Edinburgh Research Explorer

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

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

Leonidas A. A. Doumas
University of Hawaii at Manoa

John E. Hummel
University of Illinois, Urbana-Champaign

Running head: Development of the geons

Address for correspondence:
Leonidas A. A. Doumas
Department of Psychology
2430 Campus Road
Honolulu, HI 96822
leonidas@hawaii.edu

Abstract

Abecassis, et al., (2001) showed 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? According to Recognition-by-Components theory (RBC; Biederman, 1987) objects are represented as collections of categorical geometric parts ("geons") in particular categorical spatial relations. We propose that the transition from holistic to more categorical visual shape processing is a function of the development of geon-like representations via a process of progressive intersection discovery. We present an account of this transition in terms of DORA (Doumas et al., 2008), a model of the discovery of relational concepts. We demonstrate that DORA can learn representations of single geons by comparing objects composed of multiple geons. In addition, as DORA is learning it follows the same performance trajectory as children, originally generalizing shape more metrically/holistically and eventually generalizing categorically.

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

Introduction

Numerous studies have shown that both children and adults generalize labels preferentially to objects with similar shapes rather than, say, similar colors or textures (Imai & Gentner, 1997; Landau, Smith, & Jones, 1988, 1992; Smith, 1995; Woodward & Markman, 1998). That children and adults demonstrate this *shape bias* (i.e., they prefer to generalize labels to objects with shapes similar to those with which they are familiar) begs the question of what constitutes a "similar shape". The answer to this question changes with development.

Abecassis, Sera, Yonas, and Schwade (2001; Experiment 2) presented evidence that young children represent shapes more metrically and holistically than do older children and adults. For example, shown a slightly curved shape, a very curved shape and a straight shape (where the metric difference in curvature between the slightly curved shape and the straight shape is smaller than between the slightly curved shape and the very curved shape; e.g., compare the shapes in the bottom and middle vs. middle and top rows of Figure 1), adults and older children are more likely to generalize a label from the slightly curved shape to the very curved shape than to the straight shape. To adults and older children, the two curved shapes (both being categorically "curved") are more alike than the slightly curved shape is like the straight shape (one of which is "curved" and the other of which is "straight"). By contrast younger children are more likely to generalize a label from the slightly curved shape to the straight shape than to the very curved shape. This is because the slightly curved shape is metrically and holistically more similar to the straight shape than it is to the highly curved shape, and because these children are less sensitive to the visual categories "curved" and "straight". There is evidence for an analogous "relational shift" in cognitive development, in which young children appear to process objects and events holistically but, as they develop, gradually come to represent them in terms of independent objects, relations and properties (e.g., Gentner & Rattermann, 1991).

Please insert Figure 1 about here.

How does a child transition from representing objects and events as undifferentiated wholes to representing them explicitly in terms of their attributes—including categorical aspects of objects' shapes—and the relations among those attributes? This question is really two questions. The first is the question of how the categorical 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 (see Hummel & Biederman, 1992). The second is the question of how these categorical invariants come to dominate judgments of perceptual similarity.

This paper presents an effort to understand the answer to the second of these two questions. Our starting point is the hypothesis that the cognitive processes responsible for the developing reliance on categorical invariants in shape perception are the same as those responsible for the relational shift. That is, we hypothesize that the child's emerging reliance on categorical invariants in shape perception and label generalization is a manifestation of her discovery that relational invariants reliably predict other aspects of her world (for precedents, see Biederman, 1987; Garner, 1974; Hummel, 2003; Hummel & Biederman, 1992). Under this hypothesis, the question 'How do categorical invariants come to dominate judgments of perceptual similarity?' becomes 'How does the child isolate and predicate those visual

invariants that underlie the transition from a metric/holistic representation of object shape to a more categorical/analytic representation?'[1]

We present *DORA* (*Discovery Of Relations by Analogy*; Doumas & Hummel, 2005a; Doumas, Hummel & Sandhofer, 2008), a model that uses analogical mapping and intersection discovery to highlight shared abstract properties between separate systems (e.g., separate shapes, or separate views of the same shape) to discover those similarities and represent them as explicit structures (i.e., to *predicate* them). DORA was developed as a model of how structured relational representations are learned from examples (e.g., how children learn the relation *larger-than*(*x*, *y*) by observing examples of objects of various sizes). The simulations presented here suggest that the same processes DORA uses to isolate and predicate relational concepts from examples also permit it to discover and predicate categorical visual invariants (such as straight vs. curved) from metrically detailed representations of multi-part objects. DORA's metric-to-categorical transition provides a natural account of the developmental transition observed by Abecassis et al. (2001) and may therefore provide a partial account of the development of the role of geons (Biederman, 1987; see directly below) in shape perception.

