

## Chapter 4

---

# Mechanistic models of associative and rule-based category learning

Bradley C. Love and Marc Tomlinson

### Editor's Preview

Chapter 4, 'Mechanistic models of associative and rule-based category learning', reviews some of the computational models that have been used to simulate the different forms of concept learning that were introduced in Chapters 2 and 3. One class of models has been designed to capture our ability to learn rule-based categories that have defining features. Such models are well suited to learn concepts such as *triangle* in which all members have three sides, angles that add to 180°, and form a closed figure. With rule-based concepts, membership is all or none, meaning that as long as an object has the criterial features of a triangle, it is a triangle, and no triangle is any better example of a triangle than any other triangle.

A difficulty with category-learning models that can just learn rules is that many categories do not seem to be rule based. For example, one might try to define the concept of *bird* with features such as, has feathers, can fly, and lays eggs, but exceptions can be found (i.e. baby male ostriches are birds, but they do not (1) have feathers, (2) fly, or (3) lay eggs). In a more famous example, the philosopher Wittgenstein argued that a concept such as *game* has no attribute shared by virtually all of its members. Rather, members of a category bear a relationship of family resemblance where there is a cluster of attributes that characterizes the family, but hardly any attribute would hold for all members of the family.

Categories with a family-resemblance structure are learned well by models that represent prototypes. A prototype can be thought of as the central tendency of a category or as the best representative of a family in the sense of having the greatest number of attributes in common with other members of the category and the fewest number of attributes in common with members of contrast categories. Models that represent categories as prototypes can account for typicality effects in which category membership is graded with some members

judged as more representative of the category than other members, based on their greater resemblance to the prototype (e.g. robins are rated as more typical birds than are penguins).

While prototype models perform well in terms of learning, many natural categories that have a regular, similarity-based structure, models that represent concepts as sets of exemplars perform better at representing more complex, irregular category structures (i.e. large and dark things, and small and light things go into one category, and small and dark things, and large and light things go into another category, see Figure 4.3 of Chapter 4). Exemplar models represent more information about categories than do prototype models in terms of retaining data about the frequency of the exemplars, their variability, and the correlations among them. However, at the level of intuition, exemplar models are not always viewed as cognitively efficient because they store every exemplar ever encountered (e.g. is it reasonable to think that a person walking down the street will record every similar-looking bird that is experienced and store it in memory?).

Given that each of the models (rule, prototype, and exemplar) runs into one or another difficulty, investigators have more recently been implementing hybrid models where the manner in which the model represents concepts is determined by the structure of those concepts (i.e. categories that can be organized by a regular, similarity-based structure are represented as prototypes, whereas categories that have a more complex, irregular structure are represented as exemplars). There are also models that have different category-learning systems within them such that each category-learning system functions via a distinct set of operating principles. Despite the differences in the various systems of category learning, one commonality is that they all function based on experiences with exemplars. Given that human adults, human infants, and nonhuman animals all experience exemplars, one may therefore look for evidence that the various category-learning systems are continuous across development and species.

## Introduction

Judging a person as a friend or foe, a mushroom as edible or poisonous, or a sound as an *l* or *r* are examples of categorization problems. As people never encounter the same exact stimulus twice, they must develop categorization schemes that capture the useful regularities in their environment. One challenge for psychological research is to determine how humans acquire and represent categories.

The focus of this chapter is on proposed category learning mechanisms. We focus on models that attempt to explain how people acquire categories from observed


examples, as opposed to verbal instruction. Most of the models that are discussed in this chapter were developed to account for adult human performance, but many of these models have also been successfully applied to studies involving humans of all ages and to other species. Category learning is a theory and model-rich area within cognitive psychology. Models have played a prominent role in shaping our understanding of human category learning. Accordingly, proposed mechanisms are diverse, including rule-, prototype-, and exemplar-based models, as well as hybrid models and models that contain multiple systems. One general trend is towards models with increasingly sophisticated processing mechanisms that can mimic the behaviours of existing models, as well as address behaviours outside the scope of previous models.

In the course of reviewing these various models, we emphasize what the relative merits of each model reveal about the nature of human learning. When we discuss exemplar models, we devote special attention to a model of category learning that attempts to bridge work in the analogy and category-learning literatures. The model, ‘building relations through instance-driven gradient error shifting (BRIDGES), successfully accounts for findings in the child and animal learning literatures. (Tomlinson & Love, 2006) We choose to showcase this particular model because it is well matched to the overarching goals of this book, and the focus should help the reader understand the basis for all the models reviewed, which is the primary goal of this chapter.

In the remainder of this chapter, we briefly review several models of human category learning. Presentation order is organized chronologically from oldest to most recent accounts of category learning. Although more recent models offer some advantages over their ancestors, it would be a mistake to view ancestral models as being supplanted by their descendants. Each model class addresses some key aspects of human category learning and serves an important theoretical role. In fact, many older models have taken on new life as components in recently proposed multiple systems models. One common component in these multiple systems models is a rule-based system, which is the first model class that we consider.

