

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Opponent Process Control in Linked, Dynamical Agents

Permalink

<https://escholarship.org/uc/item/8sm8f941>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 30(30)

ISSN

1069-7977

Authors

Ward, Ronnie G.

Ward, Robert

Publication Date

2008

Peer reviewed

Opponent Process Control in Linked, Dynamical Agents

Ronnie G. Ward (ward@cs.tamu.edu)

Dept. of Computer Science; Texas A&M University
College Station, TX USA

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

Wolfson Centre for Clinical and Cognitive Neuroscience; Bangor University
Bangor, UK

Abstract

Previously, we showed that a minimally cognitive, visual agent demonstrated selective attention and reactive inhibition (Ward & Ward, in press). Surprisingly, we discovered the existence of an opponent-process architecture in our agent's evolved, neural-network controller. Here, we investigate how opponent processes affect response control in the agent. We scale up cognitive problem solving by evolving links between multiple copies of the visual agent to solve tasks that a single agent cannot work out alone. Opponent processing effects are demonstrated in the linked agent's response control.

Keywords: opponent process; inhibition; selection; attention; genetic algorithm; neural network; linked agents.

Introduction

Research suggests that opponent processes may be ubiquitous in cognition, and they are key inhibitory mechanisms for implementing response control (Bowman, Schlaghecken, and Eimer, 2006). As illustrated in Figure 1, implementation of opponent processes has activation of a response followed by inhibition of that response through excitatory linkage to an opponent process (Hurvich, L., & Jameson, D., 1974).

Houghton and Tipper (1994) developed a representational model of inhibitory mechanisms used in cognitive response control that captured valuable insights concerning the operational dynamics of opponent processes in realizing reactive inhibition. Burgess and Hitch (1999) demonstrated a connectionist model of the articulatory loop, which employed opponent processes to deliver decaying inhibition to a model-selected, lexical response. Bowman, et al., (2006) subsequently developed a computational model explaining how response activation from subliminal priming can be inhibited and suppressed using an analogous opponent process for response retraction.

Using an approach of artificial comparative psychology, Ward & Ward, (in press) reported opponent-process architecture in the evolved network structure of a non-representational visual agent capable of selective attention and action. This discovery was exciting because our agent's cognitive solution developed without supervised learning, or general-purpose representation as used in the connectionist models mentioned above. It emerged in a dynamical environment in the context of a genuine perception-action loop where senses guide action, and action generates new sensory inputs, and the action involves an extended

sequence of motor commands in response to sensory inputs. This development exemplifies what Clark (1999) calls "the unexpected intimacy between the brain, body, and world". Our agent was based on an environmentally situated and dynamical visual agent originally developed by Beer (1996, 2003) and Slocum, Downey & Beer (2000).




Figure 1. Opponent-process architecture (Bowman, et al., 2006). An activated response deactivates itself by exciting an opponent, which in turn inhibits the response.

The existence of opponent-process architecture in our implementation of the visual agent potentially has important implications because we also found that people exhibit similar results in equivalent "catching experiments" (Ward & Ward, in press). Hence, we show here in some detail the structure of our agent's evolved, neural-network controller, and identify the response-control mechanisms developed by a genetic algorithm. We examine the regulatory nature of opponent processes in controlling the agent behavior in the context of conflicting stimuli.

In addition, Beer (in press) raised the general question of how to "scale up" visual agents to solve more cognitively demanding problems. Our approach here is to link together multiple, fixed-agents through subsequent evolution in an extended environment and investigate what control mechanisms develop. As a preview, experimental results demonstrate the presence of opponent processes in the control network of such agents. We conclude with observations about opponent processing, and offer suggestions for further research concerning the neural organization required for successfully scaling cognitive power.

The Visual Agent

The visual agent is a tractable model that can perform interesting cognitive tasks. For example, Figure 2 shows the visual agent as presented in Ward & Ward (2006, in press).