Recognition by components

As noted by Abecassis et al. (2001), the problem of knowing when to generalize a name from one shape to another is a special case of the problem of recognizing objects: If two shapes represent different instances of the same object, then it is appropriate to generalize the name from one to the other. According to Biederman's (1987; Hummel & Biederman, 1992) Recognition-by-Components (RBC) theory of object recognition, adults visually represent object shape as 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 cylinder atop a vertical cylinder or truncated cone.

The parts, in turn, are represented as *geons*: classes of generalized cylinders[2] 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 [like the lid of a jar]; somewhat squat [like a tuna can]; neither squat

---

[1] Although the issue of metric vs. categorical appears at first blush orthogonal to, or at least different than, the issue of holistic vs. analytic/ predicate-based, in the context of representations of shape they turn out not to be (see Hummel, 2003). Holistic representations of shape— such as the representations generated in visual area V1, or the "view-based" representations endorsed by some models of object recognition (e.g., Edelman & Intrator, 2003)—are metrically precise (i.e., rather than categorical) by necessity (the reason being that they are coordinate-based, and coordinates are necessarily metrically precise). Similarly, barring exclusively feature-list-based representations of shape (which provide an extremely poor account of human shape perception; see, e.g., Hummel & Biederman, 1992), representations based on categorical invariants of the type that constitute geons are necessarily analytic, in the sense that they explicitly predicate the invariants in question and represent them independently of one another. (At the same time, however, it is perfectly possible for an analytic representation to be metrically precise, so the metric:holistic and categorical:analytic correspondences are not absolute; see Hummel & Stankiewicz, 1998.) Accordingly, in this paper, we tend to treat "metrically precise" and "holistic" as roughly synonymous and "categorical" and "analytic" as roughly synonymous. See Hummel (2003) for a much more extensive discussion of this issue.

[2] 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.

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 (and the categorical relations among those geons), then it follows that as we develop more robust representations of geons and their relations, we will transition from more holistic to more categorical shape generalization.

The DORA Model

DORA (Doumas & Hummel, 2005a; Doumas et al., 2008) is a symbolic connectionist network that learns structured (i.e., symbolic) representations of relations from unstructured (i.e., holistic) inputs. DORA derives this ability from its ability to maintain the structure of representations by dynamically binding distributed representations of relational roles to distributed representations of the objects (or geons) filling those roles.[3] 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, 2005b).

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 effects of progressive-alignment on relational learning, and relational learning in adults, among others (see Doumas et al., 2008). In the present work we used DORA to simulate the discovery of categorical representations of geons from examples of multi-geon objects and, accordingly, the development of the shape-bias in children and adults as reported by Abecassis et al. (2001). As DORA is described fully elsewhere (see Doumas et al., 2008), here we cover DORA in broad strokes focusing only on the aspects of its operation that are relevant for the current simulations.

DORA represents relational structures (such as propositions or structural descriptions of object shape) using a hierarchy of distributed and localist codes adapted from Hummel & Holyoak's (1997, 2003) LISA model (see Figure 2). At the bottom of the hierarchy "semantic" units represent the features of objects and relational roles in a distributed fashion. For example, a geon would be represented by semantic units specifying the shape of its major axis, the shape of its cross section, its aspect ratio, size, orientation, and so forth. At the next level, these distributed representations are connected to localist *predicate/object* (*PO*) units representing individual objects (or geons) and relational roles. Localist *role-binding* (*RB*) units link object and relational roles units into specific role-filler bindings. At the top of the hierarchy, localist *proposition* (*P*) units link RBs into complete propositions.

_____

[3] DORA can also represent higher-order relations (e.g., *cause*(gravity, *revolve-around*(earth, sun))) by binding relational roles to other relational structures.

Development of geons

Please insert Figure 2 about here.

