

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

An Embodied Dynamical Approach to Relational Categorization

Permalink

<https://escholarship.org/uc/item/7gb1k0rs>

Journal

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

ISSN

1069-7977

Authors

Williams, Paul L.
Beer, Randall D.
Gasser, Micheal

Publication Date

2008

Peer reviewed

An Embodied Dynamical Approach to Relational Categorization

Paul L. Williams¹ (plw@indiana.edu)
Randall D. Beer^{1,2,3} (rdbeer@indiana.edu)
Michael Gasser^{1,2} (gasser@cs.indiana.edu)

¹Cognitive Science Program, ²Dept. of Computer Science, ³Dept. of Informatics
Indiana University, Bloomington, IN 47406 USA

Abstract

This paper presents a novel approach to the study of relational categorization based on the evolution of simulated agents in a relational task. In contrast to most previous models of relational categorization, which begin by assuming abstract representations and role-filler binding mechanisms, we begin by studying relational behavior in embodied dynamical agents, which results in a wider range of possibilities for relational mechanisms. The mathematical tools of dynamical systems theory are used to analyze the relational mechanism of the best evolved agent, and we then identify some of the insights offered by this analysis.

Keywords: continuous-time recurrent neural networks; dynamical systems; genetic algorithms; relational categories

Introduction

Recent research in cognitive science has seen flourishing interest in relations and relational categories (e.g., Gentner & Kurtz, 2005; Markman & Stilwell, 2001). Relational categories are categories determined by common relational structure among category members, rather than intrinsic similarities between category members. For example, *same* and *smaller* are instances of relational categories. There is a vast body of research on relational categorization in a wide range of species, including humans (Kurtz & Boukrina, 2004), pigeons (Wills, 1999), rats (Saldanha & Bitterman, 1951), and insects (Giurfa, Zhang, Jenett, Menzel, & Srinivasan, 2001). Moreover, relational categories are of particular importance for cognitive science because they are fundamental to many topics, such as analogy and language.

In recent years, a number of authors have proposed connectionist (Gasser & Colunga, 1998; Tomlinson & Love, 2006) and hybrid symbolic-connectionist (Hummel & Holyoak, 1997) models of relational categorization. These models put forth various mechanisms for the implementation of relational categories in cognitive systems. However, because these models are largely hand-crafted, they are limited to the space of possible mechanisms considered by their designers.

In contrast, the work presented here takes an entirely different approach to the study of relational mechanisms. We use an evolutionary algorithm to evolve dynamical neural controllers for simulated agents in a relational categorization task. There are numerous benefits of using this kind of approach to study cognition (Beer, 1996; Harvey, Di Paolo, Wood, Quinn, & Tuci, 2005), two of which are central to our purposes. First, by evolving agents we are able to make minimal prior assumptions about how various behaviors should be implemented. Second, this approach allows

us to take seriously the view that cognition is situated, embodied, and dynamical. On this view, cognition is a continuous, ongoing interaction between a brain, a body, and an environment. After successfully evolving agents, we apply the tools of dynamical systems theory to analyze the resulting brain/body/environment systems.

The ability of dynamical neural circuits to perform relational tasks has been demonstrated several times in previous work. For example, in one study a simple recurrent network was trained to recognize string sequences of the form $a^n b^n$, and thus to identify a *same count* relationship between sequences of inputs (Rodriguez, Wiles, & Elman, 1999). However, in this study the relational task was disembodied and computational in nature, whereas the work here is concerned with relational behavior in situated embodied agents. In another study, a neural model was proposed that captured findings from a relational task performed by nonhuman primates (Miller, Brody, Romo, & Wang, 2003). In this case, though, the relational mechanism was hand-designed, while in the work presented here we employ evolutionary techniques, thereby attempting to minimize a priori assumptions about how the relational mechanisms should work.

