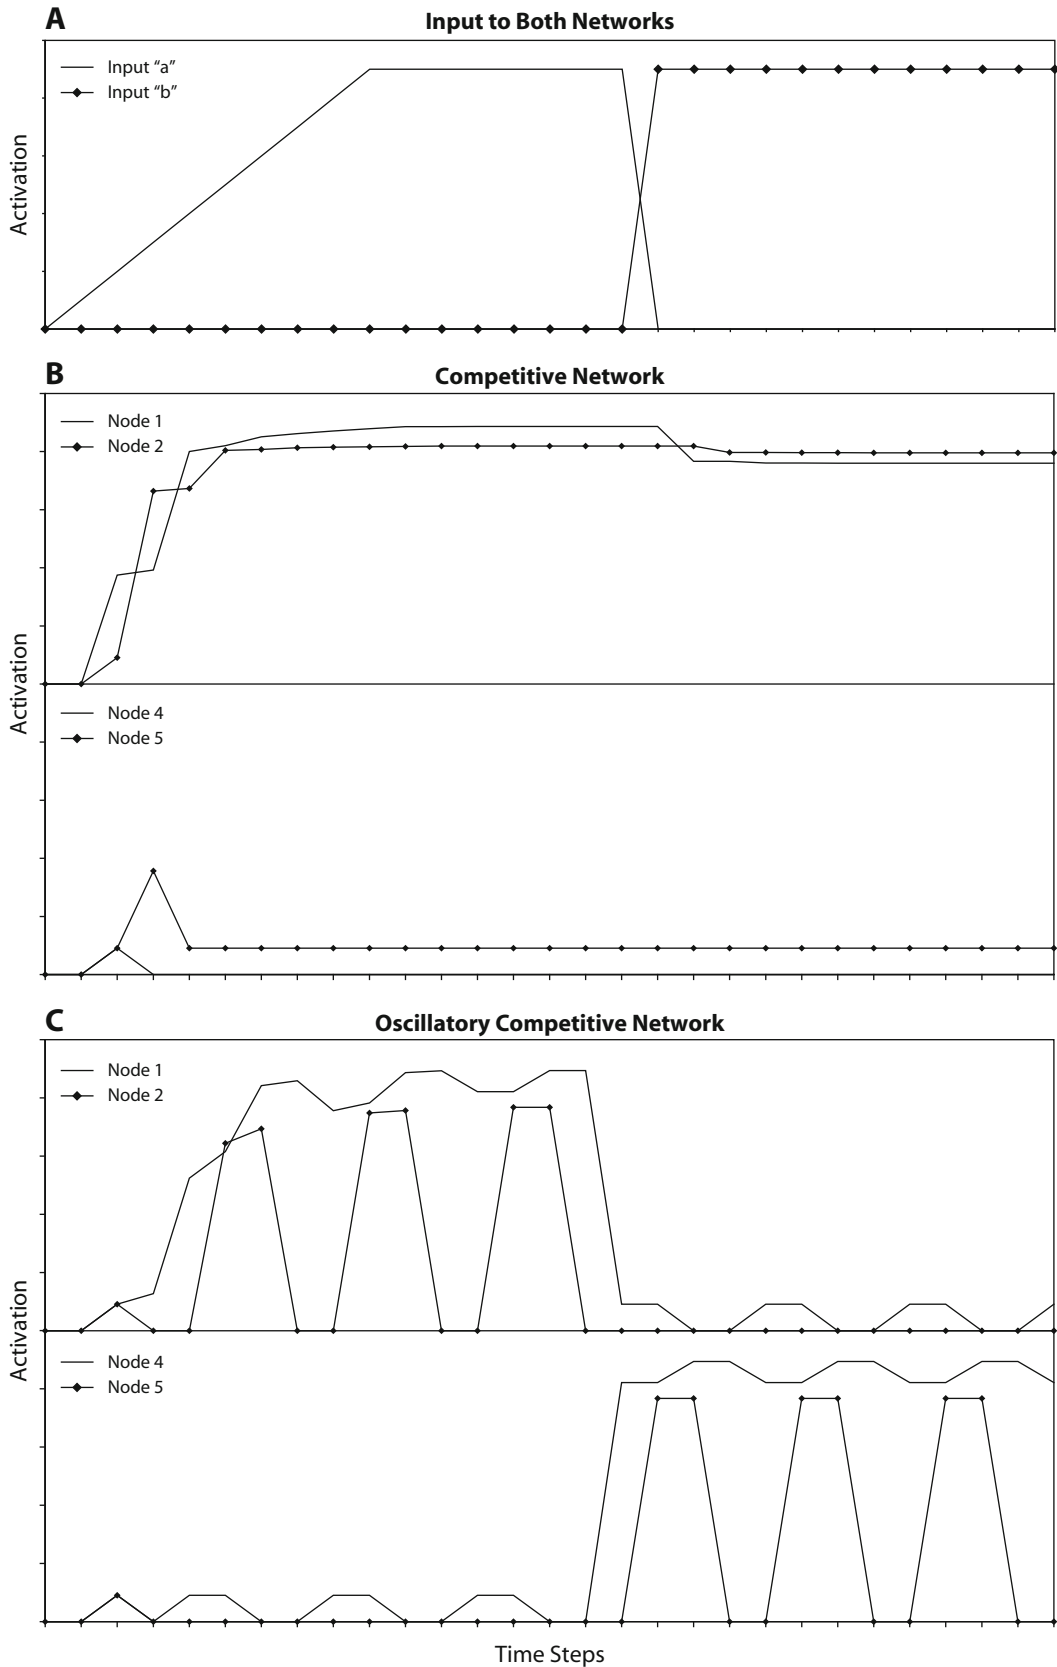


**Figure 4.** The network configuration distilled from the evolved agents that generates the oscillations. The “I” nodes provide inputs of scent “a” and scent “b.” The solid lines with the arrowheads depict excitatory connections (with weight = 5), and the dashed lines with the globule heads depict inhibiting connections (with weight = -10). Without the gray nodes  $N_3$  and  $N_6$ , the nodes  $N_1$  and  $N_2$  simply compete with nodes  $N_4$  and  $N_5$ , resulting in the attentional selection of one of these pairs. With the indicated connections to the nodes  $N_3$  and  $N_6$  added, the network oscillates, which modulates the competitive process.

For the inhibitory mechanism to successfully overrule the excitatory input, all inhibitory connections were, in this idealized model, exactly twice the strength of the excitatory connections (-10 and 5). When 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 only the most strongly activated node (i.e., the “winner”) can preserve a nonzero activation.

The switching efficacy was examined by comparing performances of oscillatory and nonoscillatory competitive networks in a switching task (see Figure 5). The nonoscillatory network had the same connection weights as the oscillatory network, but the connections involved in the generation of oscillations were lesioned (i.e., set to zero). So, both types of networks were derived from the evolved agents that showed oscillatory activations. Nonoscillatory agents that evolved in the simulations may have developed other switching mechanisms, presumably involving less mutual inhibition in the first place, but from their fitness scores, we know that these mechanisms were less effective than oscillations (see Heerebout & Phaf, 2010). For the switching task, input consisted of a monotonically increasing signal to the input node  $I_a$ , which would be replaced instantaneously by a strong signal to the  $I_b$  input node (see upper panel A of Figure 5). Panels B and C of Figure 5 show the ensuing activations of the nodes  $N_1$ ,  $N_2$ ,  $N_3$ , and  $N_4$  of the nonoscillatory and oscillatory competitive networks, respectively. It is clear that the nonoscillatory network was unable to overcome the winner-take-all mechanism after the switch and remained stuck. Because the oscillatory network never completely settled, it did not take a large change in input to tip over to another winner. 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, the oscillatory network smoothly switched in a few time steps. The abstract model demonstrates that because its winning “steady” state consists of oscillations, it is able to switch to another steady state in the periodically occurring troughs of near-zero activation.

