

Emergence of sensory selection mechanisms in Artificial Life simulations

Carolina Feher da Silva

Department of Physiology and Biophysics
Institute of Biomedical Sciences
University of São Paulo
São Paulo, SP, Brazil
Email: carolina@icb.usp.br

Nestor Caticha

Department of General Physics
Institute of Physics
University of São Paulo
São Paulo, SP, Brazil
Email: nestor@if.usp.br

Marcus Vinícius C. Baldo

Department of Physiology and Biophysics
Institute of Biomedical Sciences
University of São Paulo
São Paulo, SP, Brazil
Email: baldo@usp.br

Abstract—The evolutionary advantages of selective attention are unclear. It has been suggested that the nervous system only processes the most relevant stimuli because of its limited capacity. We tested this theory by creating Artificial Life simulations in which artificial animals had neural networks of different processing capacities. To survive and reproduce, the animals had to identify two stimuli whose relevances could be the same or different. When relevances were the same, the animals processed only one or both stimuli depending on their processing capacity. When relevances were different, the most relevant stimuli were usually processed with a higher detectability than the least relevant stimuli, regardless of processing capacity. Selection mechanisms and the ability to shift the focus of sensory processing arose in evolution depending both on the stimuli’s relevances and the animals’ processing capacity.

Index Terms—Selective attention, Artificial Life, genetic algorithms, Signal Detection Theory, evolution

I. INTRODUCTION

Only a small fraction of the stimuli collected by our sensory receptors is fully processed and becomes conscious. Our nervous system is always focusing its resources on the processing of some stimuli and inhibiting the processing of others, an ability called selective attention [1], [2], [3]. The stimuli selected for processing are the most relevant. For example, in a noisy and crowded cocktail party, we are able to concentrate on a single conversation while ignoring all the distracting noises—the so called “cocktail effect” [4].

A. Limited Capacity

Since the study of selective attention began, it has been proposed that the selection of sensory stimuli is necessary because the nervous system has a limited amount of computational resources to process them [5]. Whenever a sensory processing module is working at its full capacity, selection mechanisms are activated to control the flow of information and make sure that high-priority stimuli are processed [6]. It was even suggested that “if the brain had infinite capacity for information processing, there would be little need for attentional mechanisms” [7]. Based on these ideas, studies have been trying to identify the limiting stages in sensory processing. It was first proposed that selection occurs at the early stages (the early selection theory) [5], [8]. First the basic physical characteristics of all stimuli are extracted, then this

information is used to make the selection. For a stimulus to be further processed until the late stages of semantic analysis, selective attention is necessary. Thus this theory predicts that unattended stimuli are not identified.

However, studies suggest that sometimes unattended stimuli are identified [9]. To explain these findings, it was proposed that sensory processing does not go through any limited stage until semantic analysis and can be performed in parallel (the late selection theory). Selection only occurs later, after all the stimuli have been identified. This theory also states that early sensory processing is involuntary, i.e., not identifying a stimulus is impossible [10], [11]. However, the late selection theory is inconsistent with experiments that show that the analysis of unattended stimuli is reduced when the number of stimuli is increased [12]. Given that neither the early selection nor the late selection theory could fully explain experimental data, conciliatory proposals were made; for instance, Treisman [8] proposed that unattended information is not completely inhibited, but just degraded or attenuated.

Nevertheless, according to all of the theories above, it is the nervous system’s limited capacity that makes selective attention necessary. They differ only on the location of the “bottleneck”: if it is early, during the analysis of basic physical characteristics, or late, after the stage of semantical analysis [13]. But attention focus on what is relevant for the organism and, in order to determine what is important, some processing is necessary. If selection is made early in the process, it is not clear how the nervous system is able to select what is relevant. On the other hand, if selection is made only after the stimuli have been processed to an advanced stage, it is easy to determine what is relevant but selection is not necessary anymore, because resources have already been wasted on the processing of irrelevant information [13], [14].

