

The Development of Analogy: Working Memory in Relational Learning and Mapping

Leonidas A.A. Dumas (leonidas@hawaii.edu)
University of Hawaii at Manoa, Department Psychology
2430 Campus Rd. Honolulu, HI 96822

Robert G. Morrison (rmorrison@luc.edu)
Loyola University Chicago, Psychology Department,
6525 North Sheridan Road Chicago, IL 60626 USA

Lindsey E. Richland (lerich@uci.edu)
University of California, Irvine, Department of Education
2001 Berkeley Place, Irvine, CA 92697-5500

Abstract

Individual differences in analogical reasoning, long of interest to intelligence researchers, provide a unique window to view how changes in working memory and relational learning may jointly contribute to development. Hosenfeld, van der Maas, and van den Boom (1997) collected geometric analogy data from 6-7 year children during repetitive testing sessions over the course of one year. They identified three groups of children who showed different performance trajectories. We simulate these data in DORA/LISA and suggest that improved performance over training sessions likely results from children improving in being able to identify spatial relations, while the differences in learning trajectories across the groups of children of the same age are best explained by individual differences in working memory.

Keywords: analogy; analogical reasoning; development; computational models; individual differences; working memory.

Introduction

Developments in children's analogical reasoning are traditionally attributed either to increased working memory resources due to maturation (e.g., Halford, 2005; Richland, Morrison & Holyoak, 2006) or accretion of a knowledge base relevant to the particular analogical reasoning task (see Rattermann & Gentner, 1998; Goswami, 2001). To address these alternative claims, studies have either held knowledge constant and correlated age with success on analogy tasks with increasingly demanding working memory requirements (see Halford, 2005; Richland et al, 2006), or correlated performance on a knowledge test with performance on an analogical reasoning task (see Goswami, 2001). Most of these experiments have been cross-sectional, which has impeded the field's ability to develop a comprehensive theory of development that includes both factors.

A study by Hosenfeld et al. (1997) used a longitudinal methodology, collecting geometric analogy

data from 6-7 year old children during repetitive testing sessions over the course of one year. Most children improved over the sessions, but these authors were particularly interested in the qualitative nature of the developmental trajectories. Children's performance followed into three distinct patterns of change (see Figure 6) and the authors analyzed these patterns for qualitative insights into analogical change.

These data are also illustrative in considering the relations among processing resources and knowledge in development. We simulate these data in DORA/LISA to better understand the hypothesized contributions of resource maturation (i.e., working memory), and knowledge accretion. Specifically, we use DORA to simulate children's ability to better recognize spatial relations over repeated training sessions and then use LISA to simulate children's reasoning based on these representations of spatial relations. Importantly, we manipulate working memory in both models (via changes in lateral inhibition in the model) to simulate individual differences across groups of children.

Methods

In this section we describe the Hosenfeld, van den Boom, and Resing (1997) study, followed by a general description of the DORA and LISA models and present task simulations.

Task Description

In Hosenfeld et al.'s (1997) study children solved geometric analogy problems consisting of simple shapes in common relations such as above/below (see Figure 1). The complexity of the problems was varied by changing the number of relations needed to characterize the A:B transition. During testing, children solved A:B :: C:D problems in which they had to draw the missing D term to make a valid analogy. 6-7 year-old children were administered the task eight times over the course of one year, at three-week intervals.




Figure 1. Analogy problems varying in complexity based on those in Hosenfeld et al., (1997).

Researchers recorded accuracy rates, time to solution, and types of errors made.

Children's performance could be collapsed 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 learning trajectories of accuracy over time are shown in Figure 6.

Model Description

LISA (Hummel & Holyoak, 1997, 2003) is a symbolic-connectionist model of analogy and relational reasoning. DORA (Doumas, Hummel, & Sandhofer, 2008) is a model, based on LISA, which learns structured (i.e., symbolic) representations of properties and relations from unstructured inputs. That is, DORA provides an account of how the structured relational representations LISA uses to perform relational reasoning can be learned from examples.

DORA accounts for over 20 phenomena from the literature on relational learning, as well as its development (Doumas et al., 2008). In addition, as DORA learns representations of relations and properties it can be coupled to LISA to simulate an additional 30+ phenomena in relational thinking. The description of DORA/LISA that follows is a brief overview due to space constraints. For full details see Doumas et al. (2008) and Hummel and Holyoak (1997, 2003).

LISAese Representations In LISA (and by extension in DORA *after* it has gone through learning) relational structures are represented by a hierarchy of distributed and localist codes (see Figure 2). At the bottom, “semantic” units represent the features of objects and roles in a distributed fashion. At the next level, these

distributed representations are connected to localist units (POs) representing individual predicates (or role) and objects. Localist role-binding units (RBs; alternatively called subpropositions, SPs) link object and relational role units into specific role-filler bindings. At the top of the hierarchy, localist P units link RBs into whole relational propositions.