The different abilities of oscillatory and nonoscillatory competitive networks to switch winners clearly support our hypothesis that oscillations increase behavioral flexibility. Inhibition may have been exaggerated in the nonoscillatory networks relative to nonoscillatory networks evolved in actual simulations, but still, in the latter networks, the switching ability is inferior to that of evolved oscillatory networks (Heerebout & Phaf, 2010). In addition, the abstract network distinguishes the specific underlying mechanisms for selection and switching. In both oscillatory and nonoscillatory agents, action tendencies (i.e., approach or avoidance) compete with each other, which, according to the “selection-for-action” framework (see, e.g., Allport, 1989), can be characterized as an attentional process (cf. Desimone & Duncan, 1995; Duncan, 1996; Phaf et al., 1990). In the abstract network, the modulation of competition by oscillations is superimposed on the process of attentional selection. The oscillations thus enable the switching between winning representations without affecting the nature of the response that is being selected.



**Figure 5.** The input activations to both networks (A), resulting activations in the nonoscillatory network (B), and resulting activations 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.

In a similar evolutionary-computation study by Ward and Ward (2008), agents evolved to actively direct attention to successive falling targets in order to “catch” them. They found that in order to be able to focus on the first falling target, the activations from the second falling target, which was simultaneously present, needed to be inhibited. This inhibition subsequently caused difficulty in redirecting attention to the second target after the first one was caught. As a result, the agent showed a moment of hesitation after the first target was caught. If the targets were closer together, the inhibition required to stay focused on the first target was more powerful, and the cost in terms of hesitation was larger. Switching cost thus increased with the level of competition. These findings are remarkably consistent with our own (Heerebout & Phaf, 2010) in showing that evolution in most real-world situations needs to solve the problem of not only attentional selection but also how attention should be subsequently redirected. The artificial neural networks in Ward and Ward’s setup had bilaterally symmetric weights, which prevented the emergence of the flip-flop mechanism (i.e., requiring asymmetrical excitatory and inhibitory connections) that was responsible for oscillations. To be sure, the actual simulations should, of course, be performed, but we would guess that letting loose the constraint of bidirectional symmetry in their simulations would lead to a strengthening of mutual inhibition and competition, and to a reduction in switching costs because of the emergence of oscillations.