B. Relevance for Action

Not all theories approach selective attention as a solution for a limited capacity problem. Allport [15] proposed that action planning requires the inhibition of irrelevant stimuli, which forces the nervous system to limit its processing. According to this theory, called “selection-for-action,” it is the need for selection that causes the limitation, not the opposite. Sensory

processing is limited to relevant stimuli for action to be coherent, based only on relevant information.

Dayan, Kakade, and Montague [16] give an example of a classical conditioning experiment in which selective attention is important for action control and the notion of limited capacity can be abandoned as redundant. Classical conditioning is the learning of predictive relationships between stimuli: an animal learns what a conditioned stimulus predicts about an unconditioned stimulus. For instance, a laboratory rat learns that soon after a blue light is turned on, a food pellet is delivered. In the natural environment, few stimuli have a predictive value. Dayan, Kakade, and Montague proposed a statistical model that exemplifies how selective attention can be an advantage when different stimuli have different predictive values. According to their model, when a conditioned stimulus doesn't allow an animal to make reliable predictions about the unconditioned stimulus, it must be suppressed so that it won't affect the animal's decisions. Thus if a rat learns that a green light is turned on and off regardless of food delivery, it will not pay attention to this light when estimating the timing of food delivery. What happens is a form of selective attention that isn't caused by limited capacity but by the need for action control.

Although the inhibition of irrelevant stimuli appears to be necessary for action control, most methods of detecting the effects of selective attention in experiments rely on the subject failing to notice a stimulus or taking longer to respond to an unattended stimulus. All of these suggest that subjects would achieve a better performance in the experiment if they could process more stimuli at once, but they can't because of limited capacity.

C. Artificial Life

It was by natural evolution that all cognitive processes arose in animal history, therefore an evolutionary approach might provide additional clues to clarify the role of selective attention. One method that makes possible to model the evolution of the nervous system over many generations in a short time is Artificial Life. Artificial Life is the study of life-like behavior in computers, machines and other alternative media. Many Artificial Life experiments are computer simulations of artificial animals. To simulate cognitive processes, these animals have neural networks, which receive sensory stimuli from their environment and control their behavior.

The process of natural evolution can also be simulated with a genetic algorithm. Each artificial animal has a genome and reproduces according to how adapted they are to their environment. In studies of Artificial Life applied to cognitive neuroscience, the genome usually determines the structure of the animal's neural network. If each gene represents a synapse between two neurons, genetic algorithms can be seen as a learning algorithm for neural networks. Learning in neural networks is usually associated with synaptic changes [17]. Even though an artificial animal's neural network might remain unchanged during its life, by means of mutations and other genetic operators, the neural networks in a simulation

are gradually optimized for survival. A genetic algorithm could be able to focus the network's computational resources on the selective processing of a stimulus, in a way analogous to an attentional shift.

Selective attention is already being investigated with Artificial Life [18], [19], [20]. Bartolomeo et al [18] developed an Artificial Life simulation to study the evolutionary emergence of a central region of greater visual acuity in an animal's visual field (a "fovea"), depending on its neural network's processing capacity. They observed that animals with larger neural networks had good visual acuity in either the central and the peripheral regions of the visual field, while animals with smaller neural networks had good visual acuity only in the central region. When the latter animals had to identify a periferic stimulus, they orientated their bodies so as to centralize ("foveate") the stimulus in their visual field. These authors see the orientation movements as an embodiment of attentional processes (attention shifts) and conclude that their study is an evidence for the theory that selection arose to compensate for our nervous system's limited capacity.

