

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Conflict Monitoring in a Dynamical Agent

Permalink

<https://escholarship.org/uc/item/6zr7q5sm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 28(28)

ISSN

1069-7977

Authors

Ward, Robert

Ward, Ronnie

Publication Date

2006

Peer reviewed

Conflict Monitoring in a Dynamical Agent

Ronnie Ward (rward@acm.org)

Department of Computer Science, Texas A&M University
College Station, TX USA

Robert Ward (r.ward@bangor.ac.uk)

Centre of Cognitive Neuroscience, University of Wales
Bangor, UK

Abstract

This paper examines a dynamical agent for conflict monitoring in cognitive functions acquired through evolutionary learning processes. Existence of conflict in the agent is established using reactive inhibition theory. Agent move hesitation correlates to peak violations of a stable state equation solved by the agent. As points of cognitive conflict, the peak violations are also associated to places of disagreement in source inputs to the agent's control circuit. The control circuit is analyzed in search of an explicit conflict-monitoring mechanism. Analysis of the agent's neural connections assigned by the genetic algorithm suggests distributed conflict management as opposed to a top-down monitoring mechanism.

Keywords: cognitive conflict; evolutionary neural modeling.

Introduction

Conflict in distributed cognitive processing occurs when neural pathways associated with multiple stimuli intersect, and thus interfere with one other such that related task performance suffers (Botvinick, Braver, Barch, Carter & Cohen, 2001). For example, consider a frog eyeing two flies. By targeting a single fly for capture, the frog may succeed if it manages the conflicting input from the second fly. The case to be avoided is an intermediate response in which the frog attacks a midpoint between the flies.

To avoid such performance problems, Botvinick et al. (2001) hypothesized the existence of what they call a top-down conflict-monitoring system where conflict is detected and control invoked to activate appropriate cognitive regulatory processes. To test the sufficiency of their hypothesis, they measured conflict in terms of a rising value of the Hopfield (1984) energy function in the response layer of a discrete interactive model of the Stroop task (Cohen & Hudson, 1994). Energy increased during incongruent trials suggesting that a potential monitoring mechanism could be activated by such a signal, and subsequently causes invocation of appropriate cognitive control elements.

In discussing conflict monitoring, Botvinick et al. (2001) describe control as an active or passive mechanism, which is either on-line and modulated, or off-line based on conflict monitoring information. They modified the Cohen and Hudson (1984) Stroop model by adding recurrent links to

task control nodes back from an inserted conflict-monitoring node. The feedback links strengthened involvement of the color-naming task unit, making response on incongruent trials faster. This correlates to experimentally observed response time improvement on incongruent trials in humans. They suggest that activation of the anterior cingulate cortex in humans is a result of conflict monitoring recruiting executive control to improve performance.

Recent work by Beer (2003, 1996) proposed an artificial visual agent (VA) as an idealized model of a complete brain-body-environment system that exhibits minimally cognitive behavior. Unlike the connectionist modeling methodologies used by Botvinick et al. (2001) and others (Cohen, Dunbar & McClelland, 1990; Cohen et al., 1994), the VA is an embodied, situated and dynamical (ESD) agent operating in continuous time. ESD agents stress what Clark (1998) calls "the unexpected intimacy between the brain, body, and world". As model systems, they emphasize the contextually bound nature of solutions to cognitive problems, and allow a tractable analysis of the type of cognitive processing going on in more complex systems.

While VA is vastly different from a human, it is reasonable to investigate the nature of VA's conflict monitoring, if any, in its cognitive control. The agent has been shown to demonstrate both memory and selective attention (Slocum, Downey, & Beer, 2000; Goldenberg, Garcowski, & Beer, 2004). Moreover, Ward and Ward (2004) found that VA exhibits reactive inhibition (Houghton & Tipper, 1994) such that greater inhibition is observed for more salient targets. VA offers an opportunity to analyze an evolved conflict management mechanism in a dynamical system established through artificial evolutionary processes using genetic algorithms.

