

A COMPUTATIONAL INVESTIGATION OF  
RELATIONAL REASONING IN NONHUMAN ANIMALS

A THESIS SUBMITTED TO THE GRADUATE DIVISION OF THE  
UNIVERSITY OF HAWAI'I AT MĀNOA IN PARTIAL FULFILLMENT  
OF THE REQUIREMENTS FOR THE DEGREE

MASTER OF ARTS

IN

PSYCHOLOGY

MAY 2012

By

Margeaux F. Ciralo

Thesis Committee:

Leonidas Dumas, Chairperson  
Patricia Couvillon  
Scott Sinnett

Keywords: cognition, animal cognition, computational model, relational reasoning

## *Acknowledgements*

I would like to acknowledge Drs. Patricia Couvillon, Leonidas Doulas, Brandy Fraizer, Scott Sinnett, and Catherine Sophian for their contributions in my continued development as a cognitive scientist and the development of this project.

## ***Dedication***

To my mother and father, Donna L. & Charles R. Ciruolo; Neal Ching; and the rest of my family.

To Sif, who may not be able to reason relationally, but *loves*(Margeaux, Sif).

## *Table of Contents*

TITLE	PAGE
Acknowledgements	i
List of tables	iv
List of figures	v
<b>Introduction.....</b>	<b>1</b>
Relational Reasoning	6
Computational Models	8
Computational Models of Relational Reasoning	10
LISA/DORA Models of Relational Reasoning	14
i. Nature of representations	15
ii. Flow of control	16
iii. Relational learning	19
<b>Simulations.....</b>	<b>21</b>
Simulation 1 – Lazareva et al. (2004)	21
Simulation 2 – Cook & Wasserman (2007)	25
Simulation 3 – Call & Tomasello (1999)	29
<b>General Discussion.....</b>	<b>33</b>
<b>Future Research.....</b>	<b>36</b>
Predication	36
Comparison-Based Predication	37
Asynchrony of Firing	38
Refinement	39
Mapping	40
<b>Conclusion.....</b>	<b>42</b>
<b>Appendices.....</b>	<b>44</b>
Appendix A	44
Appendix B	45
<b>References.....</b>	<b>50</b>

## *List of Tables*

Table 1. Lazareva et al. (2004) Ordered Stimuli Proportions.....	44
Table 2. Lazareva et al. (2004) Unordered Stimuli Proportions.....	44
Table 3. Lazareva et al. (2004) Simulation File Overview – Ordered.....	45
Table 4. Lazareva et al. (2004) Simulation File Overview – Unordered...	46
Table 5. Cook & Wasserman (2007) Simulation File Overview.....	47
Table 6. Call & Tomasello (1999) Simulation File Overview.....	49

## *List of Figures*

Figure 1. LISAese Representation.....	15
Figure 2. Synchrony of firing in LISA.....	17
Figure 3. Lazareva et al. (2004) Simulation Results – Ordered.....	23
Figure 4. Lazareva et al. (2004) Simulation Results – Unordered.....	23
Figure 5. Cook & Wasserman (2007) Simulation Results.....	27
Figure 6. Call & Tomasello (1999) Simulation Results.....	31

## *Introduction*

Since Darwin (1859) suggested that all animals share a common ancestry, researchers have attempted to assess how the divergence between human and nonhuman animal biology and cognition came about. Animals have a large repertoire of physical and behavioral adaptations that have afforded them the ability to overcome myriad obstacles to their existence. On the behavioral level, these adaptations are as diverse as the species that exhibit them. One merely has to select an animal species at random and do a cursory investigation in order to find many complex behaviors which make that animal a perfect fit for the environment that it inhabits. However, it is possible that the high degree to which an animal is adapted to its environment may limit that animal's ability to employ flexible behavior. For example, the digger wasp has a highly stereotypical set of behavioral steps for readying a nest. If any one of these steps is disrupted in the process of nest creation, the digger wasp will not alter its behavioral program, even if doing so would ensure the life of its offspring (Fabre, 1919). In short, the digger wasp lacks flexibility in its behavioral routine. Indeed, many animal species display rigidity in their behavior. Humans, on the other hand, appear to have much more control of their behavior as well as a flexibility that has allowed them to manipulate and change their environment. The mechanism that underlies this ability may be at the heart of what makes human and nonhuman animal cognition so qualitatively different.

Humans are capable of manipulating their environment in such a way that they create art, science, math, and poetry. At different times, scientists have argued that complex tool use, grammatically structured language, causal-logical reasoning, and mental state attribution are what make human cognition "special." However, as researchers have examined what is involved in each of these phenomena, they have found compelling evidence that nonhuman animal

cognition may be harboring cognitive processes that initially seemed quite unique to humans (see, e.g., Lazareva et al., 2004; Gentner, T.Q., 2006; Cook & Wasserman, 2007). The key question then becomes: Are the behaviors documented in animal species and humans at all the same? At some level they must be, given that at one point humans shared a common ancestor with other animal species. In order to properly assess the degree to which human and nonhuman animal cognition is similar and different, one must establish how behavioral and cognitive changes occurred over time, then suggest a way to adequately study where those changes may have occurred in the shared history of nonhuman animals and humans, as well as propose an underlying mechanism that may account for the stark contrast between human and nonhuman animal cognition.

Darwin (1859, 1871) began to address this question with the proposal of a process known as natural selection. Darwin asserted that all species have come to be as they are today through small changes over large periods of time, and that life is a struggle where any mutation that betters an individual's probability of survival and subsequent reproduction will likely be passed down to its descendants. The adaptation that allows the cheetah to run so quickly began by some proto-cheetah having a mutation that made it just a little bit faster than rest of the members of its species, and over many generations and through the accumulation of many mutations, this speed became a unique attribute of that particular large cat species. This same logic can be applied to behavior; some proto-digger wasp possessed a mutation that compelled it to dig in the first place, and through the same processes that endowed the cheetah with its incredible speed, the complex nest creation sequence was instantiated and is thus inherited by all modern digger wasps. This is indeed an over-simplification of this process, but it is important to illustrate that the mechanism



that allows for physical adaptations is the same causal mechanism that allows for behavioral adaptations (Brown, 1975).

Natural selection has resulted in the abundance of species alive today, and much (certainly not all) is known about the ways in which animal species have physically changed over millennia of evolution. Evolutionary biologists, paleontologists, and others interested in the origin of adaptations on the biological level have many tools at their disposal, e.g., the fossil record, mitochondrial DNA, and sophisticated dating methods (see McKinney, 1991; Castro et al., 1998; Renne et al., 1998). Researchers interested in the behavioral and cognitive changes across time have much less in regards to methodology to work with. Unfortunately, most behaviors leave behind little trace, with perhaps the notable exception of early humans from whom we have some artifacts, including paintings and tools; yet there is still much debate in regards to which artifacts are reflective of modern human cognition (see Henshilwood & Marean, 2003 for a discussion). The task for comparative and cognitive psychologists interested in how behavioral differences have come about then must begin with an earnest understanding of what cognitive capabilities are shared across species. Although in isolation these data do not appear to give much insight into how these behaviors were developed and subsequently inherited over time, but in conjunction with biological evolution data, researchers may be able to infer at which points in time these cognitive differences may have appeared.

Researchers of various backgrounds have been accruing data in regards to many different species' behavioral capabilities for a long time. At first the methods for assessing the behavior of animals was methodologically flimsy and the evidence was mostly anecdotal (e.g., Romanes, 1882), but eventually these anecdotal methods were abandoned in favor of much more parsimonious methods (e.g., Thorndike, 1898). Comparative psychology has continued in this

tradition and has developed unique methodologies and has successfully investigated cognitive phenomena from low-level associative learning (e.g., Pavlov, 1927) to much more complex behaviors like animal communication (e.g., Cheney & Seyfarth, 1990). Currently, researchers have a large corpora of experimental data at their disposal and although it has led to many insights into the ways in which animal and human cognition could conceivably differ, there is still little consensus in regard to which underlying cognitive process or processes allow for the behavioral flexibility observed in human reasoning but not in that of other nonhuman animal species.

Part of what lies at the heart of the issue is that understanding differences in behavior across species requires specifying the cognitive processes that give rise to specific behaviors. Historically, however, this has not been the focus of comparative studies of cognition in animals. Shettleworth (2009) notes that the attempts to explore cognition in nonhuman animals often cast a cognitive difference between human and nonhuman animals – an ‘x,’ then ask: “Do animals have ‘x,’ yes or no?” ‘X’ is then subjected to a battery of pass fail/tests which do not end up leading to fruitful descriptions of animals’ cognitive abilities, due largely to the fact that knowing whether an animal can do a task that humans can also do says nothing about whether or not the problem is being solved in the same way. The same can be said of cases where animals cannot do a task for our hypothetical ‘x.’ Instead, Shettleworth suggests that “more progress may be made by breaking a broadly defined capability down into components, asking which are shared among species, and under what conditions, and why” (2009, p. 216). However, few cognitive processes cited earlier as potential fundamental differences are specified in enough detail to be broken down into components, or realistically differentiated between species. Therefore, an alternative tactic to addressing these issues must be developed.

The need for a better understanding of the cognitive processes being employed by nonhuman animals is shown to be especially important if we consider some examples of attempts to find evidence for complex human-like reasoning capabilities in a number of species. For example, Bergman et al. (2003) argues that free-ranging baboons have the ability to reason about social hierarchy, citing results of playback experiments in which baboons respond more intensely to calls from dominance rank reversals from between families than rank reversal calls from within families. Likewise, Dally et al. (2006) utilized the natural food caching behaviors of scrub jays to show that these birds are capable of a type of mental state attribution. In these experiments a dominant conspecific, subordinate conspecific, or preferred partner watched as another scrub jay cached a food item. The bird was then able to return and re-cache the food if it chose to. They found that scrub jays re-cache more often when a dominant conspecific or subordinate conspecific has viewed their cache site in comparison to control viewers (i.e., the scrub-jay's partner or in private).

