

From Pigeons to Humans: Grounding Relational Learning in Concrete Examples

Marc T. Tomlinson and Bradley C. Love

[mtomlinson,love]@love.psy.utexas.edu

University of Texas At Austin

1 University Station, A8000

Austin, TX 78712-0187 USA

Abstract

We present a cognitive model that bridges work in analogy and category learning. The model, Building Relations through Instance Driven Gradient Error Shifting (BRIDGES), extends ALCOVE, an exemplar-based connectionist model of human category learning (Kruschke, 1992). Unlike ALCOVE which is limited to featural or spatial representations, BRIDGES can appreciate analogical relationships between stimuli and stored predicate representations of exemplars. Like ALCOVE, BRIDGES learns to shift attention over the course of learning to reduce error and, in the process, alters its notion of similarity. A shift toward relational sources of similarity allows BRIDGES to display what appears to be an understanding of abstract domains, when in fact performance is driven by similarity-based structural alignment (i.e., analogy) to stored exemplars. Supportive simulations of animal, infant, and adult learning are provided. We end by considering possible extensions of BRIDGES suitable for computationally demanding applications.

Introduction

We live in a world of concrete experiences, yet we appreciate seemingly abstract concepts. Category contrasts like “same” or “different” and “predator” or “prey” are not based on featural regularities. Instead, these concepts are relational in nature (e.g., two identical objects of any size, color, or shape are the same). How do we acquire concepts that are not based on featural regularities?

The category learning literature has focused on how we acquire new concepts, but largely in domains in which the regularities are defined by distributions of features (e.g., red objects are in one category, whereas blue objects are in the other category). In contrast, the analogy literature has considered how people appreciate regularities defined by common relational structures, but has not leveraged work in category learning detailing how people integrate knowledge across examples and shift attention.

Here, we extend an existing model of category learning to include a more sophisticated notion of similarity that is sensitive to both featural and relational match. We propose that in many domains where performance appears to be governed

by mental rules, it is actually driven by similarity-based activation of stored exemplars. This account is supported by simulations of how young infants learn seemingly abstract grammars and how animals and people learn the distinction between same and different.

Category Learning

Proposals for category representation are diverse, ranging from exemplar-based (Medin & Schaffer, 1978) to prototype-based (Smith & Minda, 1998) and include proposals between these two extremes (Love, Medin, & Gureckis, 2004).

In the present work, we focus on exemplar-based proposals. Exemplar-models of category learning hold that all abstraction or generalization occurs through similarity-based activations of concrete examples. In exemplar models, each experienced instance is stored in memory. When a new item is encountered, the similarity between the item and each exemplar in memory is calculated. The stimulus is predicted to belong in the category with the greatest sum of pairwise similarity. Thus, exemplar models clearly link experienced events to later generalization.

The model of relational learning that we will develop here, BRIDGES, is an extension of Kruschke’s (1992) ALCOVE connectionist exemplar model of category learning. BRIDGES’s exemplar representation will support the notion that analogy to stored experiences is sufficient to appreciate abstract relationships. Furthermore, by incorporating ALCOVE’s attentional shifting mechanisms into BRIDGES, we forward an explanation of how perceived similarity can change over the course of learning as more predictive stimulus properties are accentuated. Such attentional shifts will prove useful in demonstrating how seemingly abstract understandings can arise from concrete experiences.

Finally, ALCOVE has been a very successful model of category learning and provides a strong foundation for bridging the category learning and analogy literatures. By virtue of containing ALCOVE as a special case, BRIDGES accounts for a wide variety of prior human learning studies examining how people acquire categories determined by featural regularities.




Figure 1: Example of attention on mapping

Analogy

Analogical comparison can reveal non-obvious similarities. For example, people can appreciate that the Rutherford atom is similar to the solar system because the planets revolve around the sun much like how the electrons revolve around the atom’s nucleus (Gentner, 1983). The perceived similarity between the atom and the solar system is not based on feature overlap, but on common relational structure.