To represent the proposition *contains* (shield, square) as shown in the top left stimulus in Figure 1, PO units (triangles and large circles in Figure 2) representing the relational roles *outside* and *inside*, and the fillers shield and square are connected to semantic units coding their semantic features. RB units (rectangles) then conjunctively code the connection between roles and their fillers (one RB connects shield 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*+shield and the RB representing *inside*+square, thus encoding the relational proposition *contains* (shield, square).

Propositions are divided into two mutually exclusive sets: a driver and one or more recipients. In LISA, the sequence of firing events is controlled by the driver. Specifically, one (or at most three) proposition(s) in the driver becomes active (i.e., enter working memory). 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, 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 2, in order to bind *outside* to shield and *inside* to square (and so represent *contains* (shield, square)), the units corresponding to *inside* fire directly followed by the

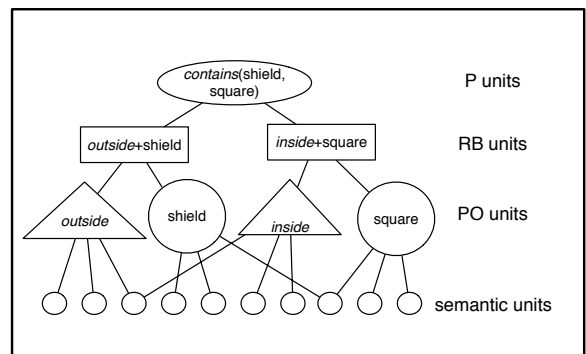


Figure 2. A proposition in LISA/DORA. Triangles denote roles and circles denote objects.




Figure 3. 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*”).

units corresponding to shield, followed by the units for coding *inside* followed by the units for square.¹

Relational Learning At a basic level, DORA uses comparison to isolate shared properties of objects and to represent them as explicit structures. DORA starts with simple feature-vector representations of objects (i.e., a node connected to set of features describing that object). When DORA compares one object to another, corresponding elements of the two representations fire simultaneously. For example, when DORA compares a square that is inside some object to a triangle that is inside some other object (e.g., the square inside the shield and triangle inside the circle in the first row of Figure 1), then the nodes representing the square and triangle fire together (Figure 3a). Any semantic features that are shared by both compared objects (i.e., features common to both the square and the triangle) receive twice as much input and thus become roughly twice as active as features connected to one but not the other (Figure 3b). DORA then recruits a new PO unit and learns connections between that unit and active semantics via Hebbian learning. Because the strength of connections learned via Hebbian learning is a function of the units’ activations, DORA learns stronger

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 (Domas et al., 2008).

connections between the new PO unit and more active semantic units (Figure 3c). The new PO thus becomes an explicit representation of the featural overlap of the compared square and triangle. In this example, DORA forms an explicit predicate representing “*inside*” (i.e., the features common to both the square and triangle). Importantly, the new PO acts as an explicit predicate representation of *inside* that can be dynamically bound to fillers.²

DORA then learns representations of multi-place relations by linking sets of constituent role-filler pairs into relational structures (see Domas et al., 2008 for details). Continuing the previous example, when DORA thinks about a triangle inside a circle, and a square inside a shield, it will map *outside* (circle) to *outside* (shield) and *inside* (triangle) to *inside* (square) (Figure 4a). This process produces a distinct pattern of firing over the units composing each set of propositions (namely, the RB units of *outside* (circle) fire out of synchrony with those of *inside* (triangle) while the RB unit of *outside* (shield) fire out of synchrony with those of *inside* (square); Figure 4b-d). The pattern serves as a reliable signal that DORA exploits to combine sets of role-filler pairs into multi-place relations. The diagnostic firing pattern signals DORA to recruit a P unit that learns connections to any active RBs in the recipient (Figure 4e). The end result is a P unit linking the RBs in the recipient into a whole relational structure (in Figure 4f-h, *contains* (shield, square)).

Mapping For the purposes of analogical mapping, DORA uses LISA’s mapping algorithm. DORA learns *mapping connections* between units of the same type (e.g., PO, RB, etc.) in the driver and recipient (e.g., between PO units in the driver and PO units in the recipient). These connections grow whenever corresponding units in the driver and recipient are active simultaneously. They permit LISA to learn the correspondences (i.e., mappings) between corresponding structures in separate analogs. They also permit correspondences learned early in mapping to influence the correspondences learned later.

Analogical inference When augmented with the capacity for self-supervised learning, LISA’s mapping algorithm (Hummel & Holyoak, 2003) naturally allows for analogical inference. To illustrate, consider how LISA solves an inference problem like one in the first row of Figure 1. LISA represents the A and B terms in the driver and the C term in the recipient. As the proposition coding for the A term, *contain* (shield, square), becomes active in the driver, it activates, and

² The new predicates DORA learns might be initially “dirty” in that they contain some extraneous features (e.g., any other features shared by the square and triangle from the above example). However, through repeated iterations of the same learning process, DORA forms progressively more refined representations (see Domas et al., 2008).

