A Computational Account of the Development of the Generalization of Shape Information

Leonidas A. A. Doumas (adoumas@indiana.edu)
Department of Psychological and Brain Science, 1101 E. Tenth Street
Bloomington, IN 47405

John E. Hummel (jehummel@cyrus.psych.uiuc.edu)
Department of Psychology, 603 E. Daniel Street
Champaign, IL 61820

Abstract

Abecassis, Sera, Yonas, and Schwade (2001) have shown that young children represent shapes more metrically, and perhaps more holistically, than do older children and adults. How does a child transition from representing objects and events as undifferentiated wholes to representing them explicitly in terms of their attributes—including invariant aspects of objects' shapes—and the relations among those attributes? According to recognition-by-components theory objects are represented as collections of arranged geons. We propose that the transition from more holistic to more categorical shape processing is a function of the development of geon-like representations. We present a model, DORA, that was originally proposed to solve the problem of discovering relations, but can also learn representations of single geons from representations of multi-geon objects. We demonstrate that DORA follows the same trajectory humans do, originally generalizing shape more holistically and eventually, after more learning, generalizing categorically.

Keywords: Shape bias, relation learning, relation discovery, development, computational modeling.

Introduction

Numerous studies have shown that both children and adults apply similar labels to objects with similar shapes (e.g., Imai & Gentner, 1997; Landau, Smith, & Jones, 1988, 1992; Smith, 1995; Woodward & Markman, 1998). This phenomenon is often referred to as the shape bias. There is considerable debate about the origins of the shape bias (see e.g., Jones & Smith, 1993; Landau, Smith, & Jones, 1988, 1992; Woodward & Markman, 1998), but there are also questions about how children and adults can and do see two shapes as similar in the first place.

Abecassis, Sera, Yonas, and Schwade (2001) have shown that young children represent shapes more metrically, and perhaps more holistically, than do older children and adults. For example, presented with a slightly curved shape, a much more curved shape and a straight shape (i.e., where the metric difference in curvature is smaller between the slightly curved shape and the straight shape than between the slightly curved shape and the more curved shape), adults and older children tend to choose the more curved shape, rather than the straight shape, as more like the slightly curved shape. That is, to adults and older children, the two curved shapes are more alike than the slightly curved shape and the straight shape. Presumably the two curved shapes are more similar because they share the visual invariant "curved". By contrast younger children are more likely to say the slightly curved shape is like the straight shape, presumably because it is metrically closer and they are insensitive (or less sensitive) to the visual invariants "curved" and "straight". There is evidence for an analogous "relational shift" in cognitive development, in which young children appear to process objects and events rather holistically but, as they develop and learn, gradually come to represent them in terms of independent objects, relations and properties (see e.g., Gentner & Rattermann, 1991; Smith, 1989).

How does a child transition from representing objects and events as undifferentiated wholes to representing them explicitly in terms of their attributes—including invariant aspects of objects' shapes—and the relations among those attributes? This question is really two questions. The first is the question of how the invariant properties (e.g., "straight" vs. "curved" regardless of the degree of curvature) come to be detected from the holistic early visual input (e.g., as in V1) in the first place. The second is the question of how the child comes to notice that these invariants remain constant across separate objects. That is, how does the child discover that the "straightness" of one shape is, in some sense the same as the "straightness" of another? In other words, how does the child predicate *straightness* as an explicit property that retains its identity across different instances? We argue that these processes, discovery and predication, are necessary precursors to the shift from reliance on metric representations of shape to representations based more on abstract visual invariants.

We present our early efforts at understanding the answer to the second of these two questions. We present a model that relies on the processes of analogical mapping and intersection discovery to highlight shared abstract properties between separate systems (e.g., separate shapes) and subsequently predicates these similarities as explicit (i.e., symbolic) properties of the systems. Simulations suggest that these basic processes may permit the discovery and predication of geon like representations from examples containing multiple goens. In addition, learning more refined representations of geons leads to the more categorical (i.e., adult) processing of shapes observed by Abecassis et al. (2001).

Recognition by components

As noted by Abecassis et al. (2001) the problem of learning to generalize shapes (i.e., understanding that two shapes are similar and so the same name should be applied to both) is similar to the problem of recognizing objects in the world.