Seth [20] came to a different conclusion by observing the emergence of behaviors similar to those associated to selective attention as a result of action control mechanisms. In his simulation, artificial animals had a nervous system consisting of direct links between receptors and two wheels; the receptors detected water, food, and traps in the environment and wheels moved the animal. They also had food and water "batteries" connected to the wheels by links. To survive and reproduce, they had to keep their batteries from running out and avoid traps. The wheel speed was determined by the sum of the links' outputs, which were themselves determined by their input and genetic parameters. At each generation, these parameters evolved so that the animals' actions became more coherent with their survival needs. Seth observed that the animals with the highest fitnesses exhibited behaviors related to action selection and selective attention, such as giving priority to behaviors according to their needs, opportunistic behavior, interruption of current behavior etc. For instance, when an animal was positioned next to a food source, it went towards it. If a trap was suddenly introduced between the animal and the food, it usually changed its course to avoid the trap, but still got closer to the food source. The frequency with which the animal avoided the trap depended on its food battery level. With low food levels, the animal was unable to avoid the trap. This was interpreted as "variable attention." When the animal needed food, it seemed to "pay less attention" to traps. Because in the model there are no interconnections between receptors and effectors, the phenomena of action selection and selective attention are equivalent. According to the author, these results suggest that these two problems are the same. But the animals' behavior of falling into the trap in the above example is not adaptive. Maybe they would have been able to avoid the trap if they had more processing units; after all, they might have a limited capacity problem.

The conflicting results obtained in these two experiments do not allow us to reach a conclusion. Our capacity seems

to be limited and we can't pay attention to many things at once. Regardless, the ability to keep the focus on relevant stimuli and block distractors from consciousness seems to be an advantage. In this study, we aim to shed more light on this question by following the evolution of artificial animals with or without limited capacity and by varying stimuli's relevances for action.

II. METHODS

In our Artificial Life simulations, artificial animals lived for 500 time units. At each time unit, they were presented two objects, which could be of two types: "food" or "non-food." The animals could choose whether to eat both objects, only one of them, or none of them, and their fitness (a measure of adaptation) changed with each object they ate. The objects had values, which were added to the fitness when the object was eaten; the food objects had positive values and the non-food objects had negative values. Not eating any objects would keep their fitness at the same level.

The two objects shall be referred to as the "left" and the "right" objects, because they were presented at these locations in the animals' visual fields and stimulated different input nodes in the animals' neural networks. The neural networks had two input nodes, one for each object: the left input node received a stimulus that represented the left object and the right input node received a stimulus that represented the right object. Two output nodes determined if the animal would eat each of the presented objects: the left output node determined if the animal would eat the left object and the right output node determined if the animal would eat the right object.

Besides having input and output nodes, the neural networks also had a variable number of hidden nodes, which determined whether or not it had enough capacity for processing the two stimuli. The evolutionary relevances of the left and the right objects could also vary depending on how they affected fitness when ingested. We compared sensory processing in animals with or without limited capacity, which evolved in simulations in which the objects had the same or different relevances.

If selection is a solution for a problem of limited capacity, we expect that selection will emerge only in artificial populations wherein the animals had limited capacity.

A. Neural networks

The animals had feedforward, totally connected neural networks with two linear input nodes, whose output was the same as their input, a variable number of hidden nodes, and two output nodes. Each input node was stimulated with a number from the interval $[-1, 1)$, which indicated a food object if the number was also an element of the interval $[-0.3, 0.3)$ and a non-food object otherwise. Each hidden or output node j was a simple perceptron, whose output S_j was given by the equation

$$S_j = \tanh(5(b_j + \sum_{i=1}^n w_{ji}S_i)) \quad (1)$$

wherein b_j is the node bias, S_i is the output of the presynaptic node i and w_{ji} is the weight of the synapse between node i

and node j . The output of the network's two output nodes determined whether the animal ate the two objects. If the left node's output S_l was greater than zero, the animal ate the left object; otherwise, that object was not eaten. The same rule was applied to the right output node and object.

Each node bias and each synaptic weight was determined by the animals' chromosomes. The synaptic weights were constant during an animal's life (no learning), but could change through mutation and crossover through generations.

B. Simulation Sets

In several simulation sets, we studied the effect of two factors on the emergence of selection mechanisms: the number of hidden nodes in the neural networks and the stimuli's relevances for action (Table I). The number of hidden nodes could be two or eight. Two hidden nodes are not sufficient for the correct identification of both stimuli all the time. These animals had to process only one stimulus or both of them partially; these neural networks had limited capacity. Neural networks with eight hidden nodes had the capacity to correctly identify both stimuli all of the time; these animals didn't have limited capacity.