A genetic algorithm evolved the agent to selectively attend and respond to targets in a 2D environment. The agent has an array of seven proximity sensors, a small hidden layer of eight units, and left and right motor units to effect its movement. Its control circuit is a continuous-time recurrent neural network (Beer, 1996), with reasonably generic connections (see Figure 2).

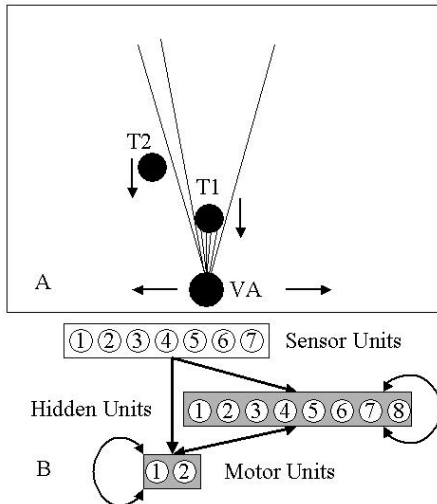


Figure 2. The visual agent (VA) catching two targets, which was originally presented in Slocum et al, (2000). Part A of the figure shows a schematic of the agent and environment. The agent was evolved to catch two targets (T1 and T2), falling from the top of the environment. VA has left and right motors, enabling it to move along the bottom of the environment to catch the targets as close as possible center-to-center. The rays emanating from VA represent proximity sensors. Part B illustrates the continuous-time recurrent neural network (CTRNN) that controls VA. A genetic algorithm evolved its parameters. The seven input units each pass on activation from an associated proximity sensor. The activation of the input units reflects where the proximity sensor intersects a target. Each input unit is connected to both motors and to each hidden unit (HU). Each of the eight HUs receives connections from every other unit in the network, and sends connections to all hidden and motor units. The two motor units (M1, M2) receive inputs from both sensor and hidden layers, and send connections back to the hidden layer. The motors are self-connected and connected to each other. A shaded box indicates fully recurrent connections between the designated nodes.

Targets fall from the top of the 2D environment, and VA catches them by aligning its center under one and then the other as the targets impact at the bottom. In our experiments, the two targets fall straight down, and with constant velocity, the first target (T1) by definition has greater velocity than the second target (T2). This two-ball catching task requires many cognitive operations, including: (1) prioritizing T1 over T2; (2) selectively focusing

responses on T1, while preventing T2 from interfering with the responses; (3) creating a memory for the unselected T2 item, so that it can be efficiently processed later; and (4) reallocating processing towards a perhaps unseen T2 after catching T1. Our evolved agent demonstrated all these abilities. Note, these processes of selection, response control, and reconfiguration following a change of targets are all important themes in current selective attention research (for an overview, see Driver, 2001).

Earlier, we analyzed our agent’s control circuit for explicit conflict-monitoring in periods of cognitive conflict defined by peak violations of its stable state equation and time course disagreements in its source inputs (Ward & Ward, 2006). Analytical and simulation results implied a distributed conflict management system rather than a top-down monitoring mechanism as suggested by Botvinick Braver, Barch, Carter & Cohen, (2001, 2004).

Response Control Mechanisms

Figure 3 shows the evolved structure of the neural units in our agent. A genetic algorithm assembled neural units into left and right move groups, and weighted the intra- and inter-layer connections in the control circuit as shown.