### Oscillations in Biological Neural Networks

The oscillations emerging in the evolutionary simulations can probably be associated with relatively high-frequency gamma oscillations (30–80 Hz) in biological neural networks (for an overview, see Buzsáki & Draguhn, 2004). From the different frequency bands, gamma oscillations have most consistently been associated with attention (Bauer, Cheadle, Parton, Müller, & Usher, 2009; Womelsdorf & Fries, 2007). The relevance of these oscillations, if any (Pareti & DePalma, 2004), remains a matter of debate. Because evolutionary simulations optimize the genotypic neural networks given particular environmental conditions, they may be in a unique position to shed light on the adaptive function of the emerging oscillations from a computational angle. Our previous work, at least, strongly argues in favor of an adaptive (i.e., fitness-enhancing) function of simulations (Heerebout & Phaf, 2010).

Neurobiological research has suggested a large variety of roles for the oscillations in different areas of the mammalian brain. Possible functions include the binding of cell assemblies (Gray, König, Engel, & Singer, 1989), biasing input selection through neuronal resonance (Hutchison & Yarom, 2000), providing a sense of time (Buhusi & Meck, 2005), selective amplification (Lengyel, Huh, & Érdi, 2005), and sequence learning (Ulanovsky & Moss, 2007). The enhanced ability of the agents to reorganize behavior in response to specific stimuli in our simulations may correspond to the suggestion of Schaefer, Angelo, Spors, and Margrie (2006) that oscillations serve to discriminate between stimuli (see also Brody & Hopfield,

2003, who showed that simple oscillating network models implemented sensory segmentation). In ground-breaking research, Sohal, Zhang, Yizhar, and Deisseroth (2009) even demonstrated in genetically modified mice that light-induced gamma oscillations enhanced information transmission in neocortex by reducing circuit noise and amplifying signals through the circuit. Increased flexibility and switching ability would indeed imply that the information carried by the neural output is more strongly related to the input signal. Thus, the mice with induced gamma oscillations seemed to be able to switch attention more easily than the mice without gamma oscillations.

According to Fries, Nikolić, and Singer (2007), stimuli can be distinguished more easily because the interaction between excitatory pyramidal neurons and inhibitory interneurons results in a time-critical competition during the gamma cycle. The few pyramidal cells that are able to spike first are the only ones to spike at all. Subsequently, the activity of these pyramidal cells is suppressed by strong inhibitory interneurons that are themselves activated by the initial activation of these cells. The mechanism closely resembles the flip-flop mechanism that we encountered in the agents’ networks from our simulations. They suggest that this mechanism enables fast processing and flexible routing of activity, supporting fast selection of responses. The correspondence between the neurobiological findings of Fries et al. and our results further strengthens the association of the oscillations emerging in our simulations with gamma band oscillations in biological neural networks.

In contrast with our hypothesis, Börgers, Epstein, and Kopell (2008) argued, on the basis of their (nonevolutionary) simulations, that oscillations increase stimulus competition, help to suppress distractors, and more narrowly focus selective attention. Their oscillations resulted from the interactions between fast-spiking interneurons, which were periodically subjected to a “bath of inhibition” due to cholinergic activation by a second class of inhibitory interneurons. During this bath of inhibition, all neural activation was suppressed. Particularly, activation from potential distractors could not reach threshold values. Subsequently, only the strongest representation in the network could peak before inhibition started all over again. In their view, only the selection resulting from competitive processes is strengthened by the oscillations. The increased focusing of attention due to oscillations in the Börgers et al. model would decrease, rather than increase, behavioral flexibility, and would reduce the amount of information transferred by a network (see, e.g., Sohal et al., 2009). In contrast with those in Börgers et al., in our simulations, a distinction between attentional selection and attentional switching emerged, with oscillations affecting only the latter. The additional finding of affect-specific oscillation frequencies, which will be discussed hereafter, provides additional support for the distinction and this specific role of oscillations. The evolutionary motivation, which can be distilled quite easily from the simulations, and the experimental evidence supporting the hypothesized link between affect and attentional switching, make it likely that a similar setup can also be found in biological neural networks.



## New Simulations Reveal Stimulus-Specific Frequency

According to LeDoux (1996), the faster, more coarsely grained, direct pathway is biased to evoke an aversive response, because the costs of hesitation are higher when confronted with a threat than with food. This evolutionary reasoning suggests that threatening stimuli should quickly capture and hold attention (cf. Öhman, Flykt, & Esteves, 2001). We wondered whether a similar bias toward focusing attention on negative stimuli could be found in the oscillating networks, even when they did not have the classical LeDoux (1996) type of dual-route architecture. While moving away from an approaching predator, there is little need to heed other stimuli. A high degree of distractibility caused by high-frequency oscillations seems to be a disadvantage in this situation.

A plant stimulus, on the other hand, should enable rapid behavioral switching if a predator appears suddenly. This increased behavioral flexibility, moreover, is supported by ample empirical research in which positive affect was found to increase cognitive and behavioral flexibility (Baumann & Kuhl, 2005; Das & Fennis, 2008; Dreisbach & Goschke, 2004; Fenske & Eastwood, 2003; Isen, 1999; Luu, Tucker, & Derryberry, 1998; Tan, Jones, & Watson, 2009). The broaden-and-build theory of Fredrickson (1998), for instance, generalizes these ideas even further by assuming that positive emotions serve to broaden an individual's momentary thought-action repertoire by expanding the attentional focus. If indeed attentional flexibility is modulated by oscillations, it is only a logical step to expect higher oscillation frequencies for foraging agents than for fleeing agents in the simulations.