Numerous accounts of how people detect these analogical similarities exist. Transformational accounts hold that one analog is transformed over a series of steps until it matches the other analog (Hahn, Chater & Richardson, 2003). The fewer the steps and the smaller their cost, the higher the resulting similarity is. The dominant account of detecting non-obvious similarities is structure mapping theory (Gentner, 1983). Although BRIDGES could utilize other frameworks to reach its ends, such as transformational approaches, we adopt the structure mapping account because of the broad support it enjoys in the analogy literature.

Structure mapping holds that people encode stimuli (e.g., objects, scenes, events) in terms of predicate representations that capture relations among entities (e.g., *Revolves(planets,sun)*). Relations can serve as arguments to other relations (e.g., *Causes(GreaterMass(sun,planets), Revolves(planets,sun))*). Structure mapping posits that people align such structured representations to find the most satisfying correspondences. In the solar system/atom example, this alignment would involve placing the sun in correspondence with the nucleus and the planets in correspondence with the electrons. Sounder mappings lead to increased perceived similarity (Gentner & Markman, 1997).

Factors that influence the soundness of a mapping include: 1) one-to-one correspondence 2) parallel connectivity and 3) systematicity. One-to-one correspondence says that corresponding entities and relations can map to at most one node in the other structure. Parallel connectivity is a preference for mapping arguments playing the same role within corresponding predicates to one another. For example, parallel connectivity is preserved for the *revolves* relation in the atom/solar system example because the two *revolves* predicates correspond and the arguments of these two relations also map to the appropriate argument in the other predicate (i.e., mapping the sun to the electrons would not preserve parallel connectivity). As Figure 1 illustrates with a simpler perceptual analogy, parallel connectivity can be at odds with featural similarity. Finally, systematicity is a preference for mappings that are complete and deep (i.e., contain embedded systems of relations). These principles have been suc-

cessfully embodied in the SME model of analogical alignment (Falkenhainer, Forbus, & Gentner, 1989).

More recent models within the structure mapping framework have demonstrated that systematicity can fall out of a system that enforces one-to-one mappings and prefers parallel connected mappings (Larkey & Love, 2003). Thus, BRIDGES’s approach to structural alignment will incorporate only one-to-one correspondence and parallel connectivity. As will be more fully discussed below, BRIDGES factors in both relational and featural matches when determining mappings. When these two forces are at odds (as in Figure 1), learned attentional weightings determine the preferred mapping.

Overview of BRIDGES

BRIDGES is an exemplar-based model of category learning that can utilize structure mapping to appreciate analogical similarities. For any given comparison between a stimulus and a stored exemplar, BRIDGES considers all possible one-to-one mappings. Models of analogical alignment avoid this exhaustive search by using heuristics to guide the mapping process. BRIDGES could be extended to incorporate these shortcuts, but we do not do so here. Instead we focus on the basic ideas underlying BRIDGES.

Out of all permissible mappings, BRIDGES chooses the mapping that minimizes a difference measure that incorporates notions of featural and relational mismatch (see Figure 1). A relational mismatch of 1 occurs when a relation does not exhibit parallel connectivity (see the right panel of Figure 1). A featural mismatch of 1 occurs when non-identical entities or entities containing mismatching features are mapped to one another (see the right panel of Figure 1).

Both types of mismatch are weighted by attention weights and the sum of these attention weighted mismatches yields an overall differences measure that is inversely proportional to similarity. The mapping that maximizes similarity (i.e., minimizes attention weighted difference) is chosen. These exemplar similarity values serve as exemplar unit activations and are passed across association weights to category units (e.g., *predator* and *prey*). The stimulus tends to be classified into the category whose unit has the greatest activation.

After feedback is provided, attention weights and association weights between exemplars and category units are adjusted to reduce error. Changes in attentional weights can lead to different future mappings.

