THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A theory of the discovery and predication of relational concepts

Citation for published version:

Doumas, LAA, Hummel, JE & Sandhofer, CM 2008, 'A theory of the discovery and predication of relational concepts', *Psychological Review*, vol. 115, no. 1, pp. 1-43. <https://doi.org/10.1037/0033-295X.115.1.1>

Digital Object Identifier (DOI):

[10.1037/0033-295X.115.1.1](https://doi.org/10.1037/0033-295X.115.1.1)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Psychological Review

Publisher Rights Statement:

© 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), 1-43. 10.1037/0033-295X.115.1.1

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Formatted manuscript copy of paper appearing as:

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.

© APA. The article does not exactly replicate the final version published in the APA journal. It includes some corrections to the original noted in DOI: 10.1037/a0032966, 2013. It is not the copy of record.

A theory of the discovery and predication of relational concepts

Leonidas A. A. Doumas

Indiana University, Bloomington

John E. Hummel

University of Illinois, Urbana-Champaign

Catherine M. Sandhofer

University of California, Los Angeles

Address for correspondence: Leonidas A. A. Doumas, Dept. of Psychological and Brain Sciences, Indiana University, 1101 E. Tenth Street, Bloomington, IN 47405. E-mail: adoumas@indiana.edu.

Abstract

Relational thinking plays a central role in human cognition. However, we do not know how children and adults acquire relational concepts and come to represent them in a form that is useful for the purposes of relational thinking (i.e., as structures that can be dynamically bound to arguments). The authors present a theory of how a psychologically and neurally plausible cognitive architecture can discover relational concepts from examples and represent them as explicit structures (predicates) that can take arguments (i.e., *predicate* them). The theory is instantiated as a computer program called *DORA* (*Discovery Of Relations by Analogy*). *DORA* is used to simulate the discovery of novel properties and relations, as well as a body of empirical phenomena from the domain of relational learning and the development of relational representations in children and adults.

Keywords: Learning relations, learning structured representations, relation discovery, cognitive development, analogy.

Relational thinking—thinking that is constrained by the relational roles things play, rather than just the literal features of those things—is a cornerstone of human perception and cognition. It underlies our ability to comprehend visual scenes (Biederman, 1987, Green & Hummel, 2004), learn and use rules (Anderson & Lebiere, 1998; Lovett & Anderson, 2005), appreciate analogies between different situations or knowledge systems (Gentner, 1983, 1989; Gick & Holyoak, 1980, 1983; Holyoak & Thagard, 1995), understand and produce language, science, art, and mathematics, and even our ability to appreciate basic perceptual similarities (Medin, Goldstone & Gentner, 1993). The ability to appreciate and use relations also underlies the transition from similarity-based to structure-based cognition in children (e.g., Gentner, 2003; Gentner & Rattermann, 1991; Halford, 2005).

In order to think about a relation, it is necessary to represent it as an explicit entity (i.e., a *predicate*) that can take novel arguments. Doing so entails solving three more basic problems (Doumas & Hummel, 2005; Hummel & Holyoak, 1997, 2003). First, it is necessary to have a representational element (e.g., a symbol, node[s] in a network, or neuron[s] in a brain) that corresponds specifically to the relation (or, more accurately, to the roles of the relation, as elaborated shortly), because in order to appreciate what one instance of a relation has in common with another, the two situations

must have something in common (specifically, the representational elements corresponding to the roles of the relation; Hummel & Biederman, 1992; Hummel & Holyoak, 1997).

Second, it is necessary to specify the bindings of relational roles to their arguments. Together the first and second requirements imply that the bindings of roles to their fillers must be *dynamic* (Hummel & Biederman, 1990, 1992; Hummel & Holyoak, 1997; Shastri & Ajenagadde, 1993). That is, it must be possible to create and destroy bindings on the fly, and the mechanism or tag that represents the bindings must be independent of the elements that represent the roles and fillers so bound. For example, in order to understand how a cone above a brick is similar to and differs from a brick above a cone, one must simultaneously appreciate that the same elements are involved in the same relation in both cases, and that, while in one case the brick is bound to the *higher* role and the cone to the *lower* role, in the other case these role bindings are reversed.

Finally, these representational elements—and the concepts they represent—must come from someplace. In particular, unless all relational concepts (e.g., above, causes, chases, loves, larger-than, ameliorates, etc.) are assumed to be innate, they must somehow be learned from examples. Although numerous cognitive and perceptual models postulate elements corresponding specifically to relations (e.g., Anderson & Lebiere, 1998;

Falkenhainer, Forbus & Gentner, 1989) or relational roles (Hummel & Biederman, 1992; Hummel & Holyoak, 1997, 2003; Shastri & Ajenagadde, 1993), and all of these models specify methods for binding relational roles to their arguments, to date no model has provided a satisfactory account of how we learn relational concepts from examples (although for efforts in this direction, see Gasser & Colunga, 2001, 2003).

Accounting for how we learn and predicate relational concepts is difficult, in part, because of the requirement that the same element(s) must represent a given relation regardless of what its arguments happen to be at the time. Ideally, a relation (such as *above*) must be represented in a way that is completely agnostic with respect to its potential arguments (so that it can be represented in the same way regardless of what is above what). However, any specific example of a relation is always instantiated with some specific set of arguments (e.g., in order to have an example of *above*, some specific thing must be above some other specific thing). It is never possible to observe an example of pure, disembodied “aboveness”. As such, it is not obvious how the cognitive architecture might learn to represent “aboveness” in a way that is argument free. Our goal in this paper is to present a theory of how the human cognitive architecture solves this problem of learning and representing novel relational concepts.

The Development of Relational Thought

An important theme that has emerged in the study of relational thinking is that the ability to reason relationally changes with development (e.g., Gentner & Rattermann, 1991; Halford, 2005). Across a variety of tasks and procedures children initially make inferences based on whole-object similarity and gradually acquire the ability to make inferences based on the relational roles to which objects are bound (e.g., Gentner, 1977; 1988; 2003; Gentner & Rattermann, 1991; Gentner & Namy, 1999; Gentner & Toupin, 1986; Goswami, 1992; Halford, 1980; Kotovsky & Gentner, 1996; Richland, Morrison, & Holyoak, 2006; Smith, 1984, 1989). For example, given a picture of a dog chasing a cat and another picture of a boy chasing a girl with a cat in the background, 3 year-old children tend to match the cat in the first picture to the cat in the second picture (based on their featural similarity), whereas 5 year-old children tend to match the cat in the first picture to the girl in the second picture based on their relational similarity (both are being chased; e.g., Richland, Morrison, & Holyoak, 2006). Gentner and Rattermann (1991) refer to this developmental trend as the relational shift.

Traditional connectionist models based on distributed representations (e.g., Colunga & Smith, 2005) provide a good account of younger children’s reasoning based on whole-object similarity. However, these systems cannot account for later relational thought (see Holyoak & Hummel, 2000; St. John, 1992). On the other hand, systems based on structured representations

(e.g., Anderson & Lebiere, 1998; Falkenhainer et al., 1989; Hummel & Holyoak, 1997, 2003) provide a good account of older children’s and adult’s reasoning based on relations, but provide no account of where the structured representations on which they rely come from in the first place. That is, although we can account for the behavior of both younger and older children on relational tasks, we cannot account for how the ability to reason relationally develops because we do not know how the kinds of representations that support relational thought are learned from the kinds of representations that support whole-object similarity-based reasoning.

That we lack an account of how people do—or even *could*—learn structured representations from unstructured examples is often cited as the most significant limitation of structure-based accounts of cognition (e.g., Munakata & O’Reilly, 2003; O’Reilly & Busby, 2002; O’Reilly, Busby, & Soto, 2003). As such, an understanding of how we learn structured relational concepts from unstructured inputs will not only contribute to our understanding of the development of relational thinking and the foundations of symbolic thought in general (see Smith, 1989), but will also address a fundamental limitation of current structure-based accounts of cognition.

The Purpose of the Current Work

Discovering a relation and representing it in a form that can support relational thinking entails solving three problems. First, there must be some basic featural *invariants* that remain constant across instances of the relation, and the perceptual/cognitive system must be able to detect them. Second, the architecture must be able to isolate these invariants from the other properties of the objects engaged in the relation to be learned. And third, it must be able to *predicate* the relational properties—that is, represent them as explicit entities that can be bound to arbitrary, novel arguments.

Detecting Featural and Relational Invariants

The most basic prerequisite to explicitly relational thought is the capacity to detect relational (and featural) invariants in the environment. For example, to learn the *above* relation, the cognitive architecture must be able to detect the invariant perceptual properties present when one object is above another regardless of the nature of the objects involved. There is ample evidence that children possess at least primitive versions of these features from even a very young age.

For example, Clearfield and Mix (1999) and Feigenson, Carey, and Spelke (2002) have shown that even infants as young as six-months are sensitive to differences such as “more” or “less” in properties like size and surface area. Similarly, Baillergeon and her colleagues have shown that very young children have an intuitive understanding of basic relational concepts such as *occludes*, *contains*, *collides-with* and *supports* (for a summary, see Baillergeon, 2004). These findings sug-

gest that the mechanisms for detecting these invariants are present in the visual system at a very early age.

In addition, some models have made progress demonstrating how basic visual invariants can be detected from early visual representations. For example, Hummel and Biederman (1992) describe a model that computes abstract invariant features, including relational features, from holistic visual representations very much like those in visual area V1 (see also Hummel, 2001; Hummel & Stankiewicz, 1996). In addition, Kellman, Burke and Hummel (1999) describe a model that can learn visual invariants from similarly V1-like inputs. Together, these models suggest solutions to the problem of detecting invariant relational features.

The findings and models summarized here do not provide a complete answer to the question of where relational invariants come from; important aspects of this problem remain largely unsolved. But this work demonstrates, minimally, that children—even very young children—can detect featural and relational invariants in the world and that such invariants can be computed from early (non-invariant) visual representations.

However, being able to detect relational invariants (i.e., solving the first problem) is not the same as being able to isolate them from other features and represent them as explicit structures that can take arguments (solving the second and third). Simply having relational “features” provides no basis for binding those relations to their arguments. It is for this reason that simply declaring that various nodes in a neural network represent “relations” does not render the network’s representations relational (see Doumas & Hummel, 2005; Halford et al., 1998; Hummel & Holyoak, 1997, 2003). What makes a representation relational is not the population of features it contains, but the capacity to compose those features into structures that permit dynamic binding of relational roles to their arguments. Thus, although it may appear at first blush that the difficulty of learning relational concepts lies in discovering their constituent “features”, in fact, the ability to compose those features into relational structures is what distinguishes a relational representation from a non-relational one.

Isolating and Predicating Object Properties and Relations

In the real world, relational invariants never appear in isolation. For example, every observable instance of the *above* relation consists of some specific object above some other specific object. In spite of this, we somehow learn to represent *above* in a way that remains the same regardless of the arguments to which it is bound. In order to learn an explicit representation of a relational concept such as *above*, the cognitive architecture must be able to isolate the relevant relational invariants from the other properties of the objects engaged in the relation (e.g., the specific shapes of those objects).

In addition, the cognitive architecture must be able to *predicate* the relational properties (i.e., represent them as explicit entities that can be bound to arbitrary, novel arguments; Doumas & Hummel, 2005; Gentner 1983, 1989, 2003; Gentner & Markman, 1997; Halford et al., 1998; Hummel & Holyoak, 1997, 2003; Markman, 1999). For example, predicating the relation *above* entails learning a representation of *above* that is independent of its arguments, and explicitly and dynamically codes the binding of these arguments to their relational roles (i.e., specifies which is above which). It is this last step that marks the transition from a feature-based representation to a genuinely relational or structured one (Doumas & Hummel, 2005; Halford, 1998; Hummel & Holyoak, 1997). As such, the difficulty of learning relational representations lies not in learning their *content* (e.g., what *above* “means”) but in learning their *format* (i.e., coming to represent them in a way that allows them to bind to novel arguments).

This paper presents a theory of how the human cognitive architecture solves the problems of isolating relational invariants and representing them as explicit structures that can take arguments. The result is a theory of how structured relational representations can be learned from unstructured examples—that is, of how relational thought can be bootstrapped from non-relational beginnings.

We begin by discussing a set of representational and processing constraints that make the problem of learning structured representations of relational concepts tractable. We then describe a theory of the discovery and predication of relational concepts based on these constraints. The theory is instantiated in a computer model called *DORA* (*Discovery Of Relations by Analogy*), which we have constrained to be both cognitively and neurally plausible (e.g., it works within intrinsic WM capacity limits, and both its knowledge representations and the operations on those representations are designed to have a transparent neural analog). We demonstrate the sufficiency of the model by using it to simulate (a) the discovery and predication of novel object properties and relations, (b) a body of empirical phenomena from the domain of relational learning and (c) the development of relational representations in children and adults. Finally, we discuss the implications and limitations of the model and suggest directions for future research.

Constraints on Relation Discovery

Knowledge Representation

One step toward constraining the problem of relation discovery and predication is to choose an appropriate form of knowledge representation. Formally, any multi-place predicate can be recast as a collection of single-place predicates (one for each role of the relation), with functions for linking them (Mints, 2001). Such representations are known as *role-filler bindings*. For example, the role-filler binding representation of *above* (ball,

table) would consist of a representation of the *higher* role bound to ball (*higher* (ball)) and the *lower* role bound to table (*lower* (table)) linked together to form the structure *higher* (ball)&*lower* (table) (where & is the linking function; see, e.g., Doumas & Hummel, 2004a, 2005).

Role-filler binding provides a natural constraint on the problem of relation discovery because it reduces the problem of learning relations to the problem of learning single-place predicates (i.e., properties or roles) and then linking them together to form complete relational structures. This approach allows us to recast the question “How do we discover and predicate relational structures from examples?” as two simpler questions: “How do we learn single-place predicates (such as object properties)?” and “How do we link them together to form multi-place relations?”

The Role of Mapping and Comparison

An important theme that has emerged in the literature on relational reasoning is that analogical mapping—the processes of discovering which elements of one system of knowledge correspond to which elements of another based on their shared relations—and the related process of comparison play a central role in all forms of relational reasoning (see Gentner, 1983, 2003; Holyoak & Thagard, 1995). For example, mapping bootstraps the induction of abstract relational schemas (e.g., Gick & Holyoak, 1983; Ratterman & Gentner, 1998; Sandhofer & Smith, 2001), and comparison assists in early category learning (e.g., Fisher & Sloutsky, 2005; Gentner & Medina, 1998; Gentner & Namy, 1999; Namy & Gentner, 2002; Namy, Smith, and Gershkoff-Stowe, 1997; Oakes & Madole, 2003), helps people appreciate which known relations might be relevant to a specific task (e.g., Bowdle & Gentner, 1997; Dixon & Bangart, 2004; Gick & Holyoak, 1980; Kottovsky & Gentner, 1996; Kurtz & Burkina, 2004; Markman & Gentner, 1993; Ross, Perkins, & Tenpenny, 1990; Spalding & Ross, 1994; Yamauchi & Markman, 1998, 2000), and aids in the discovery and predication of novel higher-order relations from known lower-order relations (Doumas & Hummel, 2004b).

Gentner (1983, 2003) and Mandler (1988, 1992, 2004) suggest that comparison plays a role in the discovery of new relations by highlighting the shared properties of objects. Similarly, we hypothesize that comparison may bootstrap the discovery of relations by leading to the discovery of shared properties, which can be linked together to form the roles of new relations. The general idea is that, during comparison, properties that objects share become more active than properties unique to one object or the other, thus highlighting the shared (invariant) properties and setting the stage for their becoming represented as explicit predicates.

As a simplified example, consider a child learning a property like “big” by comparing a truck to an elephant (Figure 1a). A key theoretical claim is that when objects

are compared, features they share receive input from both objects (i.e., because they are connected to both), whereas features unique to one object or the other receive input from only one object or the other. As a result, shared features tend to receive about twice as much input—and thus become about twice as active—as features unique to one object or the other. Thus, since trucks and elephants are both “big”, the features representing “big” should become more active than the features trucks and elephants do not share (Figure 1b). If the child can predicate these shared features (i.e., learn an explicit representation, such as a unit in a network that is connected to the features representing “big”), then she will have learned an explicit representation the property “big” (Figure 1c). If she then binds that property to the truck (or the elephant) she will have explicitly predicated the property “big” about the truck (or elephant). Applied to objects with various properties, this process can serve to highlight shared properties between compared objects and bootstrap their predication. Once representations of various object properties have been predicated in this fashion, comparison can also serve to link multiple roles together to form multi-place relational structures, as elaborated below.

One apparent limitation of this comparison-based approach to predication is that comparison (by assumption) highlights all the features two objects have in common, not just those that are relevant to the property or relation that is nominally in question. For example, in addition to being big, a truck and an elephant might also both be “in motion”. As a result, these shared features (like the features of “big”) would become highlighted and become part of the child’s (initial) explicit representation of “big”. For the same reason, the child’s earliest representations of relations (such as *above* and *larger-than*) would likewise be “corrupted” by irrelevant features that just happened to be present in the child’s earliest examples of the relations. Consistent with this prediction, Quinn and his colleagues (e.g., Quinn et al., 1996) have shown that infants do not initially represent spatial relations independently of the objects over which they hold, and Smith, Sera and Ratterman, (1988) have shown that children’s initial relational representations include properties of the objects that participate in these relations.

Shared Semantic Pools

In order to support comparison-based predication, the same representational units that represent object features must also be able to serve as features of predicates and relational roles. Consider the problem of explicitly predicating the property *red*—that is, transitioning from *red* as feature in a holistic representation of an object, to *red* as an explicit predicate that can take various objects as arguments. If object features and role features come from separate pools of units, then *red*, the object feature, will have nothing in common with the features of *red*, the explicit predicate.

Figure 1. (a) Representations of an elephant and a truck. Both are connected to a set of features. Some features, such as “big”, are shared by the elephant and the truck. (b) When the child thinks about an elephant and a truck simultaneously (e.g., during comparison), her representations of elephant and truck both become active simultaneously (denoted by the grey interior). The elephant and truck representations send activation to their constituent features. Features shared by both the elephant and truck receive twice as much input—because they receive input from both the elephant and the truck—and thus become more active than unshared features (greater activation is denoted by the thicker border). (c) A new unit learns a connection to the most active features and thus explicitly represents the featural over-lap of the truck and elephant (i.e., the property “big”).

Although this constraint seems straightforward, models of relational reasoning have traditionally represented predicates and objects as qualitatively different kinds of things (i.e., as different data types or separate pools of units; e.g., Falkenhainer, et al., 1989; Forbus, et al., 1995; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997, 2003; Keane, Ledgeway, & Duff, 1994; Kokinov & Petrov, 2001; Larkey & Love, 2003; Salvucci & Anderson, 2001). In brief, the reason for this convention is that the algorithms these models use for analogical mapping need to ensure that relations map to relations and objects map to objects, and the most straightforward way to enforce this constraint is to assume that they are simply different kinds of things. However, this convention precludes representing a property such as “red” as both an object property and as an explicit predicate that can take arguments. As such, it precludes learning predicates, or relational roles, by comparing objects to discover what they have in common. We therefore assume that both objects and relational roles (predicates) share a common pool of basic representational features.

Linking Single-place Predicates into Multi-place Relations

In a role-filler binding system, once single-place predicate-argument bindings have been learned, learning full-fledged relational structures becomes a matter of linking them together. For example, once a child has learned the predicates *high* (x) and *low* (y), she can form a primitive version of the *higher-than* (x, y) relation by linking them into a single two-place structure (i.e., *high* (x)&*low* (y)). Analogical mapping provides a mechanism by which smaller arity structures (e.g., predicates with arity-one, or single-place predicates) can be composed into larger arity structures (e.g., predicates with arity-two and above, or multi-place relations). Specifically, we assume that when sets (pairs or triads) of sin-

gle-place representations enter WM together, they can be mapped as a unit onto other such sets of representations, and that this situation serves as a signal to link the single-place predicates into a larger relational structure. For example, mapping “*high* (bird) and *low* (cat)” onto “*high* (clock) and *low* (ball)” will bootstrap the formation of the relations *higher-than* (bird, cat) and *higher-than* (clock, ball). This hypothesis is supported by the results of Dumas & Hummel (2004b), who demonstrated that mapping lower arity structures can bootstrap the formation of higher-arity representations (see also, Gentner, 2003; Gentner & Namy, 1999; Gick & Holyoak, 1983; Kotovsky & Gentner, 1996; Namy & Gentner, 2002; Yamauchi & Markman, 1998, 2000).

Summary

Our approach is based on four constraints (see Table 1). Together these constraints constitute our core theoretical claims. First, role-filler binding representations reduce the problem of learning relations to the problems of learning object properties or relational roles (single-place predicates) and linking them together to form multi-place relational structures. Second, comparison leads to the discovery and predication of shared properties, and, applied iteratively, results in progressively more refined representations of predicates, and eventually multi-place relations. Third, predicates and their arguments share a common representational basis (i.e., both predicates and their arguments are coded by a common vocabulary of representational primitives). Fourth, mapping multiple predicates of smaller arity leads to the formation of higher arity relational structures. Thus, as elaborated in the next section, our general proposal is that the same psychological mechanisms that underlie analogical inference and schema induction also underlie the discovery and predication of the object properties and relations that make analogical reasoning possible in the first place.

Table 1.
Core Theoretical Claims and Their Implementation in DORA.

Theoretical Claim	Implementation in DORA
1. Role-filler representations reduce the problem of learning relations to the problems of learning object properties (single-place predicates) and linking them together to form multi-place relational structures.	DORA represents relational structures as linked sets of role-filler pairs. This allows DORA to learn structured representations of novel relations from simple object representations.
2. Comparison can lead to the discovery and predication of shared properties.	Mapping and intersection discovery routines coupled with systematically asynchronous binding leads to the isolation and explicit predication of shared semantics.
3. A common vocabulary of representational primitives codes both predicates and their arguments.	Both predicates and objects share the same pool of semantic units.
4. Mapping predicates of smaller arity can lead to the formation of higher arity structures.	DORA forms larger arity relational structures (i.e., propositions with more RBs) when sets of propositions map consistently.

The Model

Overview

The heart of the DORA model is a set of algorithmic operations for exploiting the core theoretical constraints outlined in the previous section. The resulting model provides a mechanistic account of how structured relational representations can be learned from unstructured non-relational beginnings.

DORA performs four basic operations: 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. In conjunction, these operations allow DORA to predicate object properties by comparing examples, successively refine those predicates into progressively “purer” representations of properties (or relational roles), and combine lower-arity properties/relational roles into higher-arity relations. Importantly, these same four operations (memory retrieval, mapping, intersection discovery and self-supervised learning) are also the basic operations underlying analogical mapping, inference and schema induction (see Hummel & Holyoak, 2003a).

Elements of DORA’s operation and aspects of its knowledge representation are adopted from Hummel and Holyoak’s (1997, 2003) LISA model. Importantly however, significant aspects of DORA’s operation and knowledge representations depart sharply from those of LISA. Notably, DORA’s use of a single pool of units to represent the semantic features of both objects and predicates/relational roles, its comparison-based learning and refinement routines, its relation formation algorithm, and the manner in which its operations work together to learn and predicate properties and relations

are all unique to DORA. Whereas LISA requires predicates and relations to be hand-coded by the modeler, DORA provides an account of how these representations can be learned in the first place. Consequently, after DORA learns predicates and relations it takes LISA as a special case: DORA can simulate all the findings that have been simulated with LISA (as elaborated in the Discussion); in contrast, LISA cannot account for any of the phenomena simulated with DORA and reported in this paper. DORA’s use of a single pool of units to represent the semantic features of both objects and predicates/relational roles, its comparison-based learning and refinement routines, its relation formation algorithm, and the manner in which its operations work together to learn and predicate properties and relations are all unique to DORA. Whereas LISA requires predicates and relations to be hand-coded by the modeler, DORA provides an account of how these representations can be learned in the first place. Consequently, after DORA learns predicates and relations it takes LISA as a special case: DORA can simulate all the findings that have been simulated with LISA (as elaborated in the Discussion); in contrast, LISA cannot account for any of the phenomena simulated with DORA and reported in this paper.

