Individual Differences in Relational Learning and Analogical Reasoning:

A Computational Model of Longitudinal Change

Leonidas A. A. Doumas
Department of Psychology
University of Edinburgh

Robert G. Morrison
Department of Psychology
Neuroscience Institute, Stritch School of Medicine
Loyola University Chicago

Lindsey E. Richland
Department of Comparative Human Development
University of Chicago

Corresponding Author:

Robert G. Morrison
Department of Psychology
Loyola University Chicago
rmorrison@luc.edu
312-498-2612

Doumas, A. A., Morrison, R. G., Richland, L. E. (2018), Individual differences in relational learning and analogical reasoning: A computational model of longitudinal change, *Frontiers in Psychology: Cognitive Science*, 9:1235. doi: 10.3389/fpsyg.2018.01235

Abstract

Children's cognitive control and knowledge at school entry predict growth rates in analogical reasoning skill over time; however, the mechanisms by which these factors interact and impact learning are unclear. We propose that inhibitory control is critical for developing both the relational representations necessary to reason and the ability to use these representations in complex problem solving. We evaluate this hypothesis using computational simulations in a symbolic connectionist model of analogical thinking, DORA/LISA (Discovery Of Relations by Analogy; Doumas, Hummel, & Sandhofer, 2008). Longitudinal data from children who solved geometric analogy problems repeatedly over six months show three distinct learning trajectories though all gained somewhat: analogical reasoners throughout, non analogical reasoners throughout, and transitional - those who start nonanalogical and grew to be analogical. Varying the base level of lateral inhibition in DORA affected the ability to learn relational representations, which, in conjunction with lateral inhibition levels used in LISA during reasoning, simulated accuracy rates and error types seen in the three different learning trajectories. These simulations suggest inhibitory control may not only impact reasoning ability, but may also shape the ability to acquire relational knowledge given reasoning opportunities.

Keywords: analogical reasoning, relational knowledge, development, computational modeling, cognitive control

Analogical reasoning, the process of representing information as systems of relationships and mapping between these representations, is ubiquitous in learning and discovery throughout the lifespan, and is part of what makes humans uniquely intelligent and adaptive (Gentner, 2003; Penn, Holyoak & Povinelli, 2008). Analogical reasoning may play a crucial role in childhood, serving as a cognitive-bootstrapping mechanism that enables children to make increasingly abstract inferences and generalizations (e.g., Gentner, 2003), and supporting learning across a wide range of educational domains (Richland & Simms, 2015). The mechanisms by which children's analogical reasoning improve, however, are not well understood. In particular, little attention has been paid to the processes by which children develop the relational representations used for analogical reasoning.

Children's cognitive-control resources have been implicated as one source of individual differences in relational representation and reasoning. Also described as executive function (EF) (Diamond, 2013), these resources refer to the ability to use selective attention to manipulate the contents of working memory, and are believed to include a variety of functions including inhibitory control (IC), updating, and shifting (Bainich, 2009). Cross-sectional studies have revealed that children who can solve analogies successfully make mistakes when the requirements for cognitive control are raised, either by increasing the requirements for controlling attention in the face of distraction, or increasing the complexity of the relations (Richland, Morrison & Holyoak, 2006; Thibaut, French & Vezneva, 2010). The difficulty of controlling attention to relations in the face of distraction has been identified across children from different cultural and linguistic backgrounds (Richland, Chan, Morrison & Au, 2010). Computational work has provided support for the interpretation that these errors are due to low levels of resources for inhibitory control (Morrison, Doumas & Richland, 2011).

However, a full theory of relational reasoning development must go beyond performance accuracy to provide a mechanism for developmental change over time. There is reason to believe that cognitive-control resources not only predict performance at a single time point, but also may impact children's growth in reasoning skill. An analysis of data from a large-scale longitudinal study found that children's performance at school entry on an inhibitory control task (Children's Stroop; Gerstadt, Hong & Diamond, 1994), and an executive function task (Tower of Hanoi) both predicted distinct variance in children's analogical skill, and more interestingly, their growth in analogical skill from school entry to adolescence (Richland & Burchinal, 2013). This relationship held even when controlling for environmental factors (e.g., parental education, SES, gender), as well as short-term memory, sustained attention, knowledge measures, and analogy skill at third grade. This pattern of change suggests that early executive function skills play an important role in shaping children's trajectory of learning reasoning skills.

EF as a mechanism underpinning relational reasoning growth

The current paper reports computational simulations that test a mechanism by which early inhibitory control resources could alter the trajectory by which children's reasoning develops through the course of children's reasoning opportunities. We simulated data from one of the few longitudinal studies on the development of analogical reasoning (Hosenfeld, van der Maas, & van den Boom 1997). Our aim was to explore how relational knowledge and variations in children's inhibitory control could predict children's rate of reasoning development over a series of repeated opportunities to solve geometric analogies. We focus in particular on the interplay between the learning of relational representations and individual differences in inhibitory control.

Behavioral data on reasoning change over time. In the original study (Hosenfeld, van der Maas, & van den Boom 1997), children aged 6-7 years solved 20 geometric analogy problems consisting of five common relations between simple shapes including: adding an element, changing size, halving, doubling, and changing position repeatedly over 8 testing sessions. The measure was originally designed by randomly combining six basic geometric shapes and five transformations to create 12,150 problems. The authors used the difficulty metric (Difficulty = 0.5 x Elements + 1 x Transformations) to select problems for a large norming project (Hosenfeld, van der Boom, & Resling, 1997). Twenty of these problems were then selected for use in the longitudinal study to represent a range of difficulty both with respect to the difficulty metric and actual child performance.

During testing, children solved A: B :: C: D problems in which they had to infer the missing D term in order to construct a valid analogy. Figure 1 provides three examples of these geometric analogy items in increasing difficulty, showing duplication (top line), halving/duplication and "inside" (middle line), and an above/ below/inside set of transformations (bottom line).

On each testing occasion, the children were first given practice time. This included naming and drawing the basic geometric shapes that would be part of the relational problems. They were then told they would be solving puzzles, and completed three practice analogies with the experimenter. The following instruction was provided: "These two boxes belong together (point to A and B), and those two boxes belong together (point to C and D). These two ones (A and B) belong together in the same way as those ones (C and D) do. Do you know what the solution is?" (van der Mas & Boom, 1997, p. 375).

Twenty test items were then presented, in which children were instructed to draw the missing piece for each problem, and were provided with feedback following errors. The problems within each session varied in complexity based on changes in the number of relationships needed to characterize the A:B transition.

Insert Figure 1 about here

Children were tested eight times over the course of one year, at three-week intervals. Researchers recorded accuracy rates, time to solution, and types of errors made. These data were used to examine the trajectory of children’s analogical reasoning over the course of the study. Children's performance could be separated into three learning profiles: 1) *Non-Analogical Reasoners*, who solved the majority of problems non-analogically throughout all sessions, 2) *Transitional Reasoners*, who moved from solving problems largely non-analogically to solving problems largely analogically, and 3) *Analogical Reasoners*, who solved the majority of problems analogically throughout the treatment. The reasoning accuracy results for the three groups of children over time are shown in Figure 7.

The data from Hosenfeld et al.’s (1997) study are informative, and they raise a challenge of interpretation. One cannot fully explain these three trajectories by access to learning opportunities, since all children had access to the same number of learning opportunities. Further, children’s initial skill-based starting point is not fully predictive either, since one group started low and ended high, and another group started low and ended low. Cognitive maturation of growth in executive function (EF) capacity is a similarly an unsatisfactory explanation. While some EF growth over the period of six months might be expected, there is no reason to expect three different yet systematic patterns of EF growth that would explain these three performance trajectories.

Current simulation study aims. In the present this study we use computational simulations of these data to argue that 1) differences in inhibitory control EF resources may explain initial differences in reasoning, but 2) they also help to explain differences between the three groups in their ability to learn relational representations necessary for reasoning over repeated learning opportunities. Thus while all children received the same number of learning opportunities during the six training sessions, the level of structure they identify in the problem inputs may increase or decrease their likelihood of successfully reasoning analogically with these representations over time. Furthermore, the rate at which they learn is constrained by their inhibitory control EF resources. The interaction of processing ability and learning representations produces a more complete picture of the development of analogical reasoning then either factor independently.