This paper examines the VA for conflict monitoring in a dual-task scenario where the agent must select actions in the presence of stimuli suggesting conflicting responses. First, we demonstrate that the agent exhibits cognitive conflict. Periods of conflict are located using a stable-state equation, which the agent solves during processing. These conflict periods are equated with disagreements in the source inputs to the agent's control circuits, which are examined for explicit conflict monitoring, as are the intra- and inter-layer network connections in the agent.

Methods

Visual agents were evolved with the same connection architecture used in the selective-attention experiments of Slocum et al. (2000). As illustrated in Figure 1, our agent has 7 sensor rays of length 220 evenly spaced over a $\pi / 6$ degree visual angle. External input magnitude was 0 to 10, inversely proportional to distance to an object. Seven sensor neurons (S1-S7) were connected bilaterally symmetric to eight hidden units (H1-H8) and two motor units (M1-M2). Units H1-H8 and M1-M2 were fully interconnected in bilaterally symmetric, recurrent fashion. Units H1-H8 were also connected bilaterally symmetric to M1-M2, which in turn were recurrently connected back to H1-H8 in bilaterally symmetric fashion (see Figure 1A). Agent diameter was 30 units and target diameter was 22 units.

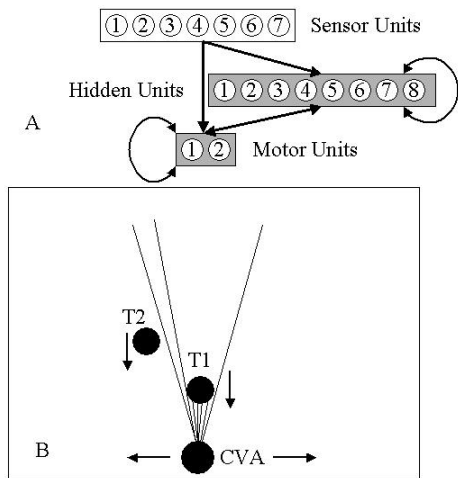


Figure 1. (A) Network layers and connections of the controlled visual agent (CVA). The top unfilled box indicates the seven-node sensor layer, which has no intra-layer connections. The middle box represents the eight-node hidden unit layer, and the lower box illustrates the two-node motor layer. The filled boxes indicate that each unit is connected to every other unit within the layer using bilaterally symmetric weights. Arrows between layers represent bilaterally symmetric connections. (B) The agent moves left and right underneath the falling targets (T1 & T2). T2 is out-of-view (OOV). The rightmost and the two leftmost sensor rays detect nothing. All other rays intersect T1. An agent viewer and associated data files are available at <http://www.psychology.bangor.ac.uk/ward>.

The 102 network parameters were encoded for genetic algorithm (GA) search using GALib (Wall, 1999). See Appendix A. Simulated Continuous-Time, Recurrent Neural Networks (CTRNNs) were based on a derivative of the CTRNN Grid Neural Simulator (Gallagher, 2000).

The agent was required to catch targets T1 and T2 falling simultaneously. T2 serves as a distracting, interfering

stimulus during T1 processing. The experimental factors include: (1) the side on which T1 and T2 appear relative to the agent—Left/Right, (2) the spatial separations between T1 & T2—24 or 48 (designated Near and Far), (3) the temporal (velocity) separation between T1 & T2—4 and 3, or 5 and 2 (designated Near and Far). The Left/Right, Near/Far designations of T1 and T2 dictate 16 two-target trials centered about the agent’s horizontal start position 200 (T1 start positions are 188, 212 and T2 start positions are 140, 164, 188, 212, 236 and 260). A range of T2 saliency exists between Near/Near and Near/Far trials, and between Far/Near and Far/Far trials. The 16 trials with only T1 or T2 were also presented while the agent was evolved. Figure 1B illustrates the agent processing a Far space/Near time trial where T2 is OOV during and after T1 catch.