We begin by discussing knowledge representation in DORA, both the representations it begins with and those that it eventually learns. We then describe how DORA performs binding and the flow of activation. Finally we describe the operations that allow DORA to learn structured relational representations from unstructured non-relational examples (namely, retrieval, mapping, predication and refinement, and relation formation). For the purposes of clarity, in this section we outline DORA’s operation in comparatively broad strokes. The full details of DORA’s operation, along with the majority of the equations and parameters, appear in Appendix A.

Figure 2. The structure of propositions in DORA. (a) An illustration of the representations that DORA begins with. An object bound to a set of features that describe it. (b) An illustration of DORA's relational (i.e., final) representations. At the bottom layer semantic units (small circles) code for the features of individual objects and relational roles. At the next layer, localist PO units code for individual predicates and objects. Although predicates and objects are not distinct "data types" in DORA (see text), we distinguish them in the figures for clarity: Circles for POs acting as objects, and triangles for POs acting as predicates. At the next layer, localist RB units code for specific role-filler bindings. At the top layer, localist P units represent collections of role-filler bindings that form complete propositions. Units in the PO, RB, and P layers are represented using different shapes for the purposes of clarity.

Knowledge Representation

As noted previously, the mental representations underlying adult (and older children's) relational thinking are characterized by two properties that make them simultaneously flexible, structured and difficult to simulate (Doumas & Hummel, 2005; Hummel & Holyoak, 1997, 2003): Relational roles are represented explicitly and independently of their fillers, and role-filler bindings are specified explicitly. By contrast, young children's representations are holistic (i.e., unstructured) in the sense that the features composing the representation are not accessible to processing independently of one another. Accordingly, DORA starts with holistic representations of objects that simply list the objects' features (Figure 2a). These kinds of representations are not capable of supporting relational thinking such as analogy or the use of variabalized rules (see Doumas & Hummel, 2005). Our goal is to provide an account of how children transition from representing the world in terms of holistic (unstructured) representations of objects, to representing the world in a structured fashion that makes relations, relational roles and their bindings to fillers explicit.

Individual Propositions

DORA begins with holistic representations of objects (Figure 2a), and learns relational representations that dynamically bind distributed representations of relational roles and objects into explicit propositional structures (Figure 2b). Figure 2a depicts the kind of representations with which DORA begins, and Figure

2b depicts the kind of representations that result from the application of its basic operations to representations like those in Figure 2a. We describe the latter representations (i.e., DORA's end state) now in order to make the rest of the model's description easier to follow.

Propositions are represented in four layers of units (see Figure 2b). At the bottom of the hierarchy *semantic* units code the features of objects and relational roles in a distributed fashion (thereby capturing their semantic content). These units might represent features such as visual invariants (e.g., "round", "square", "shiny"), relational invariants (e.g., "more", "less", "same"), dimensional properties (e.g., "size-3", "height-3", "color-red"), other perceptual properties (e.g., "sweet", "noisy", "rough"), complex perceptual/cognitive properties (e.g., "furry", "barks", "has-wheels", "fast"), category information (e.g., "apple", "dog", "fire engine") and information identifying individuals (e.g., "me", "Spot", "Jane", "mommy").

In this paper we label semantic units both in the text and figures. For example, we might label a semantic unit "big". Of course we do not claim that this unit is the "right" representation of the property *big*. The labels attached to semantics are arbitrary and mean nothing to DORA. We use them only for clarity of exposition and for interpreting the model's behavior. Rather, our claim is that some set of units code the relevant properties of attributes, relations and relational roles. The crucial aspect of these units is that they are independent of one another in the sense that they represent separable properties (e.g., a unit that codes for "red" will become active in response to any red object, not only, say, red objects in the upper-left of the visual

field). Many such features are detectable from the outputs of early visual processing (e.g., V1), as demonstrated by Hummel and Biederman's (1992) JIM model of object recognition. Indeed, many of DORA's feature units (namely, those coding for basic visual and relational invariants) correspond precisely to the kinds of features activated in the third and fifth layers of the JIM model. As noted previously, these features can serve to bootstrap the discovery and predication of relations, but are not themselves capable of supporting explicit relational thought.

At the next layer localist *predicate* and *object* (PO) units code for individual predicates (relational roles or attributes) and objects. Each PO unit serves as a token for the role or object it represents and is connected to the semantic units corresponding to its features. For example, the token for the object Fido in the proposition *bigger* (Fido, Sara) would be connected to a set of semantics corresponding to the features of Fido (e.g., "adult", "male", "dog"), while the token for the relational role *larger* would be connected to a set of semantics corresponding to its features (as would the token for Sara and the *smaller* role; see Figure 2b).

Above the PO units a layer of localist role-binding units (RBs) encode the bindings of relational roles to their fillers. For example, the proposition *bigger* (Fido, Sara) contains two RB units, one representing the binding *larger*(Fido), and one *smaller*(Sara) (Figure 2b). Each RB unit encodes one role-filler binding and shares bi-directional excitatory connections with one predicate and one argument (which can either be an object or a whole proposition, as detailed below). Above the RB units, a layer of proposition (P) units binds sets of RBs into multi-place relations. For example, the P unit encoding the relation *bigger* (Fido, Sara) has bi-directional excitatory connections with the RBs encoding *larger* (Fido) and *smaller* (Sara). All localist token units (POs, RBs, and Ps) laterally inhibit all units in the same layer (e.g., P units in the driver inhibit other driver Ps, RBs in the recipient inhibit other recipient RBs; driver and recipient are defined below).

Each level of the representational hierarchy serves an important purpose. The semantic units capture the semantically rich (i.e., distributed) nature of human mental representations. The three layers of localist units make it possible to treat each level of the hierarchy as an independent entity for the purposes of mapping and inference (Hummel & Holyoak, 1997, 2003). These localist units (i.e., POs, RBs, and Ps) are not different "kinds" of units (i.e., they do not work differently). Rather, they are simply units in different layers of the network. We use different names for the units in each layer (and different shapes in the figures) only to make them easier to distinguish.

RB units encode the bindings of specific roles to specific fillers conjunctively; likewise, P units bind col-

lections of role-filler bindings into full propositions in a conjunctive fashion. For the purposes of storage in (and retrieval from) LTM this kind of conjunctive coding is sufficient (even necessary; Hummel & Holyoak, 1997, 2003; Shastri, 1996). However, conjunctive coding violates role-filler independence (in the sense that a given role [or filler] is represented by different units depending on the filler [role] to which it happens to be bound), making it inadequate as a basis for binding in WM (Hummel & Holyoak, 1997; Hummel et al., 2004).

Dynamic Role-Filler Binding

When a proposition enters WM (i.e., when it becomes active), its role-filler bindings must be represented explicitly (and dynamically) on the units that preserve role-filler independence (i.e., the PO and semantic units). A common approach to dynamic binding in neural network models of perception and cognition is based on synchrony of role-filler firing (see Hummel et al., 2004, for a review): Units representing relational roles fire in synchrony with the arguments bound to those roles and out of synchrony with other role-filler bindings. For example, to represent *bigger* (Fido, Sara), the units representing the *larger* role fire in synchrony with those representing Fido while those representing *smaller* fire in synchrony with those representing Sara. Critically, the *larger* and Fido units must fire *out* of synchrony with the *smaller* and Sara units (Figure 3a). This approach to dynamic binding has the virtue that the bindings are represented both explicitly and independently of the units so bound: If Sara suddenly gained weight, the proposition *bigger* (Sara, Fido) could be represented by the very same semantic and PO units simply by reversing the synchrony relations. Another virtue of role-filler synchrony for binding is that it provides a natural *a priori* account of the limitations of visual attention and WM because only a limited number of RBs can be simultaneously active and still be mutually out of synchrony (Hummel & Holyoak, 2003; Morrison, Dumas & Richland, 2006; Morrison et al., 2004; Viskontas et al., 2004).

One limitation of binding by role-filler synchrony is that it implicitly assumes that predicates and objects are different "data types": Because object semantics fire at the same time as the predicate semantics to which they are bound, the only way to know whether a given unit represents an object feature or a predicate feature is to assume that the two are represented by separate (non-overlapping) pools of semantic feature units (see Hummel & Holyoak, 2003). As noted earlier, this "different data types" assumption, while convenient for many purposes, precludes learning attributes (such as "red" or "small") or relations (such as "same color" or "larger") from the properties of examples. Role-filler synchrony of firing is thus inadequate as a binding signal for our current purposes.

Figure 3. The firing of the proposition *bigger* (Fido, Sara) in DORA. (a) Binding by role-filler synchrony. The units for the *larger* (Fido) role binding fire followed by those for the *smaller* (Sara) role binding. (b) Binding by role-filler asynchrony. The units for the *larger* role fire followed by the units for Fido, then the units for the *smaller* role fire followed by those for Sara. (c) A time-series illustration of the firing of units during asynchronous binding. Each graph corresponds to one unit from (b) (specifically, the unit with the same name). The y-axis represents the unit's activation, the x-axis time. Prop = P unit.

Under binding by role-filler synchrony, roles fire in synchrony with the fillers to which they are bound and out of synchrony with other role-filler bindings (Figure 3a). Importantly, although role-filler bindings from the same proposition fire out of synchrony with one another, they fire in closer temporal proximity than bindings belonging to separate propositions.¹ This close temporal proximity plays a crucial role in disambiguating which role-filler bindings are part of the same proposition. This notion can be generalized to represent dynamic role-filler bindings while still permitting a single pool of units to represent the semantic features of both predicates and objects. The idea is that roles fire, not in synchrony with their fillers, but in close temporal prox-

imity (see Figure 3b and c). That is, role-filler binding is represented, not as temporal synchrony of firing, but as systematic *asynchrony* of firing (see Love, 1999). Although this approach may sound like the opposite of binding by synchrony, it is in fact a generalization of the same idea.

In DORA, role-filler bindings are represented by systematic asynchrony of firing (see Figure 3b and c). For example, to represent *bigger* (Fido, Sara), the PO unit for *larger*, along with its semantic features, fires followed by the units representing Fido, followed by the units for *smaller*, followed by the units for Sara. The binding of Fido to *larger* (rather than *smaller*) is carried by the fact that Fido fires in immediate temporal prox-

¹ For this reason, among others, it is necessary for a given role-filler binding to fire more than once. For example, if A represents "*larger*+Fido" and B "*smaller*+Sara", then the proposition *bigger* (Fido, Sara) would be represented by the sequence ABABAB. If C represents "*larger*+Bill" and D "*smaller*+Mary" then the propositions *bigger* (Fido, Sara) and *bigger* (Bill, Mary) would be represented by the sequence ABABABCDABCD; the propositions *bigger* (Fido, Mary) and *bigger* (Bill, Sara) would be represented ADADADCBBCB. Thus, temporal proximity signals which bindings belong together as parts of the same proposition.

imity to *larger* (as described more fully under *Establishing Asynchrony of Firing* below). For clarity of exposition, we describe this process as taking place with roles firing before their fillers (see Appendix A). However, none of DORA's learning routines (predicate learning, refinement, and multi-place relation formation) are affected by whether roles or fillers fire first.

Under asynchronous binding, information is carried by *when* units fire rather than *which* units are firing. As a result, there is no need to use different types of units (as in SME; Fagenheiner et al., 1989; Forbus et al., 1995; or STAR; Halford et al., 1998) or different sets of units (as in LISA; Hummel & Holyoak, 1997, 2003) to represent relations/relational roles and objects. Consequently, predicates and objects in DORA are coded with a common set of semantic units. The capacity to treat role and filler semantics equivalently and still specify their bindings dynamically makes all of DORA's other operations possible.

Representing higher-order relations

When a proposition takes another proposition as an argument, the P unit of the lower-order proposition serves as the argument under the appropriate RB unit of the higher-order proposition. For example, in the higher-order relation *causes* (gravity, *revolve-around* (earth, sun)), the P unit of the proposition *revolve-around* (earth, sun) serves as the argument of the *caused* role of the higher-order *cause* relation (see Figure 4). However, none of our simulations use this ability so we do not discuss it further.

Figure 4. Illustration of the structure of a higher-order relational proposition in DORA.

Flow of Control

Driver, Recipient, Emerging Recipient, and LTM

For the purposes of all the operations DORA performs, propositions are divided into four sets (Hummel

& Holyoak, 1997, 2003): a *driver*, a *recipient*, an *emerging recipient*, and *long-term memory (LTM)*. The driver contains the proposition or propositions that are the current focus of attention.² All of DORA's operations are controlled by the driver. Propositions in the driver pass activation to the semantic units. Because the semantic units are shared by propositions in all sets, activation flows from the driver to propositions in the other three sets. During retrieval, patterns of activation generated on the semantic units retrieve propositions from LTM into the recipient. Propositions in the recipient are available for mapping onto propositions in the driver. As detailed below, DORA then learns and refines explicit representations of the features common to the objects and predicates in the driver and the recipient, and can form higher arity structures from mapped lower arity structures. The new predicates, objects and propositions are encoded into the emerging recipient (equivalent to a *schema* in Hummel & Holyoak's, 2003, LISA), and may later be encoded into LTM and subsequently enter the driver or recipient.

Establishing Asynchrony of Firing

One at a time, P units (propositions) in the driver become active (if there are no P units in the driver, then the units in the next level of the hierarchy, here RBs, become active), exciting their RBs, which compete (via lateral inhibition) to become active. Similarly, RBs excite their POs, which also compete to become active. Each RB and PO is an oscillator consisting of an excitatory unit (an exciter) yoked to an inhibitory unit (an inhibitor). Exciters pass activation to their inhibitors and to any inhibitors in lower layers (i.e., RB exciters pass activation to PO inhibitors). Inhibitors (which have a long temporal integration period), in turn, inhibit their exciters. As a result, in response to a fixed excitatory input from some other unit (e.g., in the case of an RB input from a P unit), the activation of an excitatory unit will tend to oscillate (see Hummel & Holyoak, 1997). Groups of such oscillators that share mutually inhibitory connections (such as RBs and POs) will tend to oscillate out of synchrony with one another. PO inhibitors receive roughly twice the input of RB inhibitors (recall that they receive input both from their own PO and any active RBs). Therefore, POs oscillate at roughly twice the rate of RBs.

Collectively, these excitatory and inhibitory interactions have the following effects: When a P unit becomes active it excites its RBs, one of which becomes active. The active RB, in turn, excites its POs, one of which becomes active. When the inhibitor on the first PO becomes active, it inhibits that PO, allowing the second PO to become active. Likewise, once the active

² As detailed in Appendix A, within the driver, those propositions that are currently active constitute the *phase set*—the set of mutually desynchronized role-filler bindings (see Hummel & Holyoak, 1997). In general, the driver will contain too many propositions to fit all of them into the phase set simultaneously. However, in all the simulations reported here, the driver consists of only a single phase set. Therefore, for clarity, we discuss the model's operation in terms of the driver, rather than the phase set.

RB is inhibited, the next RB connected to the active P becomes active and excites its POs, one of which becomes active, followed by the second PO, and so forth.

In order to distinguish boundaries between separate RBs, DORA produces an inhibitory “refresh” signal (e.g., Horn, Sagi, & Usher, 1992; Horn & Usher, 1990; Hummel & Holyoak, 1997, 2003; von der Malsburg & Buhmann, 1992) at the level of role-bindings (RBs) when no RB exciter is active above a threshold (i.e., after a RB inhibitor fires). An analogous inhibitory signal is produced at the level of POs when no PO is active above threshold. The resulting pattern of activation on the semantic units is (for a two-role proposition): role1, /refresh/, filler1, /refresh//REFRESH/, role2, /refresh/, filler2, /refresh//REFRESH/, where /refresh/ is the inhibitory signal produced at the level of POs, and /REFRESH/ is the inhibitory signal produced at the level of RBs. (See Appendix A for details.) Although these operations may appear complex, they are all locally realizable as exchanges of excitatory and inhibitory signals between units.

Unit Activation

All token units in DORA update their activations according to the simple leaky integrator function,

$$\Delta a_i = \gamma_i(1.1 - a_i) - \delta a_i, \quad (1)$$

where Δa_i is the change in activation of unit i , γ is a growth rate, n_i is the net input to unit i , and δ is a decay constant. Activation is hard limited between zero and one. Net input, n_i , is a simple weighted sum:

$$n_i = \sum_j w_{ij} a_j, \quad (2)$$

where w_{ij} is the connection weight from unit j to unit i , and a_j is the activation of j . Semantic units compute their activations as,

$$a_i = \frac{n_i}{\max(n_i)}, \quad (3)$$

where n_i is input to unit i and $\max(n_i)$ is the maximum input to any semantic unit.

For simplicity, propositions fire in a random order. However, Hummel and Holyoak (1997; 2003) describe an algorithm that allows that network to determine the order in which propositions fire, based on constraints such as pragmatic centrality (Holyoak & Thagard, 1989), causal relations and other higher-order relations (see Hummel & Holyoak, 1997; Viskontas et al., 2004).

Retrieval from LTM

One of the most basic operations DORA performs is to retrieve a proposition or *analog* (situation, story, or event) from LTM given a driver as a cue. The model’s algorithm for retrieval also forms an essential part of its algorithm for mapping, comparison and predication. Retrieval in DORA is a form of guided pattern recognition (Hummel & Holyoak, 1997): Patterns of activation generated by the driver on the semantic units excite token units in LTM. Propositions and RBs in LTM become active to the extent that their roles and fillers are semantically similar to the patterns generated by the driver, and are retrieved into the recipient as a probabilistic function of their activation (see Appendix A).

Mapping

Analogical mapping is the processes of discovering which elements (objects, relational roles, whole propositions, etc.) of one analog correspond to which elements of another. In DORA, mapping guides the predication of new properties, the formation of new relations, and refinement of these properties and relations. Mapping is performed by the same guided pattern recognition that drives retrieval, augmented with the ability to learn mapping connections between coactive units in the driver and recipient. A collection of *mapping hypotheses*, h_{ij} , is generated for each unit, i , in the recipient and for every unit, j , of the same type in the driver (i.e., P units have mapping hypotheses for P units, RBs for RBs, etc.).³ At each instant in time, t , mapping hypothesis h_{ij} accumulates evidence that unit i corresponds to unit j using a simple Hebbian learning rule:

$$\Delta h_{ij}^t = a_i^t a_j^t, \quad (4)$$

where a_i and a_j are the activations of units i and j .

After the propositions in the driver have fired, the mapping hypotheses are used to update numerical weights, w_{ij} , on the mapping connections (Hummel & Holyoak, 1997; see Appendix A for details). The mapping connections serve to represent the correspondences between elements of the driver and recipient. Because mapping connections allow driver units to excite recipient units directly (as opposed to strictly through the semantic units), they allow mappings the model has discovered earlier in the mapping process to constrain mappings discovered later. Hummel and Holyoak (1997) demonstrated that this algorithm provides a natural account of the strengths and limitations of human analogical mapping. It also correctly predicted previ-

³ We assume that each mapping hypothesis is a unit that accumulates and compares evidence that a given unit, j , in the driver maps to a given unit, i , in the recipient, and that both mapping hypotheses and the mapping connection weights they will ultimately inform are not realized literally as synaptic weights, but rather correspond to neurons in prefrontal cortex with rapidly modifiable synapses (e.g., Asaad, Rainer & Miller, 1998; Fuster, 1997; Hummel & Holyoak, 1997, 2003).

ously unknown properties of human analogical mapping (Kubose, Holyoak & Hummel, 2002).

Mapping connection weights are constrained to take values between 0 and 1—that is, the mapping connections themselves are strictly excitatory. However, in addition to the excitatory signals transmitted over individual mapping connections, each driver unit also transmits a global inhibitory input to all recipient units of the same type (i.e., PO units to PO units, RB units to RB units; see Eq. 9 in Appendix A). This inhibitory signal, I_{ij} , is proportional to the activation of the driver unit a_i , the value of the recipient unit's largest in-coming weight, $\max(w_i)$, and the value of that unit's largest outgoing mapping weight, $\max(w_j)$:

$$I_{ij} = a_i(\max(w_i) + \max(w_j)), \quad (5)$$

As detailed shortly, this signal plays a crucial role in the model's self-supervised learning algorithm and thus in the discovery and predication of new properties and relations.

Comparison-based Predication

DORA performs comparison-based predication using a simple algorithm for intersection discovery. When DORA compares its two objects, for example, an elephant and a bear (Figure 5a), it will attempt to map them. After mapping, units in the driver will activate corresponding units in the recipient. Active units in both the driver and the recipient pass excitation to the semantic units. As a result, semantic units connected to coactive POs in both the driver and the recipient tend to get about twice as much input—and thus become about twice as active (recall Eq. 3)—as any semantics connected to only one PO or the other (see Figure 5b). The resulting heightened activation for semantics shared by both the driver and recipient POs serves to signal that those semantics are shared: That is, the semantic intersection of two POs is tagged as such by the units' heightened activations.

DORA exploits this activation-based intersection tagging in order to explicitly predicate the semantic properties shared by the coactive POs. Specifically, when DORA maps (i.e., compares) a solitary PO in the driver to a solitary PO in the recipient, it recruits a new PO unit in the recipient to learn the semantic intersection of the two compared objects. (Of course, we do not assume the brain grows a neuron to correspond to the active PO. Rather, we assume the new unit is recruited by the same process that recruits structure units at the time of encoding; see Hummel & Holyoak, 2003.) Newly inferred PO units have activation=1 and learn connections to all active semantics according to:

$$\Delta w_{ij} = a_i(a_j - w_{ij})\gamma \quad (6)$$

where Δw_{ij} is the change in weight between the new PO unit, i , and semantic unit, j , a_i and a_j are the activations of i and j , respectively, and γ is a growth rate parameter. By this equation, the weight between the new PO and a semantic unit asymptotes to that semantic's activation. Because semantics shared by the two compared POs tend to be about twice as active as semantics unique to one of the two compared POs, the new PO becomes approximately twice as strongly connected to the semantics shared by the compared POs than to semantics that are unshared (Figure 5c). The result is a new representation defined by the common features of the compared POs. The active PO units in the recipient signal DORA to recruit an empty RB unit (also in the recipient) that becomes active and learns connections to the two active PO units (the recipient object and the new PO) by simple Hebbian learning (Figure 5d). In short, comparison based learning causes DORA to represent the features shared by the compared objects as an explicit predicate (i.e., a PO unit) that is bound to one of the compared objects. If the shared features denote a categorical property (e.g., “red”, “big”, etc.), then the resulting PO will represent that property as a single-place predicate. If the shared features denote a relational property (e.g., “same-color”, “more-size”, etc.), then the resulting PO will again represent that property as a single-place (relational) predicate. As elaborated under *Forming Relations* (below), it is this latter kind of single-place predicate that is most likely to eventually serve as one role of a multi-place relation.

This algorithm recruits a new PO unit whenever two existing POs (predicates or objects) are explicitly compared to one another. Although the algorithm might appear to result in an explosion of PO units (one for every comparison ever made), it is at worst comparable to an exemplar-based model of category representation (e.g., Nosofsky, 1988). According to exemplar models, a category is represented in memory by storing, individually, all e exemplars of that category. If there are c categories, each with (on average) e exemplars, then there will be $c * e = t$ total exemplars in memory. In the very worst case, that is, if DORA compares *every* exemplar of *every* category with every other exemplar of every category, then the number of first-order (i.e., object-object) comparisons is t^2 . If it compares the results of every one of those comparisons with every other result of those comparisons, then the number of second-order (i.e., predicate-predicate) comparisons is still only $t^{2*2} = t^4$. In other words, under the very worst scenario in which DORA compares *absolutely everything* it knows to *absolutely everything else* it knows, the number of POs resulting from these comparisons grows only as a polynomial function of the number of objects, t , in its memory.