This paper has two primary aims. The first is to contribute to a growing body of research on minimally cognitive behavior, which studies simple behaviors of cognitive interest in order to develop and refine our conceptual and analytical tools for understanding cognition (Beer, 1996; Slocum, Downey, & Beer, 2000). The second is to show how the approach used here may inform the study and modeling of relational mechanisms more generally. To this end, we discuss how relational categorization is addressed by other models and identify some of the specific advantages of our approach.

The rest of this paper is organized as follows. In the next section, we discuss the basic challenge posed by relational categorization tasks and how this challenge has been addressed in previous models. The following section then describes the agent, neural network model, and evolutionary algorithm used here. Next, we briefly discuss the results of the evolutionary simulations. After, we present an analysis of the relational mechanism used by the best evolved agent. Finally, we conclude with some general remarks, and an outline of ongoing and future work.

Relational Categorization

Relational categorization tasks can be described in terms of two sets of features of the related objects (sometimes, but not

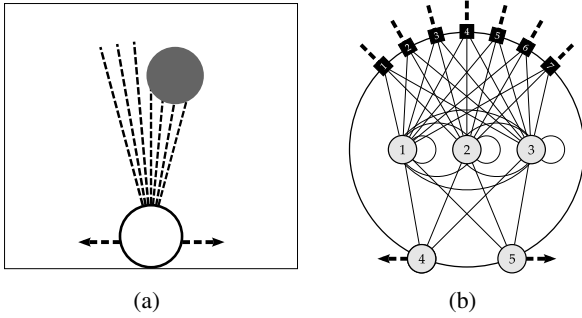


Figure 1: The agent and environment. (a) The agent moves horizontally while circles fall towards it from above. The agent’s sensory apparatus consists of an array of seven distance sensors. (b) The distance sensors fully project to a layer of fully interconnected interneurons, which in turn fully project to the two motor neurons.

always, conceived of as “roles” and “fillers”). Consider a situation in which a book is above a table. The two objects have object category features (*book*, *table*) and spatial relational features (*above*, *below*). A relational categorizer should be able to distinguish this situation from the situation in which the table is above the book. In the relational task explored here, an agent is presented with two circles, one after the other, and its task is to catch the second circle if it is smaller than the first, and to avoid it otherwise. Again, this task can be described in similar terms: each of the two circles is either *first* or *second* and either *larger* or *smaller*.

Models of relational categorization traditionally begin by taking this kind of *description* as a reflection of the *solution* to the categorization problem, by assuming explicit representations for each of the features. In symbolic models, these representations are symbols; in connectionist models, they are distributed representations. In both cases, relational categorization is based on the separation of the four features and the binding together of the appropriate features (e.g., *book* to *above*, *table* to *below*).

A natural consequence of adopting this binding approach is the so-called binding problem, which concerns how features can be associated with each other in such a way that the two possible bindings are distinguished. Diverse solutions to the binding problem have been proposed in the literature (see Gasser (1998) for an extended discussion). One common solution is to dynamically mark separate role and filler elements as belonging together (e.g., Hummel & Holyoak, 1997). That is, in addition to an activation, elements have another associated value which, when it matches that value for another element, represents a binding between them. Another solution is to incorporate role information directly into the encoding of a set of input values. For example, a feature vector and a special “role” vector may be combined—e.g. using the tensor product (Smolensky, 1990) or convolution (Plate, 1995)—to form the inputs to a connectionist network. In this way, binding is implemented through an explicit combination

mechanism. A third solution is to allocate different parts of a relational system—e.g., separate banks of units in a connectionist network—to different roles. A key point about all of these solutions is that they depend on the assumptions that information is represented as a set of abstract features and that it is processed in a stepwise computational procedure. These assumptions allow modelers to make ad hoc decisions in designing their representations or network architectures to deal with the binding problem.