According to Biederman's (1987; Hummel & Biederman, 1992) Recognition-by-Components (RBC) theory of object recognition, adults visually represent objects in terms of a structural description that specifies the categorical relations among an object's parts. For example, a coffee mug would be represented as a curved cylinder (the handle) side-attached to a vertical cylinder (the body). A bucket would be represented as a curved top-attached to a vertical cylinder or truncated cone. The parts, in turn, are represented as *geons*: classes of generalized cylinders[1] that can be discriminated from one another based on categorical contrasts in their 3-D shape (which, in turn, can be detected based on non-accidental categorical contrasts in the object's 2-D image). For example, a cylinder has a curved cross section, parallel sides and a straight major axis; a cone has a curved cross section, nonparallel sides and a straight major axis; and a curved brick has a straight cross section, parallel sides and a curved major axis. Each geon is represented in terms of its general aspect ratio (i.e., degree of elongation: very squat [e.g., like the lid of a jar]; somewhat squat [like a tuna can]; neither squat nor elongated [like a cube or ball]; somewhat elongated [like a soup can] or very elongated [like a lamp post]), but importantly, a geon's metric properties (such as the precise degree of curvature of its major axis or the precise shape of its cross section) are otherwise completely left out of the description. The resulting categorical structural descriptions are naturally robust to variations in viewpoint and variations in an object's precise 3-D shape and thus provide a natural basis for recognizing objects in novel viewpoints and for recognizing different exemplars as members of the same basic-level category (e.g., a Toyota Camry and a Mazda 626 have identical geon-based descriptions).

If what allows us to recognize two objects as members of the same category is our ability to process and represent the geons that compose those objects, it follows that as we develop more refined representations of geons and their relations we transition from more holistic to more categorical shape generalization.

[1] A generalized cylinder is the 3-dimensional (3-D) volume produced by sweeping a 2-D shape (the cross-section) along an axis in the third dimension. For example, sweeping a circle along a straight axis produces a cylinder; sweeping the same cylinder along the same axis while linearly reducing its size produces either a cone (if the circle eventually disappears into a point) or a truncated cone (if he circle never completely disappears); and sweeping a rectangle along a curved axis results in a curved brick-like shape.

The DORA Model

DORA (Doumas & Hummel, 2005; Doumas, Hummel & Sandhofer, submitted) is a symbolic connectionist network that learns structured representations of relations from unstructured inputs. DORA dynamically binds distributed (i.e., connectionist) representations of relational roles and objects into explicitly relational (i.e., symbolic) structures. The resulting representations enjoy the advantages of both connectionist and traditional symbolic approaches to knowledge representation, while suffering the limitations of neither (see Doumas & Hummel, 2005).

DORA was developed as a model of the discovery of relational concepts. It has been used to simulate a wide range of cognitive phenomena including the discovery of novel relational concepts, the trajectory of children's relation learning, the idiosyncrasies of early relational concepts the progressive-alignment effect, and adult relation learning (see Doumas & Hummel, 2005; Doumas et al., submitted). In this paper we use DORA to simulate the discovery of simple geons from multi-geon objects and the development of the shape-bias in children and adults.

DORA uses a hierarchy of distributed and localist codes to represent relational structures. This hierarchy is adapted from Hummel & Holyoak's (1997, 2003) LISA model. At the bottom, "semantic" units represent the features of objects and roles in a distributed fashion. At the next level, these distributed representations are connected to localist units (*POs*) representing individual objects and relational roles. Localist role-binding units (*RBs*) link object and relational roles units into specific role-filler bindings. At the top of the hierarchy, localist *P* units link RBs into whole relational propositions (see Figure 1).

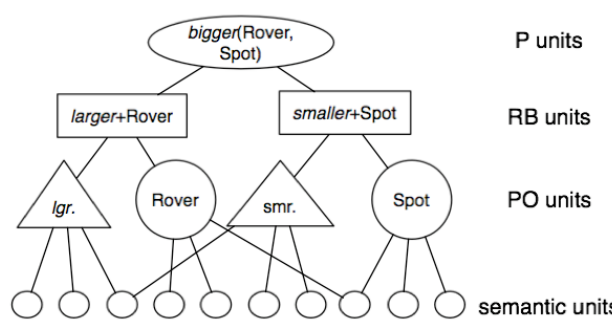

Figure 1. Example of a proposition in DORA. Triangles are used to denote roles and circles to denote objects for clarity. In DORA, the same types of units code both roles and objects.