Figure 5. Comparison based predication. (a) An elephant and a bear are compared (for the purposes of clarity, mapping connections are not depicted in the figure). (b) Semantics shared by both the elephant and bear become more active than unshared semantics. (c) A new PO (here a triangle to differentiate it from other active PO units) is recruited and learns connections to the active semantics in proportion to their activation. (d) The new unit is bound to the bear unit via an RB unit. Solid lines = strong connections, dashed lines = weaker connections, gray = more active units.

However, even this kind of polynomial growth is a gross overestimate of the number of POs produced by DORA’s comparison algorithm. The polynomial growth requires DORA to compare everything it knows to everything else (e.g., compare this chicken to that bowl of oatmeal you had last Tuesday, and to that hot air balloon, and to your toenail, etc.). Obviously, no child will compare everything she knows to every other thing she knows. As elaborated in the General Discussion, there must be strong psychological constraints on when objects (or predicates or relations) will be compared to one another, and these constraints will greatly reduce the number of comparisons that get made, and thus the number of POs that get recruited. Candidate constraints include things such as verbal labeling, and explicit direction (e.g., from an adult) to compare objects. With such constraints in hand, the proliferation of POs becomes much less problematic. Moreover, we do not assume that, once recruited, a PO becomes a permanent part of a child’s LTM. Especially if a PO is rarely used or rarely participates in additional comparisons, it is reasonable to assume that the neurons recruited to represent it will eventually be re-recruited to represent new things.

Predicate Refinement

The first time DORA learns a new predicate, that predicate will almost invariably contain a number of extraneous semantic features. For example, in the case of predicating the concept “big” by comparing an elephant to a bear (Figure 5), the resulting *big* predicate will also be connected strongly to any other properties the elephant and bear in question happen to share (e.g., “animal”), as well as weakly to all their unique features. That is, newly-predicated properties will initially be object-specific. Likewise, the second time DORA predicates “big”, for example by comparing a house to a truck, the resulting predicate will also be object-specific and contain extraneous semantics that the house and truck happen to share (e.g., “rectangular”). However, the same intersection-discovery algorithm that causes DORA to start to extract “big” from elephant and bear, and from house and truck, when coupled with the capac-

ity for self-supervised learning (SSL; Hummel & Holyoak, 2003), allows DORA to further refine its understanding of “big” by comparing its first (elephant and bear) *big* predicate to its second (house and truck) *big* predicate. When DORA compares its two “big” predicates-object pairs, by comparing, say, *big* (Bear) to *big* (House) it will first map them (Figure 6a). Consequently, units in the driver will activate corresponding units in the recipient. In a second (initially empty) emerging recipient DORA recruits new units to correspond to active units in the driver (Figure 6b). DORA uses the mapping-based global inhibition as a cue for detecting when a unit must be recruited in the emerging recipient. If a unit, i , in the recipient maps to some unit, j , in the driver, then i will be globally inhibited by all other units, $k \neq j$, in the driver (recall Eq. 5). Therefore, if some unit, k , in the driver maps to nothing in the emerging recipient (or if the emerging recipient is empty), then when k fires, it will inhibit all emerging recipient units and it will excite no units in the emerging recipient (see Hummel & Holyoak, 2003). This kind of global mapping-based inhibition, unaccompanied by any mapping-based excitation, serves as a reliable cue that nothing in the emerging recipient analog corresponds to driver unit k . It therefore causes DORA to recruit a new unit to correspond to k . As during comparison-based predication, DORA learns connections between active semantics and new POs by Eq. 6, and between active corresponding token units (i.e., between POs and RBs, and between RBs and Ps) by simple Hebbian learning (Figure 6c). This occurs for each driver PO that fires (Figure 6d and e). So, when DORA compares *big* (Bear) to *big* (House), the refined representation of the predicate *big* will have a connection weight of 1.0 to the shared semantics of the two *big* predicates (i.e., the semantic feature “big”), and a connection weight of roughly 0.5 to extraneous semantics connected to only one of the two “dirty” *big* predicates (e.g., the semantics “animal” and “rectangular”). Applied iteratively, this process produces progressively more refined representations of the compared predicates, eventually resulting in a “pure” representation of the property or relation, free of extraneous semantics.

Figure 6. Schema based predicate refinement. (a) The two propositions are mapped (semi-dashed lines indicate mapping connections; for the purposes of clarity, mapping connections are not depicted in the remaining components of the figure). (b) The two mapped *big* predicates fire. New PO and RB unit are recruited to match the firing PO and RB in the driver. (c) The new PO learns a connection to the new RB, and learns connections to active semantics in proportion to their activation. (d) The two mapped objects fire. A new PO is recruited to match the firing PO in the driver (an RB to match the firing RB was learned in (b)). (e) The new PO learns connections to the RB and to active semantics in proportion to their activation. The result is a schematized “cleaner” representation of the property *big*. Full lines = strong connections, dashed lines = weaker connections, darker gray = more active units.

Importantly, this same process works both on single role-filler sets and on multi-place relational structures composed of multiple role-filler sets (see below). Over the course of several comparisons, the algorithm retains the invariant semantic features that define a property, relational role or whole relation. For instance, if DORA compares many instances of the *bigger* relation, it will learn a representation of the invariant properties of the roles of *bigger*. Moreover, if DORA compares two relations like *bigger* and *higher* it will learn a representation of the invariant properties of the roles of those two relations, in this case the relation *more-than*.

An important prediction of this algorithm is that features will be hard to purge from a concept to the extent that they are shared by many examples of that concept. For example, if a child always compares objects in which the larger object is red and the smaller blue, then the algorithm predicts that, for that child, “red” will be a feature of “larger-than” and “blue” a feature of “smaller-than”. Fortunately, in most children’s experience, color is not systematically confounded with size. However, examples of this kind of confounding do occur in children’s experience. For example, most things that sleep also happen to have eyes. The algorithm thus predicts that “has eyes” is likely to be included, at least early on (i.e., before counter examples), in the concept “sleeps”. There is substantial evidence that children have difficulty purging widely present features from their concepts. For example, children tend to think that things that sleep have eyes, and that things that move have legs (e.g., Sheya & Smith, 2006). DORA’s inter-

section discovery algorithm thus provides a natural account of such early errors of inclusion.

The algorithm also provides a natural account of children’s ability to ignore omnipresent features—such as “has boundaries” and “is smaller than the universe”—that are not simply confounded with *some* other features (in the way that “has eyes” is confounded with “sleeps”), but instead are confounded with everything. Inasmuch as such features are ever represented at all (which we doubt they are until adulthood), they will have exactly zero effect on DORA’s behavior, precisely because they are omnipresent. The reason is that although adding a constant feature to every object will increase the pair-wise similarity of every object, x , to every other object, y , it would also increase the pair-wise similarity of every pair, (x, y) , by the same amount, for all x and all y . As a result, no feature that appears on every object will have any effect on DORA’s performance.

Forming Relations

Comparison allows DORA to discover the shared semantic features of separate objects and represent them as explicit predicates that take arguments (e.g., *big* (Bear)). However, it is one thing to learn that one object is big and another is small, but it is another thing to learn the relation *bigger-than* and predicate it over multiple arguments.

Learning a relation that takes multiple arguments entails solving two problems, each of which is related to the problems entailed in learning single-place predicates. The first problem, hinted at previously, is that the

invariants that describe relational roles differ from those that describe simple object properties. For example, “bigger-than” is not the same as “big”. If an object is bound to a relational role specifying “bigger-than”, then some other object must be bound to the complementary role specifying “smaller-than”. In other words, deciding which of two objects is bigger and which smaller requires a capacity to compare the two.

Hummel and Biederman (1992; see also Hummel & Holyoak, 2002) describe a simple neural *comparator* circuit that takes, as input, representations of numerical values along a dimension (such as size) and as output returns relative values along that dimension (such as “bigger” and “smaller”). For example, given “size-5” and “size-8” as input, Hummel and Biederman’s comparator will activate the feature for “bigger” in response to “size-8” and the feature for “smaller” in response to “size-5”. We have adapted this comparator circuit for use in DORA. Whenever DORA activates two POs in the driver simultaneously, and those POs both have features describing values along the same metric dimension (e.g., size, color, height, etc.), DORA invokes the comparator. Similar to Hummel and Biederman’s comparator, DORA’s comparator activates the semantic unit “more” in synchrony with the larger value along the dimension, the semantic unit “less” in synchrony with the smaller, and the semantic unit “same” if the two values are equal. Following Hummel and Biederman our comparator thus assumes implicit knowledge of dimensions (an assumption that is not unreasonable, see Feigensohn et al., 2002). These semantics are then connected to the active POs by Eq. 6. The resulting “more”, “less” and “same” semantics, along with the semantics describing the dimension itself (e.g., “size- x ”, etc.), provide the relational semantics for DORA’s emerging relational predicates.

Given a set of relational semantics bound to individual roles, the next problem is to link those roles together into a representation of a multi-place relation. DORA solves this problem by exploiting the temporal dynamics of binding by asynchrony. For example, if DORA thinks about a dog (size-6) and a cat (size-4) at the same time, then the comparator will attach “more” to “size-6” and “less” to “size-4”. The resulting role-argument bindings (i.e., *more+size-6* (dog) and *less+size-4* (cat)) may remind DORA of a previously experienced comparison, say, *more+size-9* (bear) and *less+size-5* (fox), permitting DORA to map the dog and the cat onto the bear and the fox, respectively, along with their relational roles. The result of this mapping is a distinct pattern of firing on the units in the recipient: The units encoding *more+size-6* (dog) and *less+size-4* (cat) oscillate out of synchrony in the driver, and, through their mapping connections, impose the same systematic oscillatory pattern on *more+size-9* (bear) and *less+size-5* (fox) in the recipient (as illustrated in Figure 7a-d). Units in the

recipient will exhibit this systematic oscillatory firing pattern only under two circumstances: When the role-filler pairs in question are already bound into a single relational structure (via a P unit), or when similar sets of role-filler pairs are in WM simultaneously and are mapped as a group (i.e., when they could be bound into a single relational structure; as in the current example). Importantly, this firing pattern is not something that DORA must learn; it is a pattern that emerges as a natural consequence of binding using time. This temporal firing pattern, coupled with DORA’s self-supervised learning algorithm, bootstraps the learning of multi-place relational structures. During relation formation, when an RB in the recipient becomes active, a P unit is recruited by SSL (also in the recipient) if no other P units are active (Figure 8a). The new P unit remains active until the Ps (or RBs if there are no Ps) in the driver have all fired, and learns connections to active RBs as during SSL (i.e., via Hebbian learning; Figure 8b-d). The result is a P unit linking the oscillating role-filler pairs in the recipient into a multi-place relation. The resulting proposition explicitly codes the relation *bigger-than* (bear, fox). In summary, DORA exploits the oscillations of the RB units in the recipient to form a new relation: RBs that fire sequentially become component roles of the new relational structure (encoded by the P unit), and the completion of firing of the driver Ps (or RBs) marks the end of the relational structure.

As during predicate learning, the resulting relational roles will initially be specific to the values from which they were learned (e.g., the roles of the resulting *bigger-than* relation will be connected to “size-9” and “size-5”). However, the predicate refinement routines described above, applied iteratively to different examples of the *bigger-than* relation, will eventually yield a value-independent representation of *bigger-than*.

Order of Operations

The order in which DORA performs the various operations described above (e.g., predication, refinement, etc.) is determined only by the state of DORA’s knowledge. If there is a proposition (or a pair of propositions) in the driver but none in the recipient, then DORA first attempts to retrieve a similar proposition (or propositions) from LTM. Given one or more propositions in both the driver and recipient, DORA next attempts to map the driver proposition(s) onto the recipient proposition(s). After mapping DORA initiates learning: If there are no mapped RBs in the driver, or if the mapped RBs are connected to only one PO unit (a situation that indicates that the active objects not bound to any relational roles), then DORA performs comparison-based predication. After one or more roles have been predicated, DORA attempts to form new relations, and refine its existing representations. DORA followed this order of operations in all the simulations reported here..

Figure 7. When one set of single-place predicates is mapped to another set of single place predicates a diagnostic firing pattern necessarily emerges. Role-filler sets in the driver oscillate out of synchrony and impose the same systematic oscillatory pattern on the recipient units to which they map (darker units = more active units). For example, when the set *more+size-6* (dog) and *less+size-4* (cat) are mapped to *more+size-9* (bear) and *less+size-5* (fox) the following firing pattern emerges: (a) First the PO unit encoding *more+size-6* in the driver and its RB unit fire, which excites the PO unit encoding *more+size-9* and its corresponding RB in the recipient. (b) Then the objects bound to those predicates (dog and bear) and their corresponding RBs fire. (c) Next the mapped POs encoding *less+size-4* and *less+size-5* and their corresponding RBs fire. (d) Then the objects bound to those predicates (cat and fox) fire and their corresponding RBs fire.

Figure 8. Forming a multi-place relational structure by mapping sets of single place predicates (darker units = more active units). DORA uses the systematic time-sharing of RBs that is imposed on units in the driver by units in the recipient (see Figure 7 and text) and recruits a new P unit in the recipient. As the mapped POs and RBs fire in sequence the new P unit learns connections to active RBs. (a-b) As the *more+size-9* PO, the bear PO and the RB encoding their binding fire, the model learns a connection between the active RB and P unit by Hebbian learning. (c-d) Similarly, as the *less+size-5* PO, the fox PO and the RB encoding their binding fire, the model learns a connection between the active RB and P unit by Hebbian learning. The result is a full relational structure encoding *bigger-than* (bear, fox).

Simulations

In this section we describe a series of simulations demonstrating DORA's ability to account for a range of empirical phenomena. Table 2 summarizes a set of core empirical phenomena that we take to be well enough established that a theory of the development of relational concepts must be able to account for them.

This section is organized as follows: First, we present a set of basic simulations demonstrating DORA's ability to learn simple structured relational representations from unstructured holistic inputs. We demonstrate that the resulting representations meet the joint requirements of structure sensitivity and semantic richness (Hummel & Holyoak, 1997) by demonstrating that they allow DORA to make relational analogies. Following this basic demonstration, we use DORA to

simulate each of the major empirical phenomena summarized in Table 2.

Basic Demonstration: Learning a Relational Concept from Examples and Evaluating the Resulting Representations

People acquire relational concepts from examples (Phenomena 1 and 2). Beginning with representations that are holistic (Phenomenon 3) children gradually learn relational representations that are both structure sensitive and semantically rich (Phenomenon 4). Here we present a series of simulations in which DORA learns relations from examples (e.g., DORA learns the relation *bigger-than* (x, y) from objects of different sizes). We then test whether the resulting representations meet the requirements of structure sensitivity and

Table 2.

Phenomena Associated With Relation Discovery*General phenomena*

1. People acquire relational concepts.
2. The ability to reason relationally develops.
3. Younger children's representations are holistic.
4. Older children and adult relational representations are both structure sensitive and semantically rich.
5. Relational concepts can be learned from interleaved examples.
6. The types of problems a reasoner can solve and the properties of the solutions depend critically on what the reasoner can and does represent.
7. Children go through a domain-specific *relational shift*, in which they transition from representing a domain in terms of its characteristic features to representing it in terms of its characteristic relations (for most domains generally around 4-5 years of age).

Specific phenomena

8. Children's relational and dimensional representations are initially holistic and tied to specific object values.
9. Over time children's relational and dimensional representations become progressively more abstract.
10. Children's initial relational representations are categorical rather than relative.
11. The trajectory of children's appreciation of similarity develops from appreciation of over-all identity, to appreciation of shared properties, to appreciation of shared relations.
12. Increased knowledge leads to more successful analogical performance.
13. Comparison helps children learn structured representations.
14. Progressively-aligned training assists in discovering relational similarities.
15. Relational learning occurs without explicit feedback.
16. Comparison helps adults learn relational rules.

semantic richness by using them to simulate some fundamental aspects of human relational reasoning (viz., the ability to solve cross-mappings, the ability to map non-identical but semantically similar relations, and the ability to violate the n-ary restriction; see Hummel & Holyoak, 1997). The purpose of these simulations is to provide a basic test of DORA's ability to learn structured relational representations from unstructured examples.

Learning Relational Concepts from Examples

Our initial simulation tested DORA's ability to learn relational representations from unstructured representations of objects. DORA started with nothing but representations of individual objects of various sizes. By applying its predication and comparison routines, it first learned to represent specific sizes as explicit (initially "dirty") predicates, then learned progressively more refined size predicates, then combined these predicates into value-specific representations of *bigger-than*, and finally compared these predicates to learn value-independent representation of *bigger-than*.

We ran two versions of these basic simulations. We ran the "perception" version as though DORA were comparing pairs or collections of objects it was "looking at". Thus, on each simulation, we gave it two or more objects to compare and allowed it to run its routines in the order described above. The advantage of this approach is that it gave us control over the comparisons DORA made. The disadvantage is that it required us to "tell" DORA which comparisons to make on each simulation, and thus effectively, which of its operations to run in what order.

In the "memory" version of the simulations, we randomly chose one or more objects for DORA to "look at" (i.e., we randomly activated objects, objects with properties predicated about them, or pairs of objects with properties or relations predicated about them) and allowed it to retrieve an object or objects from LTM to compare to the presented object(s). The advantage of the "memory" procedure is that it does not require us to tell DORA which comparisons to make, or which operations to perform in what order. Whereas the "perception" simulations entail a degree of hand-holding on the part of the modeler, the "memory" simulations allow us to observe DORA acting as an autonomous agent. The disadvantage of this approach is that it provides less control over the comparisons DORA makes and thus makes it more difficult to analyze the model's behavior. As detailed below, the model's performance was qualitatively very similar in both sets of simulations except that, not surprisingly, it required more comparisons to arrive at "clean" representations of *bigger-than* in the "memory" simulations than in the "perception" simulations. We assume that a child's experience is some combination of these approaches, with some comparisons resulting from the child viewing multiple objects simultaneously, and others resulting from the child being reminded of a comparison object by virtue of experiencing a similar object.

In both the "perception" and "memory" versions of the simulations, DORA started with objects. Each object was connected to two semantic units describing its size—specifically, "size" plus one specifying its specific size, e.g., "size-5"—plus eight other semantic units randomly chosen from a pool of 150 (see Appendix B for details).

For the purposes of analyzing DORA’s developing representations, we defined a *selectivity metric* (SM), which quantifies the degree to which a PO is selective for the semantics specifying the property or relation deemed relevant in the simulation (in this case, semantics for size and relative size). For predicates representing single-place attributes (such as specific sizes), the SM for unit i was calculated as the mean connection weight between i and all the relevant semantics to which it was connected, j , normalized by the mean connection weight between i and all the irrelevant semantics to which it was connected, k :

$$SM_i = \frac{MEAN(w_{ij})}{1 + MEAN(w_{ik})}. \quad (7)$$

One is added to the denominator to keep the SM ratio between 0 and 1: As the weights on a PO’s connections to relevant semantics approach 1, and the weights on its connections to irrelevant semantics approach 0, its SM approaches 1. For multi-place relations (such as *bigger-than*), the SM of the relation as a whole was calculated simply as the mean SM of the POs representing its roles.

For the purposes of the SM, we designated the size-related semantic units to be “relevant” (i.e., “size”, “more” and “less” were relevant for the learning of *bigger-than* as a relation) and all the other semantics to be “irrelevant.” However, it is important to stress that this designation of “relevant” and “irrelevant” semantics—as well as the naming of semantics as “size”, “more” etc.—are only for the purpose of analyzing the model’s behavior and for clarity of presentation. As noted previously, we make no claims about the actual semantic content of the concepts of size or relative size. Indeed, the details of that content, and the names of the semantic units, do not matter to DORA.

“Perception” simulations. In the “perception” simulations, DORA started with 160 (80 pairs of) objects constructed as described above. We first allowed DORA to compare objects of similar sizes and learn new predicates via comparison-based predication (as described above). Each object was randomly paired with another object of the same size. For example, DORA might compare object5 (with features “size”, “size-2”, “sweet”, “round”, “red”, “fruit”, etc.) to object37 (with features “size”, “size-2”, “alive”, “green”, “reptilian”, etc.), resulting in a predicate connected strongly to semantics the two items had in common (here “size” and “size-2”) and weakly connected to the semantics connected to only one of the objects. DORA’s predication algorithm then bound this new predicate to the objects from which it was discovered Table 3.

(recall Figure 5), producing the proposition *size-2* (object37). We ran this comparison process on 80 pairs of objects, resulting in 80 predicates coding for various object sizes, each bound to a specific object (i.e., the object in the recipient as per comparison-based predication). However, the resulting size predicates were comparatively “dirty” (i.e., still largely connected to the irrelevant—i.e., non-size-related—features of the objects from which they were predicated) with a mean selectivity of 0.33; see Table 3. We then let DORA compare these “dirty” predicate representations and refine them as described above. Each of the 80 predicate-object bindings DORA learned during the previous simulation was randomly paired with another predicate-object binding describing the same size. Each pair was then compared and refined. For example, DORA might compare the role-binding *size-2* (object7) to *size-2* (object37), resulting in a more-refined representation of *size-2*. The resulting representations had higher selectivity than the previous representations (0.40 vs. 0.33; Table 3, Row 3 vs. 2), indicating that successive comparisons allow DORA to learn progressively more selective representations of object attributes.

Once DORA had learned more selective predicates for specific sizes, we allowed it to compare them by placing pairs of predicate-object bindings in the driver at the same time and invoking the comparator (as described above). For example, to compare *size-5* (object1) to *size-8* (object2), *size-5* fires, followed by object1, followed by *size-8*, followed by object2, and the semantic patterns generated by *size-5* and *size-8* are passed to the comparator as input. As output, the comparator activates “less” in synchrony with *size-5* and “more” in synchrony with *size-8*. The PO representing *size-5* then learns a connection to the semantic “less” (by Eq. 6) and the PO for *size-8* learns a connection to “more”. The resulting POs represent, respectively, “size 5 and less size than something” and “size 8 and more size than something”.⁴ During the previous simulation DORA learned 40 refined predicates each bound to an object. We paired each of these refined predicates-object bindings with a second predicate-object binding representing a different size. For example, *size-3* (object1) might be paired with *size-7* (object2), producing 20 pairs of predicate-object bindings that DORA ran through the comparator.

Once the pairs of predicate-object bindings had been run through the comparator, DORA formed multi-place relational structures by comparing pairs of predicate-object bindings to other pairs. Each of the 20 pairs of predicate-object bindings that had been run through the comparator was randomly matched with another. For example, the predicate-object pair *size-7+more* (object1)

⁴ Recall that we are not making claims about how the relevant semantics are detected, only about how, once they are detected, they can be isolated and represented explicitly. While the comparator makes the relevant semantics active in the system, it is the predication and refinement routines that serve to organize these semantics into explicit structured relational representations.

Mean attribute and relation selectivity (see text) of the representations that resulted from each comparison during “perception” simulation. Dashes indicate undefined values.