In contrast, our approach to relational categorization makes no such assumptions. Instead of studying relational categorization in terms of abstract representations for the features belonging to a high-level description of the problem (*first*, *second*, *larger*, *smaller*) together, we focus on grounded relational behavior in embodied dynamical agents. A crucial advantage of this approach is that it significantly broadens the playing field for possible relational mechanisms. In particular, the analysis presented below will demonstrate how relational categorization may be carried out without any obvious binding mechanism, thus avoiding consideration of the binding problem altogether. Also, the analysis provides several insights about the kinds of relational mechanisms that are possible in embodied dynamical agents.

Methods

The model agent and environment used in this study are essentially the same as those used in previous work on categorical perception (Beer, 1996; Beer, 2003). The agent has a circular body with a diameter of 30, and an array of 7 distance sensors equally spaced over an angle of $\frac{\pi}{6}$ radians on the agent’s top side (Figure 1(a)). Each distance sensor has a maximum length of 220. Distance sensors take on values inversely proportional to the distance at which their corresponding rays intersect objects in the environment.

The agent is positioned along the bottom edge of a planar environment and is able to move horizontally with a maximum velocity of 5 in either direction. Circles fall towards the agent from above with a constant vertical velocity of -3. As will be elaborated later, the agent’s task is to “catch” or “avoid” various of these circles, where catching or avoidance is determined by the horizontal separation between the agent and circle when the circle completes its fall.

The agent’s behavior is controlled by a continuous-time recurrent neural network with the following state equation:

$$\tau_i \dot{s}_i = -s_i + \sum_{j=1}^N w_{ji} \sigma(s_j + \theta_j) + I_i \quad i = 1, \dots, N$$

where s is the state of each neuron, τ is the time constant, w_{ji} is the strength of the connection from the j^{th} to the i^{th} neuron, θ is a bias term, $\sigma(x) = \frac{1}{1+e^{-x}}$ is the standard logistic activation function, and I represents an external input. The output of a neuron is $o_i = \sigma(s_i + \theta_i)$. The agent’s sensors are fully connected to a layer of interneurons (Figure 1(b)), which are fully interconnected and which project fully to two motor

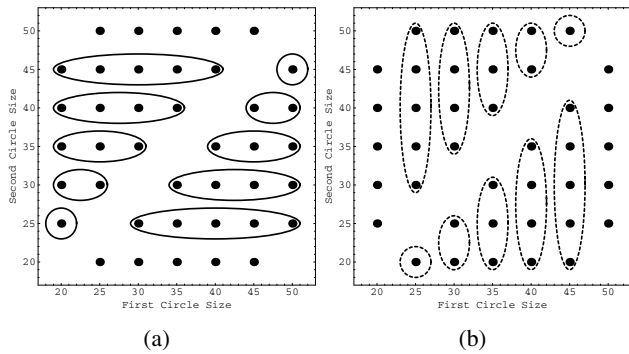


Figure 2: Procedure for fitness evaluation. Black dots indicate the 40 evaluation trials. Scores on these trials are combined to form two aggregate measures, which are then used to determine overall fitness. The first aggregate measure, dubbed *rowAgg*, is calculated by averaging the scores in each solid oval in (a), resulting in ten numbers, and then taking the average of those ten numbers. An analogous procedure is carried out for the dashed ovals in (b) to calculate *colAgg*. Overall fitness is the minimum of *rowAgg* and *colAgg*.

neurons. Horizontal velocity is proportional to the difference between the outputs of the two motor neurons.

Neural parameters are evolved using a real-valued genetic algorithm with fitness proportionate selection (see Beer (1996) for details). We used a fitness scaling multiple of 1.01 and a mutation variance of 4. The following parameters, with corresponding ranges, are evolved: time constants $\in [1, 30]$, biases $\in [-16, 16]$, and connection weights (from sensors to neurons and between neurons) $\in [-16, 16]$. Simulations are integrated using the Euler method with a step size of 0.1.

Agents are evolved for the ability to make discriminations based on the relational category *smaller*. This ability is assessed in a task wherein agents are shown pairs of circles, presented one after the other, and the objective is to catch the second circle in each pair if it is smaller than the first, and to avoid it otherwise. Crucially, the same circle may be either *smaller* or *larger* depending on the other, so the agent must attend to the size relation between the two circles. A similar task was also used elsewhere to study imprinting behavior (Izquierdo-Torres & Harvey, 2006).