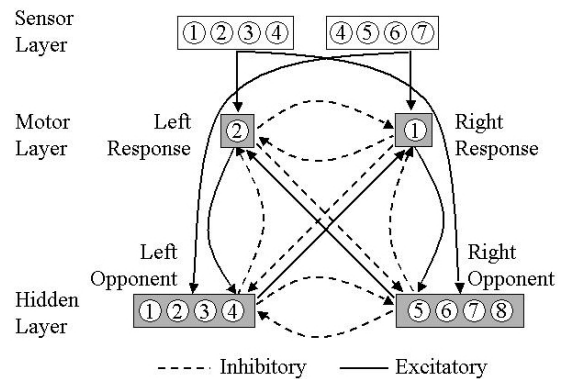


Figure 3. A genetic algorithm organized VA neural units into left/right move-bias groupings as illustrated. The summed weights from each group are labeled excitatory if positive, and inhibitory if negative. For simplicity, the middle input sensor (4) is grouped with both left/right sensor groups. As an example, activation of hidden units 1-4 (HU1-4) tends to move the agent right as they feed excitatory input to M1 and inhibitory input to HU5-8 and M2, which are biased to move the agent left. The activation of M1 excites HU5-8, which in turn inhibits M1. This structure, learned through evolutionary pressure, is an opponent architecture. For simplicity, not all connections are shown in the Figure.

We omit detailed discussion of each neuron, and present the summed connection weights of the grouped units as an approximation of group activation effect. With only feed forward connections, the left/right groupings in the sensor layer appear to have a “reactionary” role. They excite the hidden units and motors based on their biased move direction.

As illustrated in Figure 3, lateral inhibition evolved between the two hidden layer groups and between the two motors. Mutually inhibitory links such as these implement a “winner take all” control mechanism for resolving response competition (Bowman et al., 2006). The ultimate decision to move left or right is an “either/or” choice that’s won or lost on two levels. First, a winning HU group excites or inhibits the motors based on its desired movement direction. Next, a winning motor activation excites the opposite-direction HU group, which in turn inhibits the activating motor. This inter-layer competition is an implementation of opponent processing as illustrated in Figures 1 and 2. Houghton & Tipper (1996) describe this architecture as a “gain control” circuit. They argue that such a dual mechanism of excitation and inhibition is essential for observing selective attention in a cognitive agent such as VA.

Opponent Processes

Opponent processes in VA regulate activation of the motors by release of inhibition. To illustrate, Figure 4 presents the effects of interfering (or not) with opponent processing in VA. The figure illustrates a test trial while VA is processing T1 in the presence of T2 (a conflicting stimuli). T1 starts at position 203, falling at speed 4.15, just right of VA at position 200, and T2 is left of VA at position 156, speed 2.2.

The curve labeled “VA v T1” plots the horizontal, center-to-center distance between a normal, unlesioned VA and T1. VA catches T1 at its impact (timeslice 540). The curve labeled “VA v T1 OPM Cuts” is a similar plot, but for a lesioned VA. In this case, periodic cuts are made in the left-move opponent-processing circuit (when HU5-8 are acting on M2 to move the agent left, M2 excites its opponent group, HU1-4). Whenever VA moved left for 25 timeslices, the connections between HU1-4 and M2 were cut in both directions for 10 timeslices. After 10 successive left moves, the links were restored until the agent again moved left for 25 slices. Essentially, the left-move opponent links were toggled off for 10 time slices and on for 25 time slices if the agent makes a long series of left moves such as those beginning around time 300.

The regulatory effect of inhibition through an opponent process is reduced when the lesions are active. Under normal conditions (curve “VA v T1”), opponent processing suppresses the initial motor activations until sufficient target input is received through the sensor units to overcome opponent inhibition, which happens around timeslice 223 (with normal processing). However, with the lesions active, left-move inhibition is reduced, and the agent moves further left than in the normal case at times 75, 325, and 475 (contrast curves “VA v T1” and “VA v T1 OPM Cuts”). Hence, opponent processes avoid this type of over response.

Response Retraction

Once an agent starts moving in a selected direction, how is that response ever retracted? Bowman et al, (2006) argued that this function is one role of opponent processes. This function can be observed in VA in Figure 4, by comparing

“M1 Inhibition” and “M2 Inhibition” plots, which are shown for the normal processing (unlesioned) case. An activated motor excites an opposing group of HU units, which feed inhibition back to the selected motor. It also receives inhibition from the losing motor. For example, M1 activates at the point labeled A (time 225) to move the agent toward T1. This is followed by a steep magnitude increase in inhibition delivered to M2 through opponent processing (see the curve labeled “M2 Inhibition” drop below -10 around time 230). Note the small increase in M1 inhibition through opponent links (see the curve labeled “M1 Inhibition” around time 230). This inhibition lags M1’s activation, but at time 235 with reduced sensor input from T1 (the currently selected target), the agent eventually slows, and de-selects M1 around time 240. Afterwards, VA activates M2, reverses direction, and selectively attends T2. Opponent processing serves to turn the agent around based on changing sensory input from T1 and T2.