	Attribute selectivity	Relation selectivity
Initial state: holistic object representations	--	--
Compare objects to objects. Result: explicit predicates	0.33	--
Compare pairs of predicates. Result: refined predicates	0.4	--
Run comparator. Result: Values of “more” and “less”	0.71	--
Compare sets of predicate-object bindings. Result: relations	0.71	0.71
Compare pairs of relations. Result: refined relations	0.85	0.85
Compare pairs of relations again. Result: refined relations	0.91	0.91

and *size-3+less* (object2) might be matched with *size-6+more* (object3) and *size-2+less* (object4). One pair was placed in the driver and the other in the recipient and DORA mapped them. Each time DORA mapped one pair of predicate-object bindings to another, it invoked its relation discovery routine, inferring a P unitconjoining the predicate-object pairs in the recipient. That is, DORA learned explicit, two-place representations of the relation *bigger-than*. For example, if it mapped *size-7+more* (object1) and *size-3+less* (object2) in the driver to *size-6+more* (object3) and *size-2+less* (object4) in the recipient, then it would learn a P unit binding *size-6+more* (object3) to *size-2+less* (object4), forming an explicit, albeit value-specific, representation of the relation *bigger-than* (object3, object4) (with *size-6+more* as one role of the relation and *size-2+less* as the other). We ran 10 such comparisons, resulting in 10 new representations of the *bigger-than* relation. At this point, because DORA had representations of multi-argument relations, relational selectivity could be calculated. As shown in Table 3 row 5, these first relational representations are quite value-specific: Instead of representing “bigger than” as an abstract relation, DORA’s first relational representations correspond to concepts such as “size 8 bigger than size 5”.

Next, we had DORA compare sets of the relations it had learned in the previous simulation, and refine them via its predicate refinement routine. The 10 representations of the *bigger-than* relation that DORA had learned during the previous portion of the simulation were randomly paired (creating five pairs of relations). DORA then compared each pair of *bigger-than* relations. The resulting representations had greater relational selectivity than DORA’s first relational representations, indicating that the model was learning more idealized (i.e., value-independent) representations of the relations (.85 vs. .71; see Table 3 row 6). Applied iteratively, this algorithm allows DORA to learn progressively less value-laden representations. For example, one more iteration through the learning algorithm produced representations with selectivity scores of .91 (Table 3 row 7). This refinement takes place because each successive comparison reduces connection weights to features not shared by the relations (i.e., irrelevant semantics) by half, while weights to relevant semantics (i.e., the relational invariants) remain at 1.

In summary, these “perceptual” simulations demonstrated that, beginning only with holistic representations of objects of various sizes, DORA can, by comparing these objects to one another, predicate the *size* attribute in an object-dependent fashion. By comparing these predicates to one another, the model progressively refines its representations of various sizes. And by comparing different sizes to one another (using the comparator) and comparing pairs of different-sized objects to one another, DORA discovers, and gradually refines, the *bigger-than* relation. Moreover, in this simulation DORA learned representations of *bigger-than* with a very high SM with only 5 comparisons, demonstrating that DORA can learn value-independent multi-place relational representations with only a few comparisons. In short, these basic simulations demonstrate that DORA’s comparison, intersection discovery, predication and relation formation operations can work together to learn relational concepts from examples.

“Memory” simulations. Although they provide a proof of concept, the “perceptual” simulations can be criticized on the grounds that we told DORA which comparisons to make—and thus, effectively, which operations to run—in what order. We next ran the “memory” version of the same simulations to determine whether the routines DORA uses to discover predicates and relations would produce analogous results without such explicit hand-holding on the part of the modeler.

For these simulations, DORA started with 50 holistic representations of objects, constructed exactly as in the previous simulations. On each simulation we randomly chose a structure from the collection in its LTM. We then activated the randomly chosen structure in the driver (as though DORA was “looking at” or “noticing” the activated structure). We allowed DORA to run from there. Following its order of operations, DORA used the representation in the driver to retrieve representations into the recipient, then mapped the representation in its driver to those in its recipient, and then attempted to learn from the results (i.e., to perform comparison-based learning, relation formation, and refinement, as described in the *Order of Operations* section above). As in the previous simulation, DORA invoked the comparator circuit when multiple predicates describing values along the same dimension (e.g., size, color, etc.) were

active simultaneously. In other words, rather than “telling” DORA what to compare to what, we allowed it to make comparisons on its own.

As DORA runs its comparison and predication routines, the number of structures in its LTM grows, since each comparison/predication results in a new predicate-object binding or, in the case of comparing sets of predicate-object bindings, a new relation. We suggest that the same is true of a child, for whom not only every new experience results in a new structure in LTM, but every comparison of a new experience with an old experience may likewise result in a new structure in LTM. However, a child’s experiences are unlikely to all be equally salient. More recent events are likely to be more frequently thought about and more commonly retrieved than older events (Ebbinghaus, 1885/1913; Thorndike, 1914). We incorporated this fact in our “memory” simulations in a simplified manner in order to minimize the assumptions embodied therein.⁵ We manipulated the probability that DORA would notice a particular kind of thing (i.e., that an item from LTM would be activated in the driver) as a function of how recently it had been learned. DORA would notice an item learned during the previous 50 simulations with a probability of .75, and it would notice an older item with a probability of .25. Similarly, during retrieval DORA was more likely to retrieve recently learned items into the recipient (probability .75 vs. .25). In addition, because during this simulation DORA could compare items of the same size and thus learn the *same-size* relation, the “same” semantic was included as a relevant semantic for the purposes of calculating SM.

After every 50 simulations we tested the SM of the representations DORA learned during those 50 simulations. These data are presented in Table 4. Beginning with holistic representations of objects, DORA first learned single-place predicates and subsequently multi-place relations (*bigger-than*, and *same-size*) that became progressively more refined with subsequent comparisons. The “memory” simulation demonstrates that DORA’s routines for retrieval, mapping, and learning allow it to discover and predicate structured relational concepts from examples without being told which comparisons to make. Importantly, allowing DORA to run on a set of examples produces the same trajectory observed during the “perception” simulation described previously: DORA progresses from holistic objects, to “dirty” single-place predicates, to progressively more refined predicates and “dirty” relations, to progressively more refined relations. Finally, it is important to note that during the memory simulation DORA learned value-independent multi-place relations with only a few comparisons per relation. During the 300 comparisons DORA learned 50 value independent multi-place relations (with SM = .92; see Table 4).

In other words, each value independent multi-place relation took an average of 6 comparisons to learn.

Learning multiple relations from interleaved examples. For simplicity and clarity, most of the simulations reported in this paper were run in a “blocked” fashion, in the sense that DORA was charged with learning only a single relational concept at a time. Although children experience blocked practice with some concepts (e.g., at school, or with their parents), it is generally not the case that children go around the world mastering one concept (e.g., “red” or “bigger-than”) before moving on to the next (Phenomenon 5). It is therefore important to know whether DORA, like children, can discover relational concepts when different kinds of examples are presented in an interleaved fashion. From a technical perspective, it is important to know whether DORA’s learning algorithm will suffer catastrophic interference (e.g., of the type suffered under some conditions by models trained using error back-propagation; see McClosky & Cohen, 1989) when its examples are presented in interleaved fashion. If it does, then this interference would be an extremely important—even devastating—limitation of the model’s learning algorithm.

To test the model’s ability to learn relational concepts in an interleaved fashion, we gave it 100 objects whose features presented examples of four different concepts, namely: size (from which DORA learned *bigger-than* and *same-size*), width (from which DORA learned *wider-than* and *same-width*), height (from which DORA learned *higher-than* and *same-height*), and color (from which DORA learned *same-* and *different-color*).

These objects were created as in the previous simulations with the single difference that each object was attached at random to one dimension—size, width, height, or color—and with a probability of .25 to every other dimension. No object served as an example of every concept, but most objects served as examples of more than one. We then ran the model in the same way as the previous “memory” simulation (see Appendix B). Armed with its repertoire of holistic objects, we randomly activated objects (or, later, role-object bindings, pairs of role-object bindings or whole relations), and let DORA run. As in the previous “memory” simulation we manipulated the probability that DORA would “notice” a particular kind of thing as a function of how recently it had been learned ($p = .75$ for an item learned during the previous 100 simulations and $p = .25$ for an older item). Similarly, during retrieval DORA was more likely to retrieve recently learned items into the recipient (again, $p = .75$ and $p = .25$). During this simulation DORA could compare items of the same size, width, height or color, so the “same” semantic was included as

⁵ It is important to note that the model’s behavior is robust to violations in this assumption. Without this assumption the model follows exactly the same learning trajectory, it simply requires more iterations.

Mean attribute and relation selectivity (see text) of the 50 representations that resulted from each set of 50 comparisons during single-concept “memory” simulation. Dashes indicate undefined values.

	Attribute selectivity	Relation selectivity
Initial state: holistic object representations	--	--
50 Simulations	0.33	--
100 Simulations	0.63	0.63
150 Simulations	0.74	0.76
200 Simulations	0.8	0.81
250 Simulations	0.87	0.88
300 Simulations	0.91	0.92

Table 5.

Mean attribute and relation selectivity (see text) of the 100 representations that resulted from each set of 100 comparisons during multiple-concept “memory” simulation. Dashes indicate undefined values.

	Attribute selectivity	Relation selectivity
Initial state: holistic object representations	--	--
100 Simulations	0.32	--
200 Simulations	0.58	0.61
300 Simulations	0.72	0.74
400 Simulations	0.83	0.85
500 Simulations	0.89	0.9
600 Simulations	0.92	0.93

a relevant semantic for the purposes of calculating SM.

After every 100 simulations we tested the SM of the representations DORA learned during those simulations. These data are presented in Table 5. As in the previous simulations, beginning with holistic representations of objects DORA first learned single-place predicates and subsequently whole multi-place relations that became progressively more refined with subsequent comparisons.

These simulation results illustrate four very important points. First, DORA learns relational structures from interleaved examples (i.e., from training that more accurately reflects real world experience). In other words, DORA does not suffer catastrophic interference when it learns multiple concepts simultaneously. Second, DORA can learn value independent representations of multi-place relations from objects that are involved in multiple relations simultaneously. That is, DORA can learn a relation like *above* from examples of objects that are involved in the *above* relation as well as others. DORA’s learning routines isolate the features that define a specific relation (such as *above*) and learn explicit representations of that relation, even in the presence of competing relational features. Third, the same learning trajectory observed during the previous “perception” and “memory” simulations is observed when DORA learns multiple concepts simultaneously from interleaved examples: DORA progressed from holistic objects, to “dirty” single-place predicates, to progressively more refined predicates and “dirty” relations, to progressively more refined relations. The interleaved simulation (as well as those reported above) demonstrates that this trajectory is a fundamental consequence of

DORA’s learning algorithm, rather than an artifact of the manner in which we ran any particular simulation. And fourth, even when learning from interleaved examples, DORA learned value-independent multi-place relations with only a few comparisons. During the 600 comparisons DORA learned 100 value independent multi-place relations (with mean SM = .93; see Table 5), indicating that each value independent multi-place relation took an average of 6 comparisons to learn. This result demonstrates that DORA’s learning algorithm scales well in that it requires only a few comparisons to learn value-independent multi-placed relations, even when it is learning many relations concurrently.

We now turn our attention to the question of whether the representations DORA learns are genuinely relational, in the sense of supporting the basic operations necessary for relational reasoning (such as analogical mapping).

Evaluating the resulting representations

The simulations reported in this section tested whether the relations DORA learned in the previous simulations satisfy the structure sensitivity and semantic richness necessary to account for key aspects of relational reasoning (Hummel & Holyoak, 1997; see also Doumas & Hummel, 2005).

A stringent test of the structure sensitivity of a representation is its ability to support finding correct relational mappings in the face of object cross-mappings. A cross-mapping occurs when an object is mapped to a featurally less similar object rather than a featurally more similar object (because it shares a relational role

with the less similar object). For example, if dog1 is chasing cat1 and cat2 is chasing dog2, then the correct mapping places dog1 into correspondence with cat2. The ability to find such a mapping is taken as a hallmark of genuinely relational (i.e., as opposed to feature-based) processing (see, e.g., Gentner, 2003; Markman & Gentner, 1993; Halford et al., 1998; Hummel & Holyoak, 1997).

To test the *bigger-than* relations DORA learned in the previous simulations for their ability to support finding cross-mappings, we selected two of those relations at random and bound them to new objects, creating two new propositions, P1 and P2. The agent of P1 was semantically identical to the patient of P2 and patient of P1 was identical to the agent of P2. Specifically, P1 was *bigger-than1* (dog, cat) and P2 was *bigger-than2* (cat, dog). We let DORA map P1 onto P2 and observed whether it mapped the cat in P1 onto the dog in P2 (the correct relational mapping) or the cat. We repeated this procedure 10 times (each time with a different randomly-chosen pair of relations), and each time, DORA successfully mapped the cat in P1 to the dog in P2 and vice-versa. These results demonstrate that the relations DORA learned in the first series of simulations satisfy the requirement of structural sensitivity.

As a test of the semantic-richness of DORA's learned relations, we tested its ability to map those relations to similar but non-identical relations (such as *greater-than*). People can successfully map such relations (e.g., Bassok, Wu, & Olseth, 1995; Gick & Holyoak, 1980, 1983; Kubose, Holyoak, & Hummel, 2002), an ability that Hummel and Holyoak (1997, 2003), Doumas and Hummel (2004a) and others have argued depends on the semantic-richness of our relational representations.

To test DORA's ability to map non-identical predicates we had it map a new relation (R2) to one of its learned *bigger-than* relations (R1, randomly chosen). R2 was constructed to share 50% of its semantics (in each role) with R1. To assure that DORA could not "cheat" by mapping based on object similarity, the objects that served as arguments to the corresponding roles of R1 and R2 were constrained to have no semantic overlap. We repeated this procedure 10 times. Each time, DORA mapped the agent role of R1 to the agent role of R2 and the patient role of R1 to the patient role of R2; corresponding objects also always mapped to one another (by virtue of their bindings to corresponding roles) in spite of their lack of semantic overlap.

Finally, as the most stringent test of the structure sensitivity and semantic richness of DORA's learned relations, we tested the model's ability to find mappings that violate the *n-ary restriction*: the restriction that an *n*-place predicate may not map to an *m*-place predicate when $n \neq m$. The *n*-ary restriction applies to most models of analogical mapping (namely, those that represent propositions using traditional propositional notation and its isomorphs; see Doumas & Hummel, 2004a, 2005)

but does not apply to human analogical reasoning, as evidenced by our ability to find the correct role and object correspondences between *taller* (Bill, Dave) on the one hand and *short* (Fred), *tall* (James), on the other (Hummel & Holyoak, 1997).

To test DORA's ability to find such mappings, we ran a simulation in which DORA mapped one of its learned *bigger-than* relations (R1, randomly chosen) to a single place predicate (r2) that shared 50% of its semantics with the agent role of R1 and none of its semantics with the patient role. The object bound to r2 shared half of its semantics with the object in the agent role of R1 and the other half with the object in the patient role of R1. We ran this simulation 10 times, and each time DORA successfully mapped the agent role of R1 to r2, along with their arguments. We then ran the same simulation 10 more times, only on these simulations, r2 shared half its semantic content with the patient (rather than agent) role of R1. DORA successfully mapped the patient role of R1 to r2 (along with their arguments) on each run. In all these simulations, DORA overcame the *n*-ary restriction, mapping the single-place predicate r2 onto the most similar relational role of R1.

In summary, the basic simulations reported so far demonstrate that, starting with unstructured representations of object features, DORA learns relational representations, such as *bigger-than* (x, y), that meet the joint requirements of structure sensitivity and semantic richness. That is, DORA learns representations of relational concepts that support human-like relational reasoning. The simulations thus demonstrate that the representations and processes embodied in DORA can provide at least the beginnings of an account of the discovery and predication of relational concepts. In the next sections, we demonstrate that DORA can also account for several specific empirical findings in the literature on the acquisition of relational concepts in adults and children.

The relational shift

The kinds of problems a reasoner can solve depend critically on the content and form of his or her mental representations (Phenomenon 6; see Doumas & Hummel, 2005 for a review). For example, children go through a domain-specific *relational shift* in which they transition from representing a domain in terms of its characteristic features to representing it in terms of its characteristic relations as well (Phenomenon 7; see e.g., Gentner, 2003; Gentner & Rattermann, 1991). Early on children tend to appreciate similarity on a very global level (e.g., Chen, Sanchez, & Campbell, 1997; Oakes & Cohen, 1990; Smith, 1989). However, as they develop they begin to appreciate kinds of similarity. For example, they can appreciate that two items are similar because they are the same color even though they have different shapes (e.g., Smith, 1984), or that a situation in which a dog chases a cat is similar to a situation in which a police officer chases a criminal (e.g., Gentner & Rattermann, 1991; Rattermann & Gentner, 1998). Im-

portantly, this relational shift is domain-specific, in the sense that it may occur at different times for different domains of knowledge (Gentner & Ratterman, 1991).

The domain-specific nature of the relational shift has prompted Gentner and her colleagues (e.g., Gentner, 1988, 2003; Gentner & Rattermann, 1991; Rattermann & Gentner, 1998) to argue that the relational shift reflects a qualitative change in children's mental representations. Children's initial representations support the appreciation of over-all similarity (see also Smith, 1989), but as they grow older they learn more abstract representations that support reasoning based on specific properties and relations. The relational shift is an example of the manner in which a person's knowledge representations can affect their reasoning. This account of the relational shift, although very likely correct in our view, is nonetheless incomplete in that it does not provide an account of *how* this change takes place: What kinds of operations, at an algorithmic level, allow a child to make the transition from early holistic representations to progressively more relational representations that eventually support adult-like relational reasoning?

As demonstrated by the simulations presented in the previous section, DORA provides a systematic and detailed account of how this change in the quality of representations progresses. Moreover, as in the relational shift, DORA's representations progress in a domain-specific fashion. DORA's progression thus corresponds in a very intuitive way to the relational shift as a general phenomenon. As demonstrated in the following sections, the model also provides a natural account of many phenomena cited as examples of, or are otherwise related to, the relational shift.

Children's early relational representations

Phenomena 8, 9 and 10 describe patterns in children's development of relational concepts. Early in development, our relational concepts are holistic and object-based (i.e., tied to specific object values; Phenomenon 8) and over time gradually become more abstract (Phenomenon 9; e.g., Chen, Sanchez, & Campbell, 1997; Gentner, 2003; Gentner & Medina, 1998; Holyoak & Thagard, 1995; Landau, Smith, & Jones, 1988; Markman & Gentner, 1993; Rovee-Collier & Fagen, 1981; Smith, 1989; Smith, Rattermann, & Sera, 1988; Quinn et al., 1996). As a result, our early relational representations often appear more categorical (i.e., category-like) than relational (Phenomenon 10). For example, children use *bigger* to describe categorically big objects rather than objects that play the agent role in the *bigger-than* relation.

The trajectory from categorical to relational representations was demonstrated by Smith et al. (1988), Experiment 1. The experimenters tested children's ability to reason about relations between objects. Children ages 4-5 viewed pairs of toy butterflies at three different sets of heights: (1) One butterfly at five and the other at six feet above the floor (both butterflies were high); (2) one butterfly at three the other at four feet (both butter-

flies were at a medium height); (3) one butterfly at one foot the other at two feet (both butterflies were low). As the children viewed each set they were asked whether one of the two butterflies was higher (or lower) and if so which one. There were three trial types. On *consistent* trials both butterflies were high (or low) and the child was asked which butterfly was higher (or lower). On *neutral* trials both butterflies were in the middle and the child was either asked whether one was higher or lower. On *inconsistent* trials both butterflies were high (or low) and the child was asked which butterfly was lower (or higher). The experimenters found that while 4-year-olds performed well on consistent trials and above chance on neutral trials, they were at chance for inconsistent trials. However, 5-year-olds performed well above chance on all three trial types (see Figure 9).

As noted by Smith et al. (1988), these results suggest that the 4-year-olds were treating relations like categories: They treated *higher* like a category referring to high things and *lower* like a category referring to low things (i.e., only high things could be higher and only low things could be lower). However, by age 5, they understood *higher* and *lower* as relational concepts that could be true regardless of the absolute heights of the objects involved.

We simulated Smith et al. (1988) Experiment 1 in two interleaved parts. In the first, we simulated the development of the *higher-than* relation. This simulation progressed in the same way as our basic simulation of relation discovery reported above ("*Perception*" simulations), except that (a) "height" (and specific heights from 1 to 10) replaced "size" (and specific sizes), and (b) we attached category-specific values to objects (e.g., "high" for objects with height 7-10, "low" for objects with height 1-3). We let DORA make the same comparisons described in the basic simulations section.

The results of this part of the simulation were the same as the results of the relation discovery simulation reported above. As before, DORA first learned single-place representations of specific heights, which it refined through its initial comparisons. Next it ran sets of these single-place predicate-filler sets describing different heights through the comparator. Subsequently it learned relations comprised of these single-place predicate sets by comparing them to other single-place predicate sets, producing value-dependent representations of the *higher* and *lower* relations. Finally, DORA compared its value-dependent relational representations, which produced value-independent representations of the *higher* and *lower* relations (i.e., representations of *higher* and *lower* that were not attached strongly to specific heights or specific values such as "high" or "low").

It is important to emphasize that DORA's learning trajectory—from categorical (i.e., context-bound) to more context-free representations of relations—is not simply a function of the manner in which we ran the simulations, but is instead an inevitable consequence of the manner in which the model learns relations from

examples: Beginning with representations of whole objects it learns representations of specific values of height. From these it learns representations of *higher* and *lower* that are tied to specific heights or to specific categorical values (i.e., representations of *higher* that are also “high” and representations of *lower* that are also “low”). And from these it finally learns representations of *higher-than* and *lower-than* that are independent of object features and context (i.e., that are relative rather than categorical): DORA necessarily learns the context-bound representations before it learns the context-independent representations because it learns its context-independent representations by comparing sets of context-bound representations.

The second part of the simulation was the simulation of the Smith et al. (1988) experiment. Specifically, to simulate children of different ages, we ran the simulations of the Smith et al. experiment at different times during DORA’s learning of the *higher-than* and *lower-than* relations. To simulate 4 year olds, we ran DORA when it had roughly three-times as many context-dependent predicates as context-free predicates. To simulate 5 year-olds, we ran DORA after it had acquired roughly the same number of context-free predicates as it had context-dependent predicates. This procedure reflects our assumption that 5 years-olds have had more time to make the comparisons necessary to acquire context-independent predicates than have 4 year-olds. At both “ages” DORA also had roughly the same number of random propositions about butterflies as it had context-dependent predicates in its LTM (see Appendix B for details).

For each trial we placed a representation of the problem (i.e., the pair of butterflies placed in front of the child) in the driver. The same representation of the problem was used to simulate both 4 and 5 year-olds. For all three trial types, the driver contained one proposition specifying that one butterfly was higher than the other (i.e., *higher-than* (butterfly1, butterfly2)). In addition, for the consistent and inconsistent trial types we included two additional propositions. For trials on

which both butterflies were high, we included the propositions *high* (butterfly1) and *high* (butterfly2); likewise, trials on which both were low included the propositions *low* (butterfly1) and *low* (butterfly2). Neutral trials, in which both butterflies were near the middle of the display, did not include these additional statements.

We tested the model by using the propositions in the driver to retrieve propositions from LTM. Once a small number of propositions (approximately 3 or 4) had been retrieved from LTM into the recipient, DORA attempted to map the representation in its driver to the propositions in its recipient. We selected the proposition from the recipient that mapped most strongly to a driver representation as the model’s response.

We ran 40 simulations of each of the six types of trials formed by crossing the two question types (“which is higher?” vs. “which is lower?”) with the three heights (high, medium and low) for each age group. DORA’s performance across all runs is summarized in Figure 9. Just as the 4 year-olds in the Smith et al. study, DORA at “age 4” performed best on consistent trials and worst on inconsistent trials. At “age 5” DORA, just like 5 year-old children, had little difficulty with any of the conditions, and performed well above chance in all cases.