Computational Models of Analogical Reasoning

Computational models of analogical reasoning provide a unique window into the plausible cognitive underpinnings of relational reasoning, and here enable us to test correlations between the behavioral data and performance in a constrained system (see French, 2002). We use DORA (Discovery Of Relations by Analogy; Doumas, Hummel, & Sandhofer, 2008) as a model of how structured relational representations are learned from unstructured inputs, and LISA (Learning and Inference with Schemas and Analogy; Hummel & Holyoak, 1997, 2003) as a model of human relational reasoning, to simulate Hosenfeld et al.'s (1997) results, and to explore the interactions between maturation-based inhibition levels and learning opportunity cycles.

Previously, we have successfully used changes in inhibitory control in LISA's working-memory system to explain cross-sectional variations in analogical reasoning. We simulated the developmental progression (from age 3 to 14) in children's ability to handle increases in relational complexity and distraction from object similarity during analogical reasoning by varying inhibitory control (Morrison et al., 2011). In addition, we have modeled cross-cultural differences in analogy performance (Richland, Chang, Morrison & Au, 2010), considered to be the result of differences in relational knowledge accretion, via changes to the hand-coded representations used in LISA (Morrison et al., 2011).

In the current study we avoid hand coding of propositional structures. Instead, we use DORA to simulate children's ability to learn spatial relations over time, allowing relational learning patterns to be part of the investigation. We then use those representations in LISA to simulate geometric analogy accuracy and types of errors. By doing so, we are able to model the trajectory of knowledge accretion as well as reasoning ability. Importantly, we manipulate inhibitory control (via changes in lateral inhibition in both models) to simulate individual differences. We argue that inhibitory control is fundamental not only to the ability to reason relationally, but also to the ability to learn relations in the first place.

Method

Overview of LISA/DORA Model

In this section we describe the LISA (Hummel & Holyoak, 1997, 2003) and DORA (Doumas et al., 2008) models in broad strokes. Our goal is to highlight the main processing features of the models and their core theoretical claims. Knowlton, Morrison, Hummel and Holyoak (2012) provide another useful, brief description of the LISA architecture. The most complete descriptions of the models may be found in their original reports (Doumas et al., 2008;

Hummel & Holyoak, 1997, 2003).

Both LISA and DORA (DORA is a direct descendent and generalization of LISA) are connectionist models. However, unlike traditional connectionist networks (e.g., McClelland, 2010), LISA and DORA solve the binding problem, and so can process structured (i.e., symbolic) representations (see Doumas & Hummel, 2005, 2012). LISA uses structured representations of relations (represented as predicates) and their arguments to make analogies, induce schemas, and perform relational generalization. DORA provides an account of how the structured predicate representations used by LISA can be learned from unstructured representations of objects (i.e., flat feature vector representations of objects without predicates; see below).

Insert Figure 2 about here

We begin by describing the relational knowledge structures that DORA learns and that LISA uses. We then describe how DORA learns these knowledge structures from experience. Finally, we describe LISA's mapping and generalization procedures. Both DORA's learning and LISA's mapping and generalization procedures play central roles in the simulations we report in this paper.

Knowledge Structures and LISAese

DORA begins with objects represented as flat feature vectors (see Figure 3a). That is, objects are represented as in conventional distributed connectionist systems as patterns of activation in a set of units. These initial representations are holistic and unstructured (see Doumas & Hummel, 2005). Before describing how DORA learns structured predicate representations of object properties and relations from these initial representations; however, we will describe the end state of that learning. Specifically, we now describe knowledge representation in LISA (and by extension in DORA *after* it has learned).

In LISAese, relational structures are represented by a hierarchy of distributed and localist codes (see Figure 3b). At the bottom, “semantic” units represent the features of objects and roles in a distributed fashion. Semantic units don’t actually have any necessary meaning. They are simply properties of the perceptual stimulus that are detectable by the system (e.g., location on the *y*-axis, being cone shaped). As we discuss below, for the purposes of LISA and DORA, exactly what these semantic units code is not important. All that is that is necessary is that there are aspects of perceptual stimuli that are consistently detectable by the system (e.g., that when encoding two red objects, the same feature—or set of features—responds to their hue). At the next level, these distributed representations are connected to localist units termed POs (for predicate-object) that represent individual predicates (or roles) and objects. One layer up, localist role-binding units (RBs; alternatively called “subpropositions”) link object and relational role units into specific role-filler pairs. At the top of the hierarchy, localist P units link RBs into whole relational propositions.

Considering the house object containing the square in Problem 3, Term A (see Figure 1), LISA would represent the proposition *contains* (house, square) using PO units (triangles and large circles in Figure 3) to represent the relational roles *outside* and *inside*, and the role fillers house and square. Each of these PO units is connected to semantic units coding their semantic features. RB units (rectangles) then conjunctively code the connection between roles and their fillers (one RB connects house to *outside*, and one connects square to *inside*). At the top of the hierarchy, P units (oval) link sets of RBs into whole relational propositions. A P unit conjunctively codes the connection between the RBs representing *outside* (house) and the RB representing *inside* (square), thus encoding the relational proposition *contains* (house, square).

Note that all of these units are simply connectionist nodes in a layered network. While we use different names for units at different layers, and use different shapes to specify different units in our figures, we do so only for the purposes of more efficient exposition. There is nothing inherently different about PO units or RB units other than they are in different layers of a neural network (much as different units might be in the input layer or a hidden layer of a feed-forward neural network). However, just as units in a hidden layer serve a different function in relation to a network's behavior relative to units in the input layer, so units in the RB layer serve a different function than units in the semantic layer.

When a proposition enters working memory, role-filler bindings must be represented dynamically on the units that maintain role-filler independence (i.e., POs and semantic units; see Hummel & Holyoak, 1997). In DORA (and its instantiation of LISA), roles are dynamically bound to their fillers by systematic asynchrony of firing. As a proposition in the driver becomes active, bound objects and roles fire in direct sequence. Binding information is carried in the proximity of firing (e.g., with roles firing directly before their fillers). Using the example in Figure 3, in order to bind *outside* to house and *inside* to square (and so represent *contains* (house, square)), the units corresponding to *outside* fire directly followed by the units corresponding to house, followed by the units for coding *inside* followed by the units for square.[1]

Insert Figure 3 about here

---

1 Asynchrony-based binding allows role and filler to be coded by the same pool of semantic units, which allows DORA to learn representations of relations from representations of objects (Doumas et al., 2008).

Learning Structured Representations in DORA

DORA is an account of how structured LISAese representations can be learned from unstructured examples. As noted above, DORA begins with representations of objects coded by simple flat feature vectors (Figure 3a). For example, a house would be coded with a set of features describing that house, or a square would be coded with a set of features describing that square. We instantiate these representations as object token units attached to the semantic features of that object (see Figure 4a). These initial representations are holistic and unstructured (in that an object's features are active together as a mass; see, e.g., Doumas & Hummel, 2005, 2012). DORA's learning algorithm allows it to learn structured representations of specific subsets of an object's features. Vitally, these representations function like predicates in that they are explicit and can take (i.e., be dynamically bound to) arguments.

DORA uses comparison to bootstrap its learning. When DORA compares two objects, then those objects become co-active (Figure 4a). As the compared objects pass activation to their semantic features, those properties shared by both objects receive twice as much input and become roughly twice as active as unshared features (Figure 4b). DORA recruits a PO unit that learns connections to the active semantics via simple Hebbian learning. Accordingly, the new PO learns stronger connections to the more active (shared) semantics, and weaker connections to the less active (unshared) semantics (Figure 4c). DORA also recruits an RB unit at the layer above the POs, which learns connections to the active POs via Hebbian learning (Figure 4d).

Insert Figure 4 about here

The result of this learning algorithm is that DORA acquires explicit representations of the shared properties of compared objects. For example, when DORA compares two red things it will learn an explicit representation of the property *red*, and if DORA compares two objects that are containers, it will learn an explicit representation of the property *container*.[2] Importantly, these new representations function like single-place predicates: they can be bound to arguments (via asynchronous binding; see above), they specify properties of the arguments to which they are bound (see Doumas & Hummel, 2012; Doumas et al., 2008), and they support symbolic operations such as structure mapping (see Hummel & Holyoak, 1997, Doumas et al., 2008) and relational generalization (Hummel & Holyoak, 2003).