The 16 two-target trials systematically cover the spread of the sensor array, but better generalization was achieved when the 16 trials were increased to 48 by shifting all start positions left by -8 , and right by $+8$ (increasing total T1/T2 start positions to 6 and 18 respectively). The single T1 or T2 trials were offset in a similar manner to create a total of 48 T1 trials and 48 T2 trials. Thus, a total of 192 targets were processed and the objective function minimized the average agent miss distance (md) from the target at impact.

The factorial design avoids unwanted facilitation of catching T1. In trial design, it is important that the agent’s perception of T1 (or T2) not be informative about T2 (or T1) in the environment. For example, T1 catch average increased from 7 (with T2 present) to 22 (with T2 absent) in our replication of the Slocum et al. (2000) selective attention experiment. Also, the factorial trial structure includes the object passing (OP) and passing object (PO) problems described by Slocum et al. (2000). The T1-T2 Far spatial separation in combination with the T1-T2 Near temporal separation creates Far/Near trials where T2 goes out-of-view (OOV) not only when the agent is in the final stages of catching T1, but after T1 impact when T1 is removed from the environment (exhibiting OP). On Far Space trials, T2 can lie outside of T1 placing T1 between it and the agent. The VA must move past T1 to see T2 in this circumstance (exhibiting PO).

Evolved Agents

Under the above conditions, GA found agents that scored an average md of 0.096 to 0.16 units on the 192 targets. The best controlled visual agent (CVA) had an average catch accuracy of 99.7%. Catch accuracy is a measure of agent/target overlap defined by Slocum et al. (2000) as: $(VA \text{ radius} + \text{target radius} - md) / (VA \text{ radius} + \text{target radius})$. For the two-target trials used during evolution, CVA scored 99.2% on T1 and 99.8% on T2, and 99.8% on both T1 and T2 single target trials.

Generalization Testing

How well did CVA perform using random target start positions and speeds? On 500 random two-target trials with T1 and T2 positions uniformly sampled from ranges [180, 212] and [132, 268], and speeds uniformly sampled from

ranges [4, 5] and [2, 3] respectively, CVA had an average T1 catch accuracy of 95.2%. CVA never failed to catch T1, but on 36 trials, it did not catch T2 (moved away instead of toward T2 after catching T1). CVA had an average T2 catch accuracy of 89% (missed T2 *md* was capped at 26 units, the summed radii of the agent and target).

On 83 of the random trials, following T1 catch, T2 was OOV after T1 was removed from the environment. CVA scored 85.8% catching the OOV T2. When T2 was in-view after T1 removal, CVA scored 89.6% (moved away on 27 trials). Performance was close to perfect on the 1000-single target test trials corresponding to the 500 test trials. CVA scored 99.7% on T1 only trials, and 99.9% on T2 only trials.

Overall performance on the combined 1500 trials was 95.9%. Adding a uniformly distributed noise from [-1, 1] with mean of zero to each sensor input at every timeslice did not substantially alter overall performance.

Selective processing

CVA performance on T1 suggests that it is able to suppress interference from T2 and selectively process T1 (Ward & Ward, 2004). In all T1 and T2 spatial and temporal separation conditions, catch accuracy of T1 on random trials was better than 94%. The agent successfully managed interference from T2. However, the agent missed T1 on the T2 side on Near/Near test trials. In two-target trials, T1 catch accuracy was always higher than that of T2.

Verifying the Existence of Conflict

Identification of conflict periods in CVA using Hopfield's continuous-time energy function (Hopfield, 1984) was rejected since CVA weights are not symmetric. As an alternative, we examined the peak violations of a stable-state equation the agent solves while catching targets.