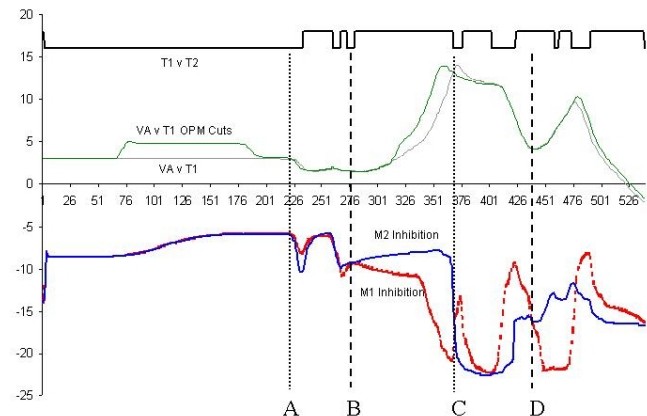


Figure 4. VA target selection and opponent-circuit inhibition release. Curve “T1 v T2” approximates VA selectively attending T1 (low) versus T2 (high) during normal processing of T1. T1 is selected if in the previous time step, VA’s change in proximity toward T1 is greater than that toward T2. Otherwise, T2 is the selected target. Curve “VA v T1” plots the horizontal separation between VA and T1 during normal processing as the agent progresses to catch T1 on its right with T2 falling on its left. Curve “VA v T1 OPM Cuts” plots the altered separation when periodic lesions are performed in the left-move opponent-process circuit. “M1 Inhibition” and “M2 Inhibition” curves graph the inhibition to the motors released through their opponent connections (see Figure 3). For example, at time 286, the agent selects T2, causing T1 to become the distractor. As predicted by Houghton & Tipper (1996), selection of one target causes a rise in inhibition of the distractor, which is reflected in VA by an increase in M1 inhibition. Examples are discussed in the text at points A, B, C and D illustrating increases in distractor inhibition as well as distractor equilibrium and inhibitory rebound (Houghton & Tipper, 1996).

Moreover, at the point labeled B (time 275) we can see examples of opponent inhibition as described by Houghton

& Tipper (1996). M2 activates to respond to the selected target T2. T1 becomes the distractor, and M1 is deactivated as it receives a slow increase in inhibition from its opponents until time 330-338 when it briefly flattens in what appears to be “distractor equilibrium” (M1 activation levels off). At time 338 a large increase of “inhibitory rebound” is released to M1 through opponent circuits (between points labeled B and C) as the agent moves increasingly away from T1 (changing sensory input from T1 and driving M1 activation to a low point). During this time, inhibition of M2 gradually decreases (smaller magnitude, negative values) until time 356. M2 is deselected when it receives a lengthy spike in opponent inhibition until point C (time 366). VA then activates M1 at time 376, changes direction to attend T1, and T2 becomes the distractor.

After point D (time 441), another instance of distractor equilibrium can be observed through M1 inhibition. The agent activates a response to T2 and sharply increases M1 inhibition until time 451 when it flattens until time 471. This suggests equilibrium in managing T1 distraction. Lack of space prevents showing the relationship between T2 salience and its level of inhibition during T1 processing.