Comparison underlies DORA's ability to learn functional single-place predicate representations, and comparison also allows DORA to learn representations of whole relational structures (see Figure 5). If multiple role-filler sets enter DORA's WM together, the model can map each set onto the other. For example, if DORA compares the circle containing the triangle in Figure 1 (Problem 3, term C) to the house containing the square (Problem 3, term A), it could map *outside* (circle) to *outside* (house) and *inside* (triangle) to *inside* (square). This process leads to a distinct pattern of firing over the units composing each set of propositions (i.e., the RB units of *outside* (circle) fire out of synchrony with those of *inside* (triangle) while the RB units of *outside* (house) fire out of synchrony with those of *inside* (square)). This pattern of oscillating activation over sets of units (with co-occurring role-filler pairs firing in sequence) acts as a signal to DORA to recruit a P unit, which learns connections to active RBs via Hebbian learning. The result is that the new P unit links co-occurring role-filler sets, and results in a rudimentary

[3] As noted above, the specific content of the units coding for a property are unimportant to DORA. So long as there is something common across the units representing a set of objects, DORA can learn an explicit representation of this commonality. For the purposes of DORA's learning algorithm, all that matters is there is something invariant across instances of a *container* (which there must be for us to learn the concept), and that the perceptual system is capable of responding to this invariance (which, again, there must be for us to respond similarly across instances of containment in the world).

representation of relations (here *contains* (object1, object2)). Importantly, this kind of relational representation, in which a relation is composed of linked sets of its roles, is a full fledged multi-place relational structure capable of the same sorts of operations and inferences as traditional multi-place relations (e.g., predicate calculus; Doumas & Hummel, 2005, 2012; Doumas et al., 2008).

Mapping and Relational generalization in LISA

In LISA/DORA, representations are divided into two mutually exclusive banks of units: a *driver* and one or more *recipients*.[3] The driver is the current focus of attention (i.e., what LISA/DORA is thinking about at the present moment), and the recipient is analogous to active memory in Cowan's (2001) terms (i.e., items primed from long-term storage, which can be potentially compared to items in the driver). The driver and recipient communicate via the semantic feature units, which are shared by both sets. Specifically, items in the driver become active and pass activation to the semantic feature units, which activate units in the recipient. Units in the recipient then compete via lateral inhibition to respond to the pattern of firing imposed on semantic units by units in the driver. Structured representations created during relational learning in DORA can be mapped using LISA's mapping algorithm (Hummel & Holyoak, 1997) with minor modifications described in Doumas et al. (2008). LISA/DORA learns which elements in the driver and recipient correspond by building mapping connections (via Hebbian learning) that keep track of when these elements are active simultaneously.

When augmented with the capacity for self-supervised learning (Doumas et al., 2008; Hummel & Holyoak, 2003; described below), LISA's mapping algorithm allows for analogical inference. To illustrate, consider how LISA/DORA solves an inference problem such as the third

[4]Mutually exclusive sets are necessary in order to perform comparison (see, e.g., Falkenhainer et al., 1989; Holyoak & Thagard, 1989). Knowlton et al. (2012) describe how such sets can be instantiated in prefrontal cortex in a neutrally-plausible fashion.

problem in Figure 1.[4] The A and B terms are in the driver and the C term is in the recipient. As the proposition coding for A term, *contains* (house, square), becomes active in the driver, it activates and consequently maps to the units coding for *contains* (circle, triangle) in the recipient. Specifically, the units coding for *outside* (house) in the driver activate and map to the units coding for *outside* (circle) in the recipient, and the units coding for *inside* (square) in the driver activate and map to the units coding for *inside* (triangle) in the recipient.

Then when the B term, *contains* (square, shield) becomes active in the driver, there are no corresponding units for it to map to in the recipient. As the representation of the C term in the recipient is already mapped to the representations of the A term in the driver (and the C term is the only item in the recipient), the representation of the B term is left with nothing to which it corresponds. This situation, in which items in the driver have no elements in the recipient that they can activate (because all recipient elements are already mapped to other driver elements), triggers the self-supervised learning algorithm in LISA/DORA. During self-supervised learning, active units in the driver prompt LISA/DORA to recruit matching units in the recipient (i.e., an active RB unit in the driver prompts recruitment of an RB unit in the recipient). Continuing the example, as units coding for *outside* (square) in the B term become active in the driver, LISA/DORA recruits RB and P units in the recipient to match the active RB and P units in the driver. The new recruited P unit in the recipient learns connections to active recipient RB units, and newly recruited RB units learn connections to active PO units via Hebbian learning. The functional result of this unit-based recruitment and Hebbian learning is that LISA/DORA infers a representation of *outside* (triangle) in the recipient, which corresponds to the representation of *outside* (square) in the driver. An analogous sequence occurs when *inside* (house) fires in the

[5]The problem is more relationally complex than the simple version we describe here; however, the same principles apply to the way LISA can solve the entire problem, including all of the nested relations.

driver and LISA/DORA infers *inside* (circle) in the recipient. Thus, LISA/DORA completes the D term in a problem via analogical inference, inferring a representation of *contains* (triangle, circle) in the recipient.

The Role of Inhibition in DORA/LISA

Of particular importance to the present simulations, inhibition plays a role in the selection of items to enter working memory because selection is a competitive process. Propositions in the driver compete to enter into working memory on the basis of several factors, including their pragmatic centrality or importance, support from other propositions that have recently fired, and the recency with which they themselves have fired. Reduced driver inhibition results in reduced competition and more random selection of RBs to fire. The selection of which RBs are chosen to fire, and in what order, can have substantial effects on DORA/LISA's ability to find a structurally consistent mapping between analogs. It follows that reduced driver inhibition, resulting in more random selection of propositions into working memory, can affect DORA/LISA's ability to discover a structurally-consistent mapping.

The role of inhibition in the activity of a recipient analog is directly analogous to its role in the activity in the driver. Recipient inhibition causes units in the recipient to compete to respond to the semantic patterns generated by activity in the driver. If DORA/LISA's capacity to inhibit units in the recipient is compromised, then the result is a loss of competition, with many units in the recipient responding to any given pattern generated by the driver. The resulting chaos hampers (in the limit, completely destroys) DORA/LISA's ability to discover which units in the recipient map to which in the driver. In short, inhibition determines DORA/LISA's working memory capacity (see Hummel & Holyoak, 2003, Appendix A; Hummel & Holyoak, 2005), controls the model's ability to select items for placement into working memory, and also regulates its ability to control the spreading of activation in the various recipient analogs. As such, inhibition is critical for the model's ability to favor relational similarity over featural similarity.

This conception is highly complementary to behavioral models suggesting inhibitory control in EF contributes to reasoning performance by enabling reasoners to inhibit rules used previously in favor of current goal requirements (e.g., Zelazo and colleagues (1998, 2003). Thus, we hypothesized that differences between the three groups of children in Hosenfeld et al.'s (1997) study were at least partially a product of differences in inhibitory control. We simulated these differences in DORA/LISA by varying levels of lateral inhibition. In DORA/LISA, inhibition is critical to the selection of information for processing in working memory. Specifically, inhibition determines the intrinsic limit on DORA/LISA's working-memory capacity (see Hummel & Holyoak, 2003, Appendix A), controls its ability to select items for placement into working memory, and also regulates its ability to control the spreading of activation in the recipient. We have previously used this approach in LISA to simulate patterns of analogy performance in a variety of populations with lesser working-memory capacity including older adults (Viskontas et al., 2004), patients with damage to prefrontal cortex (Morrison et al., 2004), and young children (Morrison, Doumas, & Richland, 2011).

Simulations

We simulated Hosenfeld et al.'s (1997) results in two steps (see Figure 2). In the first step we used DORA's relation-learning algorithm to learn representations of the transformations used in the geometric analogy problems. In our simulations, DORA began with representations of 100 objects attached to random sets of features (chosen from a pool of 100). We then defined five transformations (the same as those used by Hosenfeld et al., 1997): adding an element, changing size, halving, doubling, and changing containment). Each single-place predicate transformation (adding an element, changing size, halving, doubling) consisted of two semantic features, and the relational transformation (changing containment) consisted of two roles each with two semantic features (i.e., for the *contains* relation, both the roles inside and outside were each defined by two specific semantic units). Again, as noted above, these semantic units had no actual content. Rather, they represented our assumption that there are invariant properties of objects and transformations that are detectable by the perceptual system. Our goal in this first simulation was simply to demonstrate that DORA could isolate and learn explicit representations of invariant properties during completely unstructured training.

Each of the 100 objects was attached to the features of between two and four transformations chosen at random. If an object was part of a relational transformation, it was attached to the features of one of the roles, chosen at random. For example, object1 might be attached to the features for *doubled* (a single-place transformation) and *inside* (one role of the relational transformation, *contains*).