The agent's final behavior is determined at the motor units by integrating the non-linear inputs from the sensor neurons, hidden units and recurrent motor connections. Re-writing the non-linear CTRNN state equations given in Beer (1996) yields the following left and right (L, R) motor neuron state equations.

$$(1) \quad \begin{aligned} \tau \cdot y'_L + y_L &= \alpha \cdot M_L + \beta \cdot M_R + H_L + S_L \\ \tau \cdot y'_R + y_R &= \alpha \cdot M_R + \beta \cdot M_L + H_R + S_R \end{aligned}$$

The time constant τ , self-weight α and cross-weight β are assumed to be the same for both motor neurons; y_L and y_R are the state values and y'_L and y'_R are time-based derivatives. We define $M_L = \sigma(g(y_L + \theta))$ and $M_R = \sigma(g(y_R + \theta))$, where $\sigma(X) = \frac{1}{1 + e^{-X}}$. The gain g and the bias term θ are the same for each motor neuron. $H_L = \sum_j w_{jL} a_j$ and $H_R = \sum_j w_{jR} a_j$ are the

summed, weighted inputs to each motor from the hidden units. S_L and S_R are similarly defined for the sensor neurons, where w_{jL} and w_{jR} are the weights from neuron j to L, R and $a_j = \sigma(g_j(y_j + \theta_j))$ is the activation of neuron j feeding input to the motor neurons. Each a_j has gain and bias terms g_j and θ_j .

CVA calculates a horizontal move velocity each Euler-integration step defined as:

$$(2) \quad V = (M_L - M_R) / c,$$

where M_L and M_R are in the interval [0,1], and c is a constant. Beer (1996) selected $c = 0.2$ so that $1.0 / c = 5.0$ units per time step is the max velocity of the agent. $V < 0$ moves the agent left, $V > 0$ moves the agent right. If $V = 0$ the agent remains motionless and $M_L = M_R$. The CVA catch behavior is observed as a sequence of moves back and forth underneath a target T , using any difference in motor outputs to cause movement until T impacts.

To isolate $(M_L - M_R)$ in (2), subtract equations (1) to get the following "difference" equation, where the time constant τ is assumed to be 1.0, and $I_L = H_L + S_L$ and $I_R = H_R + S_R$.

$$(3) \quad \begin{aligned} (y'_L - y'_R) + (y_L - y_R) &= \\ k(M_L - M_R) + (I_L - I_R) \end{aligned}$$

Equation (3) has the form of a single neuron state equation $z = y_L - y_R$ with a self-connection weight $k = (\alpha - \beta) \neq 0$ and an activation function whose output $(M_L - M_R)$ is in the interval [-1,1]. The terms on the right-hand side (RHS) govern the state update for z to determine a new output value for $(M_L - M_R)$.

As CVA catches a target the network dynamics relax into a stable state, such that y'_L and y'_R in (3) become zero. Thus, as an optimization problem, the agent moves to solve:

$$(4) \quad (y_L - y_R) - [k(M_L - M_R) + I_L - I_R] = 0$$

at the T catch point. Peak violations of equation (4) are its largest absolute differences from zero. From a goal perspective, CVA should be less "happy" when its position corresponds to a peak violation of equation (4). Do these peak violations of equation (4) indicate cognitive conflict?

Peak Violations are Conflict Periods

According to the reactive inhibition model (Houghton & Tipper, 1994), inhibition should be applied most strongly during periods of greatest conflict. Hence, we make the following argument. (a) Reactive inhibition predicts that more salient T2s require greater inhibition. We have previously observed direct evidence of inhibition in VA (Ward & Ward, 2004), noticeable as a period of hesitation following the catch of T1. When T2 must be inhibited to accurately catch T1, CVA is slow to move towards T2 after catching T1. (b) If sensory input from T2 is removed for some period in the trial, less inhibition of T2 should result, and reduce hesitation. (c) The effects of removing T2 from the sensor inputs should be evident at the time of maximum interference from T2. Finally, (d) the largest peak should predict a conflict point at which removing T2 from the input has greatest impact on hesitation after catching T1.

To test these hypotheses, 1000 random trials were selected, identical in every way except for T2 speed. In this experiment, T1 and T2 were both on the left of CVA, with a spatial separation of 24 (Near). T1 speed was fixed at 4.5, and to establish a range of T2 saliency, T2 speed was chosen randomly from [2, 3.5]. Under these conditions, CVA's average hesitation was 34.78 timeslices after catching T1 (shown in Figure 2 as a dashed line). Four periods of removing T2 input were implemented around the timeslice of the greatest peak violation of equation (4). These include peak-15 to peak; peak to peak+15; peak+15 to peak+30; and peak-15 to peak+15.