To investigate stimulus-specific oscillation frequencies, we replicated the simulations with the recurrent connections of Heerebout and Phaf (2010) an additional 25 times. Doing this resulted in a total of 32 simulations. The new simulations were exact replications, but new pseudorandom numbers were used in the stochastic decisions (e.g., the weight mutations) for each simulation. We found 430 agents with oscillating networks in the last generation of 27 out of 32 simulations (totaling 547 agents). The oscillation frequencies, with both food and predator stimuli, ranged from zero to 0.5 time step<sup>-1</sup> (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 54 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 ( $n = 430$ ; average fitness = 20,112;  $SD = 17,833$  energy  $\times$  time step) reached significantly higher levels of fitness [ $F(1,546) = 7.67, p < .01, \eta_p^2 = .014$ ] than did the nonoscillating agents ( $n = 110$ ; average fitness = 15,220;  $SD = 12,493$  energy  $\times$  time step), demonstrating that the oscillations have clear functional advantages in these simulations. Over the total of 547 agents, the average frequency was significantly higher [ $F(1,547) = 25.27, p < .0001, \eta_p^2 = .044$ ] when an agent detected a plant ( $f = 0.22, SD = 0.19$  time step<sup>-1</sup>)

than when an agent detected a predator ( $f = 0.18, SD = 0.17$  time step<sup>-1</sup>). The extrapolation of these results to biological neural networks, of course, remains hazardous, but if we assume that a time step (i.e., the time needed for activation transfer by a neuron) is about 5 msec, this would result on average in 44-Hz oscillations with the food and 36-Hz oscillations with the predator (cf. the two frequency bands in Dupret et al., 2008). 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.

An anonymous reviewer to a previous version of the manuscript suggested that the dynamic properties of the plants and predators could be confounded with their valence. The plants were always stationary, but the predators actively pursued the agents. To test whether this dynamic property could cause the difference in oscillation frequency, we performed a further set of 22 simulations, in which the plants also moved through the environment, albeit in random directions. The speed (0.3 length units per time step) approximated the average speed of the predators. For the remainder, all of the conditions and parameters remained the same. If the speed difference is responsible for the differential processing of plants and predators, no frequency difference should emerge in these simulations.

The total number of agents in the last generations of all simulations equaled 377. Oscillations were again recorded in response to plants and to predators. Of the 377 agents, 113 showed oscillations only in response to plants, and 17 showed oscillations only in response to predators. The majority, 209 agents, oscillated with both plants and predators. A small group of 38 agents did not show oscillations with either stimulus. Similar to those in the previous simulations, the oscillating agents ( $n = 339$ ; average fitness = 16,977;  $SD = 14,752$  energy  $\times$  time step) reached significantly higher levels of fitness [ $F(1,375) = 10.9, p < .01, \eta_p^2 = .026$ ] than the nonoscillating agents ( $n = 38$ ; average fitness = 9,282;  $SD = 6,738$  energy  $\times$  time step). The difference in oscillation frequency was significant [ $F(1,376) = 67.98, p < .0001, \eta_p^2 = .153$ ] and even larger than it was with the previous simulations. When an agent detected a plant, the average frequency was 0.25 ( $SD = 0.18$ ) cycles per time step, and when an agent detected a predator, the average frequency was 0.16 ( $SD = 0.17$ ) cycles per time step. If we again assume that a time step lasts 5 msec, then the plant frequency corresponds to 50 Hz, and the predator frequency to 32 Hz. These results do not seem to support alternative interpretations in terms of stimulus dynamics.

## Frequency Reflects Affective and Motivational State

Specific stimuli have acquired their positive affective value because they have generally raised fitness levels throughout evolutionary history. Inversely, negative stimuli are characterized by their recurrent potential for lower fitness. Positive and negative affect indeed appear to be generated by the nervous system as a neural code "to those aspects of the environment that were a consistent benefit or

threat to gene survival in ancestral environments” (Johnston, 2003, p. 173). Affective states in the networks are thus inferred from the consequences that specific stimuli may have for fitness level. During evolution, the network “learns” to predict these consequences and to prepare for appropriate actions. Appetitive-approach tendencies toward positive stimuli and aversive-avoidance tendencies away from negative stimuli have likely developed through 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 focusing attention on the negative stimulus and by low distractibility. Fitness benefits are maximized by an approach to all positive stimuli and a high level of distractibility by other stimuli. The level of distractibility in our simulations was positively related 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 aforementioned, highly comparable simulation study by Ward and Ward (2008) also addressed attentional switching, but did not investigate affective influences on switching. Interestingly, Ward and Ward linked the reactive inhibition of the second target to the attentional blink phenomenon, and they referred to neurocognitive studies (Hommel et al., 2006; Kessler, Gross, Schmitz, & Schnitzler, 2006) that also relate attentional dwell time to reactive inhibition. The inhibition of the second target when it is sufficiently close to the first target in the Ward and Ward simulations can, of course, be easily transposed to interference in the attentional blink task. The comparison can even be taken one step further, because affective influences also have been found on the attentional blink (Olivers & Nieuwenhuis, 2006; see also Most, Chun, Johnson, & Kiehl, 2006; Todorović, 2009, recently obtained similar affective modulation of the attentional blink in our laboratory). There is substantial evidence for the hypothesized link between affect and attentional switching not only for the attentional blink task, but also for other experimental tasks (Baumann & Kuhl, 2005; Das & Fennis, 2008; Dreisbach & Goschke, 2004; Fenske & Eastwood, 2003; Isen, 1999; Luu et al., 1998; Tan et al., 2009). In comparison, the empirical evidence for a relationship between affect and oscillation frequency is much scarcer.