An agent’s performance in this task is determined based on its behavior in a number of evaluation trials. Each trial proceeds as follows. First, the agent’s neural states are initialized to zero. Then, a circle with diameter $\in [20, 50]$ begins falling from the top of the environment with a horizontal offset of zero relative to the agent. The circle falls until it reaches the top of the agent and is then removed from the environment. A second circle then begins falling from the top of the environment with the same horizontal position as the first circle. The second circle also has diameter $\in [20, 50]$ but differs from the first circle by at least 5. We impose this restriction on the size difference so that the difference will be perceivable to the agent, given the coarse spatial resolution of its sensors. The




Figure 3: Performance of the best evolved agent. A density plot of the agent’s final horizontal separation from the second circle is shown, with black indicating a perfect catch and white indicating a perfect avoidance. The gray line indicates where the two circle sizes are equal. The blue lines indicate where the size difference between the two circles is equal to the minimum size difference used in evaluation trials.

second circle falls until it reaches the agent and we record the final horizontal separation between the agent and second circle. This final separation constitutes the agent’s catch/avoid response, with a separation of zero corresponding to a perfect catch, and a separation of *MaxDistance* ($=75$) or greater corresponding to a perfect avoidance. The agent’s score on a trial is $1 - d$ if the second circle is smaller and d if the second circle is larger, where d is the final separation clipped to *MaxDistance* and normalized to run between 0 and 1.

Fitness is evaluated with 40 circle pairs (black dots in Figure 2) in order to adequately sample the stimulus space. Initially, we tried simply averaging the trial scores to produce the overall fitness measure. However, assigning fitness in this way resulted in agents adopting one of two suboptimal solutions: agents would decide to catch or avoid based on the size of either the first or second circle alone, thus ignoring the relation between the two. To compensate for this, we used a slightly more sophisticated procedure for weighting and averaging the trial scores, as described in the caption for Figure 2.

Results

Ten evolutionary runs each were performed with agents having five, four, and three interneurons, respectively. In each run, a population of 200 individuals was evolved for 1,000 generations. The best agent in each run had a fitness of at least 90% on 10,000 random circle pairs. A further attempt to evolve agents with two interneurons was unsuccessful, suggesting that three interneurons may be the necessary minimum for high performance on this task.

The behavioral strategies of the agents with three interneurons fall into two categories. One group of agents employ an active strategy, moving back and forth repeatedly to scan circles as they fall. The other group use a passive strategy, remaining mostly still as the first circle falls, and then either remaining in place to catch the second circle, or veering to one side and avoiding it.




Figure 4: Two features of the neural dynamics. (a) The activation of neuron 3 at the trial midpoint, which stores the first circle size. (b) Neuron 1 trajectories for a range of second circle sizes; whether or not neuron 1 switches off determines the catch/avoid response.

Here we focus on the passive strategy for three reasons: (1) the agent with the best performance used it; (2) the majority of agents (6 out of 10) used it; (3) the relational mechanism underlying it is more straightforward. In particular, we focus on the best evolved agent as a characteristic example of the passive strategy. The best agent had a fitness of 99.83% over 10,000 random circle pairs. A quick glance at the agent’s performance over the range of circle sizes verifies the desired relational categorization (Figure 3). The agent’s behavior is divided into distinct catch and avoid regions, coinciding with smaller and larger second circles, respectively.