Response control in combined agents

These above analyses reveal the psychological mechanisms for control of action within VA. However, VA is (intentionally) a limited agent. Let us assume for exposition’s sake that the human brain has on the order of 100 billion neurons (about 20 billion cortical neurons (Pakkenberg & Gundersen, 1997), and about 100 billion cerebellar neurons (Andersen, Korbo, & Pakkenberg, 1992), each one of which has far more sophisticated processing capabilities (Graham & van Ooyen, 2004) than the 17 simple CTRNN units we used with VA. Further, there are an estimated 240 trillion synapses in the cortex alone (Koch, 1999), compared to 170 synapses for VA. This is a big gap. So, how do we scale up from a minimally cognitive agent to a brain? Which direction do we go? With more VAs, or with bigger VAs? In other words, is the human brain’s cognitive architecture best thought of as a relatively small number of large, multi- or general-purpose networks, or as a collection of billions of small special-purpose agents? Here we investigate the later possibility by evolving a new agent we called the linked visual agent (LVA), and investigate its response control mechanisms.

Linked Visual Agents

To form LVA, two copies of our VA network (Ward & Ward, 2006) were linked together by a single set of shared motor units plus a new set of recurrent links between the HUs of each VA (see Figure 5). These “component VAs” were identical to our original VA except their sensors detect different color targets, red or green. The sensors of the green VA can only detect green targets; red targets cannot activate its sensors. The sensors of the red VA can only detect red targets, and not green ones. The sensor arrays of the

component VAs were physically aligned so as to always perceive the identical region of space.




Figure 5. The Linked Visual Agent (LVA). LVA has two identical component agents each with the same network structure and parameters as VA (see Figure 2). The sensors in each component VA can only see targets in one color—red or green. The input and hidden layers of each component VA project to a shared motor layer, where they compete for control of motor activity. Note that the proximity sensors of the red and green component VAs are exactly aligned so that they receive input from the identical region of space. Communication between the VAs is added by a set of fully recurrent links (weights) between the hidden layers of the two component agents.

If LVA processes two green targets, for example, they are invisible to the red VA. On the other hand, if a red and a green target are dropped, each of the component VAs will attempt to catch the target it can see. This puts the two VAs in a state of cognitive conflict, one pulling towards red and the other towards the green target. To resolve this conflict, recurrent links were added between the hidden units in one component agent to each of the hidden units in the other. These links consisted only of a weight matrix, which was evolved in an environment similar to that of VA except that the number of evolutionary trials increased by a factor of four (target combinations of 2 colors) over those used to develop VA (Ward & Ward, 2006). The previously developed weights and network parameters of the component VAs were fixed during the evolution of the additional inter-agent links.

LVA was able to process the mixed-color targets with high accuracy, a task neither of its component agents alone could accomplish. T1 and T2 catch accuracy were over 99%. Analyzing the new inter-agent links, we found that a genetic algorithm established mutual inhibition between the agents’ HU groups. When LVA decides which color agent controls the motors, it “activates” that agent over the other one using a winner take all strategy. The winning agent reinforces its selection by feeding inhibition via the inter-agent links to left/right HU groups in the losing agent. LVA then responds to the targets under control of the winning agent.

Opponent Processes in Linked Agents

Some results of this simulation are shown in Figure 6, which summarizes hesitation times for 500 novel same-

color trials and 500 equivalent mixed-color trials. LVA showed an overall cost for reallocating resources from T1 to T2, a pattern very similar to VA. After catching T1, LVA hesitated for a significant period before moving again to catch T2. Figure 6 illustrates hesitation to respond to T2 as a function of same vs mixed-color targets, and whether T2 was in-view or out-of-view (OOV) after T1 catch. We have previously reported evidence of reactive inhibition with VA based on this comparison. After catching T1, VA was faster to respond to an OOV T2 than an in-view one. This result might initially seem counter-intuitive but follows naturally from the idea of reactive inhibition. The OOV T2 required less inhibition during T1 processing, and so subsequent release of inhibition was faster (Ward & Ward, in press).




Figure 6. Hesitation in LVA. Hesitation after T1 catch is shown for Same-color (T1 and T2 are both red or both green), and Mixed-color (one red target, one green target) trials. For comparison, hesitation of our VA behavior is shown for OOV (dotted line) and in-view (dashed line) trials.