On the surface, the behaviors of these two species bear a striking resemblance to human reasoning about hierarchies and mental state attribution. Most would agree that humans have the ability to reason about hierarchies (e.g., how people know that the way one speaks to a co-worker is different than how they address their boss and is also different from how they might address the company CEO) or about the mental contents of another (e.g., if Sue knows that Sally has had a bad day, Sue may choose not to burden Sally with additional bad news). Therefore it is easy to fall into an interpretation where one might hypothesize that the baboons “understand” hierarchies or that scrub jays “know” what one another are thinking. However, it is important to understand that while behaviors may appear similar on the surface, the underlying mechanisms or mental processes employed in their service might be vastly different. A more parsimonious

interpretation of the cited studies might be that scrub jays and baboons have abilities that function in a similar way as hierarchical reasoning and mental state attribution in humans, but an equally plausible interpretation is that they do not. Instead, it could be the case that these animals are relying on perceptual cues (like the presence or absence of a subordinate conspecific) and associative learning mechanisms (i.e., when a subordinate conspecific is in view, move the food) to reason about those events (Penn et al., 2008).

It is reasonable to expect some degree of overlap in the cognitive architectures of humans and nonhuman animals due to their common descent. However, human thinking appears to be qualitatively different from that of even our closest evolutionary relatives. An alternative and potentially more worthwhile strategy than those outlined earlier would be to propose an underlying cognitive mechanism that could account for the species-level differences seen in reasoning capabilities, and then test for the presence or absence of those processes. Therefore, the focus of the current study is to analyze a well-known cognitive phenomenon called *relational reasoning* and empirically test whether or not humans and nonhuman animals utilize the same processes in tasks of this nature.

## 1. Relational reasoning

Relational reasoning is widely recognized as a key component of human thought (see, e.g., Gentner, 1983, 1989; Gick & Holyoak, 1980, 1983; Holyoak & Thagard, 1995). It allows one to reason about the role that an object plays rather than attending to only the physical features of an object, and allows for the flexibility and structural sensitivity required for many unique human capabilities, such as language production, art, science, and mathematics (Medin et al., 1993). Therefore, relational reasoning is an ideal candidate for an underlying mental process

that may be able to account for differences seen between human and nonhuman animal intelligence.

Relational reasoning requires a specific type of mental representation. In order to reason relationally, a system must be able to represent relations as explicit entities that can be dynamically bound to arguments (i.e., be *predicated*). For example, in order to appreciate a sentence like “the dog chases the cat,” the concept of “chaser” must be independent from the concept of “dog” and likewise, the concept of “chased” must be independent from the concept of “cat.” An explicit representation of “chaser” and “chased” will allow a system to reason regardless of to whom or what the representation is bound.<sup>1</sup> Therefore, when a spoon chases a dinner plate, or when an older brother chases a younger sister, nothing about the underlying representation is changed other than the arguments (or actors) being bound to them.

Relational reasoning has been investigated in select animal species, both explicitly and tangentially. However, relational match to sample tasks have been conducted with a wide range of species and at one point it was thought that only humans and symbol-trained apes were capable of performing well on these tasks (Thompson & Oden, 2000). The relational match to sample paradigm includes the following: a subject is presented with a sample pair that embodies some aspect of “sameness” or “differentness” (e.g., the “sameness” can range from identity to sameness on a particular dimension like shape, color, or some other level of abstraction). The subject is then faced with a choice between two pairs, one that matches on the particular dimension of interest and one that does not. However, there is widespread debate in regards to whether “same” and “different” are relational at all, given evidence that entropy may be a salient

---

<sup>1</sup> The types of relational representations detailed here are specific to role-filler binding systems. Alternatives to role-filler binding will be discussed in section 3, *Computational models of relational reasoning*.

perceptual feature that changes the nature of the task from a relational one to a perceptual one (Fagot et al., 2001; Wasserman et al., 2004).

Relational match to sample is one of few methodologies that attempt to address relational reasoning in animals explicitly; many forms of relational reasoning have been addressed in the comparative literature tangentially. Penn et al. (2008) provide an extensive overview of evidence gathered from a large assortment of tasks, including transitive inference, hierarchal relations, causal relations, and theory of mind. These tasks at their core are (or rather can be) relational in nature in that some of the successful strategies for performing well on these tasks require concepts like ‘bigger-than,’ ‘cause,’ and ‘knows’ to be explicitly predicated. Penn et al. argue that none of the cited experiments have successfully demonstrated that animals have a capacity for relational reasoning in the same sense that humans do because it is possible that the physical features alone may be enough to confer success. Penn et al. go on to suggest that perhaps an ideal way to investigate this phenomenon would be to utilize a computational model of relational reasoning and test whether or not these process models can account for the behavioral data that have already been collected. The present plan of research was inspired by this suggestion.

## 2. Computational models

Before addressing computational models of relational reasoning specifically, it is important to first understand why models like these exist and what benefits they afford researchers in understanding cognitive phenomena. The logic falls from many insights starting from the early 1950’s, this includes the rejection of radical behaviorism, Marr’s description of the levels of analysis (Marr, 1982), and Alan Turing’s description and the implications of the

Universal Turing Machine (UTM) (Turing, 1938). These events set in motion the modern cognitive tradition of understanding individuals as information-processing machines and gave researchers new tools for understanding cognition; tools that have continued to develop over time.

David Marr (1982) set forth a framework for understanding any kind of information-processing machine; this framework came to be known as “Marr’s levels of analysis”, or simply “Marr’s levels.” The levels of analysis are (i) the computational level, (ii) the representational/algorithmic level, and (iii) the implementation level. In concert, these three “casually and logically related levels” (Marr, 1982) can have a profound impact on the depth at which an information-processing system is understood.

The first level, the computational level, is an understanding of the goal of an information-processing system. A clear understanding of the goal of a system answers questions pertaining to why the system is doing what it is doing, as well as how that system is doing it. The algorithmic and representational level aims to account for the processes being employed by the system and address questions about what processes is the system doing, what representations are being utilized, and in what way are those representations used? Finally, the implementation level is concerned primarily with how the system is instantiated in some physical system.

First, it is important to make the argument that humans and animals alike are information-processing machines. The idea began to take hold in the 1950’s after the rejection of radical behaviorism (the main thesis of which was that it is not possible to study the internal components of the mind, only overt observable behaviors) that mental contents could in fact be studied. Methodologies were created that demonstrated that something happens between perception and output of behavior. In fact, the mind is transforming or processing information in

some way. It was also during this time that the computer was identified as a useful analogy for understanding the mind (Lachman et al., 1979). In conjunction with Alan Turing's proof that all problems with a solution can be solved utilizing a UTM, modern computational methodologies were created.

Modern computational models are in fact quite different from one another in that they can approach the problem of understanding cognition from one of many angles and from different levels of Marr's levels of analysis. Marr himself argued that some problems are better understood at some levels versus other levels. However, he did go on to note that an algorithmic and representational account is best understood by first understanding the nature of the problem (i.e., the computational level) rather than assessing the third level (i.e., the physical implementation) and then attempting to understand the algorithmic and representational levels (i.e., it is easier to go down the levels than it is to go up them), and finally that a complete and full understanding of an information processing machine would capture all three levels.

### 3. Computational models of relational reasoning

Marr's levels of analysis have a direct impact on how computational models of relational cognition have been designed. Successful models attempt to address relational reasoning by giving a satisfactory description of the behavior to be modeled, they specify the nature of representation in conjunction with how representations are used to generalize, reason, etc., and some are created using a neurally plausible architecture (one that can be seen as being easily realizable within the neurons of a brain). However, it is important to first address the various



ways in which relational reasoning has been addressed in models and specify what a model has to do in order to best approximate how people reason relationally.<sup>2</sup>

Computational accounts of relational reasoning typically take the format of one of three types of models: symbolic, connectionist, or symbolic-connectionist. The first two types have different strengths and weaknesses that will be discussed in the following section. The latter attempts to marry the strengths of symbolic models and connectionist models in a such a way that none of the negative implications of choosing a symbolic or connectionist method of modeling human relational reasoning come along for the ride.

Human relational representations are considered to be symbolic in nature, Fodor & Pylyshyn (1988) famously argue that this is the case because it is possible for one to think the thought “John loves Mary” and also be able to consider, similarly, that “Mary loves John.” This argument is typically cited as the “argument from systematicity.” However, to avoid the pitfalls of such a vague argument, the reasoning that human relational reasoning must be symbolic is borrowed from Doumas & Hummel (2005) – relational reasoning is symbolic as it entails the ability to represent relational roles independently of their arguments as well as specify which roles are bound to which arguments. The ability to maintain role-filler independence leads to an appreciation of what different bindings of the same relational roles and fillers have in common and how they differ, which is readily apparent in human reasoning. For example, if one considers the novel sentence, “the glarbile flarbs the vilbil.” Although *glarbile*, *flarbs*, and *vilbil* carry little semantic information it is possible to parse the sentence in such a way that the reasoner knows

---

<sup>2</sup> Although the following descriptions are based on attempts to understand *human* relational reasoning, it is important to recall that the focus of the present research is to meaningfully differentiate between nonhuman animal and human cognition. In this way, researchers might come to understand what aspects of human relational reasoning may be present/absent in animal reasoning. Therefore an understanding of what humans do when they reason relationally is necessary.

that the *flarber* is the *glarbile* and the *flarbed* is the *vibil*. One merely has to change the ordering of the sentence in order to recognize instead that the *vibil* is now doing the *flarbing* and the *glarbile* is now *flarbed*. Without a symbolic representation such structural sensitivity would not be possible.