Figure 2.

Figure 2 shows the average hesitation after T1 catch for each removal of T2 input; the onset and duration of T2 input elimination are indicated by the width of the shaded bars. The height of each bar extends one standard error above and below the mean ($n=1000$). Removing T2 input before the peak reduces hesitation, and suggests that inhibition of T2 develops during the period leading up to the peak.

Hesitation was significantly less for the two T2 input removals before the peak compared to that of the control

condition. The removals after the peak were marginally different from the control condition. This test suggests that peak violations of equation (4) mark conflict points within CVA.

Source Competition Conflict

If an input source on the RHS of equation (3) differentially activates one motor unit over another, we understand that source as acting to move the agent in a particular direction. Given this, we observe competition between input sources as illustrated in Figure 3. The sensor bias is a plot at each timeslice of $S_L - S_R$, the hidden bias is $H_L - H_R$ and the motor bias is a plot of $k(M_L - M_R)$.

The bias lines illustrate the activity of the source inputs to the motors during the T1 catch portion of a two-target, random trial. T1 starts just right of CVA at 203, speed 4.15, and T2 is Far left at 156, speed 2.2. Figure 3 shows the miss distance, or separation between CVA and T1 as it progresses to catch T1 at impact.




Figure 3. Y-axis miss distance: positive numbers for T1 right of CVA, negative for left. Y-axis bias: positive-going lines indicate the bias to move CVA right; negative-going, to the left. Motor source input disagreements are labeled at maximum differences.

Up until time 200 all the sources hold the agent relatively still. Shortly afterwards, the sensor biases CVA to move rightwards, underneath T1. The sensors tend to keep CVA aligned under the center of perceptual mass. The motor layer initiates a leftward bias toward T2, but this switches to a rightward bias around time 300 when the hidden units also become involved. Notice also that the hidden unit bias and motor bias curves are always in opposition.

At time 350 the hidden layer initiates a large leftward bias so that the agent moves away from T1 and closer to T2. The largest source competition (A) occurs between the hidden inputs and the recurrent motor bias after time 368 when CVA is located farthest from T1. The hidden units push CVA left to consider T2 while conflicting motor bias signals a right turn toward T1. The time of the peak

violation of equation (4) corresponds to the maximum disagreement of the hidden layer with the motor bias at (A).

Around time 400, CVA appears to pause between the attraction of T1 and T2. This is observed as a flattening of the separation curve. Afterwards, CVA moves sharply right toward T1 using the rightward sensor bias while the opposing hidden and motor bias competition relaxes.

At time 433 another source competition (B) appears when both sensor and hidden inputs signal a move right in opposition to the left turn signal from the motor bias. CVA later turns at 441 and moves left to again locate T2. A smaller source competition (C) occurs around time 467 as CVA turns back toward T1. It's resolved at 478. The last significant input competition (D) occurs at 492 as T2 passes OOV when CVA is in the final stages of catching T1. The rightward hidden unit bias overcomes the leftward motor bias to push the agent under and slightly past T1. Points B, C and D in the Figure 3 also correspond to peak violations of equation (4).

Analysis of Control Circuit Activations

Now that cognitive conflict is confirmed in the agent, what form of conflict control can we observe in its control circuit? The agent controls movement by regulating input to the motors [see RHS of equation (3)]. Is the agent monitoring for conflict in the sense of Botvinick et al. (2001) to activate conflict control? Since conflict is observed at turning points, we examine what activations occur as CVA changes directions to resolve conflict.

Without loss of generality, assume the agent is moving left with $k(M_L - M_R) > 0$, i.e. $(M_L - M_R) < 0$ and $k < 0$. This is the case with conflict (A) illustrated in Figure 3 at time 368. Immediately before time 368, the agent has $I_L - I_R < 0$, so the agent's leftward movement agrees with the greatest input I_R .