The extrapolation of the simulation results to biological neural networks remains a big step that surely needs to be supported by empirical evidence. There is some research to suggest that the frequency differences emerging from the evolutionary simulations may have a similar biological function. A motivational role of gamma oscillations is, for instance, suggested by a study by Rougeul-Buser and Buser (1997). They observed 40-Hz oscillations in a cat’s motor, parietal, and visual cortices when it 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, only lower frequencies of 10–15 Hz showed up. We argue

that the cat at the hole was in a positive affective state, full of expectations of a nice meal, and was prepared to quickly switch from immobility to vigorous attack. With the box, however, where the cat could not reach the mouse, both the need for this preparation and the corresponding oscillation frequency were lower. In the mouse, on the other hand, we would expect the oscillation frequency to be much lower when it is being predated by the cat than, for instance, when it detects a piece of cheese.

A causal relationship between gamma oscillations and positive affect has, to our knowledge, been demonstrated only very recently by Tsai et al. (2009). They showed with their highly innovative optogenetic method (see also Sohal et al., 2009) that trains of high-frequency light flashes (50 Hz) applied to the ventral tegmental area (i.e., containing many dopaminergic neurons) of optogenetically manipulated mice served as a strong reward signal in place conditioning. In the trial, mice were exposed to both high-frequency light trains in one room and low-frequency light trains in another. The total amount of light that the mice were exposed to was equal in both rooms. Afterward, when the mice were given the choice, they all preferred the room in which they had been exposed to the high-frequency light trains. If the high- but not the low-frequency light trains can substitute a reward signal that is normally produced by a positive stimulus, then this shows that, at least in some cases, high-frequency oscillations correspond to positive affect.

The other optogenetic study that we mentioned (Sohal et al., 2009), which was performed by the same group and published simultaneously with the study of Cardin et al. (2009), showed that the induced gamma oscillations enhanced information transmission. Our evolutionary simulations suggest that the two observations are connected and that high neural oscillation frequencies facilitate efficient attentional switching and high levels of information transmission, whereas low frequencies foster the maintenance of attentional focus to negative stimuli. Although at the outset of their research these authors have not formulated specific hypotheses about the relation among oscillation frequency, attention, and motivational state, their results provide converging evidence for our hypothesis.

Two different hypotheses were brought forward by the evolutionary simulations in the present study. First, oscillations have a modulatory role that is superimposed on attentional selection by means of competition. Second, oscillations have a higher frequency in an appetitive state that is associated with positive affect than they do in an aversive state that is associated with negative affect. Particularly, the latter hypothesis appears to be new and may provide the missing neurobiological link between affect and attention, which has often been investigated in psychological research. There are few ideas about the neurobiological mechanisms underlying affective influences on information processing, and in view of these simulation results, the oscillations seem a likely candidate. Computational simulations can show only that a particular hypothesis is more probable than another (i.e., the likelihood of its emergence is higher), but not that it necessarily holds true in biological neural networks. Evolutionary computation takes a special position

within the modeling approaches, because it can create new models and thus form the starting point of the empirical cycle. Of course, novel hypotheses can be successful only if the good vibrations generated by the computer simulations resonate in empirical work and eventually lead to a switching in neurobiological research perspective.

#### AUTHOR NOTE

Address correspondence to B. T. Heerebout, Psychonomics Department, University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands (e-mail: b.t.heerebout@uva.nl).