LVA's performance on same-color targets was almost identical to the component VAs. In particular, hesitation was longer when T2 is in-view compared to OOV, $F(1,498)=137.8$, $p<0.0005$, suggesting that inhibition is proportional to target salience. This result is expected, and simply shows that the addition of the inter-agent links did not fundamentally change the operation of the component VAs. This is evidence that the component VAs still used reactive inhibition from opponent processes to selectively respond to T1, and so were slower to reallocate processing towards salient T2s.

A similar pattern can be observed for the mixed-color trials. Again, when reallocating from T1 to T2 there was a significant hesitation, and this hesitation was reduced for the less salient OOV targets, $F(1,498)=78.5$, $p<0.0005$. Here we also see evidence for increased inhibition through opponent processes for salient T2s. In this way, selective attention and action in the LVA is coordinated using mechanisms very similar to those in the component VAs. Note that while hesitation for LVA with same-color targets was very similar to those of the component VAs, hesitation was in fact slightly reduced for LVA, $F(1,498)=7.52$, $p=0.006$.

Evidently LVA was able to reallocate more efficiently than a component VA. The better performance for LVA must then be due to the use of units in the "other" component VA, mediated by the inter-agent links. That is, on a trial with two green targets, LVA could reallocate more effectively than the component green VA could on its own. This suggests that LVA was using units in the red component VA to assist performance on trials with two green targets.

Another interesting result evident in Figure 6, is that reallocation was more efficient in the mixed than same color case, $F(1,98)=320$, $p<0.0005$. That is, it appears that cognitive conflict produced by the two targets was greater when both targets were loading on the same component VA. This suggests that T2 attracted less inhibition in the mixed than the same color case. This is to be expected since targets in the mixed case weren't competing for the resources of the same component VA. Less competition means less need for inhibition, and faster subsequent reallocation.

If both targets are of the same color, the winning agent uses its own opponent processing to manage LVA movement. But in a mixed-color target environment, how is a selected agent ever de-activated? LVA appears to leverage opponent processing in VA to fulfill this purpose. Recall, the motors are shared. An activating motor feeds excitatory input not only back to the winning agent's opponent HU group, but also to the corresponding opponent HU group in the losing agent. This acts to retract control from a selected agent as follows. From an LVA perspective, the losing agent HU group excited by the activated motor feeds inhibition back to both the winning agent HU group (through the inter-agent links) as well as the opposite motor (through its opponent links). Hence, the losing agent HU groups act as opponent units feeding inhibition to both sets of HUs in the winning agent in an effort to shut down their control. As the losing color agent's sensory perception of its color target activates its HUs, greater inhibition flows through the inter-agent links to the winning agent. LVA eventually retracts selection of the winning agent, and the other agent takes control over the motors. Since the motors are shared, excitatory input is feed back to opponent HU groups within both the winning and losing agents. An activating motor thus has two opponents--one in each agent. As it excites these opponents, both in turn inhibit it. So we observe opponent processes in LVA as well as VA.

Conclusion

The VA is particularly interesting as a research tool for investigating cognitive processes, such as response control mechanisms. In VA's neural circuit, we unexpectedly observed an opponent-process architecture (Houghton & Tipper, 1994). Importantly, VA is not a representational model designed to follow existing empirical work, so we can expect such surprises. As determined by a genetic algorithm, VA's neural circuit also embodies classical lateral inhibition (Bowman, et al., 2006) between the motors and also between competing hidden-unit groups.

Inhibition release through opponent processes was found to regulate agent movement, prevent over-response, and switch agent move direction in conjunction with changing stimuli. Such opponent processes affect selective attention and reactive inhibition, which VA has been shown to demonstrate (Ward & Ward, in press).

Here, we also observed similar capabilities in linked agents solving a task that individual agents cannot successfully process alone. In mixed-color test trials, the LVA demonstrated efficiencies not seen in the individual component agents suggesting the feasibility of “scaling up” cognitive capabilities by combining existing agents. These simulation results are indicative of opponent processes in LVA. Inhibition does not flow from a top-down “central inhibitor”, but from an organized and distributed structure. A genetic algorithm evolved lateral inhibition between the hidden layers of the linked agents to prevent move paralysis resulting from the competitive activation of both agents.