DORA's fundamental operating principle is to use analogical mapping to isolate shared properties of objects and to use self-supervised learning (Hummel & Holyoak, 2003) to represent them as explicit structures (i.e., to *predicate* them). When DORA maps one object or structure onto another, corresponding elements of the two representations fire together. As a simplified example, if DORA compares a mouse to a hummingbird, then the nodes representing the mouse will fire at the same time as those representing the hummingbird (Figure 3). Consequently, any semantic features shared by both 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 semantic units in proportion to their activation via simple Hebbian learning (i.e., DORA learns stronger connections to more active units; Figure 3b). The new PO thus becomes an explicit representation of the featural overlap between the compared objects (or relational roles; see below). In the case of comparing a hummingbird to a mouse, the network might form an explicit predicate representing "small", "animal", (and any other features they share; Figure 3).[4] Importantly, this new PO is now an explicit predicate (representing, e.g., *small animal*) that can be dynamically bound to new fillers.[5]

Please insert Figure 3 about here.

Although the new predicates DORA learns are initially "messy" in that they contain extraneous features (e.g., in the previous example the representation of "small animal" will also have weak connections to the other features of hummingbird and mouse), through repeated iterations of the same learning process, DORA forms progressively more refined representations. For example, consider what happens when DORA compares the "messy" representation of "small animal" it learned in the previous example to a representation of "small" it learned by comparing, say, a matchbook to a toy-car. Both representations of "small" contain the essential features "small", as well as various extraneous features (e.g., the representation learned by comparing a mouse and a hummingbird is also connected to the feature "animal"; Figure 4a). However, because only the essential "small" feature is common to both representations, when the two representations are compared to one another, the features they share ("small") will become most active (Figure 4b). When a new PO learns connections to the active features (as described above) it is most strongly connected to "small" and less strongly connected to the features idiosyncratic to either representation (Figure 4c). In short, through a series of progressive comparisons,

---

[4] For the current example (and the current simulations) we let DORA start with semantic representations like "small" and later "curved-cross-section" for the purposes of clarity and ease of expostion. Importantly, DORA can learn representations of semantics coding for featrues "smaller", "larger", "higher", "lower", etc., starting with semantics coding for specific metric values and a compariter circuit (see Doumas et al., 2008, for full details).

[5] DORA uses systematic asynchrony of firing to bind roles to fillers (see Doumas & Hummel, 2005a; Doumas et al., 2008). As this detail is not important for the simulations reported here, we do not discuss binding further in this paper. See Doumas et al. (2008) for the complete details of DORA's operation, including all assumptions, equations and parameter values.

DORA's learning algorithm preserves what remains invariant across examples and discards everything else.

Please insert Figure 4 about here.

In the previous (highly simplified) example DORA learned and refined an explicit representation of the property *small*. 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, via a kind of intersection discovery, isolate the properties that remain invariant across those examples and come to represent those properties as explicit structures (or symbols). 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 those properties (see Doumas et al., 2008, for many more examples).

Simulations

We ran two simulations with DORA to simulate Abecassis et al.'s (2001) Experiment 2. In the first we simulated the development of representations of single geons from representations of multi-geon objects. In the second, we used the representations DORA learned during the first simulation to simulate the behavior of Abecassis et al.'s subjects. Importantly, these two sets of simulations were interleaved, allowing us to test DORA at different "ages". A key assumption underlying these simulations is that the visual system represents an object's categorical properties independently of its metric properties. 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 size, orientation and location in the visual field. This kind of independence was first predicted by Hummel & Biederman (1992) and has since been supported by both psychophysical experiments (e.g., Stankiewicz, 2002) and single-unit recordings in the macaque visual system (e.g., Kayaert et al., 2005; Tanaka, 1994).

Simulation 1

To simulate the development of geon representations, we created 160 multi-geon objects. Each object consisted of two geons selected randomly from a pool of seven (including straight brick, curved brick, straight cone, straight wedge, curved wedge, straight cylinder, and curved cylinder; Biederman, 1987). Examples of these objects appear in Figure 5. Each object was represented in DORA as a PO unit attached to 12 features. Of these, up to six described the invariant categorical properties of the geons that composed the object. For example, an object consisting of a straight cone (*curved cross section*, *straight major axis* and *non-parallel sides*) and a straight brick (*straight cross section*, *straight major axis* and

parallel sides) was represented by a PO attached to five categorical features (the union of the six listed in the previous sentence; both the cone and the brick have a straight major axis, so their union consists of only five features); the remaining seven features, describing the geons' metric properties, were chosen at random (e.g., the object's location in the visual field or the degree of curvature; see Figure 6). Note that this representation of the object is holistic in the sense that it simply lists the cone and brick features without specifying how they go together in geon-based sets.

Please insert Figure 5 about here.

Please insert Figure 6 about here.