## Rule-based models

The classical view of concepts holds that categories are defined by logical rules. This view has a long history dating back to Aristotle. In Figure 4.1, any item that is a square is a member of category *A*. This simple rule determines category membership. According to the rule view, our concept of category *A* can be represented by this simple rule. Discovering this rule would involve a rational hypothesis testing procedure. This procedure attempts to discover a rule that is satisfied by all of the positive examples of a concept, but none of the negative examples of the concept (i.e. items that are members of other categories). In trying to come up with such a rule for category *A*, one might first try the rule *if dark, then in category A*. After rejecting this rule (because there are counterexamples), other rules would be tested (starting with simple rules and progressing towards more complex rules) until the correct rule is eventually discovered. For example, in learning about birds, one might first try the rule *if it flies, then it is a bird*. This rule works pretty well, but not perfectly (penguins do not fly and bats do). Another simple rule like *if it has feathers, then it is a bird* would not work either



**Fig. 4.1** Examples of category *A* and category *B*. A simple rule on shape discriminates between the two categories.

because a pillow filled with feathers is not a bird. Eventually, a more complex rule might be discovered like *if it has feathers and wings, then it is a bird*.

For decades, psychologists have conducted experiments to characterize the relative difficulty people have in learning various types of rules. (Bruner, Goodnow, & Austin, 1956; Shepard, Hovland, & Jenkins, 1961) These studies have provided the primary data used to develop and validate models of hypothesis testing. Some models, like RULEX, (Nosofsky et al. 1994) embody the hypothesis-testing procedure described above. RULEX starts with simple hypotheses and progresses towards more complex hypotheses until a set of rules and exceptions is discovered that properly discriminates between the categories.

As discussed in the Chapter 3 in this book, the term *rule* has various, somewhat conflicting, interpretations. Here, we focus on rule-based models, like RULEX, that engage in explicit, hypothesis testing. RULEX's mechanistic approach (i.e. algorithmic in the sense of Marr, 1982) contrasts with other approaches that aim to predict how difficult learning should be based on calculations of how complex the correct hypothesis is (Feldman, 2000). The latter approaches, which are not concerned with the actual process of learning, have more in common with measures of complexity and compression (Pothos & Chater 2002). Yet other approaches, such as general recognition theory (Maddox & Ashby, 1993), aim to assess and compactly describe people's performance rather than characterize the learning process. Unlike these more abstract approaches, mechanistic models of hypothesis testing, such as RULEX, largely implement the strategic and conscious thought processes that we feel (by introspection) that we are carrying out when solving classification problems.

Although rules can in principle provide a concise representation of a concept, often more elaborate representations would serve us better. Concept representation needs to be richer than a simple rule because we use concepts for much more than simply classifying objects we encounter. For instance, we often use concepts to support inference (e.g. a child infers that members of the category stove can be dangerously hot). Using categories to make inferences is a very important use of concepts (Markman & Ross, 2003). Knowing something is an example of a concept tells us a great deal about the item. For example, after classifying a politician from the USA as a Republican, one can readily infer the politician's position on a number of issues. The point is that our representations of concepts must include information beyond what is needed to classify items as examples of the concept. For example, the rule *if square, then in category A* correctly classifies all members of category A in Figure 4.1, but it does not capture the knowledge that all category A members are *dark*. One problem with rule representations of concepts is that potentially useful information is discarded. In fact, even when people explicitly use rules to classify items, performance is heavily influenced by rule-irrelevant information (Allen & Brooks, 1991; Lacroix, Giguere, & Larochelle, 2005; Sakamoto & Love, 2004), which is inconsistent with rules serving as the sole basis for category representations.

Perhaps the biggest problem with the rule approach to concepts is that most of our everyday categories do not seem to be describable by a tractable rule. To demonstrate this point, Wittgenstein (1953) noted that the concept game lacks a defining property. Most games are fun, but Russian roulette is not fun. Most games are competitive, but ring around the roses is not competitive. While most games have characteristics in common, there is not a rule that unifies them all. Rather, we can think of the members of the category game as being organized around a family-resemblance structure (analogous to how members of your family resemble one another). Rosch and colleagues (Rosch & Mervis, 1975) in their seminal work, demonstrated the psychological reality of many of Wittgenstein's intuitions. Even some paradigmatic examples of rule-based classification reveal a non-rule-based underbelly (see Love, Tomlinson, & Gureckis, 2008, for a review). Hahn and Ramscar (2001) offer one such example. Tigers are defined as having tiger DNA, which is a seemingly rule-based category definition. However, determining whether an animal has tiger DNA amounts to assessing the similarity of the animal's DNA to known examples of tiger DNA.

A related weakness of the rule account of concepts is that examples of a concept differ in their typicality (Barsalou, 1985; Posner & Keele, 1968; Reed, 1972; Rosch & Mervis, 1975). If all a concept consisted of was a rule that determined membership, then all examples should have equal status. According to the rule account, all that should matter is whether an item satisfies the rule. Our concepts do not seem to have this definitive flavour. For example, some games are better examples of the category game than others. Basketball is a very typical example of the category games. Children play basketball in a playground, it is competitive, there are two teams, each team consists of multiple players, you score points, etc. Basketball is a typical example of the category of games because it has many characteristics in common with other games. On the other hand, Russian roulette is not a very typical game – it requires a gun and one of the two players dies. Russian roulette does not have many properties in common

with other games. In terms of family-resemblance structure, we can think of basketball as having a central position and Russian roulette being a distant cousin to the other family members. These findings extend to categories in which a simple classification rule exists. For example, people judge the number three to be a more typical odd number than the number forty-seven even though membership in the category odd number can be defined by a simple rule (Gleitman, Gleitman, Miller, & Ostrin, 1996).