Because of its possible implications concerning the human brain's cognitive architecture, future research areas include greater scaling of linked-agent capability not only by increasing the number of fixed VAs, but by varying their kind and behavior possibilities. For example, we have successfully evolved visual agents that solve various cognitive problems using additional behaviors, such as shape recognition and target inference which employ avoidance behavior (Beer, 1996). A thorough investigation of the response control mechanisms, and the role of inhibition in such agents are warranted to determine which of perhaps several agents actually control linked-agent behavior. A cognitive architecture described by linking large numbers of fixed-function agents may be a reasonable means for scaling cognitive power (Minsky, 1986; Singh, 2003). The pros and cons of this suggestion have been debated among computer scientists, but it may also be worth serious consideration, as a *neuroscientific* hypothesis, the idea that human cognition emerges from the interaction of massive numbers of connected, small networks.

Acknowledgement

The authors are grateful for discussions with Yoonsuck Choe of Texas A&M University and Derek Harter of Texas A&M University-Commerce.

References

Andersen, B. B., Korbo, L., & Pakkenberg, B. (1992). A quantitative study of the human cerebellum with unbiased stereological techniques. *Journal of Comparative Neurology*, 326(4), 549-560.

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*, MIT Press, 421-429.

Beer, R. D. (in press). The dynamics of brain-body-environment systems: A status report. To appear in P.

Calvo and A. Gomila (Eds.), *The Elsevier Handbook of Embodied Cognitive Science*. Elsevier.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624-652.

Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends in Cognitive Science*, 8(12), 539-546.

Burgess, N. & Hitch, G. J. (1999). Memory for Serial Order: A Network Model of the Phonological Loop and its Timing. *Psychological Review*, 106(3), 551-581.

Clark, A. (1999). Where brain, body, and world collide. *Cognitive Systems Research*, 1, 5-17.

Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology*, 92(Pt 1), 53-78.

Graham, B. P. & van Ooyen, A. (2004). Transport limited effects in a model of dendritic branching. *J Theor Biol*, 230(3), 421-432.

Houghton, G. & Tipper, S. P. (1994). A model of inhibitory mechanisms in selective attention. In D. Dagenbach & T. Carr (Eds.), *Inhibitory mechanisms in attention, memory, and language*. San Diego, CA: Academic Press, 53-112.

Houghton, G. & Tipper, S. P. (1996). Inhibitory Mechanisms of Neural and Cognitive Control: Applications to Selective Attention and Sequential Action. *Brain and Cognition*, 30, 2-43.

Houghton, G., Tipper, S. P., Weaver, B., & Shore, D. I. (1996). Interference and Inhibition in Selective Attention: Some Tests of a Neural Network Model. *Visual Cognition*, 3(2), 119-164.

Hurvich, L., & Jameson, D. (1974). Opponent processes as a model of neural organization. *American Psychologist*, 29, 88-102.

Koch, C. (1999). *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press.

Minsky, M. (1986). *The society of mind*. Simon & Schuster, Inc. New York, NY, USA.

Pakkenberg, B. & Gundersen, H. J. (1997). Neocortical neuron number in humans: effect of sex and age. *Journal of Comparative Neurology*, 384(2), 312-320.

Singh, P. (2003). Examining the Society of Mind. *Computing and Informatics*, 22, 521-543.

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. Paper presented at the From animals to animats: Sixth International Conference on Simulation of Adaptive Behavior, Cambridge, MA, 430-439.

Ward, R. & Ward, R. (in press). Selective attention and control of action: Comparative psychology of an artificial, evolved agent and people. *Journal of Experimental Psychology: Human Perception and Performance*.

Ward, R. & Ward, R. (2006). Cognitive conflict without explicit conflict-monitoring in a dynamical agent. *Neural Networks*, 19(9), 1430-1436.