The Dynamics of Relational Categorization

To understand what produces the agent’s categorical behavior, we must examine the agent’s neural dynamics. In particular, two features of the dynamics turn out to be essential. The first feature has to do with how the agent stores the size of the first circle. The agent must store the first circle size in order to respond differently to the second circle when it is smaller or larger. This means that some aspect of the agent’s state—its neural activations and horizontal position—must correspond to the size of the first circle at the trial midpoint (i.e., when the second circle begins its fall). In general, this correspondence could be quite complex and nonlinear, incorporating any subset of the agent’s state variables. However, in the agent under consideration the size information is stored directly; namely, the output of one interneuron, neuron 3, stores the size. Neuron 3 linearly correlates with the first circle size at the trial midpoint (Figure 4(a)), whereas all other state variables take on basically constant values. We verified the role of neuron 3 by setting the other state variables to suitable fixed values at the trial midpoint, and confirming that performance did not decrease significantly as a result.

The second feature of the dynamics has to do with what initiates the agent’s catch/avoid response. The crucial observation here is that another interneuron, neuron 1, either switches off (i.e., its activation decays to zero) or remains on during the second circle’s fall (Figure 4(b)). Furthermore, whether or




Figure 5: The impact of neuron 3 on neuron 1. Trajectories of neurons 1 and 3 are shown over the second half of a trial for several first circle sizes. Gray trajectories are catch trials and black trajectories are avoid trials. Higher activations of neuron 3 cause neuron 1 to switch off more quickly.

not neuron 1 switches off determines the agent’s catch/avoid response. If neuron 1 switches off, the agent catches the second circle; if neuron 1 remains on, the agent avoids the second circle. Thus, we can account for the agent’s categorical behavior by understanding what determines whether or not neuron 1 switches off.

To summarize, an inspection of the neural dynamics reveals two key features: (1) the activation of neuron 3 stores the size of the first circle; (2) whether or not neuron 1 switches off determines the agent’s catch/avoid response. To understand the agent’s categorization, then, we need to explore how these two features of the dynamics interface with each other. Specifically, how does the activation of neuron 3 combine with information about the second circle size to determine the switching behavior of neuron 1?

To probe this question, we first examine the effect of neuron 3 on neuron 1. This effect corresponds to the connection from neuron 3 to neuron 1, which is negative (i.e., inhibitory). In other words, neuron 3 tends to cause neuron 1 to switch off, and greater activations of neuron 3 cause neuron 1 to switch off more quickly (Figure 5). Since neuron 3 stores the size of the first circle, larger first circle sizes cause faster decay in neuron 1. However, despite the tendency of neuron 3 to switch neuron 1 off, we know that in some cases neuron 1 remains on, and that whether or not neuron 1 switches off determines the agent’s catch/avoid response. So, what determines whether or not neuron 1 switches off? To answer this question, we must explore the underlying equilibrium structure of the neural dynamics.

The neural circuit is a nonautonomous dynamical system, since it receives time-varying inputs from the distance sensors. To analyze a nonautonomous system, we can examine the autonomous dynamics produced when the inputs are held constant, for each possible set of inputs. The nonautonomous dynamics can then be approximated by considering the sequence of autonomous systems corresponding to the particular inputs that the nonautonomous system receives.

In our case, three things determine the inputs to the agent’s neural circuit: (1) the size of the circle; (2) the vertical offset of the circle; (3) the horizontal offset of the circle. However,




Figure 6: A bistable region determines the behavior of neuron 1. Trajectories of neuron 1 are shown for a range of first circle sizes (gray and black lines) along with the equilibrium points (EPs). Blue EPs are stable and green EPs are saddles. The bistable region separates the trajectories into two basins of attraction. The dashed red lines show where the agent’s ray sensors are first broken. Other abrupt changes in the underlying dynamics are due to bifurcations. Note that there is a minor inconsistency between the trajectories and EPs due to the fact that horizontal offsets between the agent and circle were ignored.

since the agent remains still for most of each trial (i.e., it uses the passive strategy), we can simplify our analysis by ignoring changes in the input due to (3). Thus, we can explore the dynamical structure of the neural circuit by considering it as a function of two variables: the size of the circle and its vertical separation from the agent.