#### REFERENCES

- ALLPORT, A. (1989). Visual attention. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 631-682). Cambridge, MA: MIT Press.
- ARMONY, J. L., SERVAN-SCHREIBER, D., COHEN, J. D., & LEDOUX, J. E. (1995). An anatomically constrained neural network model of fear conditioning. *Behavioral Neuroscience*, **109**, 246-257. doi:10.1037/0735-7044.109.2.246
- BAUER, F., CHEADLE, S. W., PARTON, A., MÜLLER, H. J., & USHER, M. (2009). Gamma flicker triggers attentional selection without awareness. *Proceedings of the National Academy of Sciences*, **106**, 1666-1671. doi:10.1073/pnas.0810496106
- BAUMANN, N., & KUHL, J. (2005). The positive affect and flexibility: Overcoming the precedence of global over local processing of visual information. *Motivation & Emotion*, **29**, 123-134. doi:10.1007/s11031-005-7957-1
- BEER, R. D. (1990). Intelligence as adaptive behavior: An experiment in computational neuroethology. *Artificial Intelligence*, **72**, 173-215. doi:10.1016/0004-3702(94)00005-L
- BÖRGERS, C., EPSTEIN, S., & KOPELL, N. J. (2008). Gamma oscillations mediate stimulus competition and attentional selection in a cortical network model. *Proceedings of the National Academy of Sciences*, **105**, 18023-18028. doi:10.1073/pnas.0809511105
- BRAITENBERG, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- BRODY, C. D., & HOPFIELD, J. J. (2003). Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron*, **37**, 843-852. doi:10.1016/S0896-6273(03)00120-X
- BUHUSI, C. V., & MECK, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience*, **6**, 755-765. doi:10.1038/nrn1764
- BUZSÁKI, G., & DRAGHUN, A. (2004). Neuronal oscillations in cortical networks. *Science*, **304**, 1926-1929. doi:10.1126/science.1099745
- CARDIN, J. A., CARLÉN, M., MELETIS, K., KNOBLICH, U., ZHANG, F., DEISSEROTH, K., ET AL. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, **459**, 663-667. doi:10.1038/nature08002
- DAS, E., & FENNIS, B. M. (2008). In the mood to face the facts: When a positive mood promotes systematic processing of self-threatening information. *Motivation & Emotion*, **32**, 221-230. doi:10.1007/s11031-008-9093-1
- 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. doi:10.1162/089892903321208132
- DEN DULK, P., ROKERS, B., & PHAF, R. H. (1997). Connectionist simulations with a dual route model of fear conditioning. In B. Kokinov (Ed.), *Perspective on cognitive science* (Vol. 4, pp. 102-112). Sofia, Bulgaria: New Bulgarian University Press.
- DESIMONE, R., & DUNCAN, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, **18**, 193-222. doi:10.1146/annurev.ne.18.030195.001205
- DREISBACH, G., & GOSCHKE, T. (2004). How positive affect modulates cognitive control: Reduced perseveration at the cost of increased distractibility. *Journal of Experimental Psychology*, **30**, 343-353.
- DUNCAN, J. (1996). Cooperating brain systems in selective perception and action. In T. Inui & J. L. McClelland (Eds.), *Attention and performance 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. doi:10.1073/pnas.0810064105
- FENSKE, M. J., & EASTWOOD, J. D. (2003). Modulation of focused attention by faces expressing emotion: Evidence from flanker tasks. *Emotion*, **3**, 327-343. doi:10.1037/1528-3542.3.4.327
- FREDRICKSON, B. L. (1998). What good are positive emotions? *Review of General Psychology*, **2**, 300-319. doi:10.1037/1089-2680.2.3.300
- FRIES, P., NIKOLIĆ, D., & SINGER, W. (2007). The gamma cycle. *Trends in Neurosciences*, **30**, 309-316. doi:10.1016/j.tins.2007.05.005
- GAGE, P. J. (1995). *New approaches to optimization in aerospace conceptual design*. Unpublished PhD thesis, Stanford University.
- 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. doi:10.1038/338334a0
- HEEREBOUT, B. T., & PHAF, R. H. (2010). Emergent oscillations in evolutionary simulations: Oscillating networks increase switching efficacy. *Journal of Cognitive Neuroscience*, **22**, 807-823. doi:10.1162/jocn.2009.21205
- HOLLAND, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press.
- HOMMEL, B., KESSLER, K., SCHMITZ, F., GROSS, J., AKYÜREK, E., SHAPIRO, K., & SCHNITZLER, A. (2006). How the brain blinks: Towards a neurocognitive model of the attentional blink. *Psychological Research*, **70**, 425-435. doi:10.1007/s00426-005-0009-3
- HUTCHEON, B., & YAROM, Y. (2000). Resonance, oscillation, and the intrinsic frequency preferences of neurons. *Trends in Neurosciences*, **23**, 216-222. doi:10.1016/S0166-2236(00)01547-2
- ISEN, A. M. (1999). Positive affect. In T. Dalgleish & M. J. Power (Eds.), *Handbook of cognition and emotion* (pp. 521-539). New York: Wiley.
- JOHNSTON, V. S. (2003). The origin and function of pleasure. *Cognition & Emotion*, **17**, 167-179. doi:10.1080/026999303002290
- KESSLER, K., GROSS, J., SCHMITZ, F., & SCHNITZLER, A. (2006). Cortical dynamics and synchronization related to multiple target consolidation under rapid-serial-visual-presentation conditions. *Journal of Physiology*, **99**, 21-28.
- LEDoux, J. E. (1986). Sensory systems and emotion: A model of affective processing. *Integrative Psychiatry*, **4**, 237-248.
- 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. doi:10.1007/s00422-005-0567-x
- LUU, P., TUCKER, D. M., & DERRYBERRY, D. (1998). Anxiety and the motivational basis of working memory. *Cognitive Therapy & Research*, **22**, 577-594. doi:10.1023/A:1018742120255
- MCCLELLAND, J. L., & RUMELHART, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, **88**, 375-407. doi:10.1037/0033-295X.88.5.375
- MOST, S. B., CHUN, M. M., JOHNSON, M. R., & KIEHL, K. A. (2006). Attentional modulation of the amygdala varies with personality. *NeuroImage*, **31**, 934-944. doi:10.1016/j.neuroimage.2005.12.031
- ÖHMAN, A., FLYKT, A., & ESTEVES, F. (2001). Emotion drives attention: Detecting the snake in the grass. *Journal of Experimental Psychology: General*, **130**, 466-478. doi:10.1037/0096-3445.130.3.466
- OLIVERS, C. N. L., & NIEUWENHUIS, S. (2006). The beneficial effects of additional task load, positive affect, and instruction on the attentional blink. *Journal of Experimental Psychology: Human Perception & Performance*, **32**, 364-379. doi:10.1037/0096-1523.32.2.364
- PARETI, G., & DE PALMA, A. (2004). Does the brain oscillate? The dispute on neuronal synchronization. *Neurological Sciences*, **25**, 41-47. doi:10.1007/s10072-004-0235-2
- PHAF, R. H., DEN DULK, P., TIJSELING, A., & LEBERT, E. (2001). Novelty-dependent learning and topological mapping. *Connection Science*, **13**, 293-321. doi:10.1080/09540090110085666
- 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. doi:10.1016/0010-0285(90)90006-P