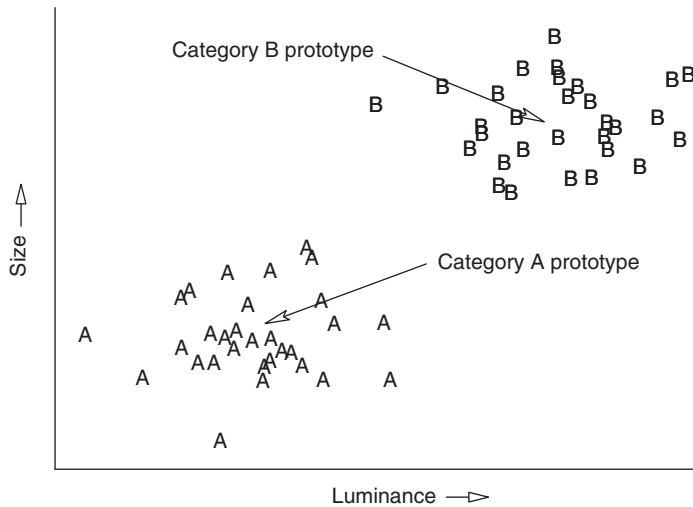
The fact that category membership follows a gradient as opposed to being all or none affords us flexibility in how we apply our concepts. Of course, this flexibility can lead to ambiguity. Consider the concept mother (see Lakoff, 1987, for a thorough analysis). It is a concept that we are all familiar with that seems straightforward – a mother is a woman who becomes pregnant and gives birth to a child. But what about a woman who adopts a neglected infant and raises it in a nurturing environment? Is the birth mother who neglected the infant a mother? What if a woman is implanted with an embryo from another woman? Court cases over maternity arise because the concept of motherhood is ambiguous. The concept exhibits greater flexibility and productivity than is even indicated above. For example, is it proper to refer to an architect as the mother of a building? All the above examples of the concept mother share a family-resemblance structure (i.e. they are organized around some commonalities), but the concept is not rule based. Some examples of the concept mother are better than others.

We do not want to imply that rule-based approaches do not have their place. For example, rule-based approaches might be viable for some socially defined categories. For example, determining whether currency is legal tender might largely involve applying a series of rules (Hampton, 2001). Moreover, as we will see later in this chapter, rule-based approaches figure prominently in multiple systems accounts. While rule-based approaches might not provide a sufficient explanation of human learning in isolation, such approaches might prove viable in certain domains or as components of multiple system models.

## Prototype-based models

The prototype approach to concept learning and representation was developed by Rosch and colleagues to address some of the shortcomings of the rule approach. Prototype models represent information about all the possible properties (i.e. stimulus dimensions), instead of focusing on only a few properties like rule models do. The prototype of a category is a summary of all of its members (Posner & Keele, 1968; Reed, 1972; Smith & Minda, 2001). Mathematically, the prototype is the average or central tendency of all category members. Figure 4.2 displays the prototypes for two categories, simply named categories *A* and *B*. Notice that all the items differ in size and luminance (i.e. there are two stimulus dimensions) and that the prototype is located amidst all of its category members. The prototype for each category has the average value on both the stimulus dimensions of size and luminance for the members of its category.

The prototype of a category is used to represent the category. According to the prototype model, a novel item is classified as a member of the category whose prototype



**Fig. 4.2** Two categories and their prototypes.

it is most similar to. For example, a large, bright item would be classified as a member of category *B* because category *B*'s prototype is large and bright (see Figure 4.2). The position of the prototype is updated when new examples of the category are encountered. For example, if one encountered a very small and dark item that is a member of category *A*, then category *A*'s prototype would move slightly towards the bottom left corner in Figure 4.2. As an outcome of learning, the position of the prototype shifts towards the newest category member in order to take it into account. A prototype can be very useful for determining category membership in domains where there are many stimulus dimensions that each provide information useful for determining category membership, but no dimension is definitive. For example, members of a family may tend to be tall, have large noses, a medium complexion, brown eyes, and good muscle tone, but no family member possesses all of these traits. Matching on some subset of these traits would provide evidence for being a family member.

Notice the economy of the prototype approach. Each cloud of examples in Figure 4.2 can be represented by just the prototype. The prototype is intended to capture the critical structure in the environment without having to encode every detail or example. It is also fairly simple to determine which category a novel item belongs to by determining which category prototype is most similar to the item.

Unlike the rule approach, the prototype model can account for typicality effects. According to the prototype model, the more typical category members should be those members that are most similar to the prototype. In Figure 4.2, similarity can be viewed in geometric terms – the closer items are together in the plot, the more similar they are. Thus, the most typical items for categories *A* and *B* are those that are closest to the appropriate prototype. Accordingly, the prototype approach can explain why robins are more typical birds than penguins. The bird prototype represents the average bird: has wings, has feathers, can fly, can sing, lives in trees, lays eggs, etc. Robins share

all of these properties with the prototype, whereas penguins differ in a number of ways (e.g. penguins cannot fly, but do swim). Extending this line of reasoning, the best example of a category should be the prototype, even if the actual prototype has never been viewed (or does not even exist). Indeed, numerous learning studies support this conjecture. After viewing a series of examples of a category, human participants are more likely to categorize the prototype as a category member (even though they never actually viewed the prototype) than they are to categorize an item they have seen before as a category member (Posner & Keele, 1968).