When attention shifts away from features and toward relations, parallel connectivity (i.e., analogical match) is stressed over featural similarity and BRIDGES demon-

strates abstract understanding of a domain. Conversely, when featural matches lead to successful predictions, attention shifts toward features and BRIDGES is governed by featural similarity. When relational information is not discriminative or present, BRIDGES reduces to the standard ALCOVE model.

BRIDGES’s Formalism

BRIDGES is a three-layer feed forward artificial neural network. The input to BRIDGES contains both featural and relational information in predicate argument form. Each possible feature or relation has an associated attention weight. The activation (i.e., similarity) of each exemplar j by stimulus s is

$$h_j = \max_{m \in M} (\exp[-c(\sum_i \alpha_i \cdot \text{Mismatch}(s_i))]) \quad (1)$$

where $m \in M$ is all possible one-to-one mappings between nodes (i.e., features, entities, and relations) forming stimulus s and exemplar j , i ranges over nodes in stimulus s , α_i is the attention weight associated with node type s_i , and c is the specificity parameter that determines the rate that similarity falls off with increasing mismatch. Mismatch is defined to be 1 for features or entities in stimulus s that map to non-identical nodes in exemplar j , otherwise 0. For relations, Mismatch is defined to be 1 for relations in stimulus s not exhibiting parallel connectivity, otherwise 0.

Activation passes from exemplars to category unit o_k :

$$o_k = \sum_j w_{kj} \cdot h_j \quad (2)$$

where w_{jk} is the association weight between exemplar j and category unit o_k . The probability of selecting the category corresponding to category unit r is

$$\text{Pr}(r) = \exp(\phi o_r) / \sum_k \exp(\phi o_k) \quad (3)$$

where ϕ is a decision parameter and k ranges over all category units.

Learning is accomplished via gradient descent error minimization. The target value for the category unit corresponding to the correct category is set to 1 and other category nodes are set to 0. A “humble teacher” scheme is used so that category unit output values in the correct direction greater than 1 or less than 0 are not penalized. The error function E minimized is

$$E = \frac{1}{2} \sum_k (t_k - o_k)^2 \quad (4)$$

where t_k is the target value for o_k . The association weight w_{kj} from exemplar j to output unit k is adjusted by

$$\Delta w_{kj} = \lambda_w (t_k - o_k) h_j \quad (5)$$

where λ_w is the learning rate for association weights. Attention weights are updated by

$$\Delta \alpha_i = -\lambda_\alpha \sum_j [\sum_k (t_k - o_k) w_{kj}] h_j \cdot c \cdot \text{Mismatch}(s_i) \quad (6)$$

where λ_α is the learning rate for attention weights. The parameter values used in the simulations reported here are $\lambda_\alpha = .007$, $\lambda_w = .001$, $c = .03$, and $\phi = 6.1$. These values were chosen to maximize R^2 in the second simulation, but both simulations are actually quite robust with regards to the actual parameter values.

Modeling Infant Grammar Learning

Marcus, Vijayan, Bandi Rao and Vishton (1999) found that infants could discriminate between abstract patterns or grammars of speech sounds. Importantly, this discrimination could not be accomplished by any weighting of phonetic features. Because featural regularities could not be leverage to discriminate between grammars, Marcus et al. proposed that infants utilized variable binding in conjunction with algebraic rules to master such learning tasks.

In Marcus et al.’s (1999) study, seven-month-old infants were exposed to sentences that followed either an AAB pattern or an ABB pattern. The sentences were made up of simple monosyllable sounds (words) such as “GA TI TI”. Each infant was trained for approximately 2 minutes on one of the grammars. The 2-minute session contained three repetitions of 16 unique sentences. There was a 250 milliseconds pause between each word in the sentences and a one second pause between each sentence. The testing phase consisted of presentation of 12 sentences containing novel words. Half of the 12 sentences were from the same grammar used in training and the other half were from the contrasting grammar. Fifteen of the 16 infants in the study had greater looking times during the presentation of the sentences from the novel grammar, indicating that the infants had habituated to the abstract pattern.