Importantly the features we use to code categorical and metric properties are features that can be detected by Hummel and Biederman's (1992) JIM model of object recognition 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 cannot *learn* which shape attributes are view-invariant, and thus does not learn 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 is hand-coded into the model's representations. As such, in the first simulation we tested whether DORA's learning algorithm could discover which features define geons simply by observing examples of multi-geon objects. For example, would 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 ran three sets of comparisons. During the first set of comparisons (CS 1), DORA compared pairs of multi-geon objects. Each pair of multi-geon objects that DORA compared contained at least one of the same geons. For example, DORA might compare a cone and brick (as in Figure 5a) to a wedge and brick (as in Figure 5b). 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). CS 1 consisted of 80 comparisons, each producing a "messy" representation of the shared geon. For example, by comparing a cone and brick to a wedge and brick, DORA learned a representation most strongly connected to the features shared by the two objects (i.e., those features defining a brick), and weakly connected to the other features of the two objects (i.e., the features of cone and wedge and any metric properties of the two objects).

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

In the third set of comparisons (CS 3) 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. CS 3 produced even more refined representations of the individual geons.[6]

After each set of comparisons a *selectivity metric* (SM) was calculated on the representations of individual geons DORA had learned during those comparisons. 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*:

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

(1)

Table 1 presents the mean SM over all POs after each set of comparisons. Over the course of the comparisons DORA learned progressively more refined representations of the six geons. The simulation thus demonstrated that DORA's learning algorithm is capable of learning representations of individual geons by comparing examples of multi-geon objects. Importantly, in these simulations DORA not only discovered which categorical properties remained invariant over different views of a given geon (thereby learning a preference for view-invariant categorical properties over view-specific metric properties), it also learned representations of individual geons from holistic examples of multi-geon objects. That is, it learned what a geon "is", both in the sense of learning which features reliably co-occur as a geon and in the sense of "pulling a single geon out" of an otherwise holistically-represented multi-geon object.

Please insert Table 1 about here.

Simulation 2

Experiment 2 of Abecassis et al. (2001) and the results are described above. To simulate the experiment we used the representations of geons DORA had learned during simulation 1 to construct multi-geon "wug" objects. The sets of comparisons performed in Simulation 1 necessarily took place in sequence (i.e., with CS 3 following CS 2, which followed CS 1). In our second set of simulations, we treated DORA's progression in this sequence as its point in development (i.e., its "age"). To simulate 4-year-olds we created all nine "wug" exemplars using the geons DORA had learned during CS 1 of the previous simulation. For example, to represent the exemplar from the middle row middle column of Figure 1 we used the representation *above* (curvedBrick1, curvedBrick2), where curvedBrick1 and curvedBrick2 were geons learned during CS 1 of Simulation 1.[8] To simulate adults we did the same

---

[6] Although we trained DORA using a blocked training design for the purpose of simplicity, DORA is very capable of learning from unstructured training sets where exemplars are not blocked (see Doumas et al., 2008, simulation 1). Importantly, the SM of the representations it learns from unstructured training follows the exact same pattern as the one it learns from blocked training.

[7] One is added to the denominator to keep the SM ratio between 0 and 1 (i.e., as weights on a PO's connections to relevant semantics approach 1, and weights on its connections to irrelevant semantics approach 0, its SM approaches 1.

[8] See Doumas et al., 2008 for a full description of how DORA learns representations of relations like *above*, *below*, and *next-to*.

thing, only we constructed the exemplars using the geons DORA had learned during CS 3 of Simulation 1.

In order to simulate Abecassis et al.'s experiment, we placed six representations of the sample items into DORA's LTM (corresponding to the training exemplars; using CS 1 geons in the case of children and CS 3 geons in the case of adults), along with six representations of randomly-chosen geons in randomly-chosen relations. On each simulation, DORA used its representation of a test exemplar (i.e., one of the objects from either the top or bottom row of Figure 1) to retrieve previously viewed exemplars from its LTM. Retrieval consists of activating the representation of the test exemplar and allowing it to pass activation, via the semantic units, to the representations in LTM and using the Luce (1959) choice axiom to retrieve these active LTM representations into working memory (WM) (see Doumas et al., 2008). After two or three exemplars had been retrieved into WM, DORA attempted to map the representation of the test exemplar onto the representations of the retrieved exemplars. During mapping, the retrieved exemplars compete (via lateral inhibition) to become active. If one of the retrieved representations matches the test item 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 it to the representation of the test exemplar. The item to which the test was mapped was taken to be DORA's "categorization" of the test: If that item was one DORA had been trained on as a "wug", then we took that response as DORA endorsing the test item as a "wug"; if it was any other item, or if no mapping was found, then we took that response as rejecting the test as a non-"wug". We ran 12 simulations of each "age", each with six trials (the three bottom row trials and the three top row trials of Figure 1).