We presented DORA with sets of objects selected at random, and allowed it to compare the objects and learn from the results (applying DORA's relation-learning algorithm). As DORA learned new representations it would use these representations to make subsequent comparisons. For example, if DORA learned an explicit representation of the property *double* by comparing two objects both attached to the features of *double*, it could use this new representation for future comparisons. On each trial we selected between two and six representations and let DORA compare them and learn from the results (i.e., perform predication, and relation-learning routines). We assume that this act of inspection and comparison is similar to what happens when children encounter the geometric analogy problems and have to consider how the various elements are related (Gentner & Smith, 2013). Importantly, this training was completely unstructured and undirected (i.e., DORA randomly selected items from memory to reason about). We have demonstrated in previous work that DORA can learn under these unstructured conditions, and that learning improves markedly with more directed or more structured training (see Doumas et al., 2008). We have no doubt that the children in Hosenfeld et al., learned from their experience with the various versions of the geometric analogy task, and that taking the test over successive sessions served to structure their training somewhat. For the current simulations, however, we wanted to make as few assumptions about the learning environments of the children in the study as possible (given this information is, very understandably, absent from Hosenfeld et al.). As such, we chose to handicap ourselves and avoid making additional assumptions that would improve our overall ability to fit the data.

Moreover, we defined three groups for the purposes of the simulation as determined by a range of lateral inhibition values. We ran 100 simulations for each group. During each simulation we chose an inhibition level from a normal distribution, with a mean of 0.4 for the low inhibition group, 0.6 for the medium inhibition group, and 0.8 for the high inhibition group (each distribution had a SD = .1). We chose to simulate groups using a distribution of inhibition scores in order to match our assumption that the learning groups from the original Hosenfeld et al. study were not completely homogeneous in their inhibitory abilities. Our decision once again served to handicap the precision of our simulations by adding some noise, but there was almost certainly some natural variability in the inhibitory abilities of the children in the initial study, and we wanted our simulations to reflect this variability.

For the low-knowledge condition, simulations were run with 800 learning trials, and we checked the quality of the representations DORA had learned after each 100 trials. Quality was calculated as the mean of connection weights to relevant features (i.e., those defining a specific transformation or role of a transformation) divided by the mean of all other connection weights + 1. (1 was added in the denominator to keep the quality metric bound between 0 and 1.) For the high-inhibition, high-knowledge condition we extended the simulations to 1000 learning trials and sampled the representations after 300 to 1000 trials. The reason for the different knowledge conditions was to test our hypothesis that children in the high performance group not only had higher inhibitory resources, but also came into the study with a higher quality of relational representations. In brief, our goal was to test whether starting at a higher knowledge state in tandem with increased inhibitory resources would provide a closer fit to the high performance group's data than increased inhibitory resources in isolation.

Figure 6 provides a summary of results from part one of the simulation. While all groups did learn, learning was obviously improved with higher levels of inhibition. In addition, learning was much faster for the higher inhibition group.

In the second part of the simulation we passed the representations DORA learned during the first part of the simulation to LISA, which then simulated solving the geometric analogy problems. Thus, unlike LISA simulations we have performed previously to account for developmental changes (e.g., Morrison et al., 2011), relational knowledge representations were not hand-coded, but rather were generated automatically by DORA. We created problems of varying difficulty to capture the range of difficulty used in the Hosenfeld et al. (1997). Thirty percent of problems were hard problems consisting of three transformations (one third had two binary transformations and one unary transformation; one third had one binary transformation and two unary transformations; and one third had three unary transformations). Thirty-five percent were medium difficulty problems consisting of two transformations (one third had two binary transformations; one third had two unary transformations; and one third had one binary and one unary transformation). Lastly, thirty-five percent were easier problems with only one transformation per problem (half had one binary transformation while the other half had one unary transformation).

We simulated all eight of the testing phases in the Hosenfeld et al. (1997) study. Each testing phase consisted of 20 trials. On each trial we presented LISA/DORA with the A and B terms in the driver and the C term in the recipient. The A, B and C terms were represented as object POs each attached to four random features, and bound to PO predicate units identifying the transformations in which they were involved. Importantly, the PO units identifying the transformations (as well as the RB units linking predicate and object POs, and P units linking RBs) were representations that DORA had learned during the first part of the simulation. For example, if the A term was a shield inside a square, we represented that with the LISEese proposition *contains* (square, shield), with a PO representing *square* bound to a PO representing *outside* (where *outside* was a PO that DORA had learned during the first part of the simulation), and a PO representing *shield* bound to a PO representing *inside* (where inside was a PO that DORA had learned during the first part of the simulation). For the first testing phase for the low-knowledge groups we used the representations DORA had learned after the first 100 learning trials, for the second testing phase we used the representations DORA had learned after the first 200 learning trials, and so on. For the high-knowledge group, we used the representations DORA had learned after the first 300 learning trials for the first testing phase, the representations learned after the first 400 learning trials for the second testing phase, and so forth. In each case we treated the level of lateral inhibition as maturational, and thus used the same levels as used for the learning phase for each group (0.4 low, 0.6 medium, 0.80 high; each with ±0.1 SD distribution).

As noted directly above, using slightly more advanced representations (the high-knowledge group) reflects the assumption that children with higher maturational inhibitory control are likely to have learned more about relations prior to beginning the study compared to children with lower maturational inhibitory control. Note that by starting testing with representations at 100 we assume that all children have some capacity for representing relations. This assumption is reflected in Hosenfeld et al.'s data, in that low- and medium-analogy group children started with similarly low scores in the first testing phase, whereas children in the high-analogy group started with much higher performance on the first testing phase.

During test trials, LISA attempted to map driver and recipient propositions and make inferences about the missing D term. For example, if LISA mapped the A term in the driver to the C term, then when the B term fired LISA inferred the D term in the recipient. We took the inferred proposition in the recipient to be LISA's answer on that trial.

As is apparent from the learning trajectories plotted in Figure 7, DORA/LISA's performance on the testing trials closely followed those of the children in Hosenfeld et al.'s (1997) study. Just like the non-analogical children, DORA/LISA with a low lateral inhibition level performed poorly throughout. Like the transitional children, DORA/LISA with a medium lateral inhibition level started slow, but slowly improved. Finally, like the analogical children DORA/LISA with high lateral inhibition levels performed well virtually from the start and maintained good performance; however, additional relational knowledge coupled with high lateral inhibition levels appears to best fit the analogical performance group.

Importantly, the types of errors that DORA/LISA makes closely follow the types of errors made by each of the performance groups (see Table 1). Specifically, like the non-analogical children, low-inhibition DORA tended to make errors based on featural association errors (e.g., objects in A, B and C copied). Like transitional children, with medium inhibition, DORA tended to make featural/ associative errors at the beginning, but these largely disappeared by the final session. Finally, like the analogical children, with high inhibition, DORA tended to make fewer errors overall, which further decreased over time, but these errors that did happen were a mix of associative and incomplete solutions.

Moreover, the kinds of problems that DORA "got wrong" at various inhibition levels seems highly in line with the kinds of problems that children seem to make errors on as they develop. While Hosenfeld et al. (1998) do not give specific data on which problems the children tended to get wrong, a good deal of previous research has been done on children's analogical development using cross sectional designs (see above; e.g., Richland, Morrison & Holyoak, 2006. Generally, children tend to develop a capacity for solving simpler analogy problems first, and solve such problems consistently before they develop the capacity for solving harder analogy problems. For example, Richland et al. (2006) found that young children around the age of three perform consistently above chance on simple analogy problems that require aligning pairs of elements across two pictures (e.g., a task that requires matching the cat in a picture of a dog chasing a cat, to the boy, in a picture of a mother chasing a boy). However, these same children perform very poorly when the task is made harder, either by adding distractor elements to one or both of the pictures, or by requiring integration across multiple relations (see above). Around age seven, children consistently solve problems either requiring relational integration or involving distractors, but perform less well on problems involving a distractor and requiring relational integration. Finally, by age 14, participants could consistently solve all the types of analogy problems tested. Similarly, DORA, across all inhibition levels, performed well on some classes of problems and less well on other classes. Specifically, low-inhibition DORA did consistently quite well on the easiest problem types, but quite poorly on the medium and hard problems. Medium inhibition DORA performed well on easy problems, made some headway on the medium problems, and did very poorly on the hard problems. High inhibition DORA, was competent across all problem difficulties, but failed most consistently on the hard problems.