At the most basic level, DORA uses analogical mapping to isolate shared properties of objects and to represent them as explicit structures. When DORA maps one object or structure onto another, corresponding elements of the two representations fire in synchrony. For example, if DORA compares a mouse and a hummingbird, then the nodes representing the mouse will fire in synchrony with those representing the hummingbird (Figure 2). Consequently, any semantic features that are shared by both compared objects (i.e., features common to both the hummingbird and the mouse) receive twice as much input as features connected to one but not the other. The network uses this firing pattern to recruit a new PO unit that learns connections to active semantics in proportion to their activation via simple Hebbian learning (i.e., DORA learns stronger connections to more active units; Figure 2b). The new PO thus becomes an explicit representation of the featural overlap of the compared hummingbird and mouse. So, in the case of comparing a hummingbird and a mouse, the network might form an explicit predicate representing "small" (and any other features they share, for example, "animal") due to their semantic overlap (Figure 2). Importantly, this new PO acts as an explicit predicate representation of the property *small* that can be dynamically bound to fillers.[2]

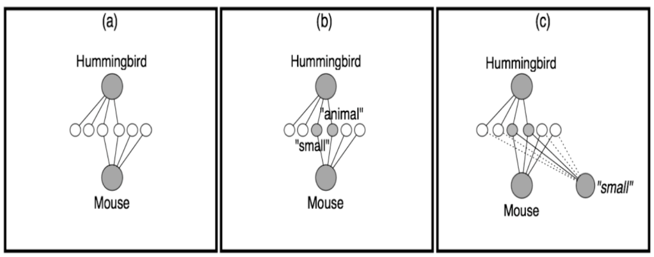

Figure 2. DORA learns a representation of "small" by comparing a hummingbird and a mouse. (a) When DORA compares a hummingbird and a mouse the units representing both become active simultaneously. (b) Feature units shared by both the hummingbird and mouse become most active (darker grey). (c) A new unit is recruited and learns connections to features in proportion to their activation (solid lines indicate stronger connection weights). The new unit codes the featural over-lap hummingbird and mouse, or a "dirty" representation of "small".

Although the new predicates DORA learns are initially "dirty" in that they contain extraneous features (e.g., in the previous example the representation of "small" also contains the feature "animal") through repeated iterations of the same learning process, DORA forms progressively more refined representations. For example, consider what happens when DORA compares the "dirty representation of "small" it learned in the previous example to another representation of "small" it learned, say, by comparing a matchbook to a playing card. Both representations of "small" contain the essential feature "small" and an extraneous feature (Figure 3a). However, because only the essential "small" feature is common to both representations of "small". When the two representations are compared the features they share will become most active (Figure 3b). When a new PO learns connections to the active features (as described above) it is most strongly connected to the feature "small" (the feature shared by both "small" representations) and less strongly connected to the features idiosyncratic to either particular representation (Figure 3c). In short, through a series of progressive comparisons DORA preserves what remains invariant across examples and discards everything else.

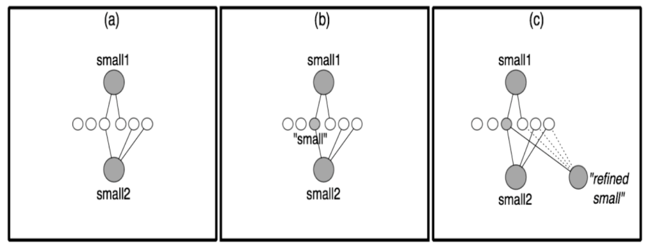

Figure 3. DORA learns a refined representation of "small" by comparing a two "dirty" representations of "small". (a) When DORA compares a the two representations of "small" the units representing both become active simultaneously. (b) Feature units shared by both representations of "small" become more active (darker grey). (c) A new unit is recruited and learns connections to features in proportion to their activation (solid lines indicate stronger connection weights). The new unit codes the featural over-lap of the compared representations, or a more refined representation of "small".

In the previous example DORA learned and refined an explicit representation of the property *small*. In the example we used a single semantic unit to code the feature "small" in order to make the example easier to follow. However, what is important about DORA's operation is not what each specific semantic unit codes, but that DORA's learning algorithm isolates and forms explicit representations of *any* features shared by compared representations, whatever those features may be. Whether "small" is coded by a single feature unit or by a set of units, when DORA compares small things it will isolate and represent the features that are invariant in small things (i.e., whatever is integral to being small) and discard other features. In other words through progressive comparisons of examples of a concept, DORA will isolate the properties that are invariant across those examples and represent those properties with an explicit predicate that can take arguments. Given that there are invariant properties in the world and the human cognitive system can detect them, DORA provides a means to learn explicit structured representations of these properties.