The results from Abecassis et al. (2001) and our simulation are presented in Figure 7. Like the younger children in Abecassis et al.'s study, DORA with earlier (CS 1) geon representations tended to generalize the name "wug" roughly equally often to both exemplars from the top and the bottom rows. By contrast, with more refined (CS 3) representations, "adult" 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. Importantly, we ran these simulations using exactly the same settings and parameters that Doumas et al. (2008) used to simulate numerous other finding in the literature. We did no parameter fitting and these results reflect DORA's first run.

Please insert Figure 7 about here.

Discussion

Through a process of iterative comparison and intersection discovery, DORA gradually comes to discover features that remain invariant over instances of an object category, relational role or concept. This process allows it to discover invariant object attributes and to form representations of geon-like structures from examples of objects composed of geons. 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 (as well as other relational representations) and to the transition from holistic to categorical and, more generally, analytic, representations of object shape. This account predicts that the rate of invariant discovery and predication should closely follow the rate of comparison: The more comparisons a child makes, the faster she should discover whatever invariant properties characterize the objects so compared. And 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 spatial proximity.

A related prediction is that the class of invariants a child discovers should be a function of the invariants present and absent in the set of objects she has compared. This property of the model can be illustrated by considering the implications of comparing different views of the same object. The simulations reported here were simplified in the sense that we only allowed DORA to compare objects to *other* objects. In the real world, however, objects move in 3-dimensional space, providing numerous opportunities to compare one view of an object to other views of the *same* object. Such within-object comparisons provide a weaker set of contrasts than do the between-object comparisons used in the simulations reported here. For example, consider comparing one view of a cone-on-a-brick object to a different view of the same object. Although this comparison makes it possible to discover which aspects of the cone and brick remain invariant over multiple views (especially if many such comparisons are made), it does not make it possible to discover that cone-like things can exist independently of brick-like things. As a result, it does not make it possible to predicate "cone", as a geon, independent of "brick": If the child always sees cones paired with bricks (and vice versa), DORA predicts she will never learn to predicate them as independent geons (see Doumas, et al., 2008).

This property of the model highlights the unusual location it occupies in the space of statistical vs. symbolic models of cognitive development. On the one hand, DORA is decidedly symbolic, possessing as it does the ability to (a) represent relational roles independently of their arguments and (b) dynamically bind arguments to whatever roles they happen to be filling at the time. This property gives it the full inductive flexibility of any relational, i.e., non-associative, computational system—a flexibility that we and others argue it shares with the human cognitive architecture (Doumas et al, 2008; Holland, Holyoak, Nisbett & Thagard, 1989; Hummel & Holyoak, 1997, 2003; Penn, Holyoak & Povinelli, 2008). In this respect, DORA stands in stark contrast to purely associationist models of cognition (e.g., Elman, 1991; O'Reilly & Busby, 2002) and purely holistic "view-based" models of object perception (e.g., Reisenhuber & Poggio, 1999; Poggio & Edelman, 1991). Even Edelman and Intrator's (2003) "Chorus of Fragments" model, which is nominally "part"- (i.e., "fragment"-) based cannot discover the kinds of invariants discovered by DORA because it cannot represent object properties independently of one another or of their location in the visual field (see Hummel, 2003).

At the same time, however, the model's reliance on intersection discovery for the discovery and predication of spatial and relational invariants makes it heavily dependent on the statistics of the examples over which the intersection discovery is performed: If DORA never sees a geon with a straight major axis and a straight cross section, for example, it will never predicate "brick" as a possible geon. The model's "innate" endowment is thus in its ability to compare examples to discover the invariants they share (and in its ability to reify these invariants as explicit predicates that can be bound to novel arguments), not in its expectations about what those invariants will be. In running the simulations reported here, we were forced, by the practical realities of modeling, to endow the model with a population of basic semantic features with which to work. In this sense, we gave it an "innate"

population of (invariant and non-invariant) features as raw material. However, beyond that, we did not endow it with any "expectations" about which of these features would either (a) remain invariant over different object views or (b) occur in clusters (e.g., in the form of geons). These latter properties it had to discover for itself through experience with the statistical properties of the objects in its world.