BRIDGES’s Simulation BRIDGES was applied to the Marcus et al. (1999) study. Each sentence (e.g., “GA TI TI”) was represented as an exemplar. BRIDGES’s exemplar representation for “GA TI TI” is shown in Table 1. Each syllable is represented as an entity. Each syllable’s position in the speech stream is encoded by a positional feature. These syllables have a number of phonetic features that are not represented in these simulations. Not including such features follows Marcus et al.’s presumption that no significant regularities exist across these features. Importantly, including uncorrelated features would only slow learning, but would not alter the overall results.

Table 1: BRIDGES’s representation of “GA TI TI.”

Entities	Features	Relations
GA_1	$Position(GA_1) = 1$	$TypeOf(GA_1, GA)$
TI_1	$Position(TI_1) = 2$	$TypeOf(TI_1, TI)$
TI_2	$Position(TI_2) = 3$	$TypeOf(TI_2, TI)$

Critically, relational information was included in BRIDGES’s representations. BRIDGES makes a distinction between tokens and types. In effect, we assume that infants in Marcus’s et al.’s study have developed categories of speech sounds (Eimas, Siqueland, Jusczyk, & Vigorito,

1971). These type relations allow for abstract patterns to be uncovered through analogy to stored exemplars as one type of sound can mapped to another.

The training and test regimen followed the original Marcus et al. (1999) study as closely as possible. Following Love et al. (2004), unsupervised learning was modeled by a network consisting of a single category output unit with a target value of 1 for all stimuli. In effect, this category unit is a familiarity detector. Association and attention weights in the model were adjusted to uncover the underlying regularities across the sentences to yield consistently high familiarity (i.e., high output values for the category unit).

During habituation, the 16 unique sentences were presented three times each to BRIDGES and stored as exemplars. On each presentation, association and attention weights were updated. Though not critical, we assumed that the saliency of positional features is sufficiently great to constrain the mapping process (i.e., words in sentences align temporally). Besides position (which does not discriminate between grammars), no regularities across features or entities existed. However, parallel connectivity was perfect for members of the same grammar. For instance, “GA TI TI” is isomorphic to “LI NA NA” in that all token and types in the type relation (see Table 1) can be mapped to one another and preserve parallel connectivity. This degree of perfect match caused BRIDGES to shift attention to the type relation. This shift makes BRIDGES sensitive to the underlying grammar, rendering novel sentences following the original grammar somewhat familiar. Sentences not following the learned grammar do analogically match stored exemplars (parallel connectivity is violated), making these items less familiar and resulting in greater looking time as infants dishabituate.

Discussion BRIDGES is able to discriminate between abstract patterns on the basis of analogical similarity. Storing concrete exemplars, shifting attention, and analogical match are sufficient to show generalization to novel items. BRIDGES’s success calls into question Marcus et al.’s (1999) claim that algebraic rules underly infant performance. However, BRIDGES’s success is attributable to its ability to bind arguments to relations, which supports Marcus et al.’s claim that infants bind variables.

Marcus (1999) has criticized other accounts (e.g. Seidenberg, & Elman 1999) of these results for including a same/different detector within a learning mechanism. The BRIDGES simulations do not explicitly label speech sounds as identical, rather the model assumes that infants can categorize speech sounds, as embodied by the type/token distinction. BRIDGES’s solution does not hinge on a same detector. In fact, the patterns that can be discriminated by analogical mapping (even in simple domains in which only the type relation is present) are more encompassing than same/different. The analogical mapping process in these simulations aligned the current stimulus to stored exemplars — BRIDGES did not label words within sentences as same or different nor did it shift attention to a same feature. Abstract responding arose through analogy to stored exemplars and attention shifting from concrete features to relations.