On the other hand, traditional connectionists reject the notion completely that relational reasoning is at all symbolic. They argue instead that mental representations are instantiated in patterns of activations of nodes in a network (which can easily be realized as neurons in the brain). Therefore a single concept is represented as pattern of activation across nodes or units and any single element, node, or unit can and will participate in the representation of many different concepts. In this way connectionist models are capable of capturing some of the semantic content that may be shared across different concepts. For example, the conceptualization of a dog and cat are indeed similar, but they share little in common at all with a fish, and even less in common with a telephone. Distributed representations like those used in connectionist models model the relative *likeness* of dogs, cats, fish, and telephones very well. That is, there will be more units in common activated when thinking about dogs and cats than when thinking about dogs and telephones. In addition to this powerful aspect of these models, they have incredibly graceful degradation. In the same way that when a human suffers brain damage, reasoning may indeed be impaired; when nodes in the network are “damaged” the model will continue to reason, albeit somewhat impaired (Doumas & Hummel, 2005).

Although this natural inferential power is highly desirable it comes at the cost of the strength of symbolic methods of modeling human relational reasoning. It is important to note that instantiating representations in this distributed fashion violates role-filler independence while also making it unclear as to which roles are bound to which fillers. Roles and fillers are

conjunctively coded in some connectionist models. This comes with incredible implications for what it means to be John in “John loves Mary” and what it means to be John in “Mary loves John.” John in the former is therefore unlike John in the latter as the *lover* and the *beloved* role is coded conjunctively. Similarly, in cases where connectionist models choose to code entire propositions conjunctively one loses information about whether John is a woman or man or which attributes belong to being loved versus loving someone, etc. (Doumas & Hummel, 2005).

Another serious negative implication of connectionist models lies in the fact that connectionist models learn to associate nodes in an input layer with an output layer, and these associations are typically between lists of features. Meaning that the performance of connectionist models is often times a function of how complete the training of the model is and which features are shared across those concepts that have been mapped. The Story Gestalt model, a model of story comprehension (St. John & McClelland, 1990), is often cited as embodying this particular drawback. St. John gave the model training on 1,000,000 examples of simple stories in the form of: <person> decided to go to <destination> so <person> drove <vehicle> to <destination>. When presented with a novel person and destination the model was incapable of inferring that the destination that individual drove to was the one in which the individual decided to go, demonstrating the model was insensitive to the structure of the story and instead was only capable of generating answers (destinations in this case) that it had previously mapped (Hummel & Holyoak, 2000; Doumas & Hummel, 2005; Marcus, 1998).

In summary, traditional connectionist models are able to generalize in a way that researchers know humans are capable, however they lack the ability to reason relationally because they are not sensitive to structure, nor are they capable of sufficiently representing roles and fillers independently. Whereas traditional symbolic models lack the ability to capture shades

of meaning and semantics in such a way that they could recognize that a cat is more like a dog than a fish, but more like a fish than a telephone. Additionally, traditional connectionist models also have a fairly transparent implementation (i.e., nodes in a network are most like neurons in the brain), whereas symbolic models do not. The alternative to these two types of model, symbolic-connectionist models, takes the flexibility of a distributed representation and marries it with the ability to represent relational roles independently of their arguments as well as specify which roles are bound to which arguments (Doumas & Hummel, 2005). The LISA/DORA models of relational reasoning (Hummel & Holyoak 1997, 2003; Doumas et al., 2008) are two symbolic-connectionist models that are relevant to the current research.

#### 4. LISA/DORA models of relational reasoning

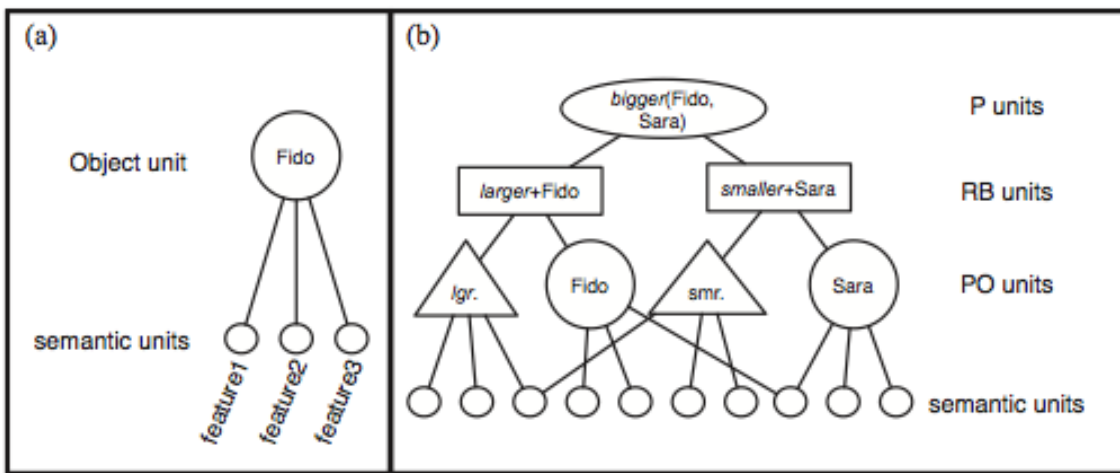
To date, there are many significant models of relational reasoning (e.g., Falkenhainer et al., 1989; Holyoak & Thagard, 1989), however of particular interest to the proposed study are two of these models of relational reasoning: LISA (Learning and Inference with Schemas and Analogies), developed by Hummel & Holyoak (1997, 2003) and DORA (Discovery of Relations by Analogy), developed by Doumas et al. (2008). These symbolic-connectionist models offer an exciting opportunity to explore the boundaries of nonhuman animal cognition. LISA/DORA collectively account for approximately 90 phenomena from the literature on human cognitive development (e.g., Hummel & Holyoak 1997, 2003; Doumas et al., 2008).

DORA was developed from LISA, in response to the criticism that the LISA model was the not able to account for where structured representations like those used in the model originate from (Munakata & O'Reilly, 2003; O'Reilly & Busby, 2002; O'Reilly et al., 2003). DORA solves this particular problem by offering a neurally plausible instantiation of how structured

representations could be learned from unstructured examples observed in the environment. For the sake of brevity the models are best understood in regards to (i) the nature of their representations, (ii) the interaction between the recipient and the driver, and (iii) how the DORA model learns relations.

#### 4.1 The nature of representations

Figure 1.



Representations within LISA/DORA exist as a hierarchy of distributed and localist units in a layered connectionist architecture (see Figure 1). On the bottom layer of the representational structure are semantic units coding for the features of objects and roles (or predicates), in a distributed fashion.<sup>3</sup>

<sup>3</sup> In DORA, semantic units are shared between predicates and objects for two important reasons. First, it is important for the meaning of some property of an object and the explicit predicate of that property to mean the same thing (Doumas et al., 2008). That is, ‘blue’ as a feature of the ocean would otherwise be unlike blue as a predicate, which can then be that cast upon any object (Hummel & Holyoak 1997, 2003; Doumas et al., 2008). In addition to these shared meanings, this shared semantic pool is important because of the role that comparison plays in learning

In the layer above the layer of semantic units are localist predicate-object (or PO) units. PO units act as tokens of individual predicates and objects and are connected to the semantic units in the layer below. Above the PO units, localist role-binding (RB) units, link predicate and object units into role-filler pairs. Proposition (P) units in the top layer, link sets of RB units together to form whole propositions (see Figure 1).

#### 4.2 Flow of Control

As propositions in LISA/DORA enter active memory they are divided into two mutually exclusive sets for the purposes of reasoning. The first set, the driver, controls the sequence of firing events within the model and is analogous to what the model is “attending to” at a given time. The activated semantic units linked to the tokens in the driver retrieve propositions (stored in long term memory) into the second set, the recipient, making them available for mapping to propositions in the driver. Activation in the model then flows from the driver into the shared pool of semantic units, which in turn causes the recipient to become active in response (Doumas et al., 2008).

When a proposition becomes active in the driver, role-filler bindings must be represented dynamically on the units that maintain role-filler independence (see, e.g., Hummel & Holyoak, 1997). In LISA binding information is carried via synchrony of firing; units representing relational roles fire at the same time as their arguments and out of synchrony with other role-filler bindings. In the relation *loves*(Neal, Margo) Neal is fired at the same time as the role “lover” and Margo is fired at the same time as the role “beloved” but out of synchrony with Neal and “lover.” Carrying binding information via synchrony allows for versatile representations,

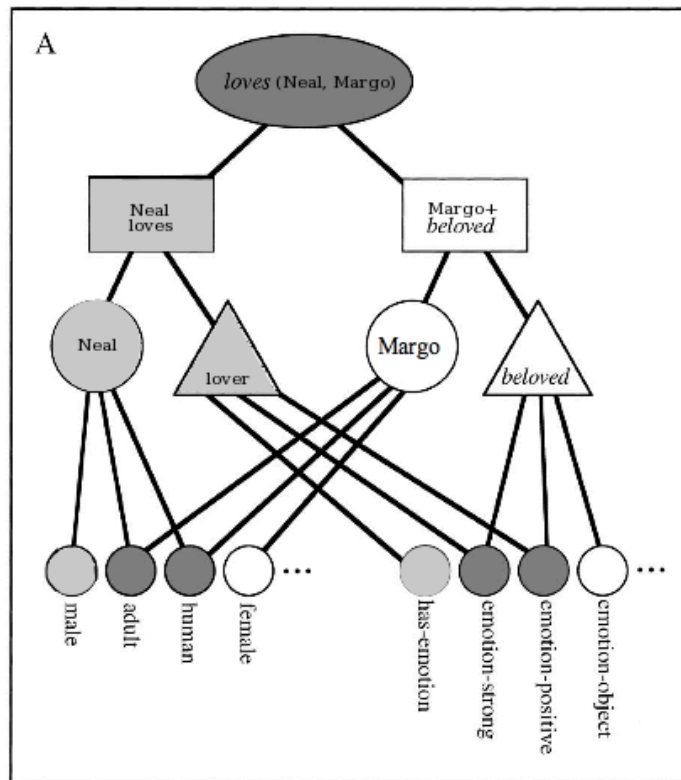
---

higher relational structures (which will be addressed in greater detail during the discussion about relational learning).

where changing the relation to *loves*(Margo, Neal) would only involve changing *when* units fire (i.e., Margo would be fired at the same time as “lover” and Neal at the same time as “beloved”).