As a model of the development of categorical representations of object shape, DORA's most fundamental claim is that adult-like representations of shape can be (indeed, are) learned in exactly this way: Adults rely on categorical shape properties precisely because those properties remain across diverse objects and object views—as pointed out forcefully by Biederman (1987)—and because the same algorithm for intersection discovery that allows us to discover abstract relational concepts (Doumas et al., 2008) also causes those invariants to emerge as the stable descriptors of object shape.

Aknowledgements

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. (2005a). 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. (2005b). 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.

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.

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

Hummel, J. E. (2003). The complementary properties of holistic and analytic representations of object shape. In G. Rhodes and M. Peterson (Eds.), *Perception of faces, objects, and scenes: Analytic and holistic processes* (pp. 212-234). Westport, CT: Greenwood.

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.

Kayaert, G., Biederman, I., Op De Beeck, H. P., & Vogels, R. (2005). Tuning for shape dimensions in macaque inferior temporal cortex. *European Journal of Neuroscience, 22,* 212-224.

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.

O'Reilly, R.C. & Busby, R.S. (2002). Generalizable Relational Binding from Coarse-coded Distributed Representations. In T. G. Dieterich, S. Becker, and Z. Ghahramani (Eds.), *Advances in Neural Information processing Systems,* 14. Cambridge, MA: MIT Press.

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.

Development of geons

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

	SM
Initial representations	.46
After CS 1	.64
After CS 2	.72
After CS 3	.84

Figure Captions

Figure 1. Example of the stimuli used in the experiment by Abecassi et al. (2001).

Figure 2. Example of the proposition *above*(cone, brick) in DORA. For clarity, triangles denote relational roles and circles to denote objects, but in the model the same types of units code both roles and objects.

Figure 3. DORA learns a representation of "small animal" by comparing a hummingbird to a mouse. (a) When DORA compares a hummingbird to a mouse the units representing both become active simultaneously. (b) Feature units shared by both become more active (small grey circles) than units belonging to one but not the other (small white circles). (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 overlap between hummingbird and mouse, effectively representing "small animal".

Figure 4. DORA learns a refined representation of "small" by comparing two "messy" representations of "small". (a) When DORA compares the two "messy" representations of "small" the units representing both become active simultaneously. (b) Feature units shared by both representations of "small" become more active (darker grey) than units connected to only one of them. (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 overlap of the compared representations, i.e., a more refined representation of "small".

Figure 5. Examples of some multi-geon objects used during simulation 1.

Figure 6. Initial representation of a straight cone and a brick in DORA. A PO coding for the shape composed by the two geons is connected to 12 semantic units. Five units jointly code for the categorical properties of the cone and brick and the remaining seven semantic units code for random metric features of the object.

