Available online at www.sciencedirect.com
ScienceDirect
Neural Networks 19 (2006) 1430–1436
Neural Networks
www.elsevier.com/locate/neunet

2006 Special Issue

Cognitive conflict without explicit conflict monitoring in a dynamical agent

Robert Ward[a,*], Ronnie Ward[b]

[a] Centre for Cognitive Neuroscience, University of Wales, Bangor, Bangor LL57 2AS, UK
[b] Department of Computer Science, Texas A&M University, College Station, TX, USA

Received 9 August 2006; accepted 16 August 2006

Abstract

We examine mechanisms for resolving cognitive conflict in an embodied, situated, and dynamic agent, developed through an evolutionary learning process. The agent was required to solve problems of response conflict in a dual-target "catching" task, focusing response on one of the targets while ignoring the other. Conflict in the agent was revealed at the behavioral level in terms of increased latencies to the second target. This behavioral interference was correlated to peak violations of the network's stable state equation. At the level of the agent's neural network, peak violations were also correlated to periods of disagreement in source inputs to the agent's motor effectors. Despite observing conflict at these numerous levels, we did not find any explicit conflict monitoring mechanisms within the agent. We instead found evidence of a distributed conflict management system, characterized by competitive sources within the network. In contrast to the conflict monitoring hypothesis [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], this agent demonstrates that resolution of cognitive conflict does not require explicit conflict monitoring. We consider the implications of our results for the conflict monitoring hypothesis.

Keywords: Conflict monitoring; Anterior cingulate; Dynamical agents; Genetic algorithms; Artificial comparative psychology

Cognitive conflict occurs when neural pathways associated with different concurrent processes or representations interfere with one another (Botvinick, Braver, Barch, Carter, & Cohen, 2001). For example, consider a frog eyeing two flies, one to the left and one to the right. Each fly might activate a specific attacking response to the fly's location. But if only one attack can be made at a time, the frog is in danger of some kind of incoherent response, like attacking a midpoint between the two targets. The frog therefore needs to manage the conflicting responses and focus on a single target. Conflict tasks have a long history of study in psychology, with the Stroop task being perhaps the best known, in which the identity of a colored word can interfere with color naming (Mari-Beffa, Estevez, & Danziger, 2000).

How might conflict be managed? An influential proposal was developed by Botvinick et al. (2001), who suggested a two-staged "evaluate–regulate" approach to the control of conflict. Botvinick et al. hypothesized a top-down conflict monitoring system, which first detects conflict in underlying neural structures, and second, invokes control mechanisms to regulate processing in a task-appropriate way. Botvinick et al. (2001) embedded this system within a number of discrete interactive models of conflict tasks, including the Stroop model of Cohen and Huston (1994). Conflict was measured by monitoring the Hopfield (1982) energy function in the response layer of the Stroop model. Energy increased during incongruent trials, suggesting that a potential monitoring mechanism could be activated by such a signal. Appropriate cognitive control could then be subsequently evoked. In the Botvinick et al. (2001) model, conflict monitoring was a localized function of a dedicated module. Botvinick et al. (2001) further speculated that in the human brain, the proposed conflict monitoring system could be localized to the Anterior Cingulate Cortex (ACC). In support of a relatively localized module for conflict monitoring, Botvinick et al. (2001) (see also Botvinick, Cohen, and Carter (2004)) considered a variety of evidence from brain imaging studies, showing ACC activation during conditions producing interference due to response conflict (Botvinick, Nystrom, Fissell, Carter, & Cohen, 1999; Carter et al., 2000; Casey et al., 2000).

In this report we investigate other possible forms of conflict monitoring in a "minimally cognitive agent" (Beer,

* Corresponding author. Tel.: +44 1248 382211.
E-mail address: r.ward@bangor.ac.uk (R. Ward).