In DORA binding information is carried via asynchrony of firing. In asynchrony of firing, roles and the arguments to which they are bound fire in direct sequence as asynchronous couplets. In the example *loves*(Neal, Margo) Neal and the role “lover” would fire in direct sequence forming the concept *lover*(Neal) and Margo and the role “beloved” would likewise fire in direct sequence, resulting in the concept *beloved*(Margo).

Figure 2.



**Figure 2:** Binding by synchrony in LISA (Hummel & Holyoak, 1997; 2003): The lightly colored Neal and *lover* are fired in synchrony with one another and out of synchrony with the white colored Margo and *beloved*. While Margo and *beloved* are fired in synchrony with one another. In DORA, Margo and *beloved* fire, instead, in close temporal proximity.

The asynchronous firing patterns that carry binding information are reliant on inhibitory signals. The first of these signals is lateral inhibition; units in the same layer connected to the

same unit above will compete to become active (e.g., PO units that are connected to RB units will compete for activation). In addition to lateral inhibition, each token unit is yoked to an inhibitor unit that integrates information from that particular token unit and active units in all the layers above. The inhibitor unit will eventually become active and turn off that token unit. A local refresh signal functions at the level of POs and becomes active any time there are no active PO units in the driver (i.e., the POs in the driver have been effectively inhibited and no other PO units are active yet), causing POs in the driver and recipient as well as the semantic units to refresh, returning all of their activations to zero. Similarly, a global refresh signal functions at the level of RBs and becomes active in the same way, but instead refreshes all units in the driver and the recipient, likewise returning all their activations to zero when there are no active RBs in the driver. The local refresh signal is important because it allows PO units in the recipient to keep pace with PO units in the driver and serves as punctuation between a role and its filler signal (e.g., Horn et al., 1992; Horn & Usher, 1990; Hummel & Holyoak 1997, 2003; von der Malsburg & Buhmann, 1992). Likewise, the global refresh signal is important because it allows units in the recipient to keep pace with units in the driver and serves as punctuation between sets of RBs in a proposition.

For example, a typical firing sequence would be: a P unit becomes active in the driver, activating connected RB units in the layer below, which compete to become active (because of noise in the system, one of the RB units will win the competition and become active). After an RB unit becomes active the unit will fire and excite PO units in the layer below. After the PO units compete to become active (once again, one will win due to noise in the system), this now active PO unit's inhibitor will eventually inhibit and turn off this unit. The local refresh signal can then take place (zeroing out all driver and recipient POs and semantics), thus allowing the



other related PO unit to become active. Similarly, once the previously activated RB unit's inhibitor becomes active the global inhibitor fires, zeroing out all units in the driver and the recipient. Note that in this typical firing sequence PO units will oscillate twice as quickly as RB units because their inhibitory information is being integrated from both the layer of PO units and layer of RB units.

### 4.3 Relational Learning

Learning in DORA is achieved primarily through comparison. Initially, DORA starts with simple feature-vector representations of objects (i.e., a node connected to a set of features describing that object). As DORA goes through the process of comparing one object to another object, the corresponding features of those objects fire at the same time. Any semantic features that the two objects have in common become highlighted, receiving twice as much activation as unshared units. DORA then recruits a new PO unit and begins learning positive connections between the new PO unit and the more active semantic units in proportion to their activation (Doumas et al., 2008). The result is an explicit representation of all the properties shared between the two objects, which may contain some extraneous information. For example, if a red apple is compared to a red fire engine, the explicit representation of "red" may also carry with it the feature of "shiny." Therefore additional examples of red are needed in order to weed out unnecessary extraneous features. After this process is applied over much iteration this new PO unit becomes an explicit and structured representation of object properties and relational roles that can be linked together to create full relational structures.

Note that this description of the DORA model is a cursory overview. There are many other capabilities (like learning whole relations) that are not addressed in this description of the

model and these processes are the only ones central to the current discussion. For additional discussion of the details of mapping and learning in the DORA model please see Doumas et al. (2008).

## *Simulations*

The general purpose of this research is to provide a computational account of relational reasoning experiments for nonhuman animals utilizing a model of *human* relational reasoning. Although the description of the DORA/LISA models that are to be used in this experiment are such that they primarily apply to humans, for the following simulations the only features that will be utilized are those that simulate associative learning (i.e., traditional connectionist learning), which involves simply creating connections between coactive units. Therefore, if simulations can be carried out successfully only utilizing connectionist learning, there is good reason to believe that associative learning mechanisms are all that are necessary to perform these tasks.<sup>4</sup> The following sections will briefly describe the behavioral data collected by the researchers who performed the studies, a description of how the study was reconceptualized for the purposes of simulations, and a brief discussion of the implications of these simulations.

### *Simulation 1 – Lazareva et al. (2004)*

Lazareva et al. (2004) utilized a familiar paradigm initially created to test transitive responding in a number of species including humans (McGonigle & Chalmers, 1977; Bryant & Trabasso, 1971) to ascertain whether or not hooded crows (*Corvus cornix*) with adequate training would be able to reason about untested pairs of stimuli. Initially crows were either trained on one of two manipulations involving linearly ordered stimuli placed on cards for which color co-varied with size (they were ordered either biggest to smallest or smallest to biggest) or unordered stimuli that only varied in color (they were all the same size). The trainings encompassed the

---


<sup>4</sup> Notice that utilizing the model in this way does not invoke any use of symbolic thought, as it is theorized that humans do. Whether or not animals invoke symbolic thought to reason is still subject to debate (Penn et al., 2008).

following pairings (+ denotes rewarded stimuli, - notes non-rewarded stimuli): A-/B+, B-/C+, C-/D+, and D-/E+. The critical test of an understanding of transitive inference was in B-/D+ pairings, as B and D had both been rewarded an equal number of times. Lazareva et al. found that hooded crows that were given linearly ordered stimuli were able to perform better on novel B-/D+ trials than crows that had unordered stimuli that varied on color alone. The authors concluded that linearly ordered stimuli were necessary for proper responding in transitive inference tasks for hooded crows.

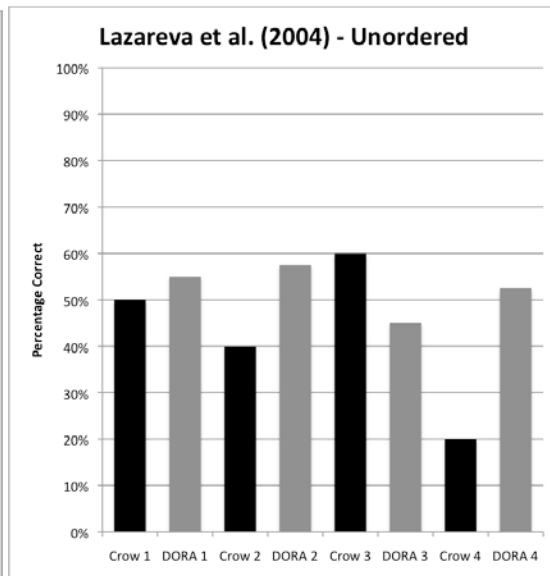
The experiment presented by Lazareva et al. (2004) was reconfigured for simulation in DORA by first taking the total number pairings of A-/B+, B-/C+, C-/D+, and D-/E+ trials and creating instances in memory proportional to the number of training trials each individual crow had with the stimuli (see Appendix A, Table 1 for the stimuli proportions for the ordered manipulation and Appendix A, Table 2 for the stimuli proportions for the unordered manipulations). The number of stimuli that corresponded with each reinforced and non-reinforced pairing was then saved to a memory set. A total of eight memory sets were created (four unordered stimuli memory sets and four ordered stimuli memory sets). Each of the eight memory sets were run for 40 trials of B-/D+ pairings. The memory sets for unordered and ordered stimuli varied from one another in regard to the features present. In ordered stimuli features involved relative surface area, color, absolute size, and whether or not the item was rewarded or non-rewarded. In unordered stimuli features included only color and whether the stimuli had been rewarded or non-rewarded (see Appendix B, Table 3 & 4 for a detailed description of these simulation files).

During simulations the model would retrieve items from memory and then attempt to map them to the current focus of attention in the driver (i.e., the novel B-/D+ pairing). Successful

mappings were those where B- was successfully mapped to non-rewarded stimuli from memory or D+ was successfully mapped to rewarded stimuli from memory. Unsuccessful retrievals from memory were not counted as misses (i.e., when the model failed to retrieve anything from memory, as there is an element of chance associated with retrieval from memory, the retrieval function would be called again until something relevant was pulled from the memory set). Incorrect B- to rewarded stimuli mappings and D+ to unrewarded stimuli mappings were counted as incorrect.



**Figure 3:** Results of the simulation of the ordered stimuli. Dark bars denote the original subjects' performance, where light bars denote the performance of the model.



**Figure 4:** Results of the simulation of the unordered stimuli. Dark bars denote the original subjects' performance, where light bars denote the performance of the model.

The results of simulation of the Lazareva et al. (2004) study are pictured in Figures 3 and 4. In the ordered manipulation (Figure 3), DORA simulated Crow 4 the best (Crow 4's performance was at 82% correct on B-/D+ trials, while DORA performance at 80% correct). Results were similarly close for Crow 3 (with the crow performing at 85% correct and DORA performing at 87.5% correct). Simulations for Crow 1 & 2 were not as close, with the model

underperforming slightly, relative to the behavioral data collected; Crow 1 scored 82% correct while the model performed at 75% correct, and Crow 2 also scored 82% correct while the model performed at 77.5%.

In the unordered manipulation (Figure 4), DORA was able to closely model the performance of Crow 1 (50% correct for the subject, 55% correct for the model). The model did a sufficient job at performing at or around chance much like Crows 2 & 3, although the difference was not as small as that of Crow 1. The last subject's performance was significantly below chance, where the model performed at chance, 52.5%.