In fact four hidden nodes are enough for the identification of both stimuli because the hidden nodes and the output nodes are perceptrons: they separate their input by a hyperplane into two categories. Two hidden nodes are necessary to identify a stimulus: one node to classify as greater than -0.3 or not and another node to classify it as greater than 0.3 or not. Thus it becomes possible to identify a food object, which was represented by a stimulus from the interval $[-0.3, 0.3)$. This is the only way an output node can respond in a different way to food and non-food objects. The left and the right stimuli had to be processed by different nodes if they were to be correctly identified, because the two stimuli were independent. As there are two objects to be identified, a neural network would then need four hidden nodes, two for each object. In fact we were able to manually design such a neural network with four hidden nodes that was able to correctly identify the two objects at all times.

Besides the number of hidden nodes, the stimuli's relevances for action could vary across simulation sets. The relevances of the left and the right stimuli were defined by how an animal's fitness changed when the animal ate the objects they represented when the objects were food. The left and the right stimuli had the same relevance when the fitness increased by the same value when either a left or a right food object was eaten. The stimuli had different relevances when the fitness increased more when a left food object was eaten than when a right one was eaten.

The simulation sets had names such as "2-9-3": the first number was the number of hidden nodes in the neural network (2 in this example), the second number was the initial value of left food objects (9 in this example), and the third number was the initial value of right food objects (3 in this example). In the second half of the simulation, the values of right and left food objects were swapped. Each simulation were run for

Variables	Simulation Sets					
	2-6-6	8-6-6	2-9-3	8-9-3	2-11-1	8-11-1
Number of hidden nodes	2	8	2	8	2	8
Initial value of left food	6	6	9	9	11	11
Initial value of right food	6	6	3	3	1	1
Value of left non-food	-3	-3	-3	-3	-3	-3
Value of right non-food	-3	-3	-3	-3	-3	-3

TABLE I

SIX SETS OF SIMULATIONS WERE PERFORMED, VARYING THE STIMULI'S RELEVANCES AND THE NUMBER OF HIDDEN NODES IN THE NEURAL NETWORK.

40,000 generations, so during the first 20,000 the left food objects were more valuable than the right food objects and in the last 20,000 generations, the right food objects became more valuable than the left food objects. All the non-food objects always had the same value of -3, regardless of their location.

Every simulation set consisted of ten simulations. They were identical except for the first generation of chromosomes, which were randomly generated. We wanted to see if the results would be consistent within a simulation set, which would indicate that the differences between simulation sets were due to different parameters instead of historical accidents during the course of the simulations.

C. Evolution

The evolution of neural networks was simulated with a genetic algorithm. Each animal lived for 500 time units and their fitnesses were computed during their lives according to the number of food and non-food objects they ate. Their chromosomes were a vector of real numbers drawn from the interval $[-1, 1)$, which determined their neural networks' biases and synaptic weights. There were a gene for each node bias and a gene for each synaptic weight.

During each simulation, 10 populations of 20 individuals evolved in parallel for 40,000 generations. To create a new population from the previous one, 5 individuals with the highest fitnesses were copied to the next generation without modification. The other 15 individuals of a population were generated by sexual reproduction with mutation (the rate of mutation was .05 per gene; a mutated gene was drawn randomly with uniform distribution from the interval $[-1, 1)$) and crossover (each gene had the same probability of coming from each parent's chromosome). The probability of an individual's being selected for reproduction was given by its fitness: the higher the fitness, the higher the probability of generating a descendant. To this end, a selection-by-rank genetic algorithm was used. The selection-by-rank algorithm sorted the individuals by fitness and attributed a rank of 20 to the fittest individual, 19 to the next fittest individual and so on. The probability of selecting an individual for reproduction was directly proportional to its rank. At each 200 generations, the individual with the highest fitness of a population migrated to a randomly selected population, always keeping the number of individuals of all populations constant.