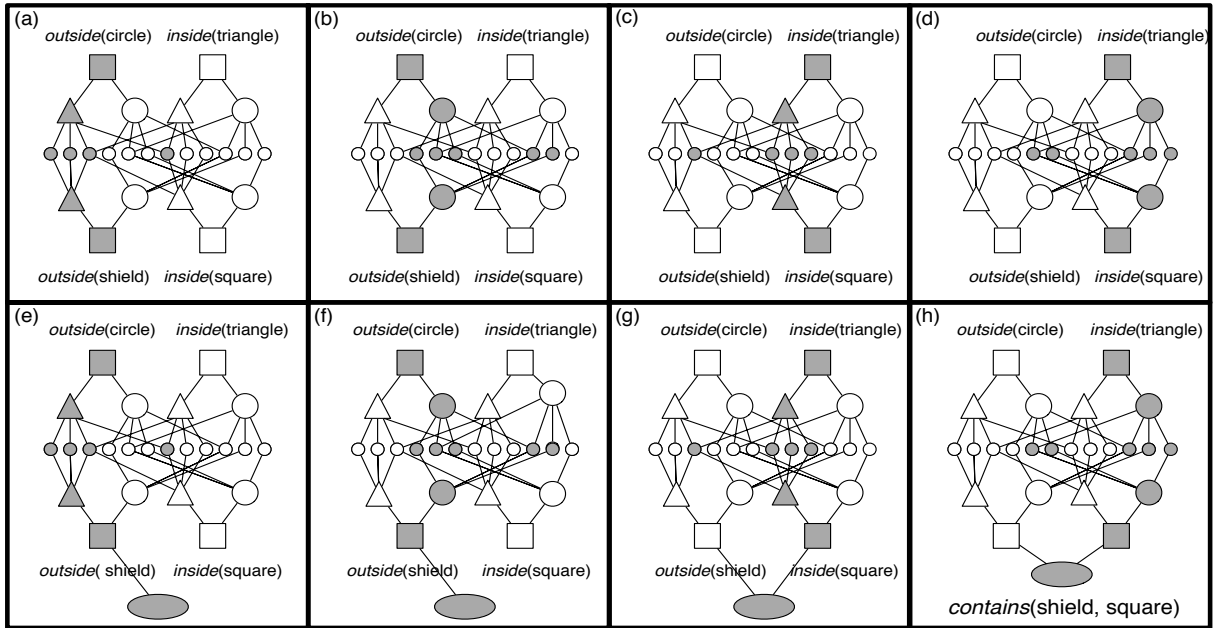


Figure 4. DORA learns a representation of the whole relation *contains*(shield, triangle) by mapping *outside*(circle) to *outside*(shield) and *inside*(triangle) to *inside*(square). (a) The units coding *outside* fire; (b) the units for circle and shield fire; (c) the units for *inside* fire; (d) finally, 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*(shield)) 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* (shield, square).

consequently maps to, the units coding for *contains* (circle, triangle) in the recipient. Specifically, the units coding for *outside* (shield) 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.

However, when the B term, *contains* (square, shield) becomes active in the driver, there are no corresponding units in the recipient for it to map to (recall the C term is already mapped to the A term). When units are active in the driver and no units are available for mapping in the recipient, LISA performs analogical inference via a self-supervised learning algorithm. During self-supervised learning, active units in the driver signal LISA to recruit matching units in the recipient. Continuing the example, as units coding for *outside* (square) become active, LISA recruits RB and P units in the recipient to match the active RB in P in the driver. Newly recruited P units in the recipient learn connections to active recipient RB units, and newly recruited RB units learn connections to active PO units (i.e., LISA learns connections between the new P and RB units and between the new RB unit and the units coding for *outside* and triangle in the recipient). In other words, LISA infers that *outside* (square) in the driver should correspond to *outside* (triangle) in the recipient. The same happens when *inside* (shield) fires in the driver and LISA infers *inside* (circle) in the

recipient. Thus, LISA completes the D term in a problem via analogical inference.

Simulation

We simulated Hosenfeld et al.'s (1997) results in two steps. In the first step we used DORA's relation learning algorithm to learn representations of the transformations used in the geometric analogy problems. We started DORA with representations of 100 objects attached to random sets of features (chosen from a pool of 100). We then defined 5 transformations (the same as 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 each relational transformation (changing containment) consisted of two roles each with two semantic features (e.g., for the *contains* relation, both the roles *inside* and *outside* were each defined by two specific semantic units). Each of the 100 objects was attached to the features of between 2 and 4 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 (as per DORA's relation learning algorithm). As DORA learned new representations it would also use these representations to make subsequent comparisons. For instance, 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 2 and 6 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.