The trajectory of the appreciation of similarity

Phenomena 11 and 12 describe the development of children’s ability to appreciate the correspondences between situations (see Sera & Smith, 1987; Smith, 1984, 1985, 1989; Smith & Sera, 1992). Early in development, children appreciate the overall similarity between objects. Later, with increased knowledge, they can reason about objects based on specific properties. Finally, they can appreciate the similarity between sets of objects based on shared relations (Phenomenon 11). This increasing relational knowledge leads to improved performance making analogical mappings (Phenomenon 12).

Figure 9. Simulation data of Smith, Rattermann, & Sera, 1988.

The trajectory of the appreciation of similarity

Phenomena 11 and 12 describe the development of children's ability to appreciate the correspondences between situations (see Sera & Smith, 1987; Smith, 1984, 1985, 1989; Smith & Sera, 1992). Early in development, children appreciate the overall similarity between objects. Later, with increased knowledge, they can reason about objects based on specific properties. Finally, they can appreciate the similarity between sets of objects based on shared relations (Phenomenon 11). This increasing relational knowledge leads to improved performance making analogical mappings (Phenomenon 12).

Smith (1984; Experiment 1) provides an example of this trajectory. This experiment tested the abilities of children aged 2-4 to match items based on overall identity, shared properties, and shared relations. In the experiment, a child and two experimenters each had three items placed in front of them. The experimenters selected items based on identity (or over-all similarity; e.g., the first experimenter selected a red house and the second experimenter selected a red house), a shared property (e.g., the first experimenter selected two red items, and the second experimenter selected two different red items), or a shared relation (e.g., the first experimenter selected two red items and the second experimenter selected two green items). The child's task was to select items from her pile that best matched the experimenters' choices. In the identity condition, the child had to select the item that matched the items selected by the two experimenters (e.g., a red house). In the shared property condition, the child had to select two items with the same property as those chosen by the experimenter (e.g., two red items). In the shared relation condition, the child had to select two items that matched the items chosen by the experimenters on the basis of a common relation (e.g., two blue items [out of, say, two blue items and a red item]). In each condition there were distracter items that shared superficial similarity with the other items in front of the child and the items chosen by the experimenters. Thus, in order to perform correctly, the child had to ignore some similarities and focus on others. Smith found that the ability to appreciate more abstract kinds of similarity increased with age. Although all the children could match items based on identity, only 3 and 4 year-olds could consistently match based on shared properties, and only 4 year-olds consistently matched based on shared relations.

DORA predicts this trajectory. Beginning with holistic object representations DORA learns representations of object properties (i.e., single place predicates). It then concatenates sets of single-place predicates to form multi-place relational structures. We simulated Smith (1984) Experiment 1 by allowing DORA to develop representations and testing it at various points along its developmental trajectory. Like the simulation Table 6.

of Smith et al. (1988), this simulation had two interleaved parts: The development of representations, followed by the use of those representations at each stage of their development. Our goal was to observe whether the kinds of problems DORA can solve as its representations develop mirror the trend demonstrated by the children in Smith's study.

We ran the first part of the simulation—the development of representations—just like our earlier basic simulations of relation discovery (*"Perception" simulations*). The only difference was that, rather than semantics describing size (or height), the objects DORA started with were connected to semantics of interest describing their color, specifically, the semantic "color", and another semantic feature indicating the specific color of the object. (As an aside, it is worth noting that many of the irrelevant semantics in these simulations described sizes and heights. Thus, the real difference between the simulations is not which semantics the objects are assumed to have attached to them; rather, it is only (a) which semantics are relevant to the task, and thus (b) which ones the objects are assumed to share.) We used 10 specific object colors (see Appendix B for details). We let DORA make exactly the same series of comparisons described in the basic simulations of relation discovery.

The results of this simulation were the same as the results of the previous relation discovery simulations. DORA first learned single-place representations of specific colors, which it refined through additional comparisons. Next it ran sets of these single-place predicate sets describing the same color through the comparator, and then learned relations comprised of these single-place predicate sets by comparing them to other single-place predicate sets, producing value-dependent representations of the *same-color* relation (i.e., *same-color* relations tied to specific color values). Finally, DORA compared its value-dependent relational representations, which produced value-independent representations of the *same-color* relation (i.e., representations of *same-color* that were not strongly attached to specific color values).

As before, this trajectory is not a function of the way we ran the specific simulations but rather is an inevitable consequence of the manner in which DORA learns relations: It must learn about specific colors before it can learn that two colors are the same, and it must learn that some colors are the same before it can abstract the relation *same-color*.

The second part of the simulation tested whether the representational trajectory DORA predicts matches the developmental trajectory demonstrated by Smith (1984). To do so, we tested DORA at various stages of its representational development during part 1 of the simulation. Specifically, to simulate 2 year-olds we tested DORA on Smith's task after it had made its initial comparisons of whole objects (i.e., when it had representations of whole-

Types of propositions and examples of these propositions used to simulate Smith (1984) Experiment 1.

Propositions	DORA “age 2”	DORA “age 3”	DORA “age 4”
<u>Identity:</u> Driver: Recipient:	<u>Unbound objects</u> P1: (Ball1) <u>Unbound objects</u> P2: (Ball3) P3: (Ball4) P4: (Ball5)	<u>Unbound objects</u> P1: (Ball1) <u>Unbound objects</u> P2: (Ball3) P3: (Ball4) P4: (Ball5)	<u>Unbound objects</u> P1: (Ball1) <u>Unbound objects</u> P2: (Ball3) P3: (Ball4) P4: (Ball5)
<u>Shared Property:</u> Driver: Recipient:	<u>“Dirty” single-place predicates</u> P1: red+size-2 (Ball1) P2: red+size-4 (Ball2) <u>“Dirty” single-place predicates</u> P3: red+size-2 (Ball3) P4: red+size-8 (Ball4) P5: green+size-4 (Ball5)	<u>Value-dependent relations</u> P1: same+color+red (Ball1, Ball2) <u>Value-dependent relations</u> P2: same+color+red (Ball3, Ball4) P3: same+size+size-5 (Ball4, Ball5)	<u>Value-independent relations</u> P1: same+color (Ball1, Ball2) <u>Value-independent relations</u> P2: same+color (Ball3, Ball4) P3: same+size (Ball4, Ball5)
<u>Shared Relation:</u> Driver: Recipient:	<u>“Dirty” single-place predicates</u> P1: red+size-2 (Ball1) P2: red+size-4 (Ball2) <u>“Dirty” single-place predicates</u> P3: green+size-2 (Ball3) P4: green+size-8 (Ball4) P5: red+size-4 (Ball5)	<u>Value-dependent relations</u> P1: same+color+red (Ball1, Ball2) <u>Value-dependent relations</u> P2: same+color+green (Ball3, Ball4) P3: more+size+red+green (Ball3, Ball5)	<u>Value-independent relations</u> P1: same+color (Ball1, Ball2) <u>Value-independent relations</u> P2: same+color (Ball3, Ball4) P3: more+size (Ball3, Ball5)

objects and the “dirty” single-place predicates). To simulate 3 year-olds we tested DORA after it had formed its initial relational representations (i.e., when it had “clean” single-place predicate representations and value-dependent relational representations). Finally, to simulate 4 year-olds we tested DORA after it had refined its relational representations (i.e., when it was first developing representations of value-independent relations).

To simulate each trial of Smith’s experiment we placed a representation of the experimenters’ choices in the driver and a representation of the three objects in front of the child in the recipient. Table 6 presents examples of the kinds of representations in DORA’s driver and recipient as a function of “age”. All the representations used in this part of the simulation were representations DORA had learned during the first part of the simulation. At “age 2” DORA had only learned “dirty” single-place predicates so it represented objects simply as objects for the identity trials (as this was sufficient to perform at ceiling), and as objects bound to “dirty” single-place predicates describing a relevant property in the shared property and shared relation conditions (Table 6 column 1). At “age 3” DORA had learned “clean” single-place predicates and value-dependent relational representations so it represented the objects as whole

objects for the identity trials, as objects bound to “clean” value-independent relations in the shared property condition, and as objects bound to value-dependent relations in the shared relation condition (Table 6 column 2). At “age 4” DORA had learned “clean” single-place predicates and value-independent relations so it represented the objects as whole objects for the identity trials, and as objects bound to value-independent relations in the shared property condition and the shared relation condition (Table 6 column 3). As in Smith (1984, Experiment 1), the proportion of relevant and distracter features was balanced across all trials (i.e., objects in the driver shared an equivalent proportion of superficial similarity, in the form of semantic features, with both the correct choice object [the object that they should match to based on the rule] and the distracter object; see Appendix B for details).

We used DORA’s ability to map representations as a metric of its ability to match them: If DORA mapped an object in the driver to an object in the recipient it was judged to have matched them. If DORA matched the objects in the driver to the objects in the recipient based on the trial rule, it was judged to have made a correct choice (e.g., in a shared relation trial it mapped the two objects in the driver to the two objects in the recipient with the same relation). If DORA mapped

either of the objects in the driver to the distracter object in the recipient, it was judged to have made an incorrect choice. If DORA failed to find any mappings, it simply

Table 7.

Probability of reaching criterion (greater-than or equal to 75% correct; 10 subjects/runs per age, 4 trials per condition per subject/run).

	Age 2 Children, DORA	Age 3 Children, DORA	Age 4 Children, DORA
Identity	1.0, 1.0	1.0, 1.0	1.0, 1.0
Shared Property	0.8, 0.7	0.9, 1.0	1.0, 1.0
Shared Relation	0.0, 0.1	0.7, 0.8	1.0, 1.0

We ran 10 simulations of each age group. Each simulation consisted of 12 trials, four of each type (i.e., identity, property, and shared relation; see Appendix B for details). DORA's performance on each type of problem at different stages of its representational development is summarized in Table 7, along with the human data.

As illustrated in Table 7, the qualitative correspondence between DORA's performance and the human data is very close. With its initial whole object representations, DORA, like the 2-year-olds in Smith (1984), could consistently match based on identity. That is, it could easily map an object in the driver to the most similar object in the recipient. At "age 3", after DORA had learned "clean" single-place predicates, it could consistently map sets of objects with similar colors predicated about them. For example, DORA would map *red* (ball) to *red* (house) because of the common red predicates even though ball and house have very different features. In fact, DORA would even map *red* (ball) to *red* (house) in the presence of another ball of a different color. In short, DORA, like the 3 year-olds in Smith's study, was able to appreciate the commonalities between items based on a shared attribute, even in the presence of distracter items. Finally, once DORA at "age 4" had learned value-independent relational representations, like *same-color*, it could consistently match pairs of items based on those relations. For instance, it would match *same-color* (ball1, ball2) to *same-color* (item1, item2) even in the presence of a distracter item.

Comparison and the development of concepts and relational representations

Phenomena 13-16 highlight the importance of comparison in the development of conceptual and relational representations in children (e.g., Gentner, 2003; Gentner & Namy, 1999; Kotovsky & Gentner, 1996; Namy & Gentner, 2002; Rattermann & Gentner, 1998; Sloutsky, 2003; Sloutsky & Fisher, 2004; Smith, 1989), and adults (e.g., Dixon & Bangart, 2004; Dumas & Hummel, 2004b; Ross & Warren, 2002; Sandhofer & Smith, 2001; Yamauchi & Markman, 1998, 2000). In DORA, comparison bootstraps the discovery of relational concepts. As a result, it may be well suited to provide an

account of the mechanisms that underlie comparison-based learning in children.

Comparison and children's learning

Phenomenon 13 refers to the central role of comparison in the formation of children's concepts, as demonstrated by many researchers and exemplified by a study by Gentner and Namy (1999). Gentner and Namy tested the effect of comparison on 4 year-old children's ability to abstract conceptual and structural commonalities from featurally dissimilar instances using a simple match-to-sample task. In a match-to-sample task the subject is presented with a sample stimulus, and two test items. Her task is to choose the test item that best matches the sample. There were two conditions in Gentner and Namy's experiment. In the No-Compare (NC) condition one of the test items was more perceptually similar to the sample and the other was a categorical match. For example, if the sample was an apple, the test items might be a red ball (a perceptual match) and a banana (a categorical match). As the authors expected, children in the NC condition more frequently chose the test item that was perceptually similar to the sample.

However, in the Compare (C) condition, Gentner and Namy elicited comparison by presenting four sample items and two test items. For example, the sample might consist of an apple, a pear, a watermelon, and a grape, with a ball and a banana as test items. Here the children had the opportunity not only to compare the sample to the test items, but also to compare sample items to one another. In this condition, although the ball, being round, is perceptually more similar to all of the sample items than is the banana, children overwhelmingly selected the test item that matched the sample categorically (in this example, the banana). Thus, comparison helped the children extract explicit representations of the categorical information from items and use this information to find correspondences between them.

In our simulation the sample and test items were generated to match the logic of Gentner and Namy's stimuli. Broadly, we generated the stimuli so that the perceptual match item was more similar to the sample item (or items) than was the conceptual match item, but

the conceptual match item and the sample item(s) did have at least one feature in common (i.e., at the very least they came from the same category; as in Gentner and Namy's experiment). To simulate the NC condition we placed a representation of the sample (as a PO unit) in the driver, and representations of the two test items (also as PO units) in the recipient. Each test item was attached to 10 distinct semantic features (i.e., the test items had no semantic features in common). The sample was attached to eight semantic units. In order to ensure that the sample had at least one feature in common with each test item, one of the semantics attached to the sample item was a semantic that was also attached to the perceptual match item (chosen at random), and a second semantic attached to the sample item was a semantic that was also attached to the perceptual match item (the categorical semantic). The remaining six semantics from the sample were chosen at random from a pool of 14. This pool was constructed by selecting all nine of the remaining semantics attached to the perceptual match test item (i.e., the semantics not already connected to the test item) and five of the nine remaining semantics (at random) that were attached to the categorical match test item. In this way we replicated the similarity of the sample and test items from Gentner and Namy: The test item generally shared two-thirds of its features with the perceptual match and one-third with the categorical match, and the sample had at least two semantics in common with the perceptual match and one semantic in common with the conceptual match. On each trial DORA attempted to map the driver representation (i.e., the sample) onto the recipient items (i.e., the test). Whichever recipient item mapped most strongly to the driver item was selected as DORA's choice on that trial.

For the C condition, we constructed sample and test items exactly as described above with the constraint that all the sample items shared the same semantic (i.e., the category label) with the categorical match item. As a result, although the perceptual match was more similar to any individual sample item, there was at least one semantic unit common to all sample items and the conceptual match, reflecting the underlying conceptual similarity between all the sample items and the conceptual match. Before mapping, we allowed DORA to make a set of comparisons. First DORA selected two sample items at random, compared them, and predicated their common properties. Then DORA compared the other two sample items and predicated their common properties (corresponding to Gentner and Namy's assumption that the four sample items elicited comparison). Finally, DORA compared the two new representations it had learned during the previous comparisons and refined them via intersection discovery. These comparisons resulted in a PO connected most strongly to the semantic (or semantics) shared by all the conceptual match items (the conceptual semantic; e.g., "fruit"). DORA used this representation in the driver and the two test

items in the recipient. Because predicates and objects share the same semantic pool, when the new predicate representation in the driver fired, it excited (and, therefore, mapped to) the conceptual match item in the recipient. Again whichever recipient object mapped most strongly to the driver item was selected as DORA's response choice on that trial.

We ran DORA for 20 simulations each consisting of five NC and five C trials. The results are presented in Figure 10. Just like the subjects in Gentner and Namy's experiment, DORA overwhelmingly chose the categorical match item in the C condition, and chose the perceptual match more frequently in the NC condition.

Figure 10. Simulation of Gentner and Namy (1999) data.

The role of progressive-alignment

Phenomenon 14 refers to the effects of a training procedure, *progressive alignment* (e.g., Gentner, 2003; Kotovsky & Gentner, 1996), on concept acquisition. Under progressive alignment, subjects begin by comparing highly similar examples of a concept and progressively compare more distant examples. This procedure is remarkable because it leads to relational learning without feedback (i.e., learning is incidental) and yields relational responding earlier (developmentally) and faster (during training) than when the examples are not progressively ordered. For example, in Experiment 2 of Kotovsky and Gentner (1996), 4 and 6 year-old children were given the task of learning to solve relational match-to-sample problems as illustrated in Figure 11. In a relational match-to-sample, the correct test item matches the sample based on a shared relation. For example, the child might be shown a sample consisting of a series of circles increasing in size, along with two test items: a series of squares increasing in size or a set of unordered squares (Figure 11a). The child's task is to choose the test item that matches the sample. This task is difficult for young children. Even harder is the problem illustrated in Figure 11b, where the child must make a cross-dimensional match (e.g., matching increasing size to increasing darkness).

Kotovsky and Gentner (1996) attempted to teach children to make these and related matches. They had

little success at the younger ages when the instances were presented in a random order. However, they had considerable success—even without providing children with any feedback (Phenomenon 15)—when the training sets were presented starting with very easy matches that could be solved by featural similarity alone, and then progressing slowly to more difficult relational matches. Each child received 16 trials. In the progressively aligned condition, the first eight were same-dimension trials (Figure 11a) arranged so that the simplest matches (those that could be solved featurally) came first. The last eight were cross-dimension trials (Figure 11b). Kotovsky and Gentner found that all the children who performed well (i.e., above chance) on the same-dimension trials also performed well on the cross-dimension trials, and that all the children who performed poorly (i.e., at chance) on the same-dimension trials also performed poorly on the cross-dimension trials.

Figure 11. Examples of stimuli used in Kotovsky and Gentner (1996). (a) Illustration of a same-dimensional match. The child must match two sets of shapes increasing in size. (b) Illustration of a cross-dimensional match. The child must match a set of items increasing in size to a set of items increasing in darkness.

There are some remarkable facts about this phenomenon: First, children discovered the relations *incidentally* (i.e. without feedback), which is important because although children receive explicit training and instruction in some circumstances (e.g., in preschool) it seems likely that the majority of children's learning takes place without feedback. At the same time, the world is not an unstructured place. Similar things often occur in similar contexts and in close temporal proximity. In this sense, the systematic ordering of the progressive alignment procedure may more closely mimic the real world than do the randomized stimulus presentations conventional in cognitive experiments. Second, each training instance is similar to the next, which may directly facilitate the comparison process. Third, the early correlation between similarity-based matches and relational matches promoted the discovery of the relation and its subsequent application without similarity support. This effect suggests that early similarity matches helped to bootstrap the discovery of the relational matches, and then faded in importance over successive examples. And fourth, familiar dimensions are introduced one at a time before the child is asked to make cross-dimensional mappings. This procedure may help the child learn to focus attention on the relevant

dimensions before requiring them to use those dimensions as the basis of mapping.

We used DORA to simulate Kotovsky and Gentner's (1996) Experiment 2. To create their stimuli, Kotovsky and Gentner crossed two relations (symmetry and monotonic increase), with two dimensions (color and size), and two trial types (same-dimension and cross-dimension). Each trial type consisted of a sample stimulus, which depicted a relation (either symmetry or monotonic increase) that held over a particular dimension (either size or color), and two match stimuli, one of which depicted the same relation holding over either the same dimension or a different dimension. This design yielded eight trial types (see Table 8). Each child received 16 trials, two of each type.

Table 8.

Summary of the eight trial types used by Kotovsky and Gentner (1996). Children received two of each trial type.

Relation	Dimension of sample	Dimension of match
Symmetry	Size	Size
Symmetry	Color	Color
Monotonic increase	Color	Color
Monotonic increase	Size	Size
Symmetry	Size	Color
Symmetry	Color	Size
Monotonic increase	Color	Size
Monotonic increase	Size	Color

To simulate each trial we placed a proposition representing the sample (P1) in the driver and a representation of two match items (P2 and P3) in the recipient. Each item was represented as three POs, one for each of the elements in a given item (e.g., one PO for each of the three circles in the sample item from Figure 11a). Each PO was attached to a unique RB. All three RBs were attached to a single P unit. This convention corresponds to our assumption that the children understood that each item (the sample and each of the two matches) was a single item composed of three elements (see Figure 12). Each PO was attached to 12 semantics: One describing its role in the predictive relation independently of the dimension over which the relation holds (e.g., "least", "medium", "most"), one describing its position in the array (i.e., "left", "middle", and "right"), two describing its shape (i.e., "circle1" and "circle2", or "square1" and "square2"), two describing its value on the relevant dimension, one specific and one abstract (e.g., "size-5" and "size" or "color-black" and "color"), and six random semantics selected from a pool of 50. We presented the trials to DORA in the progressively aligned order described by Kotovsky and Gentner (1996). Just like the children in Kotovsky and Gentner's experiment, DORA received two trials of each

type for a total of eight same-dimensional trials in the following order: two symmetrical size trials, then two symmetrical color trials, then two monotonic increase in color trials, then two monotonic increase in size trials.

Then DORA received eight cross-dimensional trials in the same order.

On each trial DORA did the following. First, it tried to map the sample item in driver to one of the two match items in the recipient. If it mapped the sample to

Figure 12. An illustration of DORA's representation of a trial from Kotovsky and Gentner (1996), Experiment 2.

one of the two match items, that item was DORA's answer on that trial. Otherwise, it picked one of the two match items at random. After DORA selected one of the two match items, it compared that match item to the sample item and learned new POs via its predicate discovery algorithm. This procedure resulted in three new POs, one for each mapped object in the match item (see Figure 12). On pairs of successive trials DORA refined the new POs it had learned via its predicate refinement algorithm (i.e., it compared and refined the new POs it learned on the first and second trials, on the third and fourth trials, on the fifth and sixth trials, and on the seventh and eighth trials). Thus, if DORA made correct choices it on the first four trials it learned refined representations of the symmetry relation (i.e., the new POs it learned were connected most strongly to the semantics describing symmetry), and if it made correct choices on the second four trials it learned refined representations of the monotonic increase relation (i.e., the new POs it learned were connected most strongly to the semantics describing monotonic increase). If DORA made incorrect choices it learned representations connected to a random set of features (those shared by the sample item and the incorrect match item).

In Kotovsky and Gentner's experiment, the same sample items were used for the same- and cross-dimension trials. That is, the triads used in the sample during the first 8 trials (the same-dimension trials) were used again in the sample for the second eight trials (the cross-dimension trials). Therefore, we allowed DORA to use the representations it had learned during the first 8 trials (the same-dimension trials) when it performed the sec-

ond 8 trials (i.e., the cross-dimension trials). For each cross-dimension symmetry trial we chose (at random) one of the sets of refined POs DORA had learned during previous same-dimension symmetry trials, and placed that representation into the driver. Similarly, for cross-dimensional monotonic increase trial we chose (at random) one of the two sets of refined POs DORA had learned during the earlier same-dimensional monotonic increase trials, and placed that representation into the driver. This procedure reflects the assumption that the children in Kotovsky and Gentner's experiment used what they learned during the earlier trials to represent subsequent trials and to make their responses. As before, on each trial DORA attempted to map the driver proposition onto the recipient items. The item in the recipient that most strongly mapped to the sample item in the driver was taken as DORA's response.

We ran 200 simulations with 16 trials per simulation (each simulation corresponds to a single subject from Kotovsky and Gentner, 1996). In Kotovsky and Gentner's study children were separated into two pseudo-groups for the purposes of analysis. Children who performed above chance on the first 8 trials were placed in the "performed well initially" (PWI) group, and those who performed at chance were placed in the "performed poorly initially" (PPI) group. Following Kotovsky and Gentner we grouped each simulation into one of these two pseudo-groups. If, on a given simulation, DORA performed well on the first 8 trials (i.e., it got 6 or more correct) it was placed in the PWI group, otherwise, it was placed in the PPI group. DORA's performance on the later 8 trials for the PWI and PPI groups, and the

human data for the children in the PWI and PPI groups from Kotovsky and Gentner, Experiment 2 are summarized in Figure 13. Like the children in Kotovsky and Gentner's study, early success with progressively aligned stimuli yielded later success on cross-dimensional matches and early failure on these trials led to performance near chance on the later trials. Importantly, like the children in Kotovsky and Gentner's experiment, DORA could discover the relevant relational categories incidentally, without any explicit feedback.