Conflict initiates at time 368 when the RHS of equation (3) becomes positive, i.e. $I_L - I_R < 0$ but $k(M_L - M_R) + (I_L - I_R) > 0$. In this scenario, the motor bias term $k(M_L - M_R)$ serves as a control "fence" on the input difference $I_L - I_R$ to affect a possible, but not guaranteed change in the agent's move direction (depending on future changes to I_L and I_R). The agent acts to make a right turn away from the greater source I_R . Based on absolute state values, it may take additional state updates to actually reverse the agent's direction.

At time 369, however, the input difference flips such that $I_L - I_R > 0$ is true. The conflict continues because the agent's movement left is in disagreement with the now greater input source I_L . The conflict over move direction is 'resolved' after four state updates at time 373 when the difference between M_L and M_R comes back into

directional agreement with the difference between I_L and I_R (i.e., when the agent changes direction and turns right toward T1).

The change to a positive update to the state of z from time 369 to 373 occurred as hidden unit H7 deactivated (activation fell from 0.95 to 0.01). All other hidden unit activations went unchanged during conflict period (A). Analysis of the other conflict points in Figure 3 shows that it is not always H7 that's deactivated. No hidden unit activation changes occur during conflicts (B) and (D). During conflict (C), hidden unit H8 was deactivated to 0.01 while the other units remained unchanged. This evidence suggests some means of distributed conflict management.

Role of Feedback Links in Conflict Monitoring

Botvinick et al. (2001) describe conflict monitoring as activating executive control to bring its regulatory power on-line, or to modulate it. To model this effect they implemented recurrent feedback links from the inserted conflict-monitoring unit back to the model control units. Similarly, CVA has feedback links from the motors to the hidden units (see Figure 1A), so we investigated what affect these links have on movement control. Consider Table 1, which shows how GA partitioned the neural units in CVA, and how GA weighted the intra- and inter-layer connections in the control circuit.

In the leftmost column, the neural units termed Groupings (S-sensors, H-hidden units, M-motors) are given along with their associated movement bias (L-left, R-right). Entries in the table indicate the connection relationships (I-inhibitory, E-excitatory) between Groupings. For simplicity, the middle sensor (S4) is grouped with both left/right sensor groups.

To read the table, consider the first row. When sensors S1-S4 are activating, they tend to move the agent left by exciting H5-H8 and M2 whose activation also moves the agent left. However, S1-S4 inhibits H1-H4 and M1 because activation of these units tends to move the agent right.

Table 1: CVA Neural Structuring by the Genetic Algorithm.

Groupings	H1-H4	H5-H8	M1	M2
S1-S4 (L)	I	E	I	E
S4-S7 (R)	E	I	E	I
H1-H4 (R)	E	I	E	I
H5-H8 (L)	I	E	I	E
M1 (R)	I	E	I	I
M2 (L)	E	I	I	I

With only feed forward links, the sensor Groupings appear to have a "reactionary" structure. That is, they excite or inhibit the hidden units and motors based only on their desired movement direction. However, the hidden unit groups have a "competitive" structure. The groupings are self-exciting and mutually inhibitory. Like the sensors, they also excite or inhibit the motors based on desired movement direction.

Most interesting is the motor activation effect, which inhibits both motors. This tends to stop movement! Moreover, motor activation inhibits the same-direction hidden-unit group, but excites the opposite-direction hidden-unit group. This inter-layer competition seems strange, but it acts to turn the agent in the opposite direction. For example, in Figure 3 around time 300 when the hidden units H5-H8 are acting to move the agent left by activating M2, M2 is activating to shut down H5-H8 and excite the opposite group H1-H4 to turn the agent around.

If the motor links to the hidden units are lesioned, and the same 500 random two-target trials described earlier are processed, T1 catch accuracy falls to 93.6% from 95.2%. This modest change indicates that CVA still manages T2 distraction to successfully select and catch T1. However, cutting the motor feedback links devastates T2 catch accuracy. It dropped from 89% with these links active to 61.3% without them. An explanation of this phenomenon is outside the scope of this paper.