Table 1 about here

Discussion

These simulations provide a mechanism by which resources for inhibitory control can account for children's analogical reasoning development (Richland & Burchinal, 2014). Like children, the model moved from preferentially attending to featural information to reasoning with relational representations after being given learning opportunities (Gentner & Rattermann, 1991). The model suggests that a child's level of inhibitory control may play an essential role in determining their learning trajectory by modulating the noise through which children identify and train their relational representations. We have previously argued that inhibitory control is an essential factor in understanding the development of analogical reasoning in children because changes in inhibitory control can explain both featural distraction and relational complexity effects during childhood (Morrison et al., 2011). However, a complete understanding of the development of analogy must also include the role of relational learning and the growth of relational knowledge over time, as evidenced by the simulations reported here.

Why does such a simple change in a single parameter have such a complex effect and thus explain so much? The effect we observe results from two factors. The first is that the greater the inhibitory control of an individual, the more they can avoid distraction from imprecise relational representations. This means that an adult can form a valid analogy between two domains about which they know relatively little, whereas a three-year-old child might know quite a bit about two domains yet still fail to inhibit a featural distractor when attempting to make an analogical mapping. Second, because relational learning also requires inhibitory control in our model (an assumption supported by recent longitudinal studies showing that executive functions can predict future analogy performance; Richland & Burchinal, 2013), children with lower inhibitory control will learn relations less efficiently. The combination of these two factors results in our complex pattern of simulations. Children low in inhibitory control have difficulty building relationally precise representations, and also are less tolerant of these “dirty” representations during reasoning. Children with high inhibitory control build relationally precise representations quickly and are also more tolerant of “dirty” representations. Our middle inhibitory control group operates at the perfect “teachable moment,” something akin to Vygotsky's zone of proximal development (Vygotsky, 1978): they possess just the right amount of inhibitory control to efficiently build relational representations, which become sufficient during the training sessions so as to yield successful analogical reasoning.

One very important limitation of our current simulations stems from the kinds of problems used in the original Hosenfeld et al. (1998) study, and in our simulations. Specifically, as has been argued previously (e.g., Thibaut & French, 2016), in the A:B::C:D tasks, the subject's goal is to find the item (D) that matches (B) in the same way that (C) matches (A). The fact that (C) corresponds to (A) is given by the structure of the problem. In many tasks people draw analogies between situations without knowing these sorts of correspondences beforehand (e.g., Markman & Gentner, 1993; Richland et al., 2006). However (as argued above), we do believe that the longitudinal nature of the data collected by Hosenfeld and colleagues has a number of merits and that any full account of the development of relational reasoning needs to account for the findings that this study reveals. Particularly, (a) Hosenfeld et al. (1998) is, to our knowledge, the only longitudinal study of analogical development with multiple repeated data collection points; (b) the difficulty of many of the problems used in the study makes solving them difficult even for adults (c) the (D) term was not given to the children to select (as is often the case with A:B::C:D analogy problems), but rather had to be generated in full. Finally, it is important to note that in our simulations we did not give DORA the (A) to (C) correspondence a priori. DORA had to discover the (A) to (C) correspondence via it's mapping algorithm, and failure to do so made the generation of the (D) term all but impossible. As such, while the our simulations are limited by the exclusive use of A:B::C:D type analogy problems in the original study, we do find the original study a very important piece in our current understanding of analogical development, and, therefore, hold that simulating the study is an important milestone for any account of the development of analogical thinking.

To conclude, while considerable effort has been directed at understanding how inhibitory control supports analogical reasoning, less attention has been given to the role of inhibitory control in its essential antecedent—relational learning. Understanding this factor constitutes an important step toward understanding how relational learning develops and how it can contribute to successful analogical reasoning in children.

Author Note

A preliminary report of these simulations was presented at the 31st Annual Conference of the

Cognitive Science in Amsterdam, Netherlands. Correspondence may be addressed to Robert G. Morrison (rmorrison@luc.edu). The research reported here was supported by grants from the National Science Foundation, SMA-1548292, and the Institute of Education Sciences, U.S. Department of Education, through R305A170488 to Lindsey Richland (PI). The opinions expressed are those of the authors and do not necessarily represent views of the Institute or the U.S. Department of Education.

References

Banich, M. T. (2009). Executive function: The search for an integrated account. *Current Directions in Psychological Science*, 18(2), 89-94.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24*, 87–185.

Diamond, A. (2013). Executive Functions. *Annual Review of Psychology*, *64*, 135–168. doi:10.1146/annurev-psych-113011-143750

Doumas, L. A. A., & Hummel, J. E. (2005). Approaches to modeling human mental representations: What works, what doesn′t and why. In K. J. Holyoak and R. Morrison (Eds.). *The Cambridge handbook of thinking and reasoning*, 73-91. Cambridge, England: Cambridge University Press.

Doumas, L. A. A., & Hummel, J. E. (2012). Computational models of higher cognition. In K.J. Holyoak & R.G. Morrison (Eds.), *The Oxford handbook of thinking and reasoning*, 52-66. New York: Oxford 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.

Falkenhainer, B., Forbus, K. Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *20*(41), 1–63.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19,* 141-205.

French, R. (2002), The computational modeling of analogy-making *TRENDS in Cognitive Sciences*, *6*(5), 200-205.

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. 195–235). Cambridge: MIT Press.

Gentner, D., & Smith, L. A. (2013). Analogical learning and reasoning. In D. Reisberg (Ed.), The Oxford handbook of Cognitive Psychology (pp. 668-681). New York, NY: Oxford University Press.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.*), Perspectives on thought and language: Interrelations in development* (pp. 225-277). London: Cambridge University Press.

Goswami, U. (2001). Analogical reasoning in children. In D. Gentner, K. J. Holyoak & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 437-470). Cambridge, MA: MIT Press.

Halford, G. S., Wilson, W. H., Guo, J., Gayler, R. W., Wiles, J., & Stewart, J. E. M. (1994). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak & J.Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2: Analogical connections* (pp. 363-415). Norwood, NJ: Ablex.

Halford, G. (2005). Development of Thinking. (pp. 529-558). In K.J. Holyoak, R. G. Morrison (Eds). *Cambridge Handbook of Thinking and Reasoning,* New York: Cambridge University Press.

Holyoak, K.J., & Thagard, P. (1989). Analogical Mapping by Constraint Satisfaction. *Cognitive Science, 13,* 295-355.

Hosenfeld, B., can der Boom, D.C., Resling, W.C.M. (1997). Constructing Geometric Analogies for the Longitudinal Testing of Elementary School Children. *Journal of Educational Measurement, 34*(4), 367-372.

Hosenfeld, B., van der Maas, H.L.J., & van den Boom, D. (1997). Indicators of discontinuous change in the development of analogical reasoning. *Journal of Experimental Child Psychology, 64,* 367-395.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427-466.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review, 110,* 220-264.

Knowlton, B.J., Morrison, R.G., Hummel, J.E., & Holyoak, K.J. (2012). A neurocomputational system for relational reasoning. *Trends in Cognitive Science, 16*(7), 373-381. doi.: 10.1016/j.tics.2012.06.002

McClelland, J. L. (2010). Emergence in cognitive science. *Topics in Cognitive Science, 2,* 751-770.

Morrison, R.G., Doumas, L.A.A., & Richland, L.E. (2011). A computational account of children's analogical reasoning: Balancing inhibitory control in working memory and relational representation. *Developmental Science, 14*(3), 516-529. doi:10.1111/j.1467-7687.2010.00999.x

Morrison, R.G., Krawczyk, D., Holyoak, K.J., Hummel, J.E., Chow, T., Miller, B., & Knowlton, B.J. (2004). A neurocomputational model of analogical reasoning and its breakdown in frontotemporal lobar degeneration. *Journal of Cognitive Neuroscience, 16,* 260-271.

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.

Rattermann, M. J., & Gentner, D. (1998). More evidence for a relational shift in the development of analogy: Children's performance on a causal-mapping task. *Cognitive Development, 13,* 453-478.

Richland, L.E., & Burchinal, M.R. (2013). Early executive function predicts reasoning development. *Psychological Science, 24*(1), 87-92.

Richland L.E., Chan, T-K., Morrison, R.G., & Au, T.K-F. (2010). Young children's analogical reasoning across cultures: Similarities and differences. *Journal of Experimental Child Psychology, 105*, 146-153.

Richland, L.E., Morrison, R.G., & Holyoak, K.J. (2006). Children's development of analogical reasoning: Insights from scene analogy problems. *Journal of Exp Child Psychology, 94,* 249–273.