- RITZ, R., & SEJNOWSKI, T. J. (1997). Synchronous oscillatory activity in sensory systems; new vistas on mechanisms. *Current Opinion in Neurobiology*, *7*, 536-546. doi:10.1016/S0959-4388(97)80034-7
- ROUGEUL-BUSER, A., & BUSER, P. (1997). Rhythms in the alpha band in cats and their behavioural correlates. *International Journal of Psychophysiology*, *26*, 191-203. doi:10.1016/S0167-8760(97)00764-2
- RUMELHART, D. E., & ZIPSER, D. (1985). Feature discovery by competitive learning. *Cognitive Science*, *9*, 75-112
- 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 Biology*, *4*, 1010-1024.
- SOHAL, V. S., ZHANG, F., YIZHAR, O., & DEISSEROTH, K. (2009). Parvalbumin neurons and gamma rhythms synergistically enhance cortical circuit performance. *Nature*, *459*, 698-702. doi:10.1038/nature07991
- TAN, H. K., JONES, G. V., & WATSON, D. G. (2009). Encouraging the perceptual underdog: Positive affective priming of nonpreferred local-global processes. *Emotion*, *9*, 238-247. doi:10.1037/a0014713
- TODOROVIĆ, A. (2009). *Affective modulation of spatial and temporal attention in early vision*. Unpublished master's thesis, University of Amsterdam, The Netherlands.
- TSAI, H., ZHANG, F., ADAMANTIDIS, A., STUBER, G. D., BONCI, A., DE LECEA, L., & DEISSEROTH, K. (2009). Phasic firing in dopaminergic neurons is sufficient for behavioral conditioning. *Science*, *324*, 1080-1084. doi:10.1126/science.1168878
- ULANOVSKY, N., & MOSS, C. F. (2007). Hippocampal cellular and network activity in freely-moving echolocating bats. *Nature Neuroscience*, *10*, 224-233. doi:10.1038/nn1829
- WARD, R., & WARD, R. (2008). Selective attention and control of action: Comparative psychology of an artificial, evolved agent and people. *Journal of Experimental Psychology: Human Perception & Performance*, *34*, 1165-1182. doi:10.1037/0096-1523.34.5.1165
- WOMELSDORF, T., & FRIES, P. (2007). The role of neuronal synchronization in selective attention. *Current Opinion in Neurobiology*, *17*, 154-160. doi:10.1016/j.conb.2007.02.002

## APPENDIX A

To create offspring, the first parent was chosen by "tournament selection." This was the agent with the highest fitness within a random subset of agents of one third of the population. To select the second parent agent, an adapted "roulette wheel selection" was applied. For each agent, a portion of a roulette wheel (i.e., proportional to the probability of selection) was allocated. The probability of an agent ( $P_i$ ) to be selected was calculated with the formula

$$P_i = 0.5 \cdot \left( \frac{\text{ReversedDistance}_i}{\text{SumReversedDistance}} + \frac{\text{ReversedReducedFitness}}{\text{SumReversedReducedFitness}} \right). \quad (\text{A1})$$

This formula combined an agent's resemblance to the first parent and the agent's reduced fitness. The resemblance to the first parent, its ReversedDistance, was measured as the inverse of the Euclidean distance between the networks, which in this context were represented as points in a multidimensional space, defined by the networks' connections. The agent's ReducedFitness was its fitness minus 90% of the lowest occurring fitness in the current generation.

The connection strengths of the parents were used to create two offspring. Half of the time, each offspring was a copy of one of the parents, and the other half, each was a mixture of both parents. In the case of a mixture, two random parts of the weights would be swapped. In addition, all of the weights, ranging between  $-10$  and  $10$ , would mutate slightly by adding  $M$ , which was calculated with the formula

$$M = -0.1 \cdot \log\left(\frac{1}{r} - 1\right), \quad (\text{A2})$$

where  $r$  was a random value drawn from a uniform distribution between 0 and 1. The offspring were added to the population, and the process was then repeated until the population size had tripled. All of the agents were considered as parents anew with the creation of each offspring pair.