We hypothesized that differences between the three groups of children from Hosenfeld et al.'s (1997) experiment were at least partially a product of differences in working memory. 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 DORA/LISA's intrinsically limited 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, Dumas, & Richland, 2006).

We defined three groups for the purposes of the simulation: (1) non-analogical, (2) transitional, and (3) analogical. We ran 100 simulations for each group. During each simulation we chose an inhibition level from a normal distribution with a mean of .4 for the non-analogical group, .6 for the transitional group, and .8 for the analogical group (each distribution had a SD = .2). For each simulation we ran 800 learning trials and checked the quality of the representations DORA had learned during the last 100 trials 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 to the mean of all other connection weights to normalize the quality measure to between 0 and 1). A higher quality denoted stronger connections to the semantics defining a specific transformation relative to all other connections (i.e., a more pristine representation of the transformation). Figure 5 shows the quality of the representations DORA learned at each level of inhibition.

In the second part of the simulation we used the representations DORA learned during the first part of the simulation in LISA to simulate solving the geometric analogy problems. We simulated all eight of Hosenfeld et al. (1997) testing phases. Each testing phase consisted of 20 trials. On each trial we presented LISA with the A and B terms in the driver and the C term in the recipient. The A, B and C term were object POs each attached to 4 random features and bound to PO predicate units identifying the transformations in which they were involved. We used representations of the transformations DORA had learned during the first simulation to represent the transformations in the testing trials. 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 we used the representations DORA had learned during the first 100 learning trials. For the second testing phase, we used the representations DORA had learned during the first 200 learning trials, and so on.

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 can be observed in Figure 6, DORA/LISA's performance on the testing trials closely followed those of the children in Hosenfeld et al. (1997). Just like the non-analogical children, DORA/LISA with a low inhibition level performed poorly throughout. Like the

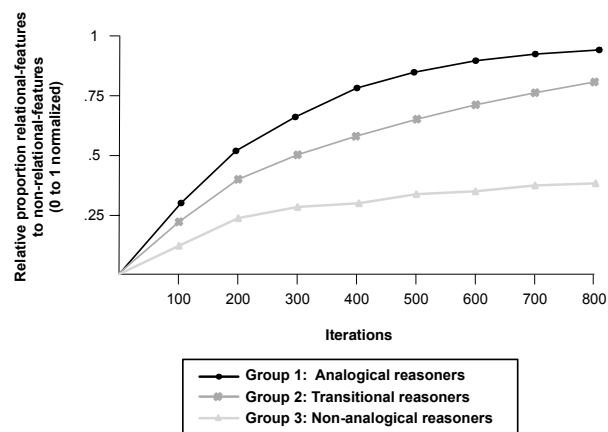


Figure 5. Simulation of relational learning in DORA. Groups were created by changing DORA's working-memory capacity via adjusting inhibition.

transitional children, DORA/LISA with a medium inhibition level started slow but improved slowly. Like the analogical children. Lastly, DORA/LISA with a high inhibition level performed well virtually from the start and maintained a good performance. It is interesting to note, however, that DORA/LISA performed poorly on the first several sessions for the Analogical group. We believe this is likely due to greater starting relational knowledge in this group of children. If for instance, we started with the DORA representations for session 3 instead of 1, the pattern of results would much more closely mirror the children's results. Thus, differences in relational knowledge may also be an important component of understanding individual differences in analogical reasoning.

Though we cannot present these data for space reasons, it is also significant to note that LISA made the same types of errors, in similar proportions, as children made in the Hosenfeld et al. (1997) study. For instance, DORA, just like children, tended to make errors by inferring a D tern solution with the correct transformations applied to the wrong objects, or simply copying all or part of the B term.

Conclusion

We assert that working-memory resources and relational knowledge each contribute differently to relational reasoning, with working-memory resources




Figure 6. Results from Hosenfeld et al., (1997b) and LISA simulations. Simulation results were the result of training in DORA followed by reasoning in LISA. Groups were created solely by changing DORA/LISA's working-memory capacity (i.e., adjusting lateral inhibition).

emerging as an important source of persistent individual differences in relational learning.

While considerable effort has been directed at understanding how working memory supports analogical reasoning, less attention has been given to looking at the role of working memory in its antecedent, relational learning. Understanding this factor will be an important step in understanding how relational learning develops and how it can contribute to successful analogical reasoning in children.

References

- 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.
- 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. (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
- 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.
- Morrison, R.G., Doumas, L.A.A., & Richland, L.E. (2006). The development of analogical reasoning in children: A computational account. *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 603-608). Mahwah, NJ: Erlbaum.
- 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.
- 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., 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.
- 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.