Modeling Same-Different Learning

BRIDGES holds that the basis for understanding abstract relations is similarity-based and therefore inherently graded. The design of the Marcus et al. study did not touch on this possibility because stimuli were either grammatical or ungrammatical. According to BRIDGES, learners can both respond abstractly (i.e., generalize to featurally novel stimuli) and show evidence for the influence of past examples. If BRIDGES is correct, category membership in relationally defined categories is graded as it is in natural categories (Rosch, & Mervis 1975).

To evaluate BRIDGES’s predictions, we will consider results from a series of studies exploring how pigeons and humans learn notions of same and different. To illustrate how BRIDGES learns the concepts same and different, we will apply BRIDGES to Young and Wasserman’s (1997) study of same/different discriminaton learning in pigeons. To foreshadow, Young and Wasserman’s results indicate that pigeons can master a notion of same and different that cannot be explained by featural similarity. At the same time, the pigeons are sensitive to the particular examples they experienced during training and display a graded notion of same and different. Although fascinating, it would be easy to dismiss these results as relevant to pigeon cognition, but not human cognition. However, work by Young and Wasserman (2001) found the same pattern of performance with human subjects. Humans as a group are slightly more deterministic than pigeons, but this group difference is within the range of individual differences. The bottom and top 20% of humans clearly bracket the mean performance of pigeons.

In Young and Wasserman (1997), pigeons learned to respond differentially to displays containing 16 identical and 16 different icons. On each trial, the 16 icons were randomly placed within a 5 X 5 grid. The pigeons were reinforced for pushing a green button when presented with a same stimulus and a red button when presented with a different stimulus. Training consisted of blocks of 16 same stimuli and 16 different stimuli in a random order. An identical set of icons was used to form stimuli for both the same and different items, making it impossible to correctly associate an icon or icon feature with a response. Training continued until the pigeons reached 85% accuracy.




Figure 2: Two examples of intermediate stimuli are shown. The numerical code below each stimulus indicates its experimental condition and is explained in the main text.

The test phase consisted of intermediate stimuli that were somewhat similar to both the same and different stim-

uli experienced in the training phase. Examples of intermediate stimuli are shown in Figure 2. These stimuli can be viewed as forming a continuum between the pure same stimuli (all 16 icons identical) and the pure different stimuli (all 16 icons different) used during the training phase. Eleven distinct conditions of intermediate stimuli were used. The 11 conditions can be characterized by their groupings of identical icons. For example, in Figure 2, the right most stimulus contains seven question marks, five dominoes, three arrows, and one magnifying glass and thus is an example of condition [7,5,3,1]. Adopting this nomenclature, the eleven intermediate conditions were [14,1,1], [8,8], [13,1,1,1], [12,1,1,1,1,1], [10,3,2,1], [7,5,3,1], [4,4,4,4], [8,1,1,1,1,1,1,1], [2,2,2,2,2,2,2,2], [4,1,1,1,1,1,1,1,1,1,1,1], [2,1,1,1,1,1,1,1,1,1,1,1,1,1]. The pigeon's performance in these intermediate conditions, as well BRIDGES's predictions, are shown in Figure 3 (with the data points ordered left to right in the order the conditions are introduced in the previous sentence).




Figure 3: The results from Young and Wasserman's (1997) studies and BRIDGES's predictions are shown. The 11 intermediate conditions, forming a continuum between pure same and pure different stimuli, are described in the main text.

BRIDGE's Simulation Like the Marcus et al. (1999) simulations, we adopted a minimal approach to stimulus representation. Each stimulus's icon was represented as an entity. Each of the 16 entities participated in a type relation as in the Marcus et al. simulations (see Table 1).