As the prototype approach does not represent concepts in terms of a logical rule that is either satisfied or not, it can explain how category membership has a graded structure that is not all or none. Some examples of a category are simply better examples than other examples. Moreover, categories do not need to be defined in terms of logical rules, but are rather defined in terms of family resemblance to the prototype. In other words, members of a category need not share a common defining thread, but rather can have many characteristic threads in common with one another.

The prototype approach, while preferable to the rule approach for the reasons just discussed, does fail to account for important aspects of human concept learning. The main problem with the prototype model is that it does not retain enough information about examples encountered in learning. For instance, prototypes do not store any information about the frequency of each category; yet people are sensitive to frequency information. If an item was about equally similar to the prototype of two different categories and one category was one hundred times larger than the other, people would be more likely to assign the item to the more common category (under most circumstances, see Kruschke, 1996). Of course, some of these concerns could be addressed by expanding the information that a prototype encodes.

However, other concerns seem fundamental to the prototype approach. Prototypes are not sensitive to the correlations and substructure within a category. For example, a prototype model would not be able to represent that spoons tend to be large and made of wood or small and made of steel. These two subgroups would simply be averaged together into one prototype. This averaging makes some categories unlearnable with a prototype model. One example of such a category structure is shown in Figure 4.3. Each category consists of two subgroups. Members of category *A* are either *small* and *dark*, or they are *large* and *light*, whereas members of category *B* are either *large* and *dark* or they are *small* and *light*. The prototypes for the two categories are both in the centre of the stimulus space (i.e. medium size and medium luminance). Items cannot be classified correctly by which prototype they are most similar because the prototypes provide little guidance.

In general, prototype models can only be used to learn category structures that are linearly separable. A learning problem involving two categories is linearly separable when a line or plane can be drawn that separates all the members of the two categories. The category structure shown in Figure 4.2 is linearly separable because, a diagonal line can be drawn that separates the category *A* and *B* members (i.e. the category *A* members fall on one side of the line and the category *B* members fall on the other side of the line). Thus, this category structure can be learned with a prototype model. The category structure illustrated in Figure 4.3 is nonlinear – no single line can be drawn





**Fig. 4.3** Two categories and their prototypes.

to segregate the category *A* and *B* members. Mathematically, a category structure is linearly separable when there exists a weighting of the feature dimensions that yields an additive rule that correctly indicates one category when the sum is below a chosen threshold and the other category when the sum is above the threshold.

The inability of the prototype model to learn nonlinear category structures detracts from its worth as a model of human concept learning because people are not biased against learning nonlinear category structures. While the extent to which natural categories deviate from linear structures is contended (Murphy, 2002), the general consensus is that people in the laboratory do not show a preference for linear structures in supervised learning (Medin & Schwanenflugel, 1981), though they might in unsupervised learning (Love, 2002). Some nonlinear category structures may actually be easier to acquire than linear category structures. For example, it seems quite natural that small birds sing, whereas large birds do not sing. Many categories have subtypes within them that we naturally pick out. One way for the prototype model to address this learnability problem is to include complex features that represent the presence of multiple simple features (e.g. large and blue). Unfortunately, this approach quickly becomes unwieldy as the number of stimulus dimensions increases (e.g. Gluck & Bower, 1988).

Related to the prototype model's inability to account for substructure within categories is its inadequacy as a model of item recognition. Unlike exemplar models considered in the following section (Medin & Schaffer, 1978; Nosofsky, 1986), prototypes models do not readily account for how people recognize specific items because the category prototype averages away item distinguishing information that people retain in some situations.

## Exemplar-based models

Exemplar models store every training example in memory instead of just the prototype (i.e. the summary) of each category. Perhaps surprising upon first consideration,

exemplar models can account for findings marshaled in support of prototype models, such as sensitivity to family-resemblance structure. At the same time, by retaining all of the information from training, exemplar models address many of the shortcomings of prototype model. Exemplar models are sensitive to the frequency, the variability, and the correlations among items. In this section, we discuss how exemplar-based models can display these behaviors.

Unlike prototype models, exemplar models can master category structures that contain substructure. For the learning problem illustrated in Figure 4.3, an exemplar model would store every training example. New items are classified by how similar they are to all items in memory (not just the prototype). For the category structure illustrated in Figure 4.3, the pairwise similarity of a novel item and every stored item would be calculated. If the novel item tended to be more similar to the category *A* members (i.e. the item was small and dark) than the category *B* members, then the novel item would be classified as a member of category *A*.

One aspect of exemplar models that seems counterintuitive is their lack of any abstraction in category representation. It seems that humans do learn something more abstract about categories than a list of examples. Surprisingly, exemplar models are capable of displaying abstraction. For instance, exemplar models can correctly predict that humans more strongly endorse the underlying prototype (even if it has not been seen) than an actual item that has been studied (a piece of evidence previously cited in favour of the prototype model). How could this be possible without the prototype actually being stored? It would be impossible if exemplar models simply functioned by retrieving the exemplar in memory that was most similar to the current item and classified the current item in the same category as the retrieved exemplar (this is essentially how processing works in a prototype model, except that a prototype is stored in memory instead of a bunch of exemplars).

Instead, exemplar models engage in more sophisticated processing and calculate the similarity between the current item (the item that is to be classified) and every item in memory. Some exemplars in memory will be very similar to the current item, whereas others will not be very similar. The current item is classified in the category in which the sum of its similarities to all the exemplars is greatest. When a previously unseen prototype is presented to an exemplar model, it can be endorsed as a category member more strongly than a previously seen item. The prototype (which is the central tendency of the category) will tend to be somewhat similar to every item in the category, whereas any given non-prototype item will tend to be very similar to some items (especially itself!) in memory, but not so similar to other items. Overall, the prototypical item can display an advantage over an item that has actually been studied. Abstraction in an exemplar model is indirect and results from processing (i.e. calculating and summing pairwise similarities), whereas abstraction in a prototype model is rather direct (i.e. prototypes are stored).