The bias of the feedback links is almost identical to the motor bias curve as shown in Figure 3. In other words, during peak conflict, the feedback links from the motors act to shut down movement in the conflicting direction. This could be viewed as an invocation of conflict control if the activations were brought on-line at that point from an off-line state. However, the feedback signal is persistent and grows, or falls in strength in tandem with the motor bias value, which we have already examined and believe not to be a mechanism of top-down conflict monitoring, but one of distributed conflict management.

Conclusion

Beer's VA is a robust model capable of demonstrating selective attention, memory, and inhibition. Its dynamical underpinnings predict conflict points as peak violations of a steady state equation associated with its behavioral goal to catch targets. According to reactive inhibition theory, these peaks correlate with agent hesitation after catching a target, and can be understood as points of conflict. These conflicts are associated with disagreement in inputs to the agent's move control circuit. Analysis of this circuit reveals no explicit conflict monitoring in the agent, such as the mechanism proposed Botvinick et al. (2001). Instead, our results suggest that the competitive groupings in VA's hidden layer provide a distributed conflict management mechanism, and give insight into its evolved motor control.

Appendix A

The GALib parameters were set as follows. The algorithm used was GASteadyState with a population sizes from 25 to 300, replacement percentage of 50-75%, cross over probability of 96%, and mutation probability of 10%. The GARealGenome (real-valued vector genome) was used with allele sets for the 102 network parameters in the following ranges: weights [-10, 10], time constants [1, 2], sensor ray biases [-10, 10], hidden unit and motor biases [-5, 5], gains [1, 5] and motor gains were set to 1.0. The default random number generator in GALib was used.

References

- Beer, R.D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences* 4(3): 91-99.
- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 421-429). MIT Press.
- Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent (with commentary and response). *Adaptive Behavior*, 11, 209-243.
- Botvinivk, M., Braver, T. Barch, D., Carter, C., & Cohen, J. (2001). Conflict Monitoring and Cognitive Control. *Psychological Review*, Vol. 108, No. 3, 624-652.
- Clark, A. (1998). Where brain, body and world collide. *Daedalus: Journal Of The American Academy Of Arts And Sciences*, 127, 257-280.
- Cohen, J. D., Dunbar, K., McClelland, J. L. (1990). On the Control of Automatic Processes: A Parallel Distributed Processing Account of the Stroop Effect. *Psychological Review*, Vol. 97, No. 3, 332-361.
- Gallagher, J. C. (2000), Programmer's Notes for CGNS.C (CTRNN Grid Neural Simulator) Library, Department of Computer Science and Engineering, Wright State University, Dayton, OH 45435.
- Goldenberg, E., Garcowski, J., & Beer, R.D. (2004). May we have your attention: Analysis of a selective attention task. *From Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior*. July 13-17, 2004, Los Angeles, CA.
- Hopfield, J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Nat'l. Acad. Sci., USA*, Vol. 81:5, 3088-3092.
- Houghton, G., & Tipper, S. P. (1994). A model of inhibitory mechanisms in selective attention. In Dagenbach, D., & Carr, T. H. (Eds.), *Inhibitory processes in attention, memory, and language*. Academic Press, San Diego, CA.
- Slocum, A.C., Downey, D.C. & Beer, R.D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat and S. Wilson (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior* (pp. 430-439). MIT Press.
- Wall, M. (1999). GALib, A C++ Library of Genetic Algorithm Components. Massachusetts Institute of Technology. Retrieved Feb. 17, 2003 from <http://lancet.mit.edu/ga/>.
- Ward, R. & Ward, R. (2004) Selective attention and action in an artificial, evolved agent: Reactive Inhibition. In Cangelosi, A., Bugmann, G., & Borisyuk, R. (Eds.). *Modelling Language, Cognition, and Action: Proceedings of the Ninth Neural Computation and Psychology Workshop*. London: World Scientific Publishing.