D. Signal Detection Theory

To learn how these animals processed the left and the right stimuli, methods from the signal detection theory were used. The signal detection theory studies detection experiments, in which weak stimuli must be distinguished from a noisy background [21]. In the simplest case, only one stimulus is present and the sensory component of a subject's performance is described by the detectability, or d' . Detectability is a function of the hit rate, the ratio of trials where the subject detects the signal and the signal is present, and the false alarm rate, the ratio of trials where the subject detects the signal but the signal is not present. Its definition is:

$$d' = z(H) - z(F) \quad (2)$$

where H is the hit rate, F is the false alarm rate and z is the inverse of the normal distribution function. A high detectability results from detecting the signal when it is present and not detecting it when it is not present. By analyzing the d' of neural networks, it was possible to determine which stimuli the networks processed and correctly identified.

The hit and the false alarm rates were calculated for the animals as follows. A hundred stimuli evenly distributed in the interval $[-1, 1)$ were generated, then combined in pairs, generating every possible pair of stimuli (totaling 10,000 pairs), and the pairs were used as input to the neural network's two input nodes. The neural network's response to each pair of stimulus was recorded and the hits and the false alarms were counted for the left and the the right objects separately. When an animal ate a food object, this was a hit. When the animal ate a non-food object, this was a false alarm. Then d' was calculated for the left and right stimuli.

III. RESULTS

In the first two simulation sets, all stimuli were equally relevant. The difference was in the number of hidden nodes in the neural networks (two or eight). The d' 's of the simulation set 2-6-6's last generation are displayed in Figure 1. About half of the animals had a high d' for the left stimuli and a zero d' for the right stimuli. The other half had a high d' for the right stimuli and a zero d' for the left stimuli. This result was expected, because these animals had only two hidden nodes in their neural networks and they were unable to identify both stimuli correctly all of the time. They could only identify one of them. As all stimuli were equally relevant for action, both animals that processed only the left stimuli and animals that processed only the right stimuli evolved. The




Fig. 1. d' for the left and the right stimuli for all the animals of the last generation of simulation set 2-6-6. All animals had two hidden nodes in their neural networks. Both stimuli were equally relevant for action.




Fig. 4. Evolution of mean d' for the left and the right stimuli for simulation set 8-6-6.

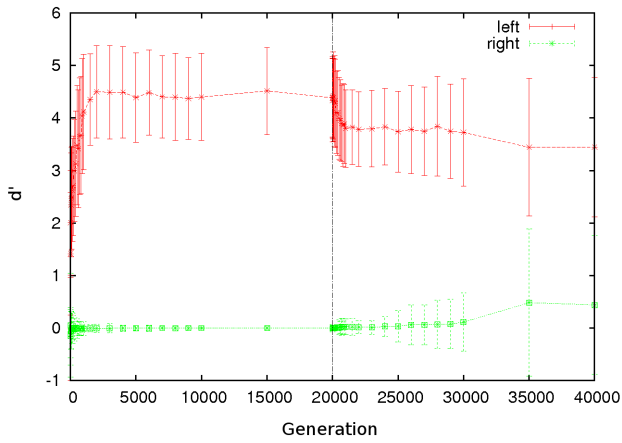


Fig. 2. Evolution of mean d' for the left and the right stimuli for simulation set 2-9-3.




Fig. 5. Evolution of mean d' for the left and the right stimuli for simulation set 8-9-3.




Fig. 3. Evolution of mean d' for the left and the right stimuli for simulation set 2-11-1.