Simulations

We ran two simulations with DORA. In the first we simulated the development of representations of single geons from representations of multi-geon objects. In the

[2] DORA uses systematic asynchrony of firing to bind roles to fillers (see Doumas & Hummel, 2005; Doumas et al., submitted). As this is not important for the simulations reported here, we do not discuss binding further in this paper.

second we simulated the findings of Abecassis et al. (2001). In these simulations we make a key assumption: We assume that metric and categorical attributes are represented by the visual system independently of one another. That is, we assume that the visual system is capable of detecting properties such as curved cross sections, straight cross-sections and parallel and non-parallel lines, and that these properties are represented independently of metric properties like location in the visual field. This assumption was predicted in the computational models of Hummel (e.g., Hummel & Biederman, 1992) and has been supported by psychophysical experimentation (e.g., Stankiewicz, 2002).

Simulation 1

To simulate the development of geon representations we created 160 multi-geon objects. These objects consisted of at 2 geons selected randomly from a pool of 7 geons (including straight brick, curved brick, straight cone, straight wedge, curved wedge, straight cylinder, and curved cylinder; see Biederman, 1987). Examples of these stimuli are presented in Figure 4. Each multi-geon object was represented in DORA as a PO unit attached to 12 features. Half of these features described invariant categorical properties of the geons that composed the object (e.g., straight cross-section, parallel sides, curved axis-of-symmetry, etc.). So, for example, the object consisting of the cone and the brick was attached to the categorical features of a brick (e.g., straight cross-section, straight axis-of-symmetry, parallel sides) and the categorical features of a cone (e.g., curved cross-section, straight axis-of-symmetry, non-parallel sides). In addition, each object was also attached to 6 features describing metric properties that were chosen at random (i.e., the object's location in the visual field or the degree of curvature).

Figure 4. Examples of some multi-geon objects used during simulation 1

Importantly the features we use to code categorical and metric properties are features that can be detected by JIM from V1-like representations of objects. For example, JIM can detect categorical features like "curved cross-section" and metric features like "x-coordinate=5". However, JIM does not *learn* which shape attributes are view-invariant, and thus form the "definition" of a geon (e.g., that straight vs. curved major axis matters, whereas the exact degree of axis curvature does not); rather this information was hand-coded into the model's operation. As such in this simulation we tested whether DORA's learning algorithm could discover which features define geons simply by observing examples of multi-geon objects. More concretely, could DORA discover that the features *straight cross section*, *straight axis* and *parallel sides* define bricks and that *curved cross section*, *straight axis* and *non-parallel sides* define cones, simply by comparing objects composed of bricks, cones and other geons.

We then allowed three sets of comparisons. During the first set of comparisons (CS 1), we allowed DORA to compare multi-geon objects. Each set of multi-geon objects that DORA compared contained at least one of the same geons. For example, DORA might compare the cone and brick in Figure 4a to the wedge and brick in Figure 4b. When DORA compared these two objects it learned a representation of what they had in common, namely, those features essential to bricks (along with some extraneous features the two objects shared by chance). That is, the first set of comparisons produced "dirty" representations of the geons.

After CS 1 we began the second set of comparisons (CS 2), during which we allowed DORA to compare the "dirty" representations of geons it had learned during CS 1 to other "dirty" representations of the same geon. For example, DORA might compare one "dirty" representation of a brick to another "dirty" representation of a brick. This produced more refined representations of the geons.

Finally, after CS 2 we began the third set of comparisons (CS 3) during which we allowed DORA to compare the more refined representations of geons it had learned during CS 2 to other refined representations of the same geon. For example, DORA might compare one representation of a cone it had learned during CS 2 to another representation of a cone it had learned during CS 2. This produced even more refined representations of the individual geons.

After each set of comparisons we tested the representations of individual geons that DORA had learned using a selectivity metric (SM). The SM was calculated for each object as the mean weight between that object and the features essential to the geon it represented (e.g., for a cone curved cross-section, straight axis-of-symmetry, non-parallel sides) divided by 1 + the mean weight between that object and all irrelevant features to which it was connected.[3] In short, the SM provided a metric of the refinement of the representation. The higher the SM of a representation of a geon the more strongly that representation is connected to relevant features and the less strongly it is connected to irrelevant features.