By and large, exemplar models can mimic all the behaviours of prototype models, but the opposite is not true. There are some subtle behaviours that the prototype model can display that versions of exemplar models cannot. For example, prototype and exemplar models predict slightly different category-endorsement gradients

(i.e. probability of membership) as one moves towards the centre of a category (see Nosofsky & Zaki, 2002; and Smith, 2002, for a recent debate).

Although exemplar models are decent models of recognition, they do have some fundamental shortcomings. Exemplar models calculate recognition strength as the sum of similarity to all items stored in memory. Thus, the pairwise similarity relations among items governs recognition. However, humans often appear to build schema-like structures in memory and store items preferentially that deviate from these structures (see Sakamoto & Love, 2004, for a review). Thus, exemplar models do not correctly predict enhanced recognition for items that violate salient rules or patterns (Palmeri & Nosofsky, 1995). Exemplar models do not capture these results because exception items that violate these patterns are not exceptional in terms of their pairwise similarity relations to other items. Exception items are exceptional in terms of violating a knowledge structure stored in memory (Sakamoto & Love, 2004, 2006).

At a more philosophical level, exemplar models seem to make some questionable assumptions. For example, exemplar models store every training example which seems excessive. Moreover, every exemplar is retrieved from memory every time an item is classified (though see Nosofsky & Palmeri, 1997, for an exception). In addition to these assumptions, one worries that the exemplar model does not make strong-enough theoretical commitments because it retains all information about training and contains a great deal of flexibility in how it processes information. In support of this conjecture, Sakamoto, Matsuka, and Love (2004) built an exemplar model that effectively built distributed knowledge structures and could account for exception recognition findings (also see Rodrigues & Murre, 2007). While their model did not explicitly build schema or exception representations, the model did learn to selectively tune exemplars (broad tunings for rule-following items and tight-tunings for exception items) and properly weight these exemplars to give rise to an exemplar model that functionally contained exception and schema-like knowledge structures. If there are no constraints on how items are processed, then in principle an exemplar model can account for any pattern of results thereby reducing the exemplar models' theoretical utility. However, in practice, exemplar models often follow previously published formalisms and serve as valuable theoretical tools.

## Exemplar-based relational learning

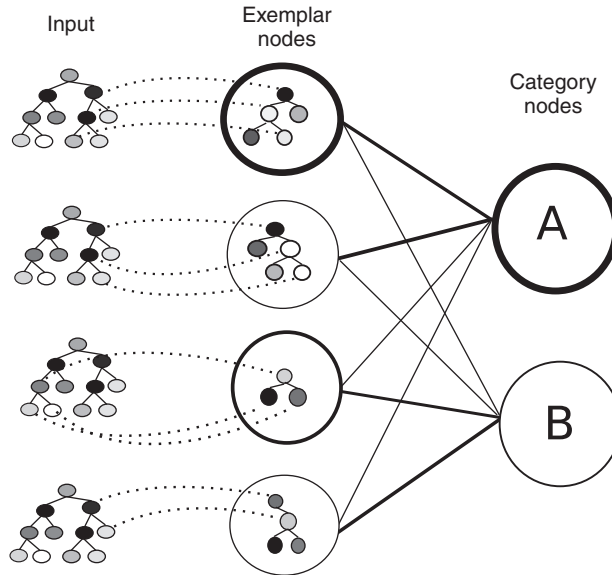
One favourable property of exemplar models is their transparency. Their predictions are purely governed by the weighting of experienced examples. This property makes them ideally suited for computational explorations of new domains, such as relational category learning. In this subsection, we consider an exemplar model of how people learn seemingly abstract concepts by analogy to exemplars. The model, BRIDGES, provides an account of how animals (and people) learn to respond relationally (Tomlinson & Love, 2006). BRIDGES differs from other exemplar models by being sensitive to relational information.

Many of our categories are relational and therefore it is important to develop models that explain how such categories are acquired. For example, membership in the category *thief* is defined by playing the appropriate relational role in the relation *steals*

rather than exhibiting some combination of concrete features (Markman & Stilwell, 2001). Differences in ability to classify relationally is often taken as a key marker of the relative mental capacities of animals, children, and adults (Thompson & Oden, 2000).

BRIDGES combines two popular approaches to cognition, exemplar-based category learning (Kruschke, 1992) and structure mapping theory (Gentner, 1983). Structure mapping theory suggests that similarity is determined between two scenes by aligning the objects and relations present within one scene with the objects and relations in the other scene (Markman & Gentner, 1993). The similarity of two scenes is then a measure of how well they align. This alignment is traditionally done using an unweighted graph-matching algorithm. Figure 4.4 provides an overview of BRIDGES.