It turns out that the only limit sets exhibited by the neural circuit are equilibrium points, whose positions and stabilities we can calculate. When the circle is very far or very near, the circuit has a single stable equilibrium point. For intermediate distances, the circuit exhibits bistability. In this case, the circuit has two stable equilibrium points, which serve as attractors for system trajectories, and a single saddle point, which separates the two basins of attraction. How does this equilibrium structure explain the switching behavior of neuron 1? To answer this question, we can superimpose representative neuron 1 trajectories over a plot of the system’s equilibrium points (Figure 6), showing how the trajectories of the system are shaped by the underlying organization of the dynamics.

Note that the bistable region separates the trajectories into those that switch off and those that remain on, corresponding to catch and avoid trials, respectively. The bistable region is produced by a bifurcation which occurs when the circle reaches a certain vertical offset. The timing and location of this bifurcation correlate with the second circle size, so that trajectories of neuron 1 are appropriately partitioned into catch and avoid regions.

With this information, we can now assemble an account of the agent’s relational mechanism as follows. During the first half of a trial, the activation of neuron 3 comes to correlate with the size of the first circle, while neuron 1 switches to fully on. Then, during the second half of a trial, neuron 1 begins to switch off. Due to the inhibitory effect that neuron 3 has on neuron 1, neuron 1 switches off more quickly for larger first circles. As the second circle falls, a bifurcation occurs whose timing depends on the size of the second circle. If the activation of neuron 1 is low enough by the time the bifurcation occurs, neuron 1 switches off completely and the second circle is caught. Otherwise, neuron 1 returns to fully on and the second circle is avoided. The inhibitory effect

of neuron 3 on neuron 1 and the location and timing of the bifurcation are coordinated such that smaller second circles are caught and larger second circles are avoided.

Discussion

One significant feature of the analysis presented above is that it demonstrates the crucial role that time can play in a dynamical mechanism. For example, the activation of neuron 3, which stores the size of the first circle, affects the subsequent network dynamics by causing neuron 1 to switch off more quickly for larger first circles. Also, the size of the second circle determines the time at which a bifurcation occurs in the system, dividing the system trajectories into catch and avoid responses. Finally, the relational mechanism as a whole relies on the coordinated timing of neuron 1 switching off and the occurrence of the bifurcation. The central importance of time runs as a common thread throughout these examples.

The analysis also shows how information can be “represented” and “processed” in a variety of ways in a dynamical mechanism. One piece of information, the first circle size, is stored in the activation level of neuron 3. This encoding of a perceptual feature by a neural activation fits with the standard approach to representation in connectionist networks. However, if we try to identify what serves to represent the second circle size, the only obvious candidate is the bifurcation that the system undergoes. This correspondence between a stimulus feature and the onset of a bifurcation is an entirely different way for information to bear on the network dynamics. Nevertheless, the analysis makes clear how these different forms of information are integrated seamlessly in producing relational categorization.

As mentioned earlier, other models of relational categorization start by assuming abstract representations and role-filler binding mechanisms. As a result, one major difficulty for these models is the binding problem, which requires modelers to make ad hoc decisions regarding the nature of representations and the structure of network architectures. In contrast, our approach begins by studying relational behavior in embodied dynamical agents, which results in a wider range of possibilities for relational mechanisms. For example, we find that information associated with different roles can manifest

itself in a number of different forms, e.g. as a neural activation level or the timing of a bifurcation. Also, our analysis demonstrates the crucial role that time can play in combining information to produce relations.

Another intriguing finding has to do with the active behavioral strategy used by other of the evolved agents. We found that these agents use their position in the environment to store information about circle sizes. Thus, the position of the agent, as well as its neural dynamics, plays a crucial role in the relational mechanism. This kind of strategy fits nicely with one popular idea that situatedness may allow a cognitive agent to offload information to its environment. One strength of the dynamical approach is that neural, bodily, and environmental variables are all represented in the same dynamical language, so cognitive processes may naturally spread across the entire brain/body/environment system.