The SM results for the representations learned during each set of comparisons are presented in Table 1. During each set of comparisons DORA learned progressively more refined representations of the six geons. Although this is, admittedly, a simplified case of learning, the simulation demonstrated that DORA's learning algorithm designed for learning relations from examples is sufficient to learn representations of individual geons from objects containing

[3] One was added to the denominator to keep the SM a ratio between 0 and 1.

multiple geons. With this in mind we proceeded to simulate the results of Abecassis et al. (2001).

Table 1. Simulation 1 results (SM = selectivity metric)

	SM
Initial representations	.5
After CS 1	.64
After CS 2	.72
After CS 3	.84

Simulation 2

In Experiment 2 of Abecassis et al. (2001) 4 year-old children and adults were presented with objects like those depicted in the middle row of Figure 5a. These sample exemplars were given a novel label, for example "wug". The participants were then given the other objects from Figure 5 one at a time and asked if these too were "wugs". While the objects in the bottom row where, on the whole, more similar to the objects in the middle row in terms of metric properties, they differed on important categorical features: The items in the middle row had curved axes of symmetry while those in the bottom row did not. On the other hand, the objects in the top row while less similar in terms of metric properties to the objects in the middle row, but were more categorically similar in that they shared categorical features such as curved axis of symmetry.

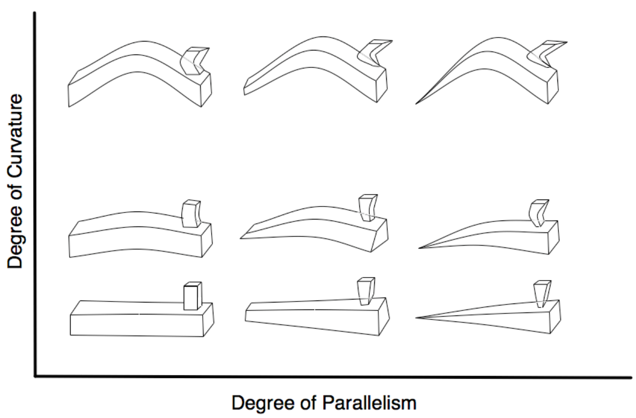

Figure 5. Example of the stimuli used in the experiment by Abecassi et al. (2001).

As noted previously, children generalized the name of the sample exemplars to the test exemplars in both the bottom row and the top row of Figure 5. Adults, on the other hand, generalized the name given to the sample exemplars much more frequently to the test exemplars from the top row. The authors concluded that as children get older they become more sensitive to invariant predictive properties (e.g., curvature) and less sensitive to over-all similarity.

We simulated the adults and children in the above experiment by varying the composition of the experimental stimuli presented to DORA. To simulate children we created all nine "wug" exemplars using the geons DORA had learned during CS 1 of the previous simulation. So, for example, to represent the exemplar from the middle row middle column of Figure 5 we used the representation *above* (curvedBrick1, curvedBrick2) where curvedBrick1 and curvedBrick2 were geons learned during simulation 1. To simulate adults we did the same thing only we constructed the exemplars using the geons that DORA had learned during CS 3 of simulation 1. In short, to simulate children we used messier representations of geons (those learned by DORA after less experience) and to simulate adults we used more refined representations of geons (those learned by DORA after more experience. To simulate children we placed 6 representations of the sample items constructed using CS 1 geons and 6 representations of random geons in random configurations into LTM. To simulate adults we placed 6 representations of the sample items constructed using CS 3 geons and 6 represenions of random geons in random configurations into LTM.

We ran 12 simulations each with 6 trials (the three bottom row trials and the three top row trials). On each trial we allowed DORA used its representation of the test exemplar to retrieve previously viewed exemplars from its LTM. During retrieval the representation of the test exemplar became active and passed activation to representations in LTM. As representations in LTM became active DORA used the Luce choice axiom to retrieve active LTM representations into working memory (WM). After two or three exemplars had been retrieved into WM DORA attempt to map the representation of the test exemplar to the representations of the retrieved exemplars. During mapping the representation of the exemplar becomes active and passes activation to the representations of the retrieved exemplars which compete (via lateral inhibition) to become active. If one of the retrieved representations matches the test items better than the others (i.e., shares a higher proportion of its semantic units with the test exemplar) then it will become most active and DORA will map the two representations. If DORA found a strong mapping correspondence, the test item was labeled a "wug", otherwise (i.e., if DORA found no strong mapping) the test item was not labeled a "wug".