BRIDGES extends the notion of similarity used in exemplar models to an attention-weighted form of structure-mapping theory. This allows relational similarity, the degree to which mapped objects play the same role in their corresponding relations (Jones & Love, 2007), to play a variable role in the alignment process. Attention can shift between the features (e.g. *red*) and the relations (e.g. *redder*). This allows for abstraction away from the features and to the relations, but only so far as the statistics of the environment warrant. Attention is updated according to a supervised or unsupervised gradient-descent algorithm. The result is that BRIDGES is able to learn to respond differentially to the presence of relations, but its response is still affected by



**Fig. 4.4** A depiction of the BRIDGES model. The structured graphs represent the input and exemplars. These graphs encode features and their relations (e.g. man biting dog and dog biting man would have different graphs\*\*\*). The luminance of the circles within the graphs represent attention to individual relations and features in the model. Node density reflects similarity-based activation of the nodes following the analogical match process. These activations are passed across connections weights to the category nodes.


AU: Please let us know what the asterisks here stand for.

the features of the stimuli. Figure 4.5 illustrates how attentional weighting can disambiguate between competing interpretations.

Previous simulations of same–different learning in pigeons (Young & Wasserman, 1997) and infant grammar learning (Marcus et al. 1999) have demonstrated that BRIDGES is capable of learning a variety of different relational behaviors without resorting to rules or symbol systems (Tomlinson & Love, 2006). Just like the participants in these experiments, BRIDGES generalizes to presentations of the relations with novel objects. Moreover, these relations are still clouded by the featural similarity of the individual stimuli because attention shifting is rarely complete, consistent with behavioural shifts seen in human development and acquired expertise (Chi, Feltovich, & Glaser, 1981; Gentner & Ratterman, 1991).

Other explanations for same–different learning centre on measures of display entropy or variability (Young, Ellefson, & Wasserman, 2003). These explanations and BRIDGES are indistinguishable with a simple goodness-of-fit measure. However, BRIDGES makes a testable prediction different from the variability model; the responses in a same–different task should not just be based on the sameness and differentness of the array but also on the featural similarity between the test array and previous arrays the animal has been trained with, because some attention should still be on the features. Gibson and Wasserman (2004) provide just such a test and confirm BRIDGE’s prediction.

In Gibson and Wasserman (2004), pigeons are trained on stimuli consisting of arrays of 16 icons drawn from one of two sets of icons, *a* and *b*. *Same* arrays always contain 16 identical *a* icons, whereas *different* arrays always contain different arrangements of the 16 unique *b* icons. When pigeons are tested with novel arrays with icons from set *c*, they behave based on the relations within the array, but when shown *different* arrays containing *a* icons, the pigeons are more likely to respond *same*, and vice-versa for *same* arrays formed with *b* icons. The pigeons learn to respond to the novel relations, but their responses are still tied to the features of the exemplars used in training.



**Fig. 4.5** An example comparison between two graphs is shown. There are two possible ways to map the elements in these corresponding relations. The example on the left preserves parallel connectivity by mapping elements that play the same role in each relation to one another. This solution is high in relational match, but low in featural match because the corresponding elements differ in shape features. The situation is reversed in the mapping shown in the right example. Attention weighting of mismatches determines which of these two possible mappings will be preferred by BRIDGES. BRIDGES chooses the mapping that minimizes attention-weighted mismatch.

These simulations provide insight into the differences among animals, infants, children, and adults. For the simulations just described, the exemplars were only represented with simple features and a type–token relationship. The type–token relationship assumes that the individual is able to recognize objects present in the input as members of the same type. In other words, when pigeons are presented with an array of shapes, they realize that all of the squares are members of an abstract type, square. This assumption is sufficient for an array of simple relational learning tasks. However, when modelling more complex behaviour, in children or adults, a more complex representation is often required. BRIDGES provides a tool to talk about these and other differences in a quantitative way.

Animals might not be able to succeed at complex relational reasoning tasks, but they can compare current examples to previous examples in a structured way, and from this respond in a manner consistent with an understanding of abstract relations. BRIDGES is a computational model of how this relation-like behaviour can be learned. By comparing concrete examples of the relations in a structured manner, one can learn to respond in a manner consistent with the relations, without true abstract knowledge. BRIDGES extends this core idea of exemplar models that all abstraction occurs as a result of online processing to relational categorization. BRIDGES's account serves to highlight the transparency and clarity of exemplar-based explanations.

## Hybrid models

Prototype and exemplar models can be seen as opposite ends of a continuum of category representation. On one extreme, prototype models store every category member together in memory. At the other extreme, exemplar models store every category member separately in memory. Between these two extremes lie a wealth of possibilities. Categories in the real world contain multiple subtypes and exceptions. For example, the category mammals contains subcategories like cats, dogs, horses, and bats. Ideally, our mental representations would reflect this structure. Both prototype and exemplar models are inflexible in that they treat the structure of each category as predetermined. These models do not let the distribution of category members influence the form category representations take. For example, prototype models assume that categories are always represented by one node (i.e. the prototype) in memory, whereas exemplar models assume that categories are always represented by one node in memory for every category example encountered.

One reasonable intuition is that similar items should cluster together in memory (Anderson, 1991; Love, Medin, & Gureckis, 2004; Vampaemel & Storms, in press). For example, a person walking down Congress Avenue in Austin in the fall will encounter thousands of seemingly identical grackles. The rationale for storing each of these birds separately in memory is unclear. At the same time, someone walking down the street probably would mentally note unusual or otherwise surprising birds.

Hybrid models embody these intuitions about memory. For example, Anderson's (1991) rational model computes the probability that an item belongs to an existing cluster (a prototype can be thought of as a cluster that encodes all category members). If this probability is sufficiently high, the cluster is updated to reflect its new member.