Figure 13. Simulation data of Kotovsky and Gentner (1996).

Comparison and adult relation learning

Comparison plays a central role in the formation of adults' relational concepts (Phenomenon 16; see e.g., Dixon & Bangart, 2004; Doumas & Hummel, 2004b; Kurtz & Bourkina, 2004). An experiment by Dixon and Bangart (2004) demonstrated how comparison facilitates adults' learning of relational concepts—specifically, the relation between parity (i.e., odd or even number) and the behavior of chains or circuits of gears. Dixon and Bangart had adults solve a series of gear problems. In each problem the subject was shown a set of gears (see Figure 14). Each gear set was either single-pathway (i.e., a set of gears arranged to form a line so there is a single pathway between any two gears; Figure 14a) or two-pathway (i.e., a set of gears arranged to form a closed circuit so there are two pathways between any two gears; Figure 14b). On each trial the subject was shown a source gear that always spun clockwise. The subject's task was to determine which direction a target gear would spin, and, in the case of two-path sets, whether it would jam. The parity of the number of gears (i.e., whether the number was odd or even) separating the source and the target gears governs the direction in which the target gear will spin. If there is an odd number of gears between the source and target, then the target will turn the same direction as the source. In addition, for two-pathway problems, if there is an even number of gears overall (or, equivalently, if the two paths from the source to the target have the

same parity), then the set will not jam. But if there is an odd number of gears (or the two paths have unequal parity) then it will jam.

The goal of this study was to determine the conditions that lead subjects to discover the relationship between parity and (a) the direction of the target gear and (b) whether a circuit will jam. Dixon and Bangart (2004) hypothesized that two-pathway trials with the same parity but different numbers of gears in each pathway (e.g., one pathway with 2 gears, the other with 4) would provide subjects with the most opportunity to extract the parity rule because these trials provided the subject with the opportunity to compare two pathways (both were present in front of the subject simultaneously) and thus notice what they had in common (i.e., parity). Therefore, repeatedly solving two-pathway problems with the same parity but different numbers of gears should increase the likelihood of the subject discovering the parity rule. Consistent with this prediction, the authors found that the probability of discovering the parity rule increased with each consecutive two-pathway same parity trial the subject encountered.

We used DORA to simulate the findings of Dixon and Bangart (2004). On each trial there were a number of features that the subject could have noticed about one or both pathways. We assumed that on each trial the subject attended to some subset of the possible features of each pathway and based their comparison of the two pathways on this feature set. For each two-pathway trial, DORA, like Dixon and Bangart's subjects, "saw" two pathways. We placed a representation of one pathway in the driver and a representation of the other pathway in the recipient. Each pathway was represented as a PO unit attached to 10 semantic features. The semantic features were selected randomly from a pool of 400.

To simulate same-parity trials, the pools of semantics used to build both pathway representations contained either the semantic "parity-even" or the semantic "parity-odd". To simulate the different-parity trials, the pools of semantics used to build one of the pathway representations contained the "parity-even" semantic, and the other pool contained the "parity-odd" semantic. On each trial DORA compared the representation of the pathway in the driver to the representation of the pathway in the recipient and predicated their common properties via its comparison-based predication routine.

Each simulation consisted of five trials. Of these, between one and five were two-pathway same-parity trials, and the remainder were two-pathway different-parity trials. This resulted in five types of simulations (i.e., one same-parity and four different-parity trials; two same-parity and three different-parity trials, etc., up to five same-parity and zero different-parity trials). We ran 1000 simulations of each type. On each simulation, we measured the probability that DORA would build an explicit predicate representing parity (i.e., a new PO connected most strongly to either "parity-odd" or "parity-even"). The results are presented in Figure 15. Like

the subjects in Dixon and Bangart's experiment, the probability that DORA discovered the parity rule

Figure 14. Examples of the stimuli used in Dixon and Bangart (2004). (a) A one-pathway gear set. (b) Two-pathway gear sets. The gear set on the left has an even number of gears and will not jam. The gear set on the right has an odd number of gears and will jam. Each set has a driving gear (marked by the black center circle that turns clockwise). The subject's task is to identify what direction the target gear (marked by a center triangle) will turn, and for two-pathway gears whether or not the gear set will jam.

Figure 15. Simulation of Dixon and Bangart's (2004) data.

increased with the number of same-parity trials in the set: Like the human subjects, the greater the number of useful comparisons DORA could make, the more likely it was to extract the rule.

In addition, these simulations along with the “perception” and “interleaved” simulations described under *Learning relational concepts from examples* demonstrate that DORA provides a natural account of learning relations in both naturalistic and experimental settings. In naturalistic settings children and adults take longer to master relational concepts (e.g., taking a year to learn a refined concept of *bigger-than*); by contrast when children and adults learn from blocked examples, they learn relational concepts much more quickly (e.g., learning a concept like *monotonic-increase* in a single experimental session). When DORA learns new relational concepts from interleaved examples, it learns comparatively slowly. However, when it learns new relational concepts from blocked training (as in the simulations of Kotovsky & Gentner, 1996, and Dixon & Bangart, 2004) it learns more quickly. Thus, DORA accounts for

both the slower naturalistic learning of children and the faster blocked learning of children and adults.

General Discussion

Summary and Overview

The question of how people acquire relational concepts is important because the relations a person can predicate tightly constrain the kinds of thoughts that person can have and the kinds of problems they can solve. Little is known, however, about how people acquire relational concepts. Accordingly, the question of how—or indeed, whether—people could learn structured (i.e., relational) representations from examples has been cited as a fundamental limitation of relational accounts of cognition (e.g., Munakata & O'Reilly, 2003; O'Reilly & Busby, 2002; O'Reilly, Busby, & Soto, 2003). We have presented a theory of how structure-sensitive and semantically-rich representations of relations are discovered and predicated from unstructured

(holistic) examples. This theory rests on a few key tenets. First, adopting a role-filler binding representational system reduces the problem of discovering relations to the problem of discovering single-place predicates (object properties or relational roles) and linking these together to form multi-place relational structures. Second, predicates and their arguments share a common representational basis (i.e., a common vocabulary of semantic primitives). Third, comparison can lead to the discovery, predication and gradual refinement of invariant properties, and the formation of multi-place relational structures. And fourth, mapping multiple predicates of smaller arity can lead to the predication of higher-arity structures.

We have instantiated the theory in a computer simulation, DORA (Discovery Of Relations by Analogy). DORA, like its predecessor, LISA (Hummel & Holyoak, 1997, 2003), is a symbolic-connectionist model—a system built from traditional connectionist computing elements that, by virtue of its distributed representations and its solution to the dynamic binding problem, represents relational knowledge in a way that is simultaneously semantically rich and meaningfully symbolic (i.e., structure sensitive).

Starting with unstructured (holistic) representations of objects, DORA learns structured and semantically rich representations of relational concepts. As such, DORA serves as an existence proof that relational representations can be learned from examples, thereby addressing one of the fundamental problems facing symbolic models of cognition. More importantly, DORA provides an integrated theory of the origins of complex mental representations and the discovery of structured representations of relational concepts. Although many researchers have hypothesized that comparison may play a vital role in learning new representations (e.g., Doumas & Hummel, 2004b; Gentner, 2003; Gentner & Medina, 1998; Sandhofer & Smith, 2001; Smith, 1989; Yamauchi & Markman, 1998, 2000), DORA is the first detailed, computationally instantiated account of how comparison can serve to bootstrap the discovery and predication of structured relational concepts.

We have shown that DORA accounts for a number of key phenomena in human cognitive development and relation discovery. Specifically, we used DORA to simulate the discovery of relational representations that support analogical thinking (i.e., representations that are both structure sensitive and semantically rich), children and adult's learning of dimensions and relational representations, and the role of comparison and progressive alignment in children's and adults' relation learning. In so doing, we have demonstrated how a system can exploit the tools of statistical learning to discover representations that allow the system to overcome the limitations of statistical learning.

Developmental Mechanism

In terms of developmental mechanism our general proposal is that the same psychological mechanisms that underlie analogical inference and schema induction—namely, analogical mapping, self-supervised learning and intersection discovery—also underlie the discovery and predication of the object properties (such as size) and relations (such as *bigger-than*) that make analogical reasoning possible in the first place. Armed with a basic vocabulary of perceptual and relational invariants (which may either be present at birth, the result of specific computing modules, or some combination of both), DORA discovers relations through general learning processes, and develops as a result of experience. Its development reflects a cascading process in which, through learning, initially holistic features become represented as explicit predicates, which then become more refined and get combined into relations, which themselves become more refined. The resulting qualitative changes in the model's representations—and thus in its ability to reason—reflect the operation of basic learning processes generating more sophisticated representations by building on previously acquired representations.

The model assumes that memory and perception are present at the start of learning. It also requires a capacity for comparison, mapping and self-supervised learning, as well as the ability to flexibly treat a feature either as a (holistic) feature of an object or as a feature of an explicit predicate. These same domain general processes are also implicated in discovering non-relational categories (i.e., properties, objects and perhaps and action categories). Importantly, the developmental changes that occur in DORA are not driven by changes in architecture (although depending on the ages being modeled changes in memory and/or perception might be appropriate). Rather, developmental change is a product of the experiences the learner has in the world. Specifically as a learner has more opportunities to compare things, she becomes less bound to the immediate situation and better able to extract context-independent relations. Our simulations demonstrate that this approach provides a powerful account of many phenomena both in cognitive development and in relational learning in adulthood. Although we by no means deny the importance of maturation in cognitive development (see, e.g., Halford, 1993; Halford et al., 1998), DORA provides an illustration of how far it is possible to go with learning alone.

Relation to Other Theories

Our theoretical proposal is broadly consistent with those of Gentner and colleagues (e.g., Gentner & Rattermann, 1991; Rattermann & Gentner, 1998) and of Goswami (1992, 2001). Like Gentner and colleagues, we argue that changes in knowledge and representation underlie changes in children's relational thinking. As children learn to represent the relations that characterize specific domains, their thinking becomes more relation-

al in these domains. Like Goswami we argue that analogy is a tool that drives knowledge acquisition in children. In DORA comparison (of which analogical mapping is a major component) is the process that drives learning novel relational concepts. Perhaps the most important difference between our account and those of Gentner and Goswami is that ours is specified in much greater algorithmic detail. Our account therefore more explicitly links the cognitive operations underlying representational change during development with those that later exploit those changes (e.g., in the service of analogy, reasoning, etc.). Our account also arguably makes more detailed predictions about the nature and course of representational change in cognitive development. In starker contrast to Goswami, we argue that early in development object similarity, without regard to relational similarity, drives analogy and the discovery of relational concepts. Goswami, in contrast, argues that relational and object-based (featural) responding compete, and thus must both be in place, from the very beginning.

Our proposal is largely orthogonal to Halford's *relational complexity theory* (Halford, 1993; Halford et al., 1998). Relational complexity theory is concerned primarily with the formal complexity of reasoning problems (i.e., in terms of the number of role bindings that must be considered simultaneously in order to solve a problem) and with the implications of complexity for children's ability to solve various problems as a function of maturational changes in their WM capacity (i.e., the number of role bindings they can hold in WM simultaneously). For example, tasks that require representing binary relations (two role-bindings), such as *chases* (John, Mary), place fewer demands on WM—and should therefore be easier—than tasks that require ternary relations (three role bindings), such as *gives* (John, Mary, book). According to relational complexity theory, as children mature, their WM capacity increases, allowing them to represent, and thus reason about, more complex relations. However, relational complexity theory does not speak to the question of where these relations come from in the first place. The DORA model, as presented here, is silent on the question of whether WM capacity changes during development. (However, it is worth noting that Hummel and Holyoak, 1997, demonstrated that it is possible to simulate the developmental changes reported by Halford & Wilson, 1980, Experiment 1, by manipulating LISA's WM capacity.) Rather, the focus of the present work is the manner in which comparison promotes the kinds of representational changes that take place during cognitive development. At the same time, it is important to acknowledge that, as argued by Halford and colleagues, the complexity of the relational concepts that can be learned by experience (e.g., as in DORA), will necessarily be sharply constrained by the learner's WM capacity.

Implications of binding by asynchrony of firing

DORA's ability to discover relational concepts from unstructured examples stems from two more basic abilities—both of which derive from its use of systematic asynchrony of firing to represent role-filler bindings. The first is its ability to represent objects, attributes and relational roles using a single pool of semantic features, and simultaneously specify which features are acting as roles and which as fillers at any given time. This ability makes it possible for DORA to learn explicit predicates from examples of objects (e.g., to transition from “red” as an implicit object feature to “*red* (*x*)” as an explicit predicate). The second is DORA's ability to exploit the temporal regularity of temporal binding to link role-filler pairs into complete relations (e.g., to transition from *more+size* (*x*) and *less+size* (*y*) to *bigger-than* (*x*, *y*)).

Hummel and Holyoak's (1997, 2003a) LISA model demonstrated that, with the right architecture and knowledge representation, a connectionist system that can perform dynamic binding can account for a wide range of phenomena in relational reasoning (see Hummel & Holyoak, 2003, for a thorough review). DORA is a generalization and extension of LISA that represents role-filler bindings, not by role-filler synchrony of firing (as in LISA), but by role-filler asynchrony. If level-of-asynchrony (i.e., at the level of role-filler bindings, as in LISA, or the level of individual roles and fillers, as in DORA) is assumed to be a function of attentional focus, then DORA takes LISA as a special case: With attention directed to role-filler bindings (i.e., so that separate RBs fire out of synchrony with one another, but within RBs, roles fire in synchrony with their fillers), DORA becomes LISA. (Although we have not discussed these simulations here, we have used DORA to simulate the same phenomena in analogical reasoning and schema induction for which LISA was designed.) And with attention directed to the level of individual roles and fillers, as in the simulations reported here, DORA becomes a model of relational learning and cognitive development.

This distinction between role-filler synchrony vs. role-filler asynchrony as a function of attention makes a broad class of predictions that, to our knowledge, remain untested. Although DORA's use of role-filler asynchrony has the advantage that it makes it possible to use a common pool of features to represent both objects and relational roles—and thus makes relations learnable from examples—it has the disadvantage that it effectively cuts WM capacity in half, relative to role-filler synchrony. If WM contains roughly four or five “slots” (see, e.g., Cowan, 2001), and if role-filler bindings are represented by role-filler synchrony (i.e., with each synchronized set of neurons occupying one “slot”), then the capacity of WM should be about four to five role-filler bindings (Hummel & Holyoak, 1997, 2003). However, if role-filler bindings are represented by role-filler asynchrony (i.e., such that each role or filler must

occupy its own “slot” in WM), then the capacity of WM drops to two or two-and-a-half role bindings.

Although LISA successfully simulates many cognitive phenomena using role-filler synchrony, the theoretical considerations presented here suggest that learning relations from examples requires role-filler asynchrony. This divide between analogy and analogical inference on the one hand (the domain of LISA), and relation learning and relational concept acquisition on the other (the domain of DORA) suggests that some cognitive operations (such as relation learning) may require twice the WM resources of some others (such as analogy-making). This issue remains open for empirical investigation, although, consistent with this general prediction, Saiki (2003) has demonstrated that when subjects have to update representations of multiple object properties (e.g., to detect color changes while tracking motion), visual WM capacity is reduced to 1-2 role-bindings (i.e., cut in half as predicted by DORA). This result, in combination with other findings suggesting the capacity of visual WM is closer to four or five (e.g., Luck & Vogel, 1997), is strikingly consistent with DORA’s prediction that different tasks might impose different capacity demands on WM.

Finally, the fact that DORA uses a common representational basis to represent both roles and fillers is broadly consistent with our ability to appreciate that a property is the same regardless of whether it is represented as a property of an object (as in *red* (ball)), or as an element of a more general class (as in *color-of* (ball, red)). This ability is fundamentally beyond the reach of any model that requires objects (or object semantics) and predicates (or predicate semantics) to be fundamentally different “data types”—namely, all other models of relational reasoning of which we are aware.

Additional Novel Predictions

DORA makes several novel predictions, some of which we have noted previously. Here we note a few additional predictions.

Trajectory of Learning. DORA predicts that a child must learn the single place predicates/roles that compose a relation prior to learning the relation. This predicate-then-relation trajectory should hold for all kinds of relations, both categorical relations over metric dimensions, as emphasized in the simulations reported here, and more abstract relations such as *chases*, *loves*, and *ameliorates*. In addition, children’s early role and relation representations should initially be “dirty” (context-bound, and attached to many irrelevant features) and only later become more refined as a result of additional comparisons.

Forming Relations. Mapping role-filler sets should produce relations composed of those role-filler sets even if they are odd pairings. That is, if DORA compares strange sets of predicates (e.g., it compares *big* (ball1) and *light-colored* (ball2) to *big* (block1) and *light-colored* (block2)) and these sets of predicate-object pairs

have sufficient over-lap, then it will form an odd relation composed of these single-place predicate sets (e.g., *big-light-colored* (ball1, ball2)). This is a novel prediction of the model that is consistent with prior work on relation learning (e.g., Sera & Smith, 1987; Smith & Sera, 1992), and which we are currently testing with children. As elaborated below, this property of the model also suggests that there must be cognitive constraints on the conditions under which multiple roles will be placed into WM at the same time for the process of forming new relations.

Confounded Features. Because of the manner in which DORA learns relations, if two or more properties consistently covary with one another, then they should not be represented separately. In other words, if, in DORA’s training corpus, a feature such as “shiny” always co-occurs with the feature “red”, then DORA will assume that all red things are shiny. DORA’s representations will reflect the statistical structure of the particular instances experienced by an individual. There is some developmental evidence in support of this prediction (e.g., Sera & Smith, 1987; Smith & Sera, 1992).

Limitations and Open Problems

Although we have demonstrated that DORA can simulate a number of empirical phenomena, especially in cognitive development, there remain many open problems for which the model does not provide an account. We now consider some of the limitations of DORA and some open problems in the study of relation discovery.

Learning Different Kinds of Relations

Most of the relations DORA learned in the simulations reported here were relations with an underlying metric dimension (e.g., *bigger-than*). In the case of a relation like *bigger-than*, the semantic content of the relational roles is intuitive: The larger role entails more of the dimension size and the smaller role entails less of that dimension. In the case of a relation like *chases*, however, it is not just that one object is running in one direction and another object is running in the same direction (multiple objects going the same direction happens in many circumstances, such as road races, flocks of birds, etc., that involve no chasing whatsoever). Rather, the object bound to the *chased* role is being chased *precisely because* the object bound to the *chaser* role is chasing it, and vice-versa: each role seems to somehow refer to the other.

This property of *chases* seems, at first blush, to give it a qualitatively different character than a relation, such as *bigger*, defined over a metric dimension. However, as noted previously, any multi-place relation is formally equivalent to a linked set of single-place predicates (Mints, 2001). Thus, although *chases* seems qualitatively different from *bigger-than*, formally speaking, they are not as different as they first appear. The only neces-

sary difference between such relations resides in the semantic features composing them. Armed with the right set of semantic invariants, the routines DORA uses to learn *bigger-than* or *same-color* are equally able to learn a relation such as *chase*. When a child chases a dog or a butterfly, she has the opportunity to compare these experiences and predicate what it “feels like” to be chaser. The semantic features attached to the feeling are likely to include those of running and motion, but also excitement and a sense of pursuit. In this way, the *chaser* role acquires semantic content. Similarly, when she is chased a dog or her older sister, she will have the opportunity to compare those experiences with other experiences of being chased and predicate their shared features: again motion and excitement, but this time, a desire to evade and escape. And when she eventually compares sets of these experiences together, she has the opportunity to learn that *chase* is a two-role relation in which both objects are in motion, both experience excitement, but one wishes to catch while the other wishes to escape.

Importantly, DORA predicts that relations like *chase* and *bigger-than* should follow identical learning trajectories. Specifically, it predicts that a child should understand what it is to be a *chaser* or a *chased* before she understands the full-blown *chase* relation. Similarly, DORA predicts that the child might understand her feelings toward her cat (i.e., that she loves the cat) and her parents’ feelings toward her (that she is loved by her parents) before she has the full-blown *loves* relation. There is already some developmental evidence that children do not learn whole relations from scratch, but rather learn individual roles and put them together to form whole relations (e.g., Smith, 1984; Smith et al., 1988), and that children comprehend relational roles before they comprehend the full relations to which these relations belong (Smith, 1984; Sera & Smith, 1987; Smith & Sera, 1992; Smith et al., 1988).

The model, as it stands, does not speak to where the semantic invariants underlying *chases* or *loves* (or *ameliorates*) come from, but it does speak to the question of how they eventually become predicates that can take arguments—and can therefore eventually support complex relational reasoning—and to the question of how and why these abilities develop as they do. Our claim is that the human cognitive architecture, starting with whatever regularities it is given or can calculate from its environment as invariants, isolates those invariants in the objects and situations it is exposed to, and composes them into relational structures with which it describes and understands its world. Given any set of such invariants (or a means to calculate them), DORA’s learning mechanisms can convert those invariants into the explicit predicates and relations that make relational reasoning possible.

Constraints on relation discovery

As noted previously, left to run unchecked, DORA’s comparison and predication routines might generate a proliferation of predicates describing random object properties and relations. In order to prevent this kind of run-away re-representation, there need to be constraints on when these learning routines are invoked. At this point, we can only speculate on what these constraints might be, but the prior literature on concept acquisition and some general theoretical considerations can provide some guidance.

The process of comparison (i.e., analogical mapping) figures centrally in all the operations DORA performs. DORA cannot learn from two situations if it does not, or cannot, align them. As such, DORA predicts that any factors that promote explicit comparison and alignment should promote relation discovery and refinement, and any factors that reduce the likelihood of comparison or alignment should reduce the likelihood of relation discovery and refinement. Certainly, explicit direction to compare two or more objects (e.g., by a parent or teacher) is likely to encourage comparison.

A second factor that might encourage comparison is joint attention and, especially, language. Gentner and her colleagues (e.g., Gentner, 2003; Gentner & Namy, 1999; Gentner & Loewenstein, 2002; Gentner & Medina, 1998; Gentner & Ratterman, 1991; Loewenstein & Gentner, in press; Namy & Gentner, 2002) have argued that relational language plays a key role in relational learning, helping us to detect and retain relational patterns.

There are at least four ways in which language may help to guide and constrain predication. First, common labels may simply encourage comparison. Namy and Gentner (2002) found that providing common labels invites explicit comparison, while providing conflicting labels deters it. Second, using common labels for similar relational concepts also increases the likelihood of retrieving one instance of a relational concept given another as a cue (e.g., Gentner, et al., 1993; Ross, 1989), thereby increasing the likelihood that they will be compared. Third, common labels seem to increase the perceived similarity between instances (e.g., Sloutsky, 2003; Sloutsky & Fisher, 2004; Sloutsky & Lo, 1999; Sloutsky, Lo, & Fisher, 2001; Sloutsky & Napolitano, 2003; Smith, 2003), which might encourage comparison, alignment, and the formation of relational concepts. Finally, a common label can serve as an invariant property across instances that are structurally similar but with little featural similarity.

There are, no doubt, a number of constraints beyond those that we mention here. In future work all of these open problems should be addressed more fully. In the mean time, however, DORA provides a framework within which it is possible to discuss, at a detailed algorithmic level, how these constraints might manifest themselves in the processes that enable the discovery of

relational representations from simple holistic beginnings.