In the simulation of Lazareva et al. (2004), DORA was able to reliably account for four out of four crows in the linearly ordered stimuli manipulation. Lazareva et al. argue that linearly ordered stimuli are necessary components for hooded crows to perform transitive inference tasks. Few would deny that crows are capable of some sensitivity to surface area, or relative size of two objects, and would be able to learn to discern between two stimuli of different sizes. The real question is, had the crows learned a concept of "bigger-than?" Or rephrased, had the crows successfully predicated a "bigger" relation and were they able to use it in service of reasoning about the novel B-/D+ pairing? The simulations demonstrate that associative mechanisms are enough to achieve above chance performance. Couvillon & Bitterman (1992) further demonstrated the idea that associative learning mechanisms are sufficient for transitive inference tasks by utilizing Rescorla & Wagner's (1972) mathematical model to simulate similar data compiled by Fersen et al. (1991). Couvillon & Bitterman's conclusion is further validated by DORA's results.

Lazareva et al.'s (2004) unordered transitive inference task was also replicated successfully by DORA. Recall that in the unordered manipulation color was the only salient

feature associated with reward and non-reward, by having one less systematic feature in the system associated with reward, it was much more difficult for the model to reliably retrieve from memory and map appropriately. Although, as stated, the model does not punish objects for having less features – one less feature that co-varies reliably with reward does impact how often that object will be retrieved from memory and does impact how readily it will be able to be mapped to another object. This can account for the differential results between conditions, all other things held equal. In sum, the evidence seems to suggest that instead of a predicated “bigger-than” or “more-than” concept, the crows were reasoning with associative mechanisms in conjunction with the extra boost to memory afforded by the relative disparities across surface areas on each of the cards.

#### *Simulation 2 - Cook & Wasserman (2007)*

Relational match to sample (RMTS) tasks have been used to assess a number of species’ understanding of the relationship between items, and at one point it was thought that only humans and language-trained apes were capable of succeeding in these tasks. In RMTS tasks, a sample pair that is either same or different is pitted against two choices: a same pair that is not a perceptual match to the sample pair or a different pair. Furthermore, putting a member of the same pair in the different pair increases the level of difficulty because a subject can then reason based on either “same pair” or “same item.” Cook & Wasserman (2007) demonstrate that in sample and test arrays that consist of 16 multi-element arrays, pigeons are capable of being successful on RMTS tasks.

The simulation of the Cook & Wasserman (2007) study was done under the assumption that entropy is a meaningful perceptual feature of arrays with 16 icons. Likewise, it is the

assumption of the simulation that at first the pigeons in the study adopted a rote memorization tactic, and over each phase slowly abandoned this in favor of entropy detection. Across baseline trials, the pigeons consistently achieved above chance performance ranging from 71% (Phase 4) to 81.8% (Phase 2), but dropped to below chance performance on novel stimuli and novel stimuli configurations, with increasing improvement to transfer tasks across each phase. This falls in line with the notion that rote memorization, which was an unsuccessful tactic because the transfer tasks were novel, was the preliminary tactic and was slowly abandoned for entropy or, similarly, the pigeons eventually learned that entropy was the relevant perceptual feature over each phase.


Ten simulation files were created for each phase; five were “different” samples and five were “same” samples. In phase 1 simulations, the different objects were created by randomly selecting three of twenty features with two noise features appended to the object. Same objects were created by randomly selecting one of twenty features with two noise features appended to the object.<sup>5</sup> In phase 2 simulations, one of the features in the different objects was replaced with a high entropy feature; in the same objects no features were removed, but a low entropy feature was added. In phase 3.1 simulations, one noise object was removed and replaced with another high entropy feature in the different objects and in the same objects nothing was removed, but a second low entropy feature was added. In phase 3.2 simulations, another feature was removed from different objects and replaced with a third high entropy feature, while no features were removed from same objects and a third low entropy feature was added. In the final manipulation, the arrays had their noise features reintroduced to simulate the novelty of the test sets in phase 4 (see Cook & Wasserman, 2007). A detailed table that depicts the simulation files for this experiment can be found in Appendix B, Table 5.

---

<sup>5</sup> In DORA objects with more features than another object do not get preferential treatment during mapping.



In each simulation, the sample (either same or different) was loaded into the driver with the choices (in this case both a “same” object and “different” object) loaded into the recipient and then mapped. A successful mapping was one in which the model mapped the same choice to the same sample, or different choice to a different sample. Incorrect mappings were ones in which the model mapped a same sample or choice to a different sample or choice. Twenty total trials were run per phase (each stimuli configuration was run twice), for a total of one hundred trials across all phases.



**Figure 5:** Results of the simulation of Cook & Wasserman (2007) study. Dark bars denote the original subjects' performance, where light bars denote the performance of the model.

The results of the simulation of the Cook & Wasserman (2007) relational match to sample are summarized in Figure 5. In the original data sample, after Phase 1, pigeons' transfer to probes was 45.3% correct; the simulation of Phase 1 in DORA was able to closely replicate Phase 1 trials, performing at 40% correct. During phase 2, the transfer for all three pigeons to probes was 59.2% correct which, once again, was closely approximated by the model's simulation, which performed at 55% correct. When novel samples were once again presented to

the pigeons in phase 3.1, the performance of the subjects was at 70.1% correct while the model performed almost identically, with 70% correct. In phase 3.2, the pigeon's pooled correct responses was 76% (this is also the best the pigeons performed on the task in all) while the model predicted their performance to be at 80%. In the final phase the model predicted that the animals would perform at about 70% correct while the subjects performed at 63%. Overall, the model was able to account for the data in these tests.

Once again, DORA reliably modeled the data of the Cook & Wasserman RMTS task with retrieval from memory and mapping. In this case, the variability of performance was captured by introducing noise and entropy detection over the course of the trials. This simulation does make the explicit assumption that the pigeons in the RMTS task were, at first, relying on rote memorization in order to learn the baseline set and then, after multiple transfer tasks, eventually abandoned that method in favor of entropy-related selection. Transfer stimuli (as noted by Cook & Wasserman, 2007) were differentially rewarded, so it is important to address why the pigeons might have abandoned rote memorization in favor of entropy detection. This could be in large part due to the sheer volume of stimuli to be memorized. For example, Set1/Set2 (i.e., Phase 1) contained 160 stimuli – 80 same arrays and 80 different arrays. Similarly, in Phase 1 Set 1 was the sample stimuli and Set 2 encompassed the choice or test stimuli, the transfer task simply involved making Set 2 the sample stimuli and Set 1 the choice stimuli. A method in which the animal memorizes the Set 1 samples and the Set 2 choices would then prove difficult to generalize to Set 2 samples with Set 1 choices (this was demonstrated in the data provided by Cook & Wasserman; see Figure 5.).

Fagot et al. (2001) established that entropy could be a salient perceptual feature associated with relational match to sample, and there has been growing concern that RMTS tasks

lack the power to demonstrate relational reasoning ever since. In this simulation the model has demonstrated that, indeed, it is not necessary to have a predicated representation of same or different in order to perform above chance; instead a subject merely needs to be able to perceive entropy. The argument has been made that same/different discriminations in humans are likewise, due to sensitivity to entropy. However it is important to note that predicated sameness and differentness manifests itself in other tasks of human cognition. For example, it is unlikely that entropy has anything to do with understanding that *loves* and *hates* are the same in some sense (i.e., they are both very strong emotions).

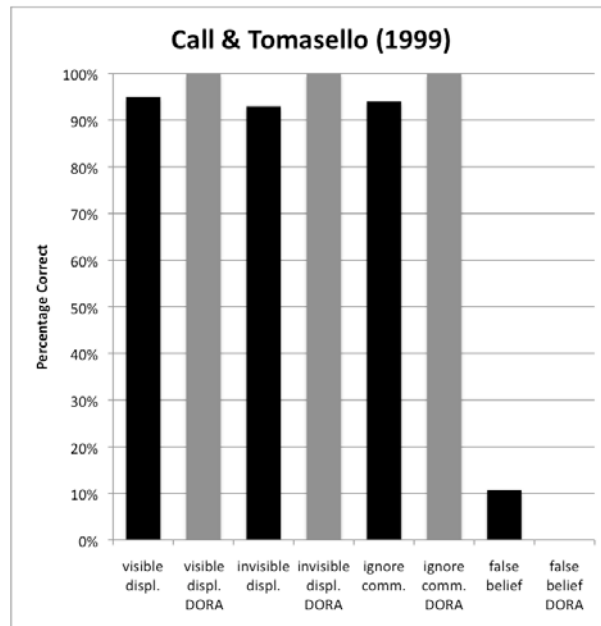
The DORA simulations of the Cook & Wasserman (2007) relational match to sample task lends support to the notion that entropy may be the only necessary feature of multi-element relational match to sample arrays. In conjunction with retrieval from memory and mapping, there may be no need to understand anything like a concept of sameness or differentness in order to perform well on these tasks.

### *Simulation 3 - Call & Tomasello (1999)*

Theory of mind tasks are often thought to be relational in nature, as it requires a subject to reason based upon the mental contents of another, effectively casting a belief state on some proposition, making the relational structure higher-order (e.g., *knows*(John, *loves*(Mary, Paul))). Reasoning in this manner is sophisticated and is present in children by age five (Call & Tomasello 1999). Theory of mind is a hotly debated topic within the comparative literature (see Penn & Povinelli 2007 for a discussion). Call & Tomasello (1999) attempted to test whether or not chimpanzees (*Pan troglodytes*) and orangutans (*Pongo pygmaeus*) are capable of reasoning about the false beliefs of an observer. In this experiment the apes were tested on all the

component capabilities one would need in order to understand a false belief task, which involved following a marker placed by an observer, tracking the movement of food between boxes, as well as tracking the movement of a box that had food inside of it. In the final manipulation of the task the observer's mental contents became critical to selecting the box containing the food reward and all apes performed far below chance on these false belief trials (although they maintained above chance performance on the other non-false belief component tasks).