The training regimen mimicked the procedure used in the original study as closely as possible. BRIDGES, like the pigeons, was trained to an 85% accuracy threshold before initiating the test phase. BRIDGES reached this level of performance by discovering analogical mappings among presented stimuli and exemplars stored in memory. For instance, a same stimulus aligns perfectly with same exemplars stored in memory. For example, consider aligning a stimulus containing 16 squares to another stimulus containing 16 triangles. Each triangle entity is put into correspondence with a square entity. This results in a perfect feature mismatch, but parallel connectivity is preserved. Within each type relation, the type triangle maps to the type square. This alignment leads to attention shifting toward the type relation and away from the entities. In contrast, only 1 out of 16 type relations

will exhibit parallel connectivity when aligning a different stimulus with a same stimulus. Thus, it is straightforward for BRIDGES to discriminate between same and different stimuli in the absence of featural support. In this regard, BRIDGES's solution for these simulations is the supervised learning analog to BRIDGE's discrimination of grammatical and ungrammatical sentences in the Marcus et al. (1999) simulations.

Strong support for BRIDGES's similarity-based discovery of the same/different relation is found in its fit of the test phase. BRIDGES correctly orders the intermediate conditions (see Figure 3). BRIDGES correctly predicts the probability that pigeons respond "different" in the intermediate conditions. Similarity-based activations are not all or none and these intermediate cases activate stored exemplars to varying degrees, leading to the successful fit. For example, an intermediate stimulus containing 12 triangles and four squares is somewhat analogous to a stimulus containing 16 circles — mapping the triangle type to the circle type preserves parallel connectivity in 12 out of 16 relations. Along similar lines, an item with two matching icons and 14 icons that are all different from one another better matches a pure different exemplar than a pure same exemplar.

Discussion

BRIDGES learned the same/different relation and achieved an excellent fit ($R^2 = .95$) of the test results involving intermediate stimuli. The simulations demonstrate how abstract concepts can be acquired through storage and analogy to concrete examples. BRIDGES's excellent fit of the intermediate conditions is a natural consequence of similarity-based processing. Like natural categories, BRIDGES predicts relational categories have a graded structure.

Young and Wasserman (1997) offered an entropy explanation of their results. Pure different displays will have maximum entropy (4 bits) whereas pure same displays have minimal entropy (0 bits). Like BRIDGES, the entropy explanation provides an excellent fit of the test results ($R^2 = .94$). Advantages of BRIDGES's account lie within its generality and ability to model learning data. BRIDGES can also predict additional phenomena that naturally follow from exemplar representations of categories: 1) When pigeons are transferred to new icons, performance remains above chance, but significantly declines, and 2) Increasing the discriminability of stimuli improves overall performance (Young & Wasserman, 1997).

General Discussion

By combining insights from the category learning and analogy literatures, BRIDGES provides an account of how people and animals can gain abstract understandings of domains based solely on experience with concrete instances. BRIDGES's power arises from using a notion of similarity informed by work in both analogy and category learning. Structural alignment processes allow BRIDGES to appreciate analogical similarities, while attention shifting modifies BRIDGES notion of similarity over the course of learning. Integrating these mechanisms allows BRIDGES to grasp ab-

stract patterns by shifting attention to relations which drive the alignment process.

In the supportive simulations, BRIDGES offered an explanation of how infants become sensitive to abstract grammars and how people and pigeons develop the concepts of same and different irrespective of a stimulus's features. Consistent with BRIDGES's stance that abstract concepts are similarity-based, the relational concepts same and different displayed graded membership like natural categories.

BRIDGES is not the first model to use analogical alignment to support category learning. SEQL can acquire category structures through a process of repeated abstraction of a structured category representation and has been successfully applied to the infant grammar learning studies considered here (Kuehne, Gentner & Forbus, 2003). While SEQL stresses building abstract representations, abstraction in BRIDGES arises from shifting attention. Some relative strengths of BRIDGES are that it extends an existing model of category learning (ALCOVE is a special case of BRIDGES) and incorporates attentional mechanisms.