Acknowledgements

We are grateful to Keith Holyoak, Bob Morrison, Derek Devnich, Linda Smith, Steve Engel, Phil Kellman, Barbara Knowlton, Ed Stabler, and Tom Wickens for very helpful discussions of this work. We thank three reviewers for the helpful comments on previous versions of this manuscript. The first author was supported in part by the UCLA Department of Psychology Dissertation Year Grant and the National Institute for Child Health and Human Development (T32 HD07475). Portions of this research were presented at the 2005 meeting of the Cognitive Science Society.

References

- Albrecht, D. G., & Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Journal of Neurophysiology*, 7, 531-546.
- Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: LEA.
- Asaad, W. F., Rainer, G., & Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron*, 21, 1399-1407.
- Baillargeon, R. (2004). Infants' Physical World. *Current Directions in Psychological Science*, 13, 9-94.
- Bassok, M., Wu, L., & Olseth, K. L. (1995). Judging a book by its cover: Interpretive effects of content on problem-solving transfer. *Memory and Cognition*, 23, 354-367.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, 2, 41-55.
- Bowdle, B., & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology*, 34, 244-286.
- Chen, Z., Sanchez, R., & Campbell, T. (1997). From beyond to within their grasp: Analogical problem solving in 10- and 13-month-olds. *Developmental Psychology*, 33, 790-801.
- Clearfield, M. W., & Mix, K. S. (1999). Number versus contour length in infants' discrimination of small visual sets. *Psychological Science*, 10, 408-411.
- Colunga, E., Smith, L. B. (2005) From the Lexicon to Expectations About Kinds: A Role for Associative Learning. *Psychological Review*, 112, 347-382.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87-114.
- Doumas, L. A. A., & Hummel, J.E. (2004a). A fundamental limitation of symbol-argument-argument notation as a model of human relational representations. In *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, 327-332.
- Doumas, L. A. A., & Hummel, J.E. (2004b). Structure mapping and relational predication. In *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, 333-338.
- Doumas, L. A. A., & Hummel, J. E. (2005). 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*, 73-94. Cambridge: Cambridge University Press.
- Dixon, J. A., & Bangert, A. S. (2004). On the spontaneous discovery of a mathematical relation during problem solving. *Cognitive Science*, 28, 433-449.
- Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology* (H. A. Ruger & C. E. Bussenius, Trans.). New York: Teachers College Press.
- Faulkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Feigenson, L., Carey, S. & Spelke, E. S. (2002). Infants' discrimination of number vs. continuous extent. *Cognitive Psychology*, 44, 33-66.
- Fisher, A. V., & Sloutsky, V. M. (2005). When induction meets memory: Evidence for gradual transition from similarity-based to category-based induction. *Child Development*, 76, 583-597.
- Foley, J. M. (1994). Human luminance pattern-vision mechanisms: Masking experiments require a new model. *Journal of the Optical Society of America*, 11, 1710-1719.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19, 141-205.
- Fuster, J. M. (1997). *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe* (3rd ed.). Philadelphia, PA: Lippincott-Raven.
- Gasser, M. & Colunga, E. (2001). Learning relational correlations. *International Conference on Cognitive Modeling*, 4, 91-96.
- Gasser, M. & Colunga, E. (2003). Pattern learning in infants and neural networks. In P. Quinlan (Ed.), *Cognitive approaches to developmental psychology*. Psychology Press.
- Gentner, D. (1977). If a tree had a knee, where would it be? Children's performance on simple spatial metaphors. *Papers and Reports on Child Language Development*, 13, 157-164.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, 59, 47-59.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*, 199-241. London: Cambridge University Press.
- Gentner, D. (2003). Why we're so smart. In D. Gentner and S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought*, 95-235. Cambridge, MA: MIT Press.
- Gentner, D., & Loewenstein, J. (2002). Relational language and relational thought. In J. Byrnes & E. Amsel (Eds.), *Language, Literacy, and Cognitive Development*, 87-120. Mahwah, NJ: LEA.
- Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52, 45-56.
- Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition*, 65, 263-297.
- Gentner, D., & Namy, L. (1999). Comparison in the development of categories. *Cognitive Development*, 14, 487-513.
- 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*, 225-277. London: Cambridge University Press.
- Gentner, D., Ratterman, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability

- from inferential soundness. *Cognitive Psychology*, 25, 524-575.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 277-300.
- 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.
- Goswami, U. (2001). Analogical reasoning in children. In Gentner, D., Holyoak, K. J., and Kokinov, B. N. (Eds.) *The analogical mind: Perspectives from cognitive science*, 437-470. Cambridge, MA: MIT Press.
- Goswami, U. (1992). *Analogical reasoning in children*. Hillsdale, NJ: LEA.
- Green, C. B., & Hummel, J. E. (2004). Relational perception and cognition: Implications for cognitive architecture and the perceptual-cognitive interface. To appear in B. H. Ross (Ed.) *The psychology of learning and motivation*. Vol 44. San Diego: Academic Press.
- Halford, G. S. (1993). *Children's understanding: The development of mental models*. Hillsdale, NJ: Erlbaum.
- Halford, Graeme S. (2005) Development of thinking. In Holyoak, K. J. & Morrison, R. G. (Eds.). *The Cambridge handbook of thinking and reasoning*, 529-558. New York, NY: Cambridge University Press.
- Halford, G. S., & Wilson, W. H. (1980). A category theory approach to cognitive development. *Cognitive Psychology*, 12, 356-411.
- Halford, G. S., Wilson, W. H., & Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Brain and Behavioral Sciences*, 21, 803-864.
- Heeger, D. J. (1992). Normalization of cell responses in cat visual cortex. *Visual Neuroscience*, 9, 181-197.
- Holyoak, K. J., & Hummel, J. E. (2000). The proper treatment of symbols in a connectionist architecture. In E. Dietrich and A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines*, 229-263, Hillsdale, NJ: 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, MA: MIT Press.
- Horn, D., & Usher, M. (1990). Dynamics of excitatory-inhibitory networks. *International Journal of Neural Systems*, 1, 249-257.
- Horn, D., Sagi, D., & Usher, M. (1992). Segmentation, binding and illusory conjunctions. *Neural Computation*, 3, 509-524.
- Hummel, J. E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. In E. Dietrich and A. Markman (Eds.), *Cognitive Dynamics: Conceptual Change in Humans and Machines*. Hillsdale, NJ: Erlbaum.
- Hummel, J. E. (2001). Complementary solutions to the binding problem in vision: Implications for shape perception and object recognition. *Visual Cognition*, 8, 489-517.
- Hummel, J. E., Biederman, I. (1990). Dynamic binding: A basis for the representation of shape by neural networks. *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, 614-621.
- Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, 99, 480-517.
- Hummel, J. E., Burns, B., & Holyoak, K. J. (1994). Analogical mapping by dynamic binding: Preliminary investigations. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory*, Vol. 2: *Analogical connections*. (pp. 416-445). Norwood, NJ: Ablex.
- 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. (2002). Analogy and creativity: Schema induction in a structure-sensitive connectionist model. In T. Dartnall (Ed.), *Creativity, Cognition and Knowledge: An Interaction*, 181-210. Westport, CT: Praeger.
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, 110, 220-263.
- Hummel, J.E., Holyoak, K.J., Green, C., Doumas, L.A.A., Devnich, D., Kittur, A., & Kalar, D.J. (2004). A Solution to the Binding Problem for Compositional Connectionism. In S.D. Levy & R. Gayler (Eds.), *Compositional Connectionism in Cognitive Science: Papers from the AAAI Fall Symposium*, 31-34. Menlo Park, CA: AAAI Press.
- Hummel, J. E., & Stankiewicz, B. J. (1996). An architecture for rapid, hierarchical structural description. In T. Inui and J. McClelland (Eds.), *Attention and Performance XVI: Information Integration in Perception and Communication* (pp. 93-121). Cambridge, MA: MIT Press.
- Keane, M. T., Ledgeway, T., & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, 18, 387-438.
- Kellman, P. J., Burke, T. & Hummel, J. E. (1999). Modeling perceptual learning of abstract invariants. *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society*, 264-269. Mahwah, NJ: Erlbaum.
- Kokinov, B. N., & Petrov, A. A. (2001). Integration of memory and reasoning in analogy-making: The AMBR model. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science*, 59-124. Cambridge, MA: MIT Press.
- Kotovskiy, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, 67, 2797-2822.
- Kubose, T. T., Holyoak, K. J., & Hummel, J. E. (2002). The role of textual coherence in incremental analogical mapping. *Journal of Memory and Language*, 47, 407-435.
- Kurtz, K. J. & Burkina O. (2004). Learning Relational Categories by Comparison of Paired Examples. In *Proceedings of the Twenty Second Annual Conference of the Cognitive Science Society*.
- Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 299-321.
- Larkey, L. & Love, B. (2003). CAB: Connectionist Analogy Builder. *Cognitive Science*, 27, 781-794.
- Loewenstein, J., & Gentner, D. (in press). Relational language and the development of relational mapping. *Cognitive Psychology*.
- Love, B. C. (1999). Utilizing time: Asynchronous binding. *Advances in Neural Information Processing Systems*, 11, 38-44.
- Lovett, M. C. & Anderson, J. R. (2005) Thinking as a production system. In K. J. Holyoak & R. Morrison (Eds.) *Cambridge Handbook of Thinking and Reasoning*.

- Luck, S. J. & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279-281.
- 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*, 99, 587-604.
- Mandler, J. M. (2004). *The Foundations of Mind: Origins of Conceptual Thought*. Oxford University Press.
- Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Erlbaum.
- Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, 25, 431-467.
- Marshall, J. A. (1995). Adaptive pattern recognition by self-organizing neural networks: Context, uncertainty, multiplicity, and scale. *Neural Networks*, 8, 335-362.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation: Advances in Research and Theory*, 24, 109-165. San Diego: Academic Press.
- Medin, D. L., Goldstone, R. L., & Gentner, D. (1994). Respects for similarity. *Psychological Review*, 100, 254-278.
- Mints, G. E. (2001). Arithmetic, Formal. In M. Hazewinkel (Ed.) *Encyclopedia of Mathematics* (pp. 63-64). Berlin: Springer.
- Morrison, R. G., Doumas, L. A. A., & Richland L. E. (2006). The a unified account of the dual constraints of capacity and complexity on children's relational reasoning. *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*.
- Morrison, R. G., Krawczyk, D. C., Holyoak, K. J., Hummel, J. E., Chow, T. W., Bruce L. Miller, B. L., & Knowlton, B. J. (2004). A Neurocomputational Model of Analogical Reasoning and its Breakdown in Frontotemporal Lobar Degeneration. *Journal of Cognitive Neuroscience*, 16, 260-271.
- Munakata, Y. & O'Reilly, R.C. (2003). Developmental and Computational Neuroscience Approaches to Cognition: The Case of Generalization. *Cognitive Studies*, 10, 76-92.
- Namy, L. L., & Gentner, D. (2002). Making a silk purse out of two sow's ears: Young children's use of comparison in category learning. *Journal of Experimental Psychology: General*, 131, 5-15.
- Namy, L. L., Smith, L. B., & Gershkoff-Stowe, L. (1997). Young children's discovery of spatial classification. *Cognitive Development*, 12, 163-184.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 14(4), 700-708.
- Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development*, 5, 193-207.
- Oakes, L. M. & Madole, K. L. (2003). Principles of developmental change in infants' category formation. In D. H. Rakison & L. M. Oakes (Eds), *Early Category and Concept Development*. Oxford: Oxford University Press.
- O'Reilly, R.C. & Busby, R.S. (2002). Generalizable Relational Binding from Coarse-coded Distributed Representations. In T. G. Dietterich, S. Becker, and Z. Ghahramani (Eds.), *Advances in Neural Information processing Systems (NIPS) 14*. Cambridge, MA: MIT Press.
- O'Reilly, R.C., Busby, R. S. & Soto, R. (2003). Three forms of binding and their neural substrates: Alternatives to temporal synchrony. In A. Cleeremans (Ed.), *The Unity of Consciousness: Binding, Integration, and Dissociation*, 168-192. Oxford: Oxford University Press.
- Quinn, P., Cummins, M., Kase, J., Martin, E., & Weissman, S. (1996). Development of categorical representations for above and below spatial relations in 3- to 7-month old infants. *Developmental Psychology*, 32, 942-950.
- 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 Experimental Child Psychology*, 94, 249-273.
- Ross, B. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 456-468.
- Ross, B. H., Perkins, S. J., & Tenpenny, P. L. (1990). Reminding-based category learning. *Cognitive Psychology*, 22, 260-492.
- Ross, B. H., & Warren, J. L. (2002). Learning abstract relations from using categories. *Memory & Cognition*, 30, 657-665.
- Rovee-Collier, C. K., & Fagen, J. W. (1981). The retrieval of memory in early infancy. In L. P. Lipsitt (Ed.), *Advances in infancy research (Vol. 1)*, 225-254. Norwood, NJ: Ablex.
- Saiki, J. (2003). Feature binding in object-file representations of multiple moving objects. *Journal of Vision*, 3, 6-21.
- Salvucci, D. D., & Anderson, J. R. (2001). Integrating analogical mapping and general problem solving: The path mapping theory. *Cognitive Science*, 25, 67-110.
- Sandhofer, C. M., & Smith, L. B., (2001). Why children learn color and size words so differently: Evidence from adults' learning of artificial terms. *Journal of Experimental Psychology: General*, 130, 600-617.
- Sera, M. & Smith, L. B. (1987). Big and little: "Nominal" and relative uses. *Cognitive Development*, 2, 89-111.
- Shastri, L. (1996). Temporal Synchrony, Dynamic Bindings, and SHRUTI: A Representational but Non-Classical Model of Reflexive Reasoning. *Behavioral and Brain Sciences*, 19, 1996, 331-337.
- Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, 16, 417-494.
- Sheya, A., and Smith, L. B. (2006). Perceptual features and the development of conceptual knowledge. *Journal of Cognition and Development*, 7, 455-476.
- Sloutsky, V.M. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences*, 7, 246-251.
- Sloutsky, V. M., & Fisher, A. V. (2004). Induction and categorization in young children: A similarity-based model. *Journal of Experimental Psychology: General*, 133, 166-188.
- Sloutsky, V. M., & Lo, Y.-F. (1999). How much does a shared name make things similar? Part 1: Linguistic labels and the development of similarity judgment. *Developmental Psychology*, 35, 1478-1492.
- Sloutsky, V. M., Lo, Y.-F., & Fisher, A. (2001). How much does a shared name make things similar? Linguistic Labels,

- Similarity and the Development of Inductive Inference. *Child Development*, 72, 1695-1709.
- Sloutsky, V. M., & Napolitano, A.C. (2003). Is a picture worth a thousand words? Preference for auditory modality in young children. *Child Development*, 74, 822-833.
- Smith, L. B. (1984). Young children's understanding of attributes and dimensions. *Child Development*, 55, 363-380.
- Smith, L. B. (1985). Young children's attention to global magnitude: Evidence from classification tasks. *Journal of Experimental Child Psychology*, 39, 472-491.
- Smith, L. B. (1989). From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning*, 147-177. Cambridge: Cambridge University Press.
- Smith, L. B. (2003) Learning to Recognize Objects. *Psychological Science*, 14 (3).
- Smith, L. B., Rattermann, M. J., & Sera, M. (1988). "Higher" and "lower": Comaparative interpretations by children. *Cognitive Development*, 3, 341-357.
- Smith, L. B., & Sera, M. D. (1992). A developmental analysis of the polar structure of dimensions. *Cognitive Psychology*, 24, 99-142.
- Spalding, T. L., & Ross, B. H. (1994). Comparison-based learning: Effects of comparing instances during category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 1251-1263.
- St. John, M. F. (1992). The Story Gestalt: A model of knowledge-intensive processes in text comprehension. *Cognitive Science*, 16, 271-302.
- Thomas, J. P., & Olzak, L. A. (1997). Contrast gain control and fine spatial discriminations. *Journal of the Optical Society of America*, 14, 2392-2405.
- Thorndike, E. L. (1914). *The psychology of learning*. New York: Teachers College Press.
- Usher, M. & Niebur, E. (1996). Modeling the temporal dynamics of IT neurons in visual search: a mechanism for top-down selective attention. *Journal of Cognitive Neuroscience*, 8, 311-327.
- 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.
- von der Malsburg, C., & Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67, 233-242.
- Yamauchi, T. & Markman, A. B. (1998). Category learning by inference and classification. *Journal of Memory and Language*, 39, 124-148.
- Yamauchi, T., & Markman, A. B. (2000). Learning categories composed of varying instances: The effect of classification, inference and structural alignment. *Memory & Cognition*, 28, 64-78.

Appendix A: Details of DORA's Operation

In DORA firing is organized into *phase sets*. The phase set is the set of units in the driver that is currently active and firing out of synchrony with one another—i.e., it is the set of things DORA is currently “thinking about”. A single phase set runs until each RB in the set has fired three times. All the routines described below and in the text (e.g., retrieval, mapping, refinement, etc.) are allowed to run for one phase set. The general sequence of events in DORA's operation is outlined below. The details of these steps, along with the relevant equations and parameter values, are provided in the subsections that follow. Within reasonable ranges, DORA is very robust to the values of the parameters. Throughout the equations in this Appendix, we will use the variable a to denote a unit's activation, n its (net) input, and w_{ij} to denote the connection from unit j to unit i .

1. Bring a prop or a set of props into active memory in the driver, D , (as designated by the user).
2. Initialize the activations of all units in the network to 0.
3. Select the firing order of propositions in D to become active. (In all the simulations described here, firing order is either set by the user or at random. However see Hummel & Holyoak, 2003, for a detailed description of how a system like DORA can set its own firing order according to the constraints of pragmatic centrality and text coherence.)
4. Run the phase set. Repeat the following until each RB in D has fired three times:
 - 4.1. Select the proposition, P , in D that is currently at the head of the firing order.
 - 4.2. Select the RB, RB_C , in P that is at the head of the firing order (chosen at random).
 - 4.3. Update the network in discrete time steps until the global inhibitor fires. On each time step t do:
 - 4.3.1. Set input to RB_C to 1.
 - 4.3.2. Update modes of all P units in R (the recipient set). (Although we do not use higher-order relations in any of the simulations described in the text and therefore the mode of P units is always set to 1, we include this step for completeness)
 - 4.3.3. Update inputs to all token units in P_S (i.e., all P, RB, and PO units in P).
 - 4.3.4. Update input to the PO inhibitors.
 - 4.3.5. Update input to the RB inhibitors.
 - 4.3.6. Update input to the local inhibitor.
 - 4.3.7. Update input to the global inhibitor.
 - 4.3.8. Update input to all semantic units.
 - 4.3.9. Update input to all token units in the recipient, R , and the emerging recipient, N .
 - 4.3.10. Update activations of all units in the network.
 - 4.3.11. Update all mapping hypotheses (if mapping is licensed).
 - 4.3.12. Run retrieval routine (if retrieval is licensed).
 - 4.3.13. Run comparison based learning (if learning is licensed).
 - 4.3.13.1. If learning from objects not already bound to predicates (i.e., if the RB that is currently most active is connected to only 1 PO) run comparison-based predication.
 - 4.3.13.2. Otherwise (i.e., if the active RB in D is connected to two POs [an object and a role]) run refinement learning:

- 4.3.13.2.1. Run relation formation.
- 4.3.13.2.2. Run predicate refinement.
- 5. Update mapping connections.

A complete *analog* (i.e., story, event, situation, etc.) is represented by the collection of token (P, RB and PO) units that together represent the propositions in that analog. Token units are not duplicated within an analog (e.g., within an analog, each proposition that refers to Fido connects to the same “Fido” unit), but separate analogs have non-overlapping sets of token units (e.g., Fido will be represented by one PO unit in one analog and by a different PO in another analog). However, all analogs are connected to the same pool of semantic units. The semantic units thus represent general types (e.g., dogs, large things, red things, etc.) and token units represent instantiations (i.e., tokens) of those things in specific analogs (Hummel & Holyoak, 1997, 2003). For example, if in some analog, the token (PO) unit “Fido” is connected to the semantics “animal”, “dog”, “furry” and “Fido”, then it is a token of an animal, a dog, a furry thing and of the particular dog Fido.

Step 4.3.2. Update P unit modes for recipient P units

P units in all propositions operate in one of three modes: Parent, child, and neutral, as described by Hummel and Holyoak (1997, 2003). Although the idea of units firing in “modes” sounds “non-neural”, Hummel, Burns & Holyoak (1994) describe how it can be accomplished with two or more auxiliary nodes with multiplicative synapses. A P unit in parent mode is operating as the overarching structure of a proposition. Parent P units excite and are excited by RBs below them to which they are connected. In child mode, a P unit is acting as the argument of a higher-order proposition. Child P units excite and are excited only by RBs to which they are upwardly connected. In neutral mode, P units take input from all RBs to which they are upwardly connected. The mode of P units in the driver are set at the beginning of each run by the rule given in the order of operations outline above. Each P unit i in R updates its mode, m_i , according to:

$$m_i = \begin{cases} \text{Parent}(1), & RB_{above} < RB_{below} \\ \text{Child}(-1) & RB_{above} > RB_{below} \\ \text{Neutral}(0) & \text{otherwise} \end{cases}, \quad (A1)$$

where RB_{above} is the summed input from all RB units to which i is upwardly connected (i.e., relative to which, i serves as an argument) and RB_{below} is the summed input from all RB units to which it is downwardly connected.

Steps 4.3.3. Updating input to driver token units

P units

Each P unit, i , in D in parent mode updates its input as:

$$n_i = \sum_j a_j - \sum_k 3a_k, \quad (A2)$$

where j are all RB units below P unit i to which i is connected and k are all other P units in D that are currently in parent mode. P units in D in child mode update their inputs by:

$$n_i = \sum_j a_j - \sum_k a_k - s \sum_l a_l - s \sum_m 3a_m, \quad (A3)$$

where j are RB units to which i is upwardly connected, k are other P units in the driver that are currently in child mode, l are all PO units in the driver that are not connected to the same RB as i , and m are all PO units that are connected to the same RB (or RBs) as i . When DORA is operating in binding-by-asynchrony mode, $s = 1$; when it is operating in binding-by-synchrony mode (i.e., like LISA), $s = 0$.

RB units

RB units in the driver update their inputs by:

$$n_i = \sum_j a_j + \sum_k a_k - \sum_l 3a_l - 10I_i, \quad (A4)$$

where j are all P units in parent mode to which RB unit i is upwardly connected, k are all PO units connected to i , l are all other RB units in D , and I_i is the activation of the RB inhibitor yoked to i .

PO units

PO units in the driver update their input by:

$$n_i = \sum_j a_j G - \sum_k a_k - \sum_l a_l - \sum_m 3a_m - \sum_n a_n - 10I_i, \quad (A5)$$

where j are all RB units to which PO unit i is connected, G is a gain parameter attached to the weight between the RB and its POs (POs learned via DORA's comparison based predication algorithm have $G=2$ and 1 otherwise), k are P units in D that are currently in child mode and not connected to the same RB as i , l are all PO units in the driver that are not connected to the same RB as i , m are PO units that are connected to the same RB (or RBs) as i , and I_i is the activation of the PO inhibitor yoked to i .