## APPENDIX B

The entities had a round shape with a radius of 10 length units. The environment was a torus-shaped square whose sides measured 400 length units. Smells were emitted by all of the entities in the environment and dissipated in all directions equally with geometrically decaying intensity ( $S$ ) according to the formula

$$\delta < \delta_{\text{MAX}} : S = \frac{1}{1 + \delta} \cdot S_{\text{MAX}} \cdot \left( 1 - \frac{\delta}{\delta_{\text{MAX}}} \right) \quad (\text{B1})$$

$$\delta \geq \delta_{\text{MAX}} : S = 0,$$

with  $S_{\text{MAX}}$  = maximal smell intensity at the source (set to 25),  $\delta$  = distance to the source (in length units), and  $\delta_{\text{MAX}}$  = maximal distance at which the source was smelled (set to 100). The entities emitted multiple smells in different ratios. Plants emitted both smells "a" and "b," with the intensity of "a" being half as strong as that of "b." With predators, smell "b" was half as strong as smell "a."

The agent's motor actuators exerted a force against the environment's surface that propelled the agent straight ahead if both forces were equal, or in a curve if the one was greater than the other. Each time step that force was exerted, the agent's energy was reduced by an amount  $\text{Energy}_t$  according to the formula

$$\text{Energy}_t = \frac{\text{MaxEnergy}}{1 - (\text{MotorAct}_t - \text{MaxAct})} \cdot \frac{\text{MotorAct}_t}{\text{MaxAct}}, \quad (\text{B2})$$

---

**APPENDIX B (Continued)**


---

where  $\text{MotorAct}_t$  = the activation of the actuator at time  $t$ ,  $\text{MaxAct}$  = the actuator's maximum activation (set to 1.0), and  $\text{MaxEnergy}$  = maximum consumption per actuator per time step (set to 0.001). In addition, the agent used a static energy cost of 0.001 per time step.

---



---

**APPENDIX C**


---

In the third stage, agents are removed from the population. The probability of staying in the population for another generation would be determined in three steps. First, a relative factor ( $RF$ ) was calculated for each agent with the formula

$$RF = \frac{b \cdot (F_i - F_{\min})}{F_{\max} - F_{\min}} + o, \quad (C1)$$

with  $F_i$  as the fitness of agent  $i$ ,  $F_{\min}$  as the lowest fitness in the current population,  $F_{\max}$  as the highest fitness in the current population,  $o$  as the lower limit of  $RF$  (set to 0.15), and  $b$  as the range of  $RF$  values (set to  $1 - o = 0.85$ ). In the second step, the relative factor was used to calculate the population factor, which accounted for the influence of the current population size ( $N$ ) with respect to the initial population size ( $N_0$ ). If the current population was large, the agents' probability to be removed increased, and vice versa. The population factor ( $PF$ ) was calculated according to the formula

$$\begin{aligned} N \geq N_0 : \quad PF &= RF \cdot \frac{N_0}{N} \\ N < N_0 : \quad PF &= 1 - \left\{ (1 - RF) \cdot \left[ 1 - \left( 1 - \frac{N}{N_0} \right)^2 \right] \right\}. \end{aligned} \quad (C2)$$

In the third step, the probability that the agent would remain in the population for another generation was calculated by multiplying the  $PF$  by 0.99. This ensured that even the fittest individual could sometimes be removed from the population.

---



---

**APPENDIX D**


---

The weights of the network's connection determine an agent's behavior. The network is symmetrical in the sense that all nodes have a mirror node and thus, for each connection running from node  $i$  to node  $j$  there is a mirror connection from mirror node  $i'$  to mirror node  $j'$ . Activations are of course unique; otherwise, the output nodes' activation would always be equal, and the agent could only move forward. The symmetry greatly reduces the number of mutations required to evolve a functional behavior and is correct under the assumption that the agent should give an exact mirrored response in an exactly mirrored environment. The activation ( $y_i$ ) of all of the nodes was calculated with the formula

$$y_i = \sigma \left( \sum_{j=1}^N \omega_{ij} y_j + \theta_i + S_i \right). \quad (D1)$$

The formula of the sigmoid was

$$\begin{aligned} \xi \geq 0 : \quad \sigma(\xi) &= \frac{\xi}{1 + \xi} \\ \xi < 0 : \quad \sigma(\xi) &= 0, \end{aligned} \quad (D2)$$

with  $\omega_{ij}$  = the weight of the connection from unit  $j$  to node  $i$ ,  $N$  = the number of connections to node  $i$ ,  $\theta_i$  = the bias of node  $i$  (0.1 in the simulations), and  $S_i$  = the smell input to node  $i$  from the corresponding sensor (only for input layer). In order to create a temporal difference between processing over long neural pathways and short neural pathways, a time delay was introduced by storing a node's activation for one time step. Hence, the activation that was calculated for a node at time  $t$  was propagated over its sending connections at time  $t+1$ . As a consequence, input signals were processed to the output nodes at  $t+1$  via the direct projections and at  $t+2$  via the indirect connections, through the hidden layer.

---