This task was simulated by first making the assumption that instead of reasoning based upon the actions of an observer, the apes in the study were instead reasoning based on some form of the following logic: if the food is seen to be in a particular location, choose that location; or if nothing has been seen, instead choose the box with the marker on it. All four control test situations as well as the false belief task were then coded with a box1 and box2 object that had feature vectors that consisted of features of boxes as well as features like having a mark or having seen food put into the particular box. These simulation file configurations can be seen in Appendix B, Table 6. Notice that the various manipulations do not appear to differ from each other meaningfully; the objects were coded based upon what was viewed by the subject, and not by taking into consideration all aspects of the task. In addition to the five task situations, two objects were created and placed in the driver. These objects can be thought of as representations of selection criteria as they are not predicated, but instead just offer a way for the model to map onto the task trials and for us to decipher what those mappings in turn mean. The selection criterion was then loaded into the driver while the boxes were loaded into the recipient and the model was to map between the two. A successful mapping was counted when the proper selection criteria were mapped to the proper object, other mappings were considered misses. A total of ten trials per manipulation were simulated.



**Figure 6:** Results of the simulation of the Call & Tomasello (1999) study. Dark bars denote the original subjects' performance, where light bars denote the performance of the model.

The results from the simulation of the Call & Tomasello (1999) non-verbal false belief task are pictured in Figure 6. In all three manipulations of the control task (visible displacement, invisible displacement, and ignore communicator conditions) the model was 100% accurate. In comparison, the apes' scored 95% correct in the visible displacement task, 92% correct on the invisible displacement task, and 93% correct on the ignore communicator task. Similarly, the model performed at 0% on the false belief manipulation, while the apes performed at 10.7% correct. The DORA model was able to perform flawlessly on all of Call & Tomasello's control tasks, resembling the behavioral data collected. The apes performed above 90% on all of these control tasks. Furthermore, the apes performed at significantly below chance for the false-belief task, whereas the simulation of the false-belief task proved to be impossible for the model. The small differences that were observed between the model and the behavioral data could likely be simulated by adding some noise into the simulations, as the model does not have to worry about

monitoring its own attention. The simulations assume perfect attention, however it could be the case that the apes were not paying full attention to the task at all times, and this may account for the less than 100% performance on their part.

These simulations imply that instead of predicating anything in regard to the knowledge of the mental contents of another, the apes were reasoning by retrieving memories of receiving food rewards and the associated perceptual features of the task configuration, and mapping those features onto the test configurations. Call & Tomasello (1999) admit that the marker proved to be such a salient predictor of the reward that the apes appeared to be unwilling to abandon this strategy in favor of reasoning about the cues from the observer. Furthermore, Call & Tomasello go on to discuss the possibility that the task may have been too difficult as success would have involved coordinating many different small pieces of evidence. In support of this notion, the results were simulated without any kind of information from the observer or from the hider. Instead, the only necessary features were the two boxes, whether they had been marked or not, and whether the animal saw the food enter a particular box. Therefore it is unlikely that the observer's behavior had any impact on the apes' reasoning. These results further validate Call & Tomasello's conclusion that the apes were not capable of utilizing the mental contents of the observer to reason successfully on the false-belief task. Although it by no means demonstrates definitively that theory of mind capabilities do not exist in apes, it does show that future tasks developed to test for theory of mind need to be cautious about what salient perceptual features exist in their tasks, such that alternative methods for reasoning may prove to be successful. As Call & Tomasello demonstrate, it is important that in theory of mind tasks that the presence or an absence of a theory of mind must be the only possible explanation of the outcome.

## *General Discussion*

The goal of the present research was to demonstrate that relational reasoning might be able to account for the divergence between the reasoning capabilities of human and nonhuman animals. In humans we observe flexible reasoning capabilities that allow for the technological advances seen in culture and society today. On the other hand, animals appear to be limited in their reasoning capabilities. Three behavioral studies that attempted to assess relational reasoning capabilities within animal species (two of which claimed to have properly demonstrated such abilities) were simulated using DORA, a model of human relational reasoning, by only using the retrieval from memory and mapping functions on perceptual features. These two processes in isolation demonstrate that all three studies could be simulated with only associative learning mechanisms (i.e., that the ability to map the perceptual features of one item to another item in memory or to an item that is the focus of attention).

The results of these experiments, however, do not definitively demonstrate that relational reasoning is absent in nonhuman animals, only that attempts thus far are incapable, in principle, of demonstrating relational reasoning, for two very important reasons. First, the relevant perceptual features in any task that may lend themselves to associative mechanisms like perceptual mapping are sometimes non-obvious. It is certainly the case that entropy is a non-obvious feature of multi-element same/different arrays. In tasks of transitive inference, especially those that employ linearly ordered stimuli, confounding perceptual features do exist. Many species have demonstrated sensitivity to surface areas, and it could be argued that species would be greatly disadvantaged if they were incapable of discerning between which is “less” and which is “more.” However, this does not provide proof positive that the animal has obtained anything

resembling a concept like “bigger-than” or “more-than.” Instead, it only shows that surface area is a salient perceptual feature in discrimination tasks.

The second issue with tasks of relational reasoning as presented to animals is they typically involve *interpolation* and not *extrapolation*. Any predicated concept extrapolates with little effort. Recall that part of what makes relational reasoning so powerful is that it is functionalized and therefore these concepts can be applied to novel cases, and these novel cases can be reasoned about with no training necessary at all. However, it is common within the comparative literature to demonstrate only interpolation. In the case of Lazareva et al. (2004) the B-/D+ pairing is a case of interpolation as the crows had already learned reward/no reward mappings for both of these cases. A better test of whether or not a “bigger-than” concept had been predicated could have been demonstrated by an X-/Y+ pairing., If the concept has been learned, it should generalize to X-/Y+ pairings without training.

Consider the following scenario: One is invited to guess the rules to a game by observing the behavior of game host and a contestant. The game host says, “One”; the contestant replies, “One.” The game host says, “Two,” and the contestant replies, “Two.” Then the game host says “Flower,” and the contestant replies, “Flower.” What would the contestant likely say if the game host says “bouncy?” It is likely that your answer would be “bouncy”. This is an example of the identity function (formally:  $f(x) = x$ ), or simply the “you say what I say” game. Notice that since the rule has been functionalized, regardless of what takes the role of  $x$ , the output will always be  $x$ . This function extrapolates without any additional training. Now imagine a possible world in which one could not generate an output without previous experience with the input. In order to perform properly on the “you say what I say” game, one would first have to learn all possible mappings; A would have to be mapped to A, B to B, and so on. In the case one was presented



with the novel “bouncy” example, the answer would be unknown because a mapping from bouncy to bouncy has never been learned (Doumas & Hummel, 2003; Hummel & Holyoak 2000). In the context of the Lazareva et al. (2004) ordered stimuli experiment, mappings had been generated to B-/D+ in the context of B-/C+ and C-/D+ pairings.

In sum, although the simulations contained in this research are not enough in and of themselves to disprove the hypothesis that animals are capable of relational reasoning, they do establish that current tests of relational reasoning are not sufficient to demonstrate this phenomenon in nonhuman animals, echoing the conclusions drawn by Penn et al. (2008). Relational reasoning is yet to be successfully demonstrated in nonhuman animals, lending some credence to the notion that aspects of relational reasoning as observed in humans may be diminished or even missing in their entirety in animals. Understanding which aspects and underlying mechanisms are and are not present in nonhuman animals is still a primary goal, and should be investigated in future research.

## *Future Research*

The following descriptions are of manipulations that could be made to the DORA/LISA models of relational reasoning in future research. The goal would be to develop the models of relational reasoning to account for the animal data across different species by establishing which mechanisms may or may not be present by removing aspects of relational reasoning from the human version of the model.

### 1. Predication

Predication is broadly used throughout the description of the DORA/LISA models to mean “the ability to represent something as an explicit entity” (Gentner, 1983). Within the models, predication is the outcome of four very important processes: retrieval of propositions from long-term memory (LTM), analogical mapping of propositions currently in working memory, intersection discovery for predication and refinement, and linking of role-filler sets into higher arity structures via self-supervised learning (SSL) (for the sake of brevity SSL will not be discussed here, instead please see Doumas et al. (2008) for details). Therefore, one way to begin asking questions about the nature of relational reasoning in nonhuman animals may be to begin with an understanding of how predication may come to be in different species. For example, we can begin by asking whether or not the cognitive architecture of nonhuman animals allows for the predication of relational properties. Admittedly, how DORA goes through the process of predication is partly due to the way in which DORA carries binding information. For the purposes of separating our potential manipulations dealing with predication, which is the focus of this section, manipulations associated with binding are covered in the section *Asynchrony of Firing* below.

Within DORA, predication is done throughout the process of building a representation allowing for higher arity structures and is done in the following way: first, DORA begins with a feature vector (that is, nodes that code for features in a network that are associated with a particular object). The features in this vector are then compared to get PO units that code explicitly for role information. This process is important because these new PO representations can then be bound (i.e., predicated about) objects. PO units are then linked together to create role-filler pairs for the purpose of long term storage.<sup>6</sup> RB (role-binding) units are then linked together in the layer above into full relational propositions, known as P units. This comparison-based learning process is then applied iteratively, refining PO and propositional representations. The final outcome of this process is a full LISAese representation (see Figure 1) that supports the high-level relational cognition seen in humans.

Therefore, truncating the representational structure at these different levels is ripe for exploration. Perhaps it is the case that nonhuman animals are incapable of representing role-filler pairs, leaving them with mental representations that only code for roles and predicates with no clear way to link them together. Alternatively, it could be the case that nonhuman animals are capable of representing role-filler pairs, but are unable to represent entire propositions.

## 2. Comparison-Based Predication

The aforementioned process of predicating higher arity structures by recruiting new nodes within the model is driven by DORA's ability to compare (Doumas et al., 2008).