Thibaut, J.-P., French, R. M., & Vezneva, M. (2010a). Analogy-making in children: The importance of processing constraints. *Journal of Experimental Child Psychology, 1*, 1-19.

Thibaut, J.-P., French, R. M., & Vezneva, M. (2010b). Cognitive load and semantic analogies: Searching semantic space, *Psychonomic Bulletin & Review*, 17 (4), 569-574 doi:10.3758/PBR.17.4.569.

Thibaut, J. P., French, R. M. (2016), Analogical reasoning, control and executive functions: a developmental investigation with eye-tracking, *Cognitive Development. 38*, 10 - 26.

Viskontas, I.V., Morrison, R.G., Holyoak, K.J., Hummel, J.E., & Knowlton, B.J., (2004) Relational integration, inhibition and analogical reasoning in older adults. *Psychology and Aging, 19,* 581-591.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, Mass.: Harvard University Press.

Zelazo, P. D., & Frye, D. (1998). Cognitive complexity and control: The development of executive function. *Current Directions in Psychological Science, 7,* 121–126.

Zelazo, P.D., Müller, U., Frye, D., & Marcovitch, S. (2003). The development of executive function in early childhood. *Monographs of the Society for Research in Child Development, 68*.

Table 1

Solution Patterns in Children's and DORA/LISA Simulations for the Three Learning Trajectories

	Non-Analogical		Transitional		Analogical		
	Children	DORA/LISA (Low Knowledge)	Children	DORA/LISA (Low Knowledge)	Children	DORA/LISA (Low Knowledge)	DORA/LISA (High Knowledge)
Analogical Solution	21	25	56	57	78	77	85
Incomplete Solution	21	17	28	26	18	17	13
Associative Solution	58	58	16	17	3	6	2

Note: Children's results based on data presented in Table 5 (p. 384) of Hosenfeld and colleagues (1997).

Figure Captions

Figure 1. Hosenfeld et al. (1997) developed a geometric analogy task with problems of varying complexity created using relations familiar to children (e.g., above/below, inside, halving or duplication, rotation). Examples of open-ended geometric analogy items of low (Problem 1), intermediate (Problem 2), and high level of difficulty (Problem 3). The D-term has to be filled in by the subjects (figure adapted from the Hosenfeld et al., 1997).

Figure 2. Schematic illustration of how DORA and LISA work together to enable relational learning and reasoning.

Figure 3. Representation of a proposition in LISA/DORA.

Figure 4. DORA learns a representation of *inside* by comparing a square that is inside some object to a triangle inside some object. (a) DORA compares square and triangle and units representing both become active. (b) Feature units shared by the square and the triangle become more active than unshared features (darker grey). (c) A new unit learns connections to features in proportion to their activation (solid lines indicate stronger connection weights). The new unit codes the featural overlap of the square and triangle (i.e., the role "*inside*").

Figure 5. DORA learns a representation of the whole relation *contains* (house, square) by mapping *outside* (circle) to *outside* (house) and *inside* (triangle) to *inside* (square). (a) The units coding *outside* fire; (b) the units for circle and house fire; (c) the units for *inside* fire; (d) finally, the units for triangle and square fire. (e-f) DORA recruits a P unit that learns connections to the

active RB unit (the RB coding for *outside* (house)) in the recipient. (g-h) The P unit learns connections to the active RB unit in the (the RB coding for *inside* (square)). The result is a structure coding for *contains* (house, square).

Figure 6. Simulation of relational learning in DORA. DORA's relational learning algorithm was run at either low (0.4), medium (0.6), or high (0.8) lateral inhibition levels for 100 to 800 iterations to generate representations used in LISA for the low-knowledge condition. For the high-knowledge version a high (0.8) lateral inhibition level was used for 300 to 1100 iterations.

Figure 7. Results from children (Hosenfeld et al., 1997) and LISA simulations. Simulation results were obtained by allowing LISA to make analogical inferences using the representations generated in DORA (see Figure 6). The three performance groups of children were simulated by using three different levels of lateral inhibition in both DORA and LISA.