Next, we would like to preempt a potential criticism of our approach regarding the issue of generality. It could be argued that most models of relational categorization are general, meaning that they apply across a wide range of inputs and relations. In contrast, the relational mechanism of our simulated agent has no obvious generalization beyond the particular stimuli and task that we used. This is a valid point, and our response is twofold. On the one hand, it is interesting to consider the possibility of more general relational mechanisms. Indeed, this possibility motivates our ongoing research, as we discuss further below. On the other hand, there need not be a general mechanism for relations. It is quite possible that natural evolution has discovered a variety of solutions which are just good enough to work in particular tasks, but nothing more. In any case, we feel that it is a strength of our approach that it leaves the necessity of general solutions as empirical questions, instead of assuming them a priori.

Current work extends the approach presented here in several ways. In one line of work, we are attempting to evolve agents in more complex versions of the task explored here by, for example, varying the horizontal positions of the circles and the speeds at which they fall, in an effort to evolve a more general relational ability in the model agents. In another line of work, we are analyzing a number of other agents in order to explore the extent to which a more general story about their relational mechanisms may hold. Other work is concerned with exploring different kinds of relational categorization. One set of experiments has successfully evolved agents to make relative size judgements of two circles falling simultaneously, rather than in sequence. Another project examines the more abstract relations of sameness and difference. The goal for all of these projects is to explore the possibility of deriving general principles for relational categorization in a situated, embodied, dynamical framework.

Acknowledgments

This investigation was supported by the National Institute of Health and National Research Service Award HD007475-13.

References

- Beer, R. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From Animals to Animats 4*, 421–429.
- Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4), 209–243.
- Gasser, M., & Colunga, E. (1998). *Where do relations come from?* (Tech. Rep. No. 221). Indiana University.
- Gentner, D., & Kurtz, K. (2005). Relational categories. In *Categorization inside and outside the laboratory* (pp. 151–175). Washington, DC: APA.
- Giurfa, M., Zhang, S., Jenett, A., Menzel, R., & Srinivasan, M. (2001). The concepts of 'sameness' and 'difference' in an insect. *Nature*, 6831, 930–933.
- Harvey, I., Di Paolo, E., Wood, R., Quinn, M., & Tuci, E. (2005). Evolutionary robotics: a new scientific tool for studying cognition. *Artificial Life*, 11, 79–98.
- Hummel, J., & Holyoak, K. (1997). Distributed representations of structure: a theory of analogical access and mapping. *Psychological Review*, 104(3), 427–466.
- Izquierdo-Torres, E., & Harvey, I. (2006). Learning on a continuum in evolved dynamical node networks. *Proceedings of Artificial Life X*, 507–512.
- Kurtz, K., & Boukrina, O. (2004). Learning relational categories by comparison of paired examples. *Proceedings of the Conference of the Cognitive Science Society*, 756–761.
- Markman, A., & Stilwell, C. (2001). Role-governed categories. *Journal of Experimental & Theoretical Artificial Intelligence*, 13(4), 329–358.
- Miller, P., Brody, C. D., Romo, R., & Wang, X. J. (2003). A recurrent network model of somatosensory parametric working memory in the prefrontal cortex. *Cerebral Cortex*, 13(11), 1208–1218.
- Plate, T. (1995). Holographic reduced representations. *Neural Networks, IEEE Transactions on*, 6(3), 623–641.
- Rodriguez, P., Wiles, J., & Elman, J. L. (1999). A recurrent neural network that learns to count. *Connection Science*, 11(1), 5–40.
- Saldanha, E., & Bitterman, M. (1951). Relational learning in the rat. *The American Journal of Psychology*, 64(1), 37–53.
- Slocum, A., Downey, D., & Beer, R. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. *From Animals to Animats 6*, 430–439.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46(1), 159–216.
- Tomlinson, M., & Love, B. (2006). From pigeons to humans: Grounding relational learning in concrete examples. *Proceedings of the National Conference on Artificial Intelligence*, 199–204.
- Wills, S. (1999). Relational learning in pigeons? *The Quarterly Journal of Experimental Psychology: Section B*, 52(1), 31–52.