However, if the item is more likely from a new cluster, then a new cluster is created. The overarching goal of Anderson's model is to create clusters that are maximally predictive.

Love et al.'s SUSTAIN model operates along similar lines in that it incrementally adds clusters as it learns, but its recruitment process is somewhat different from that of the the rational model. SUSTAIN recruits new clusters in response to surprising events. What counts as a surprising event depends on the learner's current goals. When the learner's goals are somewhat diffuse as in unsupervised learning, SUSTAIN's operation is very similar to that of the rational model. In such cases, items that are dissimilar from existing clusters result in a new cluster being recruited to encode the item. However, in supervised learning situations, such as in classification learning (the learner's goal is to properly name the stimulus's category), items are recruited when a surprising error results. For example, upon encountering a bat for the first time and being asked to name it, a child surprised to learn that a bat is not a bird would recruit a new cluster to capture this example. If the child activates this cluster in the future to successfully classify other bats, then the cluster would come to resemble a bat prototype.

Both the rational model and SUSTAIN can be viewed as multiple prototype models in which the number of prototypes is determined by the complexity of the category structure. When categories are very regular, these models will function like prototype models. When categories are very irregular (i.e. there is no discernable pattern linking members to one another), these models will tend to function like exemplar models. SUSTAIN's sensitivity to a learner's goal allows it to capture performance differences across different induction tasks. For example, people learning through inference (e.g. *This is a mammal. Does it have fur?*) tend to focus on the internal structure of categories, whereas people learning through classification (e.g. *This has fur. Is it a mammal?*) tend to focus on information that discriminates between categories (see Markman & Ross, 2003, for a review).

Hybrid models, like exemplar and prototype models, can be coupled with selective-attention mechanisms that can learn to emphasize critical-stimulus properties. For example, in learning to classify car makes, SUSTAIN would learn to weight shape more than colour because shape reliably indicates model-type whereas colour varies idiosyncratically. The motivation for selective attention comes from the observation that people can only process a limited number of stimulus properties simultaneously. Selective-attention mechanisms have been developed through consideration of human and animal learning data (see Kruscke, 2003, for a review). In tasks that require people to actively sample stimulus dimensions, selective-attention mechanisms predict which dimensions are fixated (Rehder & Hoffman, 2005).

Importantly, selective-attention mechanisms allow non-rule models to display rule-like behaviours (see Chapter 3 in this book). When a prototype, exemplar, or hybrid model places all of its attention on one stimulus dimension, a model's operation is indistinguishable from the application of a simple rule. In terms of accounting for human data, SUSTAIN outperforms RULEX in some respects on learning problems that require acquiring a simple rule and storing exceptions to these rules (Sakamoto & Love, 2004). SUSTAIN creates a small set of clusters that encode items that follow the

rules and stores exception in their own clusters. Attention is heavily biased to the rule-relevant dimensions. This allows SUSTAIN to show enhanced recognition for exceptions and rule-like behaviour for rule-following items, while maintaining some sensitivity to non-rule relevant dimensions like human subjects do. In our review of exemplar models, we discussed how selective-attention mechanisms allow BRIDGES to achieve similar ends in terms of balancing the importance of featural and relational match.

The incorporation of selective-attention mechanisms into non-rule models invites a number of theoretical questions. It is not entirely clear whether these selective-attention mechanisms should be viewed as an integral part of non-rule models or as rule mechanisms grafted onto non-rule models. One possibility is that people are relying on rule and non-rule systems, thus necessitating the need for selective-attention mechanisms in non-rule models.

## Multiple systems models

Determining the best psychological model can be difficult as one model may perform well in one situation but be bested by a competing model in a different situation. One possibility is that there is not a single *true* model. In category learning, this line of reasoning has led to the development of models containing multiple learning systems. These more complex models hold that category learning behaviour reflects the contributions of different systems organized around discrepant principles that utilize qualitatively distinct representations. The idea that multiple learning systems support category learning behaviour enjoys widespread support in the cognitive neuroscience of category learning (see Ashby and O'Brien, 2005, for a review and Nosofsky and Zaki, 1998, for a dissenting opinion).

Multiple system models of category learning detail the relative contributions of the component learning systems. Some multiple system models combine the outputs of the individual systems together (Ashby et al. 1998). Over time, one system might prove more useful and dominate responding. Alternatively, the modeler can predetermine the timing of the shift from one system to another. This is sensible in cases where there is good evidence for predictable shifts, such as the shift from rule-based to exemplar-based responding in classification learning (Johansen & Palmeri, 2002).

Both of these multiple system approaches are somewhat inadequate in that they do not allow the current situation to dictate which system is operable. For example, when trying to learn how to operate a new piece of machinery, a person might use a hypothesis (i.e. rule) system, but when riding a bicycle, a more procedural system might govern responding and be updated. In some models, like ATRIUM (Erickson & Kruschke, 1998), the relative contributions of divergent systems can depend on the circumstances (cf., Yang & Lewandowsky, 2004). ATRIUM contains a rule and exemplar learning system. Which system is operable is determined by a gating system, allowing different classification procedures to be applied to different parts of the stimulus space. For example, familiar items could be classified by the exemplar system whereas rules could be applied to unfamiliar items. The power to apply qualitatively different procedures to different stimuli is the hallmark of multiple systems models.