The results from Abecassis et al. (2001) and our simulation are presented in Figure 6. Like the children in Abecassis et al.'s study, DORA with messier geon representations tended to generalize the name "wug" roughly equally often to both exemplars from the top and the bottom row. On the other hand, with more refined representations, DORA generalized the name "wug" much more often to items from the top row than those from the bottom. In short, with more experience DORA tended to generalize a name to more categorically similar objects than to more holistically similar objects, as people do. These simulation run using exactly the same settings and parameters that we used to simulate several other finding in the literature (e.g., Dixon & Banart, 2003; Gentner & Namy, 1999; Kotovsky & Gentner, 1996; Smith (1984); Smith et al., 1988; see Doumas et al., submitted). We did no parameter fitting and these results reflect DORA's first run.

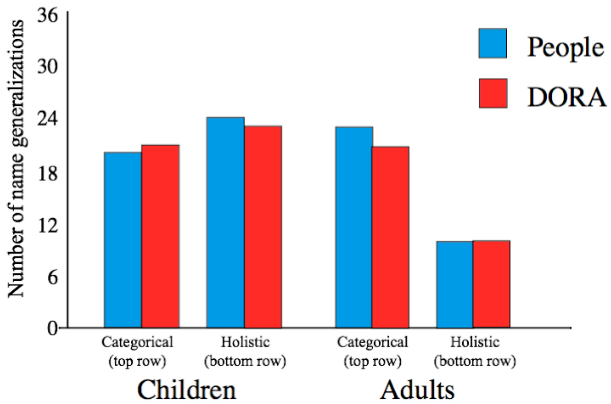

Figure 6. The experimental data from children and adults in Abecassis et al. (2001) and from DORA.

Discussion

Through a process of iterative comparison, DORA gradually comes to discover features that remain invariant over instances of an object category (or concept). This process allows it to discover invariant object attributes and, to form representations of geon like structures. The resulting representations provide a natural account of the developmental shift in the shape bias described by Abecassis et al. (2001). This process may also provide a basis for understanding how geons—clusters of co-occuring invariant features—are discovered by exposure to multi-geon objects.

An important implication of the DORA model is that comparison is central to the development of representations of geons and the transition from holistic to categorical representations of shape. Thus, DORA predicts that situations that invite comparison will provide rich contexts for developing categorical representations of shape. Such situations might include when two items share the same label, when the child is directed by an adult to compare, or when items are in close spacial proximity.

References

Abecassis, M., Sera, M. D., Yonas, A., & Schwade, J. (2001) What's in a Shape? Children represent shape variability differently than adults when naming objects. *Journal of Experimental Child Psychology*, 78, 303-326.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94* (2), 115-147.

Dixon, J. A., & Bangert, A. S. (2004). On the spontaneous discovery of a mathematical relation during problem solving. *Cognitive Science*, 28, 433-449.

Doumas, L. A. A., & Hummel, J. E. (2005). A symbolic-connectionist model of relation discovery. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 606-611. Mahwah NJ: LEA.

Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (submitted). *A theory of the discovery and predication of relational concepts*.

Gentner, D., & Namy, L. (1999). Comparison in the development of categories. *Cognitive Development*, 14, 487-513.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*, 480-517.

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

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

Imai, M., & Gentner, D. (1997). A cross-linguistic study of early word meaning: Univerasl ontology and linguistic influence. *Cognition*, 62, 169-200.

Jones, S. , & Smith, L.B. (1993). The place of perception in children's concepts. *Cognitive Development, 8*, 113-139.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, 67, 2797-2822.

Landau, B., Smith, L., & Jones, S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, 3, 299-321.

Landau, B., Smith, L., & Jones, S. (1992). Syntactic content and the shape bias in children's and adult's lexical learning. *Journal of Memory and Language*, 31, 807-825.

Smith, L. B. (1984). Young children's understanding of attributes and dimensions. *Child Development*, 55, 363-380.

Smith, L. B. (1989). From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou and A. Ortoney (Eds.), *Similarity and analogical reasoning* (pp. 147-177). Cambridge: Cambridge University Press.

Smith, L. B. (1995). Self-organizing processes in learning to learn words: Development is not induction. In The Minnesota Symposium of Child Psychology, Vol. 28, Basic and applied perspectives on learning, cognition, and development, 1-32. Mahwah, NJ: Erlbaum

Smith, L. B., Rattermann, M. J., & Sera, M. (1988). "Higher" and "lower": Comaprative interpretations by children. *Cognitive Development*, 3, 341-357.