Comparison plays a key role in how relations are initially learned by highlighting shared

---

<sup>6</sup> RB units are essential for establishing the asynchrony patterns that carry binding information in working memory. This prevents features of one role-filler pair from becoming conflated with features in the other role-filler pair.

properties between objects and roles (Gentner, 1983, 2003; Mandler, 1988, 1992, 2004). These shared properties make up what information is being encoded into the recruited nodes that in turn make up the more complicated structures in DORA's mental representations. It could then be the case that the ability to compare may have some fundamental limitations in nonhuman animals. That is to say that perhaps what things are highlighted during the process of comparison are unlike those things being highlighted when a human compares objects and roles. It could also be the case that co-activation in PO units is not enough, in and of itself, to warrant recruitment of the new nodes to code for more complicated aspects of the representational structure. This leaves comparison-based predication as a worthwhile variable to manipulate.

### 3. Asynchrony of Firing

A reoccurring theme throughout the cognitive modeling of the relational reasoning literature is that an adequate symbolic account of human relational reasoning requires dynamic binding. Binding, in and of itself, is simply a way to link together representations while maintaining independence of those representations. Frequently referred to as the "binding problem," this problem is addressed in the DORA model by systematic asynchrony of firing. Asynchrony of firing, as described earlier, involves roles and the arguments that are bound together to fire in close temporal proximity to each other (Doumas et al., 2008). Therefore, timing of firing signals is subject to manipulation within the model.

One way in which the firing signals can be manipulated is via the inhibitory refresh signals that separate RBs and POs (see Horn et al, 1992; Horn & Usher 1990; Hummel & Holyoak 1997, 2003), which create a predictable pattern of activation on the semantic units. Global inhibitory signals are easy candidates for variable manipulation; the timing of which can

be changed or they could be removed completely. Inhibitory signals are also used between nodes at the same level; these types of inhibitory signals are a type of lateral inhibition (e.g., where firing a node at the level of RBs inhibitory signals are sent to competing RB units which in turn excite the PO units in the layer below). In this case, lateral inhibition can be removed and the ability of the model to reason can then be assessed.

#### 4. Refinement

Due to the way in which DORA learns relational representations from unstructured examples (via a comparison algorithm that looks for featural overlap that in turn recruits nodes to code for those featural overlaps), initial representations will code unnecessary additional features as part of the initial structure. For example when attempting to learn to predicate the concept *big* using the exemplars ‘robot’ and ‘cactus,’ the initial representation may be confounded with the feature *green*, as the exemplars offered to the model for learning may both coincidentally be green (Doumas et al., 2008). However, as the model runs through multiple examples of *big*, *big* becomes properly predicated and featural overlaps that are extraneous are weeded out. This process is done through a self-supervised learning (SSL) algorithm (see Doumas et al. (2008) for a discussion of how SSL is used to refine predicates). The details of this process are not important in the limit, but the fact that SSL is allowed to occur for new predicates is in and of itself an interesting point. It could be the case that while attempting to predicate parts of the relational structure in nonhuman animals, the cognitive processes that allow refinement to take place may not be present, leaving animals with ‘dirty’ representations.

Within the comparative psychology literature the ability for animals to solve any particular task seems relatively task specific. For example, in the Visalberghi & Limongelli

(1994) test of causal reasoning in capuchin monkeys, subjects attempted to retrieve a piece of food placed in a transparent tube using a straight stick. There was a trap in the middle of the tube in which the food would fall if it was pushed over the hole. The idea was to push the food away from the hole in order to retrieve the food. After 90 trials, only one of the four subjects was able to successfully perform the task. Furthermore when the experimenters put the hole above the food (removing the hole as an obstacle), the one successful subject still treated the hole as if it was an obstacle. These results led Penn et al. (2008) to conclude that the causal knowledge used by nonhuman primates is “tightly coupled to specific task parameters and bodily movements” (2008, p. 119).

This type of task specificity may be the byproduct of mental representations that, once created, do not go through any type of refinement or updating and therefore are not generalizable to other tasks. Alternatively, animal mental representations may go through some refinement, but the refinement process may asymptote much earlier than those of the SSL presently implemented within the DORA model.

## 5. Mapping

Generally speaking, mapping creates opportunities for DORA to predicate new properties by linking together units that are active at the same time. When units in the driver and the recipient map together this signals that units share some correspondence with each other. Mapping is involved in many of the aforementioned processes, specifically, predication of new properties, formation of new relations, as well as refinement (Doumas et al., 2008). Changing the way the model maps will then also have implications for those processes. Similarly, changing the

way the model maps in any combination of those processes may also hold the key to some of the gradients in behavior observed across species and between human and nonhuman animals.

## *Conclusion*

Understanding the nature of nonhuman animal cognition is an incredibly important task for comparative and cognitive psychologists alike. Although the applied aspects may not be directly evident, this research does have the potential to cause researchers to reflect on the nature of the questions they posit to nonhuman animal species. In addition, computational modelers of human cognition have enjoyed some great successes in understanding the way that the human mind works. This work could prove to be a launching point for the creation of neurally plausible cognitive models of animal mental processes as well. It could also in the future provide an excellent basis for evolutionary biologists to search for the structural differences that support the process level differences that may be observed (Penn et al., 2008).

Broader implications of such research may require us to reanalyze the use of nonhuman animals as behavioral models of genetic and psychological diseases. For example, developing drug treatments for a spectrum disease like autism with the aide of transgenic mice may prove to be difficult if the way in which mice come to represent their world is drastically different to the way in which humans come to represent and manipulate those representations of their world. If the mental world of animals and humans differ at some important level it is important to realize that drugs tested on and used to treat human-like symptoms in animals may prove to have limited effect on humans.

The simulations and data collected for these experiments are only step one in the process of understanding the underlying differences that may account for the divergence between animal and nonhuman animal cognition by establishing that present research is not sufficient to draw claims about relational reasoning in nonhuman animals. Further demonstrating the need for formal models of mental processes. The next steps for understanding what underlying processes



(or lack thereof) seemingly preclude nonhuman animals from this ability have already been described in the *Future Research* section. Making functional changes to the architecture of cognitive models of relational reasoning in an attempt to better understand what underlying mechanisms may or may not be present in animal cognition is an important next step for researchers interested in an understanding of relational reasoning in nonhuman animals that satisfies all of Marr's levels of analysis.

APPENDIX A

Table 1.

*Lazareva et al. (2004) Ordered Stimuli Proportions*

Raw Trials	Crow 1	Crow 2	Crow 3	Crow 4
A-/B+	131	114	96	127
B-/C+	153	83	122	96
C-/D+	71	256	145	98
D-/E+	60	151	50	39
<b>Total</b>	<b>415</b>	<b>604</b>	<b>413</b>	<b>360</b>
<b>% Trials</b>				
A-/B+	31.5%	18.8%	25%	35.2%
B-/C+	36.8%	13.7%	31.8%	26.6%
C-/D+	17.1%	42.3%	37.8%	27.2%
D-/E+	14.4%	25%	13%	10.8%
<b>Trials in memory</b>				
A-/B+	16	9	12	18
B-/C+	18	7	16	13
C-/D+	9	22	17	14
D-/E+	7	12	6	5

Table 2.

*Lazareva et al. (2004) Unordered Stimuli Proportions*

Raw Trials	Crow 1	Crow 2	Crow 3	Crow 4
A-/B+	170	96	273	364
B-/C+	72	76	117	64
C-/D+	591	224	269	622
D-/E+	217	93	86	89
<b>Total</b>	<b>1050</b>	<b>489</b>	<b>745</b>	<b>1139</b>
<b>% Trials</b>				
A-/B+	16.2%	19.6%	36.6%	32%
B-/C+	6.9%	15.5%	15.7%	5.6%
C-/D+	56.2%	45%	36.1%	54.6%
D-/E+	23.5%	19%	11.5%	7.8%
<b>Trials in memory</b>				
A-/B+	8	10	18	16
B-/C+	4	8	8	3
C-/D+	28	23	18	27
D-/E+	10	9	6	4

APPENDIX B

Table 3.

*Lazareva et al. (2004) Simulation File Overview – Ordered Stimuli*

Ordered Stimuli – A-/B+		Ordered Stimuli B-/C+	
Object	Features	Object	Features
A-	- color1 size1 noReward	B-	- color2 size2 noReward
B+	+ color2 size2 reward	C+	+ color3 size3 reward
Ordered Stimuli C-/D+		Ordered Stimuli D-/E+	
Object	Features	Object	Features
C-	- color3 size3 noReward	D-	- color4 size4 noReward
D+	+ color4 size4 reward	E+	+ color5 size5 reward
Ordered Stimuli B-/D+			
Object	Features		
testB-	- color2 size2		
testD+	+ color4 size4		

Table 4.