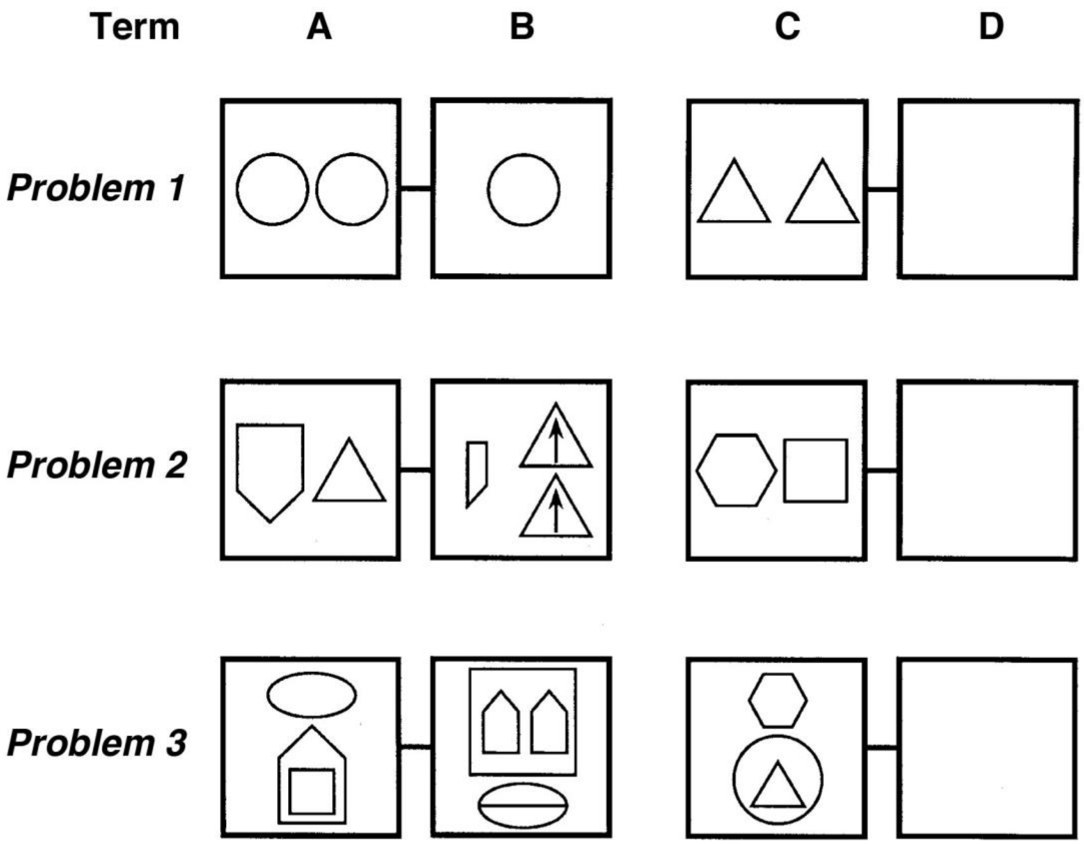

Figure 1

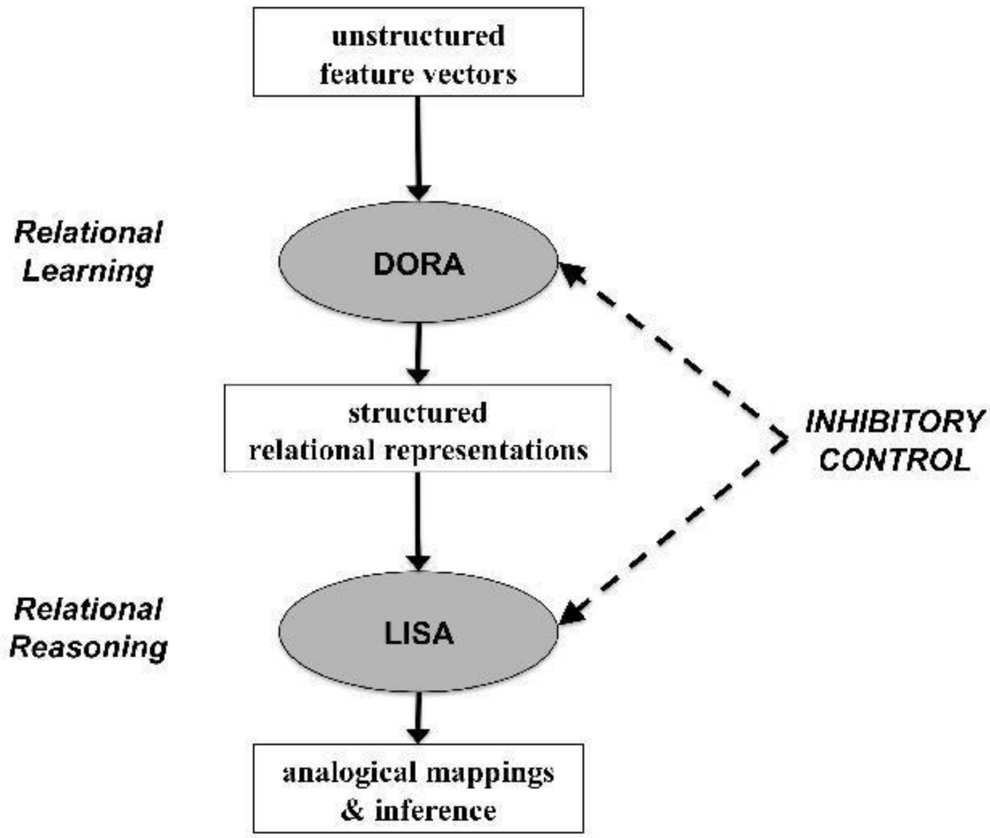
unstructured
feature vectors
Relational
Learning
DORA
structured
relational representations
INHIBITORY
CONTROL
Relational
Reasoning
LISA
analogical mappings
& inference

Figure 2

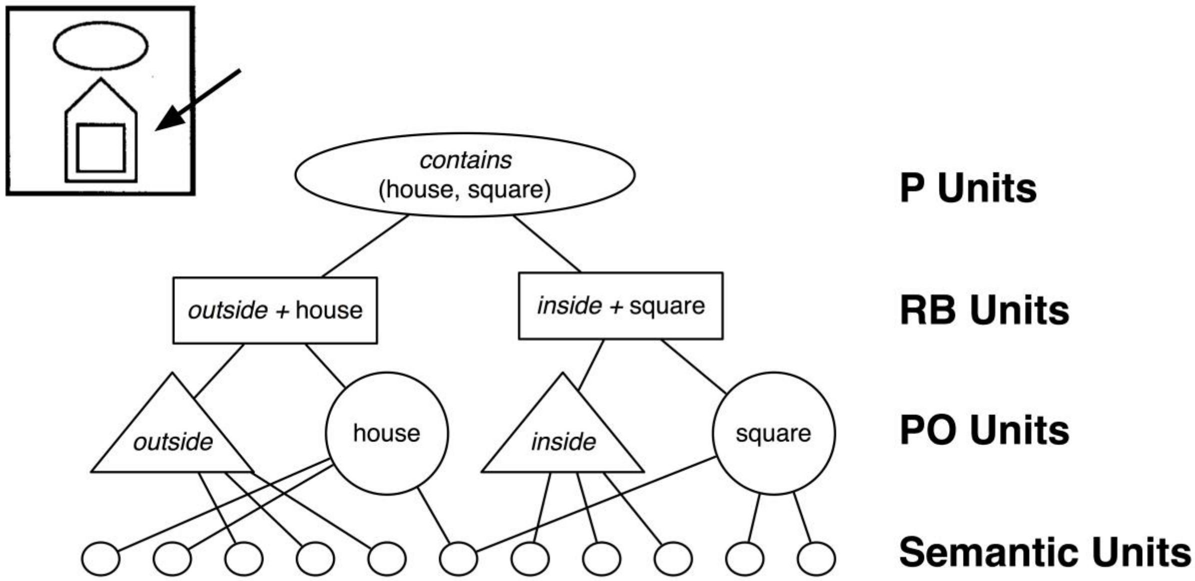
contains
(house, square)
P Units
outside + house
inside + square
RB Units
outside
house
inside
square
PO Units
Semantic Units

Figure 3

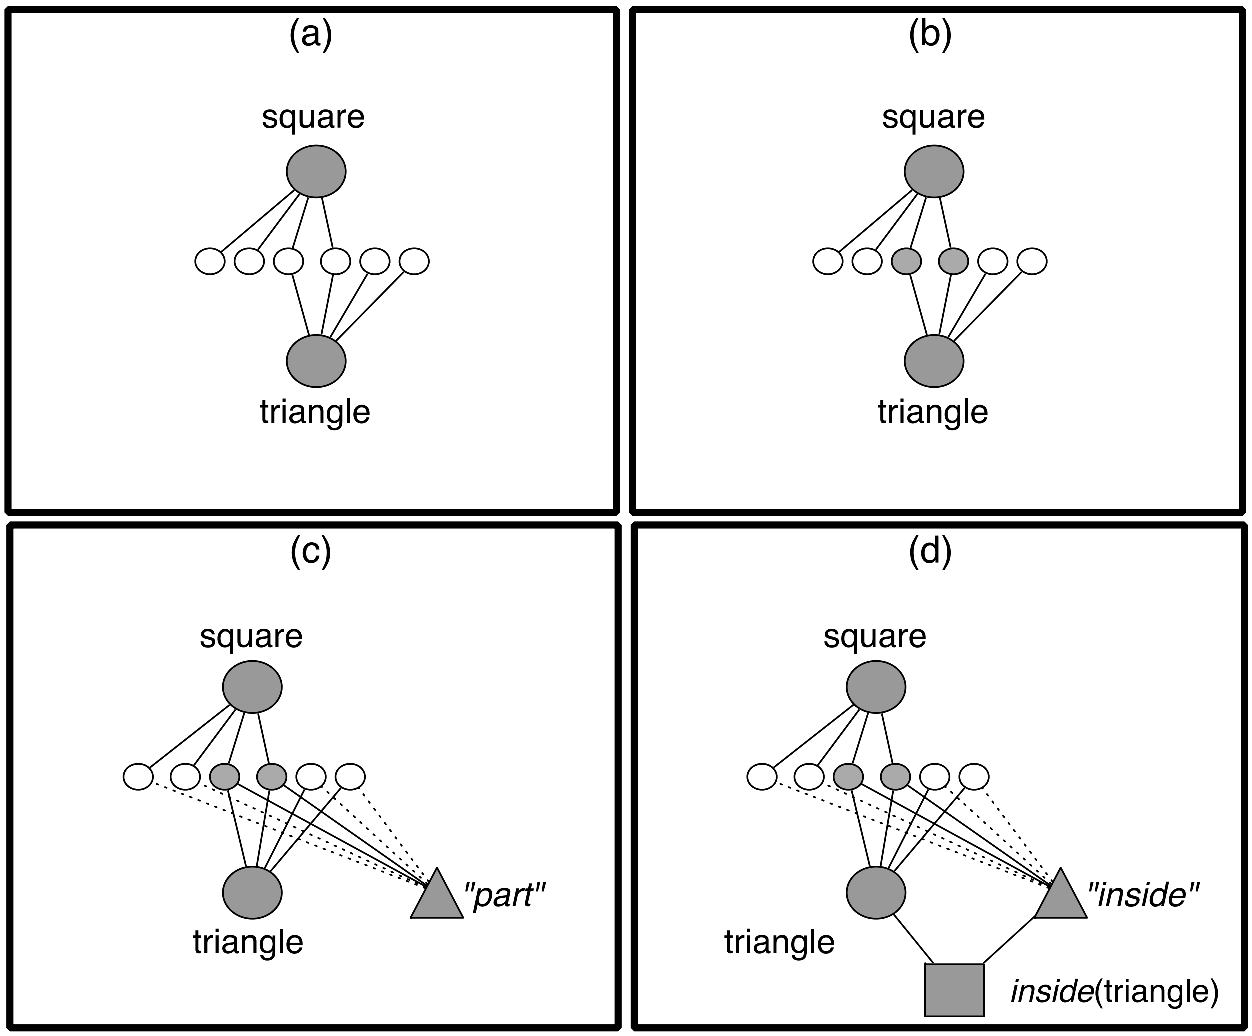
(a)
square
triangle
(b)
square
triangle
(c)
square
"part"
triangle
(d)
square
"inside"
triangle
inside(triangle)

Figure 4.