Steps 4.3.4 and 4.3.5. Update input to the RB and PO inhibitors

Every RB and PO unit is yoked to an inhibitor. Both RB and PO inhibitors integrate input over time as:

$$n_i^{(t+1)} = n_i^{(t)} + \sum_j a_j w_{ij}, \quad (A6)$$

where t refers to the current iteration, j is the RB or PO unit yoked to inhibitor unit i , and w_{ij} is the weight between RB or PO inhibitor i and its yoked RB or PO unit. RB inhibitors are yoked only to their corresponding RB. However, PO inhibitors are yoked both to their corresponding PO and all RB units in the same analog. As a result, at any given instant, PO inhibitors receive twice as much input as RB inhibitors, and so reach their activation threshold twice as fast. POs therefore oscillate twice as fast as RBs. For the current instantiation of DORA the connection weight between all POs and RBs and their inhibitors is set to 1. The purpose of the PO and RB inhibitors is to establish the time-sharing that carries role-filler binding information and allows DORA to dynamically bind roles to fillers. All PO and RB inhibitors become refreshed ($n_i = 0$) when the global inhibitor (Γ_G ; described below) fires.

Steps 4.3.6 and 4.3.7. Update the local and global inhibitors

The local and global inhibitors, Γ_L and Γ_G respectively (see e.g., Horn and Usher, 1990; Horn et al., 1992; Usher and Nieber, 1996; von der Malsburg and Buhman, 1992), serve to coordinate the activity of units in the driver and recipient sets. The local inhibitor is inhibited to inactivity ($\Gamma_L = 0$) by any PO in the driver with activation above Θ_L ($= 0.5$), and becomes active ($\Gamma_L = 10$) when no PO in the driver has an activity above Θ_L . During asynchronous binding, the predicate and object POs time-share. There is a period during the firing of each role-filler pair after the one PO fires and before the other PO becomes active when no PO in the driver is very active. During this time the local inhibitor becomes active, inhibiting all PO units in the recipient to inactivity. Effectively, Γ_L serves as a local refresh signal, punctuating the change from predicate to object or object to predicate firing in the driver, and allowing the units in the recipient to keep pace with units in the driver.

The global inhibitor works similarly. It is inhibited to inactivity ($\Gamma_G = 0$) by any RB in the driver with activation above Θ_G ($= 0.5$), and becomes active ($\Gamma_G = 10$) when no RB in the driver is active above threshold. During the transition between RBs in the driver there is a brief period when no driver RB are active above Θ_G . During this time Γ_G inhibits all units in the recipient to inactivity, allowing units in the recipient to keep pace with those in the driver.

Step 4.3.8. Update input to semantic units

The input to semantic unit i is:

$$n_i = \sum_{j \in \sigma \in (DR)} (a_j w_{ij}), \quad (A7)$$

where j is a PO unit in the driver or recipient.

Step 4.3.9. Updating input to recipient token units

Input to all token units in the recipient and emergent recipient are not updated for the first 5 iterations after the global or local inhibitor fires. This is done in order to allow units in the recipient and emergent recipient to respond to the pattern of activation imposed on the semantic units by the driver PO unit that wins the competition to become active after an inhibitor fires.

P units

P units in parent mode in the recipient update their inputs by:

$$n_i = \sum_j a_j + M_i - \sum_k 3a_k - \Gamma_G, \quad (A8)$$

where j are all RB units to which P unit i is downwardly connected, k are all other P units in the recipient currently

in parent mode and M_i is the mapping input to i :

$$M_i = \sum_j (a_i (3w_{ij} - \text{Max}(\text{Map}(i)) - \text{Max}(\text{Map}(j)))) \quad (A9)$$

where j are token units of the same type as i in the driver (e.g., if i is a RB unit, j is all RB units in the driver), $\text{Max}(\text{Map}(i))$ is the highest of all unit i 's mapping connections, and $\text{Max}(\text{Map}(j))$ is the highest of all unit j 's mapping connections. When a token unit in the driver fires, it excites all units to which it maps, and inhibits all units of the same type to which it does not map.

P units in child mode in the recipient update their inputs by:

$$n_i = \sum_j a_j + M_i - \sum_k a_k - \sum_l a_l - \sum_m 3a_m - \Gamma_G, \quad (A10)$$

where j are all RB units to which i is upwardly connected, M_i is the mapping input to i , k are all other P units in the recipient currently in child mode, l are POs in the recipient that are not connected to the same RB (or RBs if i is connected to multiple RBs) as i , and m are PO units connected to the same RB (or RBs) as i .

RB units

RB units in the recipient update their input by:

$$n_i = \sum_j a_j + \sum_k a_k + \sum_l a_l + M_i - \sum_m 3a_m - \Gamma_G, \quad (A11)$$

where j are P units currently in parent to which RB unit i is upwardly connected, k are P units currently in child mode to which i is downwardly connected, l are PO units to which unit i is connected, M_i is the mapping input to i , and m are other RB units in the recipient.

PO units

PO units in the recipient update their input by:

$$n_i = \sum_j a_j + SEM_i + M_i - \sum_k a_k - \sum_l a_l - \sum_m 3a_m - \sum_n a_n - \Gamma_G - \Gamma_L, \quad (A12)$$

where j is RB units to which PO unit i is connected, SEM_i is the semantic input to unit i , M_i is the mapping input to unit i , k is all PO units in the recipient that are not connected to the same RB (or RBs if unit i is connected to multiple RBs) as i , l is all other P units in the recipient currently in child mode that are not connected to the same RB (or RBs) as i , m is PO units connected to the same RB (or RBs) as i , and n is RB units in the recipient to which unit i is not connected (input from j is only included on phase sets beyond the first). SEM_i , the semantic input to i , is calculated as:

$$SEM_i = \frac{\sum_j a_j w_{ij}}{1 + \text{num}(j)} \quad (A13)$$

where j are semantic units, w_{ij} is the weight between semantic unit j and PO unit i , and $\text{num}(j)$ is the total number of semantic units i is connected to with a weight above θ ($=0.1$). Semantic input to POs is normalized by a Weber fraction so that the PO unit that best matches the current pattern of semantic activation takes the most semantic input, and semantic input is not biased by the raw number of semantic features that any given PO is connected to (see Hummel & Holyoak, 1997, 2003; Marshall, 1995).

Step 4.3.10. Update activations of all units in the network

All token units in DORA update their activation by the simple leaky integrator function detailed in Eq. 1 in the main text. The value of the growth parameter, γ , is 0.3, and the value of the decay parameter, δ , is 0.1. Semantic units do not inhibit one-another the way that token units do. However, in order to keep their activations bounded, their activations are divisively normalized: The activation of a semantic unit is equal to its input divided by the maximum input to any semantic unit (see Eq. 3 in the main text). There is physiological evidence for divisive normalization in the feline visual system (e.g., Albrecht & Geisler, 1991; Bonds, 1989; Heeger, 1992) and psychophysical evidence for divisive normalization in human vision (e.g., Foley, 1994; Thomas & Olzak, 1997).

RB and PO inhibitors, i , update their activations according to a simple threshold function:

$$a_i = \begin{cases} 1, & n_i \geq \Theta_{IN} \\ 0, & \text{otherwise} \end{cases} \quad (\text{A14})$$

where $\Theta_{IN} = 220$.

Step 4.3.11. Run update mapping hypotheses

The mapping algorithm used by DORA is adopted from Hummel and Holyoak (1997, 2003). During the phase set DORA learns mapping hypotheses between all token units in the driver and token units of the same type in the recipient (i.e., between P units, between RB units and between PO units in the same mode [described below]). All mapping hypotheses are initialized to zero at the beginning of a phase set. The mapping hypothesis between an active driver unit and a recipient unit of the same type is updated by Eq. 4 in the main text (i.e., by a simple Hebbian learning rule).

Step 4.3.12 Run retrieval

DORA uses a variant of the retrieval routine described by Hummel and Holyoak (1997). During retrieval propositions in the driver fire as described above for one phase set. Units in the dormant/LTM set become active in response to the patterns of activation imposed on the semantics by active driver POs. After all RBs in the driver have fired once, DORA retrieves propositions from LTM probabilistically using the Luce choice axiom:

$$L_i = \frac{R_i}{\sum_j R_j} \quad (\text{A15})$$

where L_i is the probability that P unit i will be retrieved into working memory, R_i is the maximum activation P unit i reached while during the retrieval phase set and j are all other P units in LTM. If a P unit is retrieved from LTM, the entire structure of tokens (i.e., RBs, POs, and P units that serve as arguments of the retrieved P unit) are retrieved into working memory.

Step 4.3.13. Run learning routines

In the current version of the model, learning is licensed whenever 70% of the driver token units map to recipient items (this 70% criterion is arbitrary, and in practice 100% of the units nearly always map; see the main text for a discussion of the limiting constraints on DORA's learning routines). If learning is licensed DORA invokes either its comparison-based-predication routine or its refinement learning routine. Comparison based predication is licensed when the driver contains single objects, not yet bound to any predicates (i.e., each RB in the driver is bound only to a single PO. Otherwise, DORA licenses refinement learning.

Step 4.3.13.1. Comparison-based predication

As detailed in the text, during comparison-based predication (CBP) for each PO in the driver that is currently active, and maps to a unit in the recipient with a mapping connection above the threshold Θ_{MAP} ($=0.5$), DORA infers an empty PO unit (i.e., a PO connected to no semantic features) in the recipient. The mode of the existing PO units in both the driver and recipient is set to 0 and the mode of the newly inferred PO is set to 1. While the mode of PO units is not important for the simulations described in the main text, it is important for assuring mappings from predicates to other predicates and from objects and other objects when DORA using synchronous binding (as when it is behaving like LISA). We mention it here and implement it in our code for the purposes of completeness. DORA learns connections between the new PO and all active semantics by the Eq. 6 in the main text. During CBP, DORA also infers a new RB unit in the recipient. The activation of each inferred unit is set to 1, and remains at 1 until Γ_G or Γ_L fires. DORA learns a connection with a weight of 1 between corresponding active token units (i.e., between P and RB units, and between RB and PO units) that are not already connected.

Step 4.3.13.2. Refinement Learning

During refinement learning DORA first runs its relation formation routine then its predicate refinement routine.

Step 4.3.13.2.1. When DORA successfully maps sets of role-filler bindings in the driver to sets of role-filler bindings in the recipient, the resulting pattern of firing on the recipient RB units is exactly like what would emerge from RB units joined by a common P unit (i.e., the RBs fire out of synchrony but in close temporal proximity, and within each RB, the POs fire out of synchrony but in close temporal proximity; as detailed in the main text). During relation formation DORA exploits these temporal patterns to join the recipient RBs (along with their respective POs) into a new proposition—i.e., a new relation. This process is accomplished as a case of SSL. When an RB in the recipient becomes active, if no P units are active in the

recipient, then a P unit is recruited in the recipient via SSL. The P unit remains active (activation=1) until the end of the phase set and learns connections to active RBs: A connection (weight=1) is formed between the new P and any active RB in the recipient (to which that P unit is not already connected). When the phase set ends, connection weights between the new P and any RBs to which it has connections are updated by the equation:

$$w_{ij} = \begin{cases} 1, & \left(\sum_k w_{ik} \right) \geq 2 \\ 0, & \text{otherwise} \end{cases} \quad (\text{A16})$$

where w_{ij} is the connection weight between P unit i and RB unit j , and w_{ik} is the connection weight between P unit i , and RB unit k where k is all RB units in the recipient. Essentially, if the new P has at least two connections to RB units (and the sum over k of w_{ik} is therefore ≥ 2), then DORA retains the connection weights between the recruited P and all RBs to which it has connections; if the sum is less than two, then it discards the connection (along with the P unit). This convention ensures that DORA does not learn P units that connect only to a single RB.

Step 4.3.13.2.2. As detailed in the text, during predicate refinement DORA learns a refined representation of mapped propositions or role-filler sets. For each PO in the driver that is currently active, and maps to a unit in the recipient with a mapping connection above the threshold Θ_{MAP} ($=0.5$), DORA infers an empty PO unit (i.e., a PO connected to no semantic features) in the emerging recipient. DORA learns connections between the new PO and all active semantics by Eq. 6 in the main text. In addition, DORA licenses self-supervised learning (SSL). During SSL, DORA infers token units in the emerging recipient to match active tokens in D (the driver). DORA will infer a structure unit in the emerging recipient in response to any unmapped structure unit in D . Specifically, as detailed in the text, if unit j in D maps to nothing in the emerging re-

cipient, then when j fires, it will send a global inhibitory signal to all units in the emerging recipient. This uniform inhibition, unaccompanied by any excitation, signals DORA to infer a unit of the same type (i.e., P, RB, PO) in the emerging recipient. Inferred PO units in the emerging recipient have the same mode as the active PO in the driver. The activation of each inferred unit in the emerging recipient is set to 1. DORA learns connections (weight=1) between corresponding active tokens (i.e., between P and RB units, and between RB and PO units) in the emerging recipient. To keep DORA's representations manageable (and decrease the runtime of the simulations), at the end of the phase set, we discard any connections between semantic units and POs whose weights are less than 0.1.

Step 5. Update mapping connections

At the end of every phase set mapping connections are updated. First, all mapping hypotheses are normalized divisively: Each mapping hypothesis, h_{ij} between units j and i , is divided by the largest hypothesis involving either unit i or j . Next it is normalized subtractively: The value of the largest hypothesis involving either i or j (not including h_{ij} itself) is subtracted from h_{ij} . The divisive normalization keeps the mapping hypotheses bounded between zero and one, and the subtractive normalization implements the one-to-one mapping constraint by forcing mapping hypotheses involving the same i or j to compete with one another (see Hummel & Holyoak, 1997). Finally, the mapping weights between each unit in the driver and the token units in the recipient of the same type are updated by the equation:

$$\Delta w_{ij} = \eta(1.1 - w_{ij})h_{ij}^0 \quad (\text{A17})$$

where Δw_{ij} is the change in the mapping connection weight between driver unit i and recipient unit j and η is a growth parameter set to 0.9. Δw_{ij} is truncated for values below 0 and above 1.

Appendix B: Details of Simulations

The details of the simulations reported in the main text are presented below. In the descriptions of the simulations we use the following notation. Propositions are labeled P# (e.g., P1 to denote the first proposition, P2 to denote the second, and so forth). Propositions are listed in a form of propositional notation with the predicate term in *italics* and the arguments listed in the parentheses that follow (e.g., *bigger* (object1, object2) for the proposition object1 is bigger than object2). For objects not bound to predicates we simply list the objects in parentheses without a predicate term in front of them. The semantics of each role of a relation and each object are listed under the Semantics subheading of each simulation. They are listed in the following form: name-of-role-or-object (semantics-attached-to-that-role-or-object). Names of roles are in *italics*, names of objects are not. For example, object1 (sem1, sem2, sem3) indicates that the object object1 is attached to the semantics sem1, sem2, and sem3.

General relation discovery

“Perception” simulation. P1 – P160: (object1) – (object160).

Semantics: All POs (object1 – object160) attached to “size” and one of the following (size-1, size-2, size-3, size-4, size-5, size-6, size-7, size-8, size-9, size-10), plus 10 additional semantics (sem1 – sem150).

“Memory” simulation: Learning a single relational concept. P1 – P50: (object1) – (object50).

Semantics: All POs (object1 – object50) attached to “size” and one of the following (size-1, size-2, size-3, size-4, size-5, size-6, size-7, size-8, size-9, size-10), plus 10 additional semantics (sem1 – sem500).

“Memory” simulations: Learning multiple relations from interleaved examples. P1 – P100: (object1) – (object100).

Semantics: All POs (object1 – object100) attached to one dimension at random. Dimensions include “size” and one of the following (size-1, size-2, size-3, size-4, size-5, size-6, size-7, size-8, size-9, size-10), “width” and one of the following

(width-1, width-2, width-3, width-4, width-5, width-6, width-7, width-8, width-9, width-10), “color” and one of the following (color-1, color-2, color-3, color-4, color-5, color-6, color-7, color-8, color-9, color-10), or “height” and one of the following (height-1, height-2, height-3, height-4, height-5, height-6, height-7, height-8, height-9, height-10). Also attached to each other dimension with a probability of .25. In addition, connected to 10 semantics (sem1 – sem500).

For “memory” simulations each simulation occurred as follows: First, randomly select an item from LTM. Seventy-five percent of the time that item is selected from the last 50 items DORA has learned (in the single relational concept simulation) or from the last 100 items DORA has learned (in the interleaved examples simulation). Twenty-five percent of the time that item is selected from the remainder of items in LTM. Second, place that item in the driver. Third, run DORA’s normal order of operations. First, run retrieval. During retrieval, more recently learned items are more likely to be retrieved. For the single relational concept simulation for each 50 simulations, the items learned during the previous 50 simulations and higher are weighted by .75 during retrieval, and all other items are weighted by .25. For the interleaved examples simulation for each 100 simulations, the items learned during the previous 100 simulations and higher are weighted by .75 during retrieval, and all other items are weighted by .25. If anything is retrieved, run mapping. If anything is mapped, run learning (comparison-based predication, relation formation, and refinement).

Cross-mapping

P1: *bigger* (produced during *general relation discovery* simulation) (dog, cat).

P2: *bigger* (cat, dog).

Semantics: *bigger-role*: (size, more); *smaller-role* (size, less); dog (dog1, dog2, dog3, dog4, dog5, dog6); cat (cat1, cat2, cat3, cat4, cat5, cat6).

Non-identical predicate mapping

P1: *bigger* (produced during *general relation discovery* simulation) (object1, object2).

P2: *bigger* (cat, dog).

Semantics: P2-*bigger-role* (more, x); P2-*smaller-role* (less, y); object1 (object1.1 – object1.6); object2 (object2.1 – object2.6); dog (dog1, dog2, dog3, dog4, dog5, dog6); cat (cat1, cat2, cat3, cat4, cat5, cat6).

N-ary restriction

P1: *bigger* (produced during *general relation discovery* simulation) (object1, object2).

P2: *bigger-role* (object3).

Semantics: P1-*bigger-agent* (more, x); P2-*bigger-patient* (less, y); object1 (object1.1 – object1.6); object2 (object2.1 – object2.6); object3 (object1.1 – object1.3, and object2.4 – object2.6).

Smith, Rattermann, & Sera (1988)

Simulation part 1. P1 – P160: (object1) – (object160).

Semantics: All POs (object1 – object160) attached to “height” and one of the following (height-1, height-2, height-3, height-4, height-5, height-6, height-7, height-8, height-9, height-10), plus 10 additional semantics (sem1 – sem150). In addition, POs attached to height-1 – height-3 are

also attached to (small), and POs attached to height-7 – height-10 are also attached to (big).

Simulation part 2. Driver P1: *more+height* (butterfly1, butterfly2).

Semantics: High trial: *more+height-role* (more, height); *less+height-role* (less, height); butterfly1 (butterfly1 flies fly buzzes high); butterfly2 (butterflyfly2 flies fly pretty high).

Low trial: *more+height-role* (more, height); *less+height-role* (less, height); butterfly1 (butterfly1 flies fly buzzes low); butterfly2 (butterflyfly2 flies fly pretty low)

Neutral trial: *more+height-role* (more, height); *less+height-role* (less, height); butterfly1 (butterfly1 flies fly buzzes); butterfly2 (butterflyfly2 flies fly pretty)

See text for details of propositions in LTM.

Smith (1984)

Simulation part 1. P1 – P160: (object1) – (object160).

Semantics: All POs (object1 – object160) attached to “color” and one of the following (red, yellow, blue, white, black, green, orange, purple, grey, pink), plus 10 additional semantics (sem1 – sem150).

Simulation part 2. Driver: Unbound objects: P1: (ball1)

Semantics: ball1 (ball1, ball, round, red, sphere, bouncy, medium, size-5).

Value-dependent relation: P1: *same-color+red* (ball1, ball2).

Semantics: *same-color+red-1* (same, color, red, 1); *same-color+red-2* (same, color, red, 2); ball1 (ball1, ball, round, red, sphere, bouncy, medium, size-5); ball2 (ball2, ball, round, red, sphere, bouncy, small, size-2).

Value-independent relations: P1: *same-color* (ball1, ball2).

Semantics: *same-color-1* (same, color, 1); *same-color-2* (same, color, 2); ball1 (ball1, ball, round, red, sphere, bouncy, medium, size-5); ball2 (ball2, ball, round, red, sphere, bouncy, small, size-2).

Recipient: Various proposition created during simulation part 1. To differentiate the roles of the same color relation, add semantic “1” to one of the roles and “2” to the other. This serves simply to differentiate the two roles of the *same-color* relation for the purposes of mapping. See text for additional details.

Gentner & Namy (1999)

Condition NC: P1: (apple).

P2: (ball).

P3: (banana).

Semantics: apple (round, fruit, + 6 semantics from semantic sem1 – sem15); ball (round + sem1 – sem10); banana (fruit + sem10 – sem15 + sem16 – sem20).

Condition C: P1: (apple).

P2: (watermelon).

P3: (orange).

P4: (grapes).

P5: (ball).

P6: (banana).

Semantics: apple (round, fruit, + 6 semantics from semantic sem1 – sem15); watermelon (round, fruit, + 5 semantics from semantic sem1 – sem15); orange (round, fruit, + 5 semantics from semantic sem1 – sem15); grapes (round, fruit, + 5 semantics from semantic sem1 – sem15); ball (round + sem1 – sem10); banana (fruit + sem10 – sem15 + sem16 – sem20).

Kotovsky & Gentner (1996)

Symmetry trial size correct: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, size-2, smallest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, size-5, biggest, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, size-8, smallest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, size-3, smallest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, size-6, biggest, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, size-9, smallest, black, square1, square2, + 4 random from sem1-sem50).

Symmetry trial color correct: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, grey, lightest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, lightest, darkest, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, lightest, lightest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, black, lightest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, black, darkest, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, black, lightest, black, square1, square2, + 4 random from sem1-sem50).

Monotonic-increase size trial correct: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, size-2, smallest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, size-5, medium-size, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, size-8, biggest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, size-3, smallest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, size-6, medium-size, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, size-9, biggest, black, square1, square2, + 4 random from sem1-sem50).

Monotonic-increase color trial correct: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, grey, lightest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, lightest, medium-dark, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, lightest, darkest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, black, lightest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, black, medium-dark, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, black, darkest, black, square1, square2, + 4 random from sem1-sem50).

Symmetry trial size incorrect: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, size-2, smallest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, size-5, biggest, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, size-8, smallest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, size-3, smallest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, size-6, smallest, square1, square2, + 4 random from sem1-sem50); object6 (right-side,

size, size-9, biggest, black, square1, square2, + 4 random from sem1-sem50).

Symmetry trial color incorrect: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, grey, lightest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, lightest, darkest, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, lightest, lightest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, black, darkest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, black, darkest, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, black, lightest, black, square1, square2, + 4 random from sem1-sem50).

Monotonic-increase size trial incorrect: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, size-2, smallest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, size-5, medium-size, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, size-8, biggest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, size-3, biggest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, size-6, smallest, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, size-9, medium-size, black, square1, square2, + 4 random from sem1-sem50).

Monotonic-increase color trial incorrect: P1: (object1, object2, object3).

P2: (object4, object5, object6).

Semantics: object1 (left-side, size, grey, lightest, circle1, circle2, + 4 random from sem1-sem50); object2 (middle, size, lightest, medium-dark, circle1, circle2, + 4 random from sem1-sem50); object3 (right-side, size, lightest, darkest, circle1, circle2, + 4 random from sem1-sem50); object4 (left-side, size, black, darkest, square1, square2, + 4 random from sem1-sem50); object5 (middle, size, black, lightest, square1, square2, + 4 random from sem1-sem50); object6 (right-side, size, black, medium-dark, black, square1, square2, + 4 random from sem1-sem50).

Dixon & Bangart (2004)

P1: (gear1).

P2: (gear2).

Semantics: gear1 (randomly pick 10 from a pool containing either parity-even or parity-odd and sem2-sem400); gear2 (randomly pick 10 from a pool containing either parity-even or parity-odd and sem2-sem400).