Related work in machine learning extends support vector machines (SVMs) for application to structured representations like those processed by BRIDGES (Joachims, 1993). Applying principles from SVMs to BRIDGES would offer a rigorous notion of optimality, whereas BRIDGES demonstrates the promise of adapting SVMs' kernels during learning through attentional shifts.

One challenge for BRIDGES is incorporating new relational information into its exemplar representations. Although BRIDGES can learn relational concepts, BRIDGES is not yet able to incorporate acquired relations directly into its exemplar representations (see Doumas and Hummel, 2005, for an example of a predicate discovery system).

Another challenge facing BRIDGES is the computational complexity of an exemplar-based structural alignment system. The MAC/FAC model (Forbus, Gentner & Law, 1994) uses an interesting approach to reduce the complexity of the problem. MAC/FAC involves a two stage process: first, it compares content vectors, simplified representations of the exemplars, selecting only the most similar candidates; then, it calculates the full structural alignment only on those exemplars. Further savings in processing can be achieved by using a general search procedure to discover the optimal (or near optimal) mappings in the alignment stage.

Acknowledgments

We thank Mike Young for providing access to his data. This work was supported by AFOSR grant FA9550-04-1-0226 and NSF CAREER grant 0349101 to B. C. Love.

References

Doumas, L.A.A., & Hummel, J.E. (2005). A Symbolic-Connectionist Model of Relational Discovery. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Eimas, P.D., Siqueland, E.R., Jusczyk, P., & Vigorito, J. (1971). Infant Speech Perception. *Science*, 171(3968),

303-306

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

Forbus, K.D., Gentner, D., Law, K. (1994). MAC/FAC: A Model of Similarity-based Retrieval. *Cognitive Science*, 19, 141-205

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170

Gentner, D.; & Markman, A.B. (1997). Structure Mapping in Analogy and Similarity. *American Psychologist*, 52(1), 45-56

Hahn, U., Chater, N., & Richardson, L.B. (2003). Similarity as Transformation. *Cognition*. 87, 1-32

Joachims, T. (2003). Learning to Align Sequences: A Maximum Margin Approach. Technical Report.

Kuehne, S.E.; Gentner, D.; and Forbus, K.D. (2003). Modeling Infant Learning via Symbolic Structural Alignment. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kruschke, J.K. (1992) ALCOVE: An Exemplar-Based Connectionist Model of Category Learning. *Psychological Review* 99(1), 22-44

Larkey, L.B.; & Love, B.C. (2003). CAB: Connectionist Analogy Builder. *Cognitive Science* 27, 781-794

Love, B.C.; Medin, D.L.; & Gureckis, T.M. (2004) SUSTAIN: A Network Model Of Category Learning. *Psychological Review* 111(2), 309-332

Marcus, G.F.; Vijayan, S.; Bandi Rao, S.; and Vishton, V.E. (1999) Rule Learning by Seven-Month-Old Infants. *Science* 283,5398, 77-80

Marcus, G. F. (1999). Do infants learn grammar with algebra or statistics? Response to Seidenberg & Elman, Eimas, and Negishi. *Science*. 284, 436-437.

Medin, D.L., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*. 85, 207-238

Rosch, E. & Mervis, C.B. (1975). Family Resemblance: Studies in the Internal Structure of Categories. *Cognitive Psychology*. 7, 573-605

Seidenberg, M.S., & Elman, J.L. (1999). Do Infants Learn Grammars with Algebra or Statistics? *Science*. 284, 5413, 433

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 24, 1411-1436.

Young, M.E., & Wasserman, E.A. (1997). Entropy detection by pigeons: Response to mixed visual displays after same-different discrimination training. *Journal of Experimental Psychology: Animal Behavior Processes*, 23, 157-170

Young, M.E., & Wasserman, E.A. (2001). Entropy and Variability Discrimination. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 27(1), 278-293