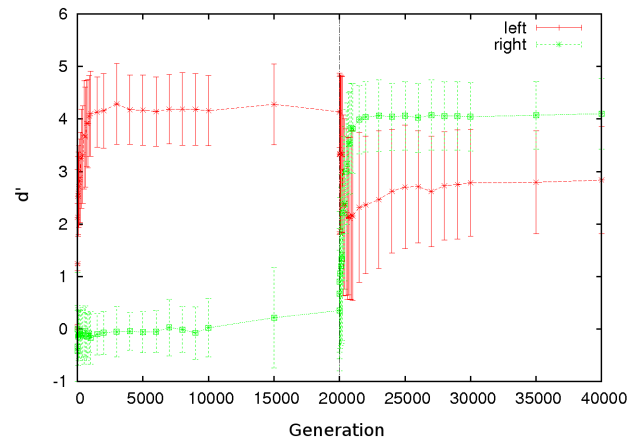


Fig. 6. Evolution of mean d' for the left and the right stimuli for simulation set 8-11-1.

animals of simulation set 8-6-6 were able to identify both stimuli correctly, because they didn't have limited capacity. The d' evolution during simulation set 8-6-6 is shown in Figure 4.

In simulation sets wherein the stimuli were of different relevances, the result could not be predicted based on the capacity of neural networks alone. All animals with only two hidden nodes evolved the ability to identify only one of the stimuli (Figures 2 and 3). But the animals that had eight hidden neurons also achieved a higher d' for the most relevant stimuli than for the least relevant ones (Figures 5 and 6).

IV. DISCUSSION

Theoretically, we had been able to demonstrate that neural networks with only two hidden nodes wouldn't be able to identify two objects correctly all the time. No simulations were needed to know that these animals wouldn't achieve a high d' for both stimuli—it would have been impossible. Which stimuli were initially processed and which were initially ignored depended on their relevances. When the left stimuli were the most relevant, they were processed. When both stimuli had the same relevances, two distinct subpopulation arose, which processed only the left or only the right stimuli. They had something akin to hemispatial neglect in biological animals. When the stimuli's relevances were swapped in generation 20,000, the networks in 2-9-3 kept on processing the left stimuli, even though they weren't the most relevant ones anymore. They were less plastic than the networks in 2-11-1, which were able to change which stimuli were processed back to the most relevant ones.

Some of the other simulations had unpredicted results. Based on the theory of limited capacity, we predicted that in simulation sets 8-9-3 and 8-11-1, animals would be able to identify both stimuli correctly, because they had enough capacity to process all their input. But while the d' for the left stimuli, which were initially the most relevant, reached its maximum value in a few generations, the d' for the right stimuli rose slowly and didn't reach as high a value. When the relevances were swapped, the d' for the left stimuli actually decreased. The larger networks were more plastic than the smaller networks, because the d' for the right stimuli increased in both simulation set 8-9-3 and 8-11-1 when the relevances were swapped in generation 20,000.

The processing of right stimuli was reduced not only when the neural networks had limited capacity, but also when those stimuli had little relevance. A higher d' for the right stimuli wasn't an important evolutionary advantage and mutations that improved their processing frequently got lost during reproduction. With enough capacity for information processing, there may have been little need for selection, but there is also little need for circuits capable of identifying less relevant stimuli. With little selective pressure, complex structures don't evolve. Animals are selected for their actions and perception is important only when it affects action. Infinite capacity for information processing wouldn't be any use without the ability to turn this information into adaptive action. For instance, a

mouse wouldn't take any advantage from paying attention to the sound of humans talking if it doesn't have the ability to understand language and act on this understanding. New mechanisms for information processing can only coevolve with the ability to make better decisions based on this information. If a set of stimuli is completely irrelevant, there is no evolutionary pressure to process them and all animals have the same expected fitness regardless what their neural networks do with these stimuli. Then, according to the theory of evolution, no structure to process them can arise except by chance.

Our results couldn't have been predicted by the theory of selection-for-action either, because the processing of the right stimuli didn't have to be inhibited for action. Only the nodes that processed the left stimuli had to inhibit the right stimuli to generate adaptive action, because the right stimuli didn't carry any information about the left objects. The remaining nodes, which were not necessary to process the left stimuli, were free to process the right stimuli.