Figure 7. The experimental data from children and adults in Abecassis et al. (2001) and from DORA.
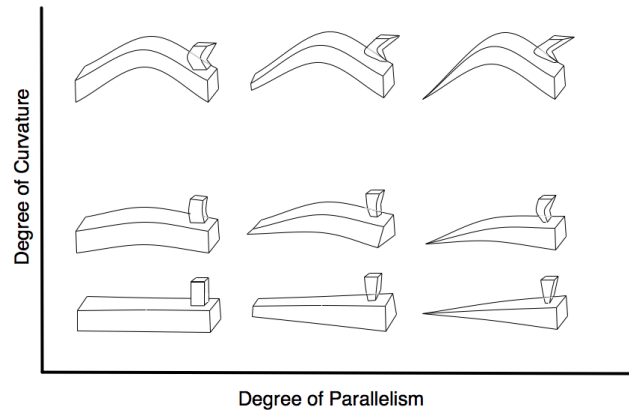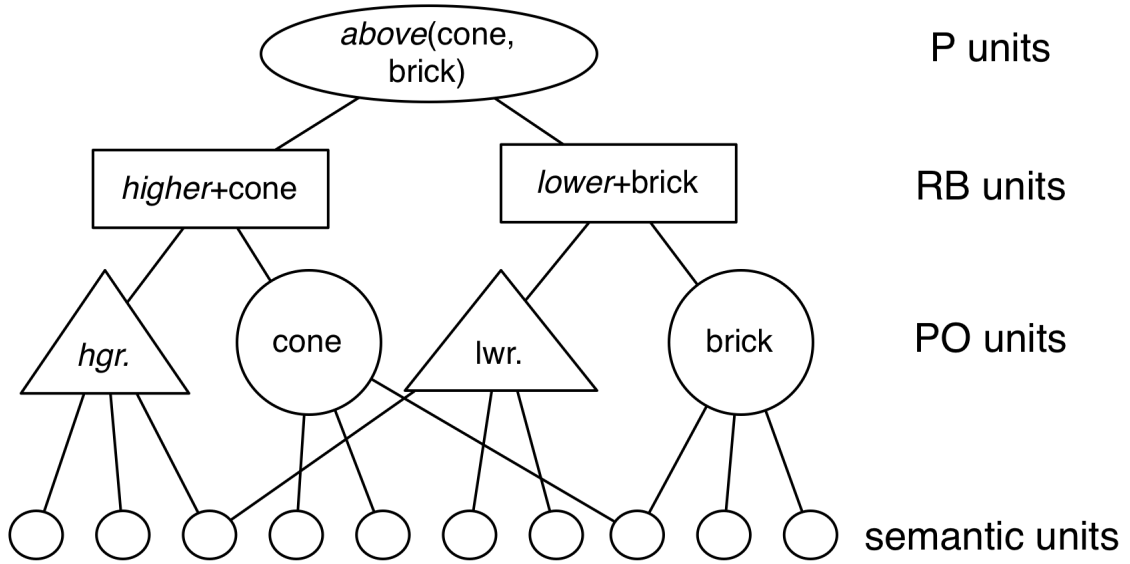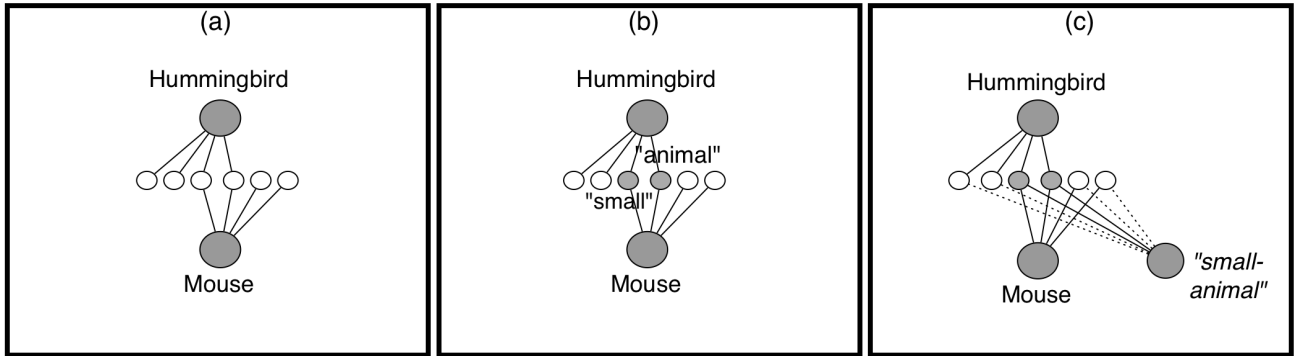
Development of geons

Figure 1.

Figure 2.

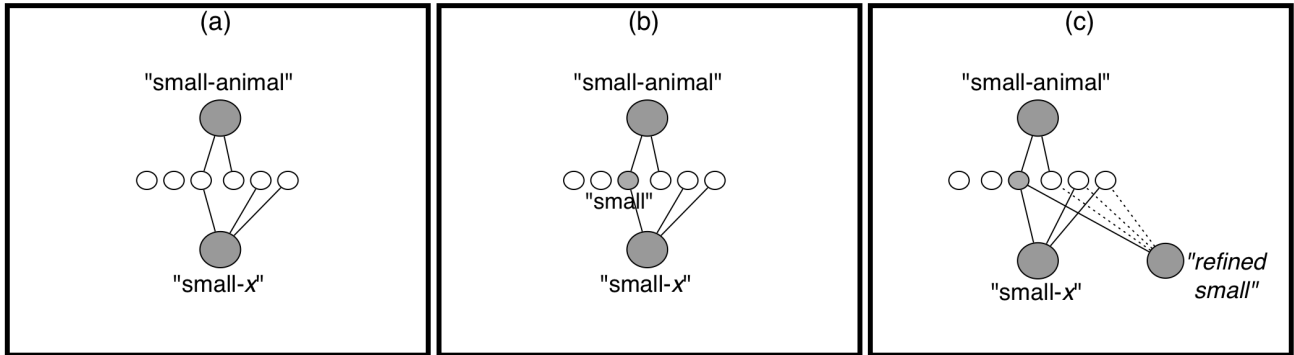Development of geons
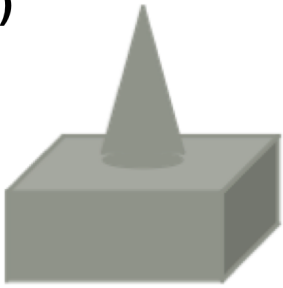
Figure 3.