0893-6080/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2006.08.003

2003). Beer (1996, 2003) suggests the use of an idealized "Visual Agent" (VA). Unlike connectionist approaches in which action is represented by the activation of an output node (e.g. Cohen, Dunbar, and McClelland (1990)), a VA is an Embodied, Situated and Dynamic (ESD) agent operating in continuous time. ESD agents stress what Clark (1999) 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. Action in an ESD agent is significantly more sophisticated than in a disembodied network. Rather than activate a single node to represent action, ESD agents must use their effectors within the context of a perception–action loop, in which actions change perceptions, and perceptions guide action. Previous work has shown that sophisticated cognitive processes can occur even within a small network, including memory, selective attention (Slocum, Downey, & Beer, 2000), and the use of reactive inhibition (Houghton, Tipper, Weaver, & Shore, 1996) in the control of selective action (Ward & Ward, submitted for publication). We suggest that exploration of minimally cognitive agents must be valuable for psychology, as long as the agents are doing genuinely interesting tasks. Either (1) the agent will use mechanisms already described in the literature, allowing for a tractable computational analysis of those mechanisms, or (2) the agent will use some entirely novel approach to the problem, suggesting new approaches.

We investigate conflict monitoring in a dual-task in which VA must select actions in the presence of stimuli suggesting conflicting responses. We used the dual-target task developed by Slocum et al. (2000), in which agents were constrained to run along the bottom of a 2D environment, moving left and right to catch two targets, T1 and T2, falling from the top of the environment. Let us briefly consider the cognitive demands of this task. First, we can see that catching a single target is not an interesting cognitive task: the sensors need only direct the motors towards the side with the greater stimulus input. In this way, the agent would track the "center of mass" of the perceptual input. However, things become much more complicated when we introduce a second target, and therefore potential response conflict: the agent can be pulled in opposite directions by the two targets. In terms of response conflict, we suggest this task has elements of both 'undetermined responding' and 'response override' (Botvinick et al., 2004). Multiple cognitive processes are required for success in this task. One of the two targets must be prioritized and selected for action. Responses must then be tied to the movement of the selected target, and insulated from the sensor activation of the other. After the first target is caught, a reallocation of processing is necessary: the second target, previously insulated from response mechanisms, must now be allowed to control them. Prioritization, selection for response control, and reconfiguration following a change of targets are all vital topics in current work on selective attention.

In our analysis, we first demonstrate that our agent does indeed exhibit 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. We then examine how the agent resolves this conflict.

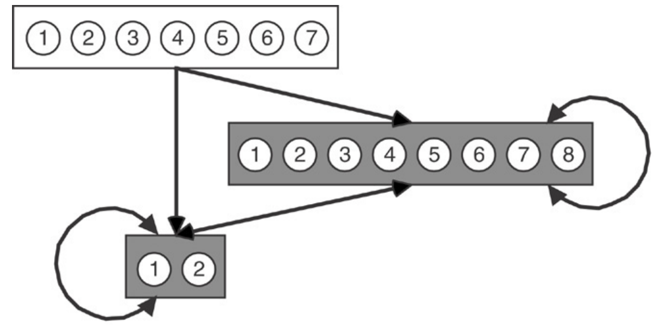

Fig. 1. Network layers and connections of the visual agent (VA). The top unfilled box indicates the seven-node sensor layer, which has no intra-layer connections. The middle box represents the eight-node hidden unit layer, and the lower box illustrates the two-node motor layer. The filled boxes indicate that each unit is connected to every other unit within the layer using bilaterally symmetric weights. Arrows between layers represent bilaterally symmetric connections.

1. Methods

Agents were created with the same connection architecture used in the selective attention experiments of Slocum et al. (2000). Agent diameter was 30 units and target diameter was 22 units, and the environment was 400 units wide by 275 units high. The agent had 7 sensor rays of length 220 evenly spaced over a visual angle of $\pi/6$ degrees. External input magnitude varied from 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 a 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 Fig. 1). Like Slocum et al. (2000) we searched for network parameters using a genetic algorithm, although in principle other unsupervised learning algorithms could also be used. In our case, the 102 network parameters were encoded for genetic algorithm search using GAlib (Wall, 1999).

The agent was required to catch targets T1 and T2. After T1 impacted, it was removed from the environment so that it no longer triggered the sensors. We will call the trials used to evolve the agent the "training" trials (see Fig. 2). The experimental factors defining these trials were: (1) the side on which T1 appeared relative to the agent (Left or Right); (2) the position of T2 relative to T1 (Left or Right); (3) the spatial separations between T1 and T2 (24 units in Near; 48 in Far conditions); (4) the velocities of T1 and T2 (relative velocities 4 and 3 in Near; 5 and 2 in Far). This factorial design creates 16 trial types. For greater generalization, these 16 trials were each presented in three epochs, offset 8 units to the left, 8 to the right, and with 0 offset. After catching T1, the agent had to travel at most 55% of its maximum velocity to catch T2. In addition