V. CONCLUSION

Selections mechanisms arose in our simulations depending on the size of the neuron networks but also on the stimuli's relevances for action. Our model is very simple, but it led us to consider selective attention from an evolutionary point of view and reach a different conclusion than what has been reported in literature. This evolutionary point of view might be Artificial Life's greatest contribution to biology.

REFERENCES

- [1] M. I. Posner, C. R. R. Snyder, and B. J. Davidson, "Attention and the detection of signals," *Journal of Experimental Psychology: General*, vol. 109, no. 2, pp. 160–174, 1980.
- [2] C. W. Eriksen and J. D. S. James, "Visual attention within and around the field of focal attention: a zoom lens model," *Perception & Psychophysics*, vol. 40, pp. 225–240, 1986.
- [3] M. I. Posner, "Exploiting cognitive brain maps," *Brain and Cognition*, vol. 42, no. 1, pp. 64–67, 2000.
- [4] E. C. Cherry and W. K. Taylor, "Some further experiments on the recognition of speech with one and two ears," *Journal of the Acoustical Society of America*, vol. 26, p. 554–559, 1954.
- [5] D. E. Broadbent, *Perception and communication*. London: Pergamon Press, 1958.
- [6] R. Desimone and J. Duncan, "Neural mechanisms of selective visual attention," *Annual Review of Neuroscience*, vol. 18, pp. 193–222, 1995.
- [7] M. M. Mesulam, *Principles of behavioral neurology*. Philadelphia, PA: F. A. Davis, 1985, ch. Attention, confusional states, and neglect, pp. 125–168.
- [8] A. M. Treisman, "Strategies and models of selective attention," *Psychological Review*, vol. 76, no. 3, pp. 282–299, 1969.
- [9] S. P. Tipper and J. Driver, "Negative priming between pictures and words in a selective attention task: evidence for semantic processing of ignored stimuli," *Memory & Cognition*, vol. 16, no. 1, pp. 64–70, 1988.
- [10] J. A. Deutsch and D. Deutsch, "Attention: Some theoretical considerations," *Psychological Review*, vol. 70, no. 1, pp. 80–90, 1963.
- [11] J. Duncan, "The locus of interference in the perception of simultaneous stimuli," *Psychological Review*, vol. 87, no. 3, pp. 272–300, 1980.
- [12] H. E. Pashler, *The Psychology of Attention*. Cambridge: MIT Press, 1998.
- [13] M. S. Gazzaniga, R. B. Ivry, and G. R. Mangun, *Cognitive Neuroscience*. New York: W. W. Norton & Company, 1998.
- [14] S. E. Palmer, *Vision science*. Cambridge: MIT Press, 1999.
- [15] A. Allport, "Selection for action: Some behavioral and neurophysiological considerations of attention and action," in *Perspectives on perception and action*, H. Heuer and A. Sanders, Eds. Erlbaum, 1987.

- [16] P. Dayan, S. Kakade, and P. R. Montague, "Learning and selective attention," *Nature Neuroscience*, vol. 3, pp. 1218–1223, 2000.
- [17] S. Haykin, *Neural Networks: A Comprehensive Foundation*. Prentice Hall, 1998.
- [18] P. Bartolomeo, L. Pagliarini, and D. Parisi, "Emergence of orienting behavior in ecological neural networks," *Neural Processing Letters*, vol. 15, pp. 69–76, 2002.
- [19] E. Goldenberg, J. Garcowski, and R. Beer, "May we have your attention: Analysis of a selective attention task," in *From Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior*, S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J. A. Meyer, Eds. MIT Press, 2004, pp. 49–56.
- [20] A. K. Seth, "Evolving action selection and selective attention without actions, attention, or selection," in *Proceedings of the fifth international conference on simulation of adaptive behavior on From animals to animats 5*. Cambridge, MA, USA: MIT Press, 1998, pp. 139–146.
- [21] N. A. Macmillan and C. D. Creelman, *Detection Theory: A User's Guide*, 2nd ed. Lawrence Erlbaum Associates, 2004.