*Lazareva et al. (2004) Simulation File Overview – Unordered Stimuli*

Unordered Stimuli – A-/B+		Unordered Stimuli B-/C+	
Object	Features	Object	Features
A-	noReward color1	B-	noReward color2
B+	reward color2	C+	reward color3

Unordered Stimuli C-/D+		Unordered Stimuli D-/E+	
Object	Features	Object	Features
C-	noReward color3	D-	noReward color4
D+	reward color4	E+	reward color5

Unordered Stimuli B-/D+	
Object	Features
testB-	color2
testD+	color4

Table 5.  
*Cook & Wasserman (2007) Simulation File Overview*

Stimuli Configuration 1 – Phase 1			
Object	Features	Object	Features
sampleDiff1	feature11	diffChoice1	feature20
	feature13		feature15
	feature7		feature14
	noise1		noise1
	noise2	sameChoice1	noise2
			feature4
			noise1
			noise2
Stimuli Configuration 1 – Phase 2			
Object	Features	Object	Features
sampleDiff1	feature11	diffChoice1	feature20
	feature13		feature15
	entropyHi1		entropyHi1
	noise1		noise1
	noise2	sameChoice1	noise2
			feature4
			entropyLo1
			noise1
			noise2
Stimuli Configuration 1 – Phase 3			
Object	Features	Object	Features
sampleDiff1	feature11	diffChoice1	feature20
	feature13		feature15
	entropyHi1		entropyHi1
	entropyHi2		entropyHi2
	noise1		noise1
		sameChoice1	feature4
			entropyLo1
			entropyLo2
			noise1

---

Stimuli Configuration 1 – Phase 3.2

Object	Features	Object	Features
sampleDiff1	feature11 entropyHi1 entropyHi2 entropyHi3 noise1	diffChoice1	feature20 entropyHi1 entropyHi2 entropyHi3 noise1
		sameChoice1	feature4 entropyLo1 entropyLo2 entropyLo3 noise1

---

Stimuli Configuration 1 – Phase 5

Object	Features	Object	Features
sampleDiff1	feature11 entropyHi1 entropyHi2 entropyHi3 noise1 noise2	diffChoice1	feature20 entropyHi1 entropyHi2 entropyHi3 noise1 noise2
		sameChoice1	feature4 entropyLo1 entropyLo2 entropyLo3 noise1 noise2

Table 6.  
*Call & Tomasello (1999) Simulation File Overview*

Control Task 1 – Marker Tracking		Control Task 2 – Visible Displacement	
Object	Features	Object	Features
0box1	boxf1	1box1	boxf1
	boxf2		boxf2
	boxf3		boxf3
	hasMarker1		sawFood1
	hasMarker2		sawFood2
0box2	boxf1	1box2	boxf1
	boxf2		boxf2
	boxf3		boxf4
			hasMarker1
Control Task 3 – Invisible Displacement		Control Task 4 – Ignore Observer	
Object	Features	Object	Features
2box1	boxf1	3box1	boxf1
	boxf2		boxf2
	boxf3		boxf3
	hasMarker1		sawFood1
	hasMarker2		sawFood2
2box2	boxf1	3box2	boxf1
	boxf2		boxf2
	boxf4		boxf4
			hasMarker1
False Belief Task		Memory Mappings	
Object	Features	Object	Features
4box1	boxf1	markerBox	boxf1
	boxf2		boxf2
	boxf3		hasMarker1
	hasMarker1		hasMarker2
	hasMarker2		
4box2	boxf1	sawBox	boxf1
	boxf2		boxf2
	boxf4		boxf4
	boxf4		sawFood1
			sawFood2

## References

- Bergman, T. J., Beehner, J. C., Cheney, D. L., & Seyfarth, R. M. (2003). Hierarchical Classification by Rank and Kinship in Baboons. *Science*, 302 (5648), 1234-1236.
- Brown, J.L. (1975). *The evolution of behavior*. New York: Norton.
- Bryant, P. E., & Trabasso T. (1971). Transitive inference and memory in young children. *Nature*, 232, 456-458.
- Call, J., & Tomasello, M. (1999). A Nonverbal False Belief Task: The Performance of Children and Great Apes. *Child Development*, 70 (2), 381-395.
- Castro, J.A., Picornell, A., & Ramon, M. (1998). Mitochondrial DNA: a tool for population genetic studies. *International Microbiology*, 1 (4), 327-332.
- Cheney, D. L., & Seyfarth, R. M. (1990). *How Monkeys See the World*. Chicago: University of Chicago Press.
- Cook, R. G., & Wasserman, E. A. (2007). Learning and transfer of relational matching-to-sample by pigeons. *Psychonomic Bulletin & Review*, 14 (6), 1107-1114.
- Couvillon, P. A., & Bitterman, M. E. (1992). A conventional conditioning analysis of “transitive inference” in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 18 (3), 308-310.
- Dally, J. M., Emery, N. J. & Clayton, N. S. (2006) Food-caching western scrub-jays keep track of who was watching when. *Science*, 312 (5780), 1662-1665.
- Darwin, C. R. (1859). *On the origin of species by natural selection*. London: J. Murray.
- Darwin, C. R. (1871). *The descent of man, and selection in relation to sex*. London: J. Murray.
- Doumas, L. A. A., & Hummel, J. E. (2005). Approaches to modeling human mental representations: What works, what doesn't, and why. In Holyoak, K. J., & Morrison, R.



- G. (Eds.). *The Cambridge Handbook of thinking and reasoning* (73-94). London: Cambridge University Press.
- Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, 115, 1-43.
- Fabre, J. H. (1919). *The hunting wasps*. London: Hodder & Stoughton.
- Fagot, J., Wasserman, E. A., & Young M. E. (2001). Discriminating the relation between relations: The role of entropy in abstract conceptualization by baboons (*Papio papio*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes*, 27 (4), 316-328.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Fodor, J. A. & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.
- Gentner, D. (1983). Structure-mapping: A theoretical framework of analogy. *Cognitive Science*, 7, 155-70.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. London: Cambridge University Press.
- Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp.95-235). Cambridge, MA: MIT Press.
- Gentner, T. Q., Fenn, K.M., Margoliash, D., & Nusbaum, H.C. (2006). Recursive syntactic pattern learning in songbirds. *Science*, 440, 1204-1207.

- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, *12*, 306–355.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1–38.
- Henshilwood, C. S., & Marean, C. W. (2003). The origin of modern human behavior: Critique of the models and their test implications. *Current Anthropology*, *44* (5), 627-651.
- Holyoak, K. J., & Hummel, J. E. (2000). The proper treatment of symbols in a connectionist architecture. In Dietrich, E. & Markman, A. (Eds.). *Cognitive dynamics: Conceptual change in humans and machines* (157-185). New Jersey: Erlbaum.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295–355.
- Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge: MIT Press.
- Horn, D., Sagi, D., & Usher, M. (1992). Segmentation, binding and illusory conjunctions. *Neural Computation*, *3*, 509-524.
- Horn, D., & Usher, M. (1990). Dynamics of excitatory-inhibitory networks. *International Journal of Neural Systems*, *1*, 249-257.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review* *104* (3), 427-466.
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review* *110* (2), 220-264.
- Lachman, R., Lachman J. L., & Butterfield, E. C. *Cognitive psychology and information processing: An introduction*. New Jersey: L. Erlbaum.

- Lazareva, O. F., Smirnova, A. A., Bagozkaja, M. S., Zorina, Z. A., Rayevsky, V. V., & Wasserman, E.A. (2004). Transitive responding in hooded crows requires linearly ordered stimuli. *Journal of the Experimental Analysis of Behavior*, 82 (1), 1-19.
- Mandler, J. M. (1988). How to build a baby: On the development of an accessible representational system. *Cognitive Development*, 3, 113-136.
- Mandler, J. M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review*, 9, 587-604.
- Mandler, J.M. (2004). *The foundations of mind: Origins of conceptual thought*. Oxford, England: Oxford University Press.
- Marcus, G. F. (1998) Rethinking eliminative connectionism. *Cognitive Psychology*, 37, 243-282.
- Marr, D. (1982). *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.
- McGonigle, B. O., & Chalmers, M. (1977). Are monkeys logical? *Nature* (267) 694-696.
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100, 254-278.
- McKinney, M.L. (1991). Completeness of the fossil record: an overview. In Donovan, S.K. (Ed.), *The process of fossilization* (66-83). New York: Columbia University Press.
- Munakata, Y., & O'Reilly, R. C. (2003). Developmental and computational neuroscience approaches to cognition: The case of generalization. *Cognitive Studies*, 10, 76-92.
- O'Reilly, R.C., & Bubsy, R. S. (2002) Generalizable relational binding from coarse-coded distributed representations. In T.G. Dietterich, S. Becker, & Z Ghahramani (eds.), *Advances in neural information processing systems (NIPS)* (Vol. 14, pp. 75-82). Cambridge, MA: MIT Press.

- O'Reilly, R. C., Bubsy, R. S., Soto, R. (2003). Three forms of binding and their neural substrates: Alternative to temporal synchrony. In A Cleeremans (ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 168-192). Oxford, England: Oxford University Press.
- Pavlov, I. P., (1927). *Conditioned Reflexes: An Investigation of Physiological Activity of the Cerebral Cortex*. London: Oxford University Press.
- Penn, D. C., Holyoak, K. J. & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31, 109-178.
- Penn, D. C., & Povinelli, D. J. (2007). On the lack of evidence that non-human animals possess anything remotely resembling a 'theory of mind.' *Philosophical Transactions: Biological Sciences*, 362, 731-744.
- Renne, P.R., Ludwig, K.R., & Karner, D.B. (1998). Progress and challenges in geochronology. *Science Progress*, 83, 107-121.
- Romanes, G. J. (1882). *Animal Intelligence*. London: K. Paul, Trench.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of classical conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In Black, A. H. & Prokasy, W. F. (Eds.). *Classical conditioning II: Current research and theory*, 64-99. New York: Appleton-Centry-Crofts.
- Shettleworth, S.J. (2009) The evolution of comparative cognition: Is the snark still a boojum? *Behavioural Processes*, 80, 210-217.
- St. John, M. F., & McClelland, J. L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, 46, 217-257.

- Thomson, R. K. R., & Oden, D. L. (2000). Categorical perception and conceptual judgments by nonhuman primates: The paleological monkey and the analogical ape. *Cognitive Science*, 24 (3), 363-396.
- Thorndike, E.L. (1998). *Animal intelligence: an experimental study of associative processes in animals*. New York: Macmillan.
- Turing, A. (1938). Systems of logic based on ordinals. *Proceedings of the London Mathematical Society*, s2-45, 161-228.
- Wasserman, E. A., Young M. E., Cook, R. G. (2004) Variability discrimination in humans and animals: Implications for adaptive action. *American Psychologist*, 59 (9), 879-90.
- Visalberghi, E. & Limongelli, L. (1994). Lack of comprehension of cause-effect relations in tool-using capuchin monkeys (*Cebus apella*). *Journal of Comparative Psychology*, 108, 15-22.
- von der Malsburg, C., & Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 39, 124-148.
- von Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. R. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 33-341.