Development of geons

Figure 4.

Development of geons

Figure 5.

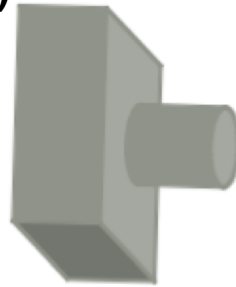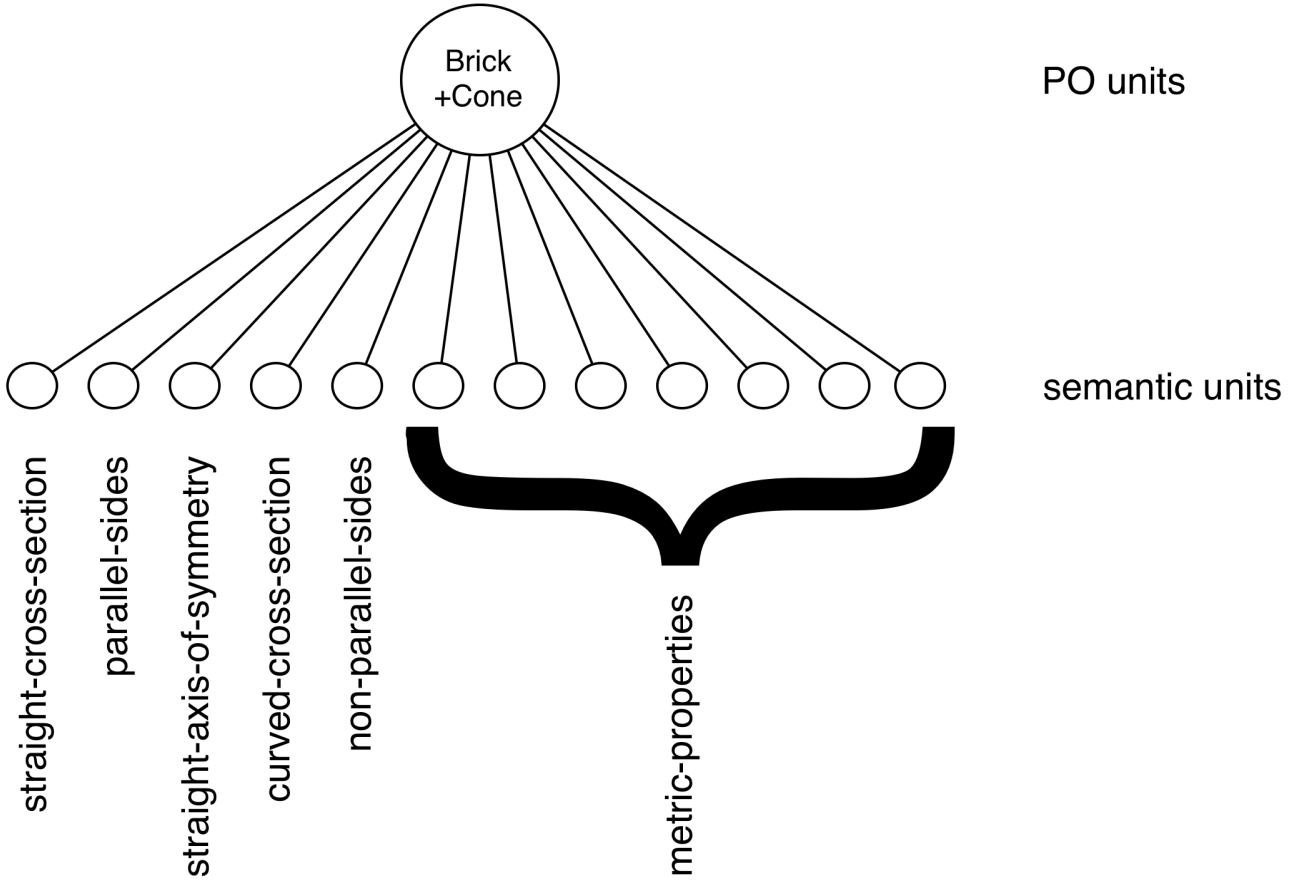(a) Straight Cone and Brick

(b) Curved Wedge and Brick

(c) Straight Cylinder and Brick

Development of geons

Figure 6.

Development of geons

Figure 7.