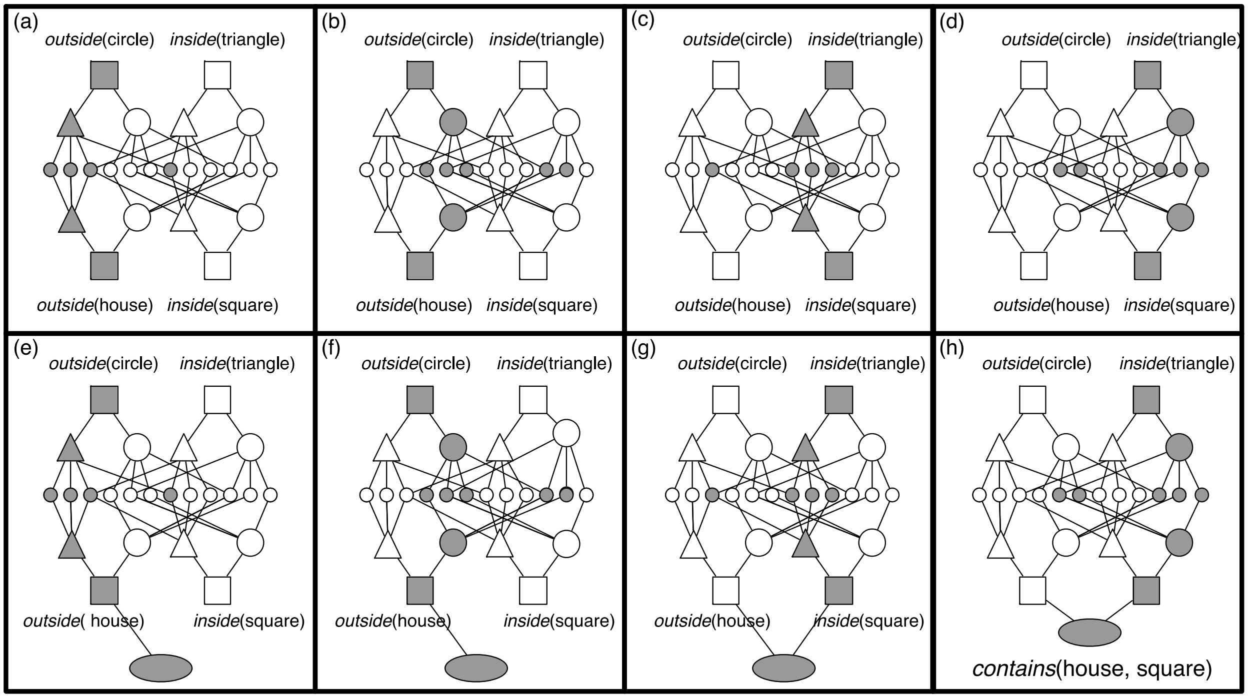
(a)
outside(circle) inside(triangle)
outside(house) inside(square)
(b)
outside(circle) inside(triangle)
outside(house) inside(square)
(c)
outside(circle) inside(triangle)
outside(house) inside(square)
(d)
outside(circle) inside(triangle)
outside(house) inside(square)
(e)
outside(circle) inside(triangle)
outside(house) inside(square)
(f)
outside(circle) inside(triangle)
outside(house) inside(square)
(g)
outside(circle) inside(triangle)
outside(house) inside(square)
(h)
outside(circle) inside(triangle)
contains(house, square)

Figure 5

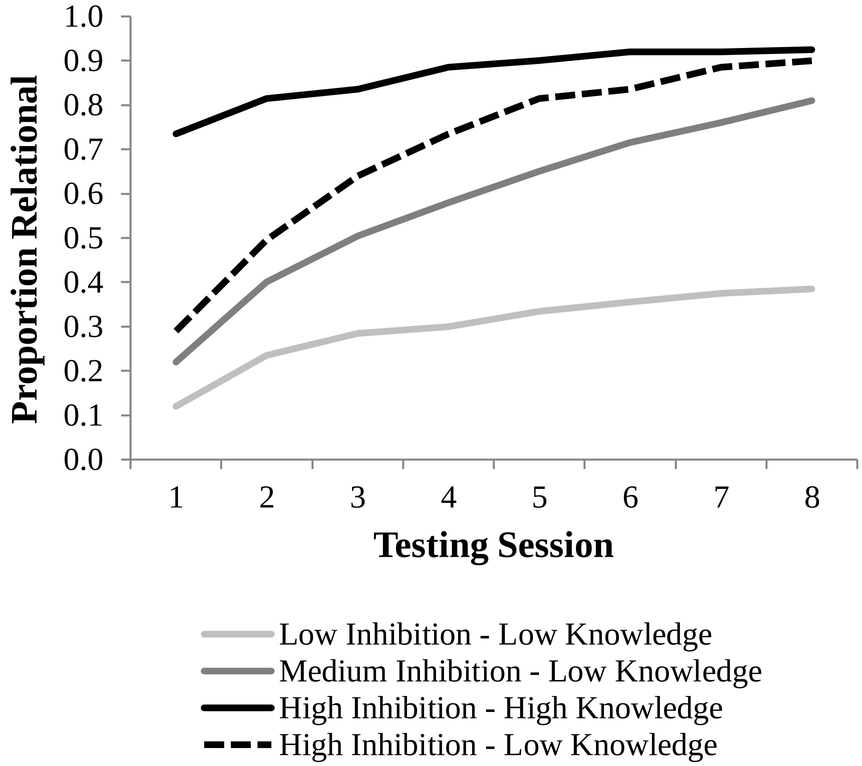
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
Proportion Relational
1
2
3
4
5
6
7
8
Testing Session
Low Inhibition - Low Knowledge
Medium Inhibition - Low Knowledge
High Inhibition - High Knowledge
High Inhibition - Low Knowledge

Figure 6

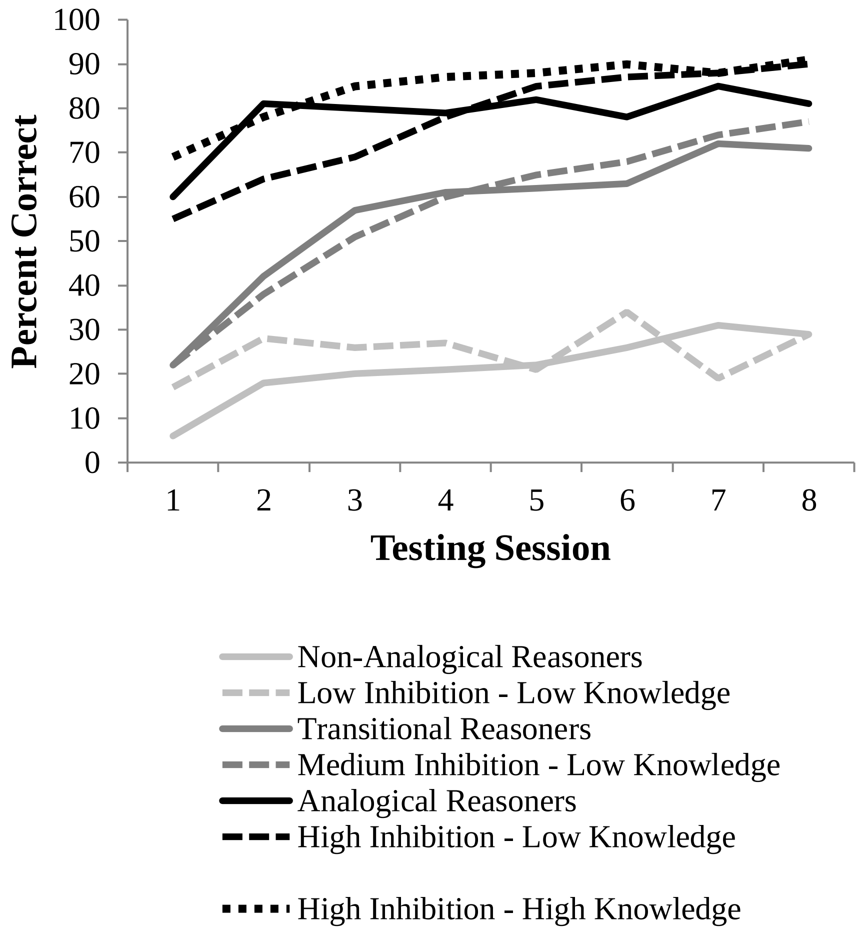
Percent Correct
100
90
80
70
60
50
40
30
20
10
0
1
2
3
4
5
6
7
8
Testing Session
Non-Analogical Reasoners
Low Inhibition - Low Knowledge
Transitional Reasoners
Medium Inhibition - Low Knowledge
Analogical Reasoners
High Inhibition - Low Knowledge
High Inhibition - High Knowledge

Figure 7