Somewhat muddying the waters, ostensibly single-system models have been developed that also manifest this ability. In CLUSTER (Love & Jones, 2006), clusters can tune themselves (i.e. attend) to different stimulus properties and encode concepts at various levels of granularity. This allows CLUSTER to apply different procedures to different parts of the stimulus space, like ATRIUM does. For example, clusters would heavily weight colour in the domain of clothing and processor type in the domain of laptops. This tuning is accomplished by minimizing an error term that reflects the model's predictive accuracy, a technique commonly used in connectionist modelling. Tunable parameters that encode each cluster's specificity and attentional weighting of different properties are shaped by experience.

Models like CLUSTER are very rich. Consideration of such models leads to the question of what constitutes or defines a system. As previously discussed, one could even construe the selective-attention mechanism of various models as being a separate system. Fortunately, models are mathematically well specified and allow researchers to make predictions and state their theories clearly without having to be overly concerned with the semantics of what constitutes a system. The mathematical specification of models can free researchers from some potentially thorny debates.

The notion of a system perhaps takes on greater significance when considered in the context of the brain (see Chapter 5 in this book). Within cognitive neuroscience, it is generally accepted that there is a hypothesis testing system that relies on frontal circuitry (Ashby et al. 1998), a dopamine-mediated procedural learning system that involves the striatum (Ashby et al. 1998), a repetition priming system that involves early visual areas (Reber et al. 2003), and a hippocampal learning system that maps onto exemplar- or cluster-based learning (Love & Gureckis, 2007). For each system, there are behavioural manipulations that tend to emphasize the system over the other systems. Lesion, patient, and imaging studies provide compelling evidence for the multiple systems view.

## General discussion

In this chapter, we reviewed the relative merits of a variety of category learning models, including rule-, prototype-, and exemplar-based models, as well as hybrid models and multiple system models that combine two or more of the aforementioned model types. We also considered how inclusion of selective-attention mechanisms can increase the capabilities of these models by endowing them with the ability to manifest rule-following behaviour.

To review briefly each model family's merits, rule-based models conform to our intuition that we effortfully search for patterns that we can verbally communicate to others. In contrast to rule models, prototype models successfully reflect the graded nature of category membership. Exemplar models address deficiencies in the prototype model and can **model** that categories often have internal correlations. Exemplar models also capture aspects of recognition memory performance. Hybrid models successfully transition between prototype- and exemplar-like representations depending on the complexity of the category structure. Finally, multiple system models align with emerging findings from cognitive neuroscience and intuitions that there are multiple paths or mechanisms available to categorize stimuli.

AU: Can this be changed to 'exemplify' instead of model, here?

All of these models have played a critical role in driving advances in theory and in the design of key experiments. The development of new models is informed by the failings of preceding models. The history of model development is marked by the arrival of models with increasingly sophisticated processing mechanisms that can manifest the behaviours of previous models as well as additional human behaviours beyond the reach of existing models. Of course, the value in models lie more in predicting unanticipated behaviours than in simply accounting for observed behaviours. Thus, it is important for models to be somewhat constrained to have theoretical value.

We would like to end by encouraging researchers to consider conducting model-directed research. All experimenters are driven by theoretical considerations and models are ideally suited to bring these issues into focus and unite seemingly discrepant findings. In the exemplar-based model section, we described how a relatively simple model brought together powerful ideas from the analogy and category learning literatures to address learning across species and development. These kinds of advances and connections can be facilitated through model-driven explorations.

To return to one of the main themes of this book, the work reviewed in this chapter largely supports the notion that cognition is continuous across development and species. The basic mechanisms discussed in this chapter appear to apply equally well across species and development. For example, the SUSTAIN model of category learning has been successfully applied to infants and older adults (Love & Gureckis, 2007). In fact, for certain tasks, data collected from 10-month-old infants and young adults are successfully fit with the same parameter values (Gureckis & Love, 2004). Likewise, in situations where one looks at performance across species, one will likely find evidence for common representations and processes. In our discussion of BRIDGES, the same exemplar-based relational learning mechanism captures performance data from pigeons, human infants, and adults.

Does this mean there are no differences in learning across development and species? In our opinion, the answer to this question is clearly no. The convergence we observe is likely a function of the task domain. The tasks considered here can all be modelled as a process of activating past experiences in memory (represented as exemplars or clusters). It may well be that differences in such tasks are minimal across development and species. Even for this family of tasks, differences can be observed across species, particularly for cases in which task factors emphasize strategic or attentional processes, which likely rely on prefrontal regions of the brain that are most prominent in adult humans (Smith, Minda, & Washburn, 2004). In such cases, models like SUSTAIN and BRIDGES can still prove useful in quantifying differences between populations in terms of best-fitting attentional parameters. Finally, it may very well be that we underestimate the extent to which our behaviours are governed by memory-based processes. As the work reviewed in this chapter indicates, such processes can support a variety of behaviours, especially when the learning and retrieval processes are sensitive to relational match (Tomlinson & Love, 2008)

## References

- Allen, S. W. & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, *120*, 3–19.
- Anderson, J. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429.
- Ashby, F., Alfonso-Reese, L., Turken, A. & Waldron, E. (1998). A neuropsychological theory of multiple-systems in category learning. *Psychological Review*, *105*, 442–481.
- Ashby, F. & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences*, *9*, 83–89.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure of categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *11*, 629–654.
- Bruner, J. S., Goodnow, J. J. & Austin, G. A. (1956). *A Study of Thinking* New York: Wiley.
- Chi, M., Feltovich, P. & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.
- Erickson, M. A. & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107–140.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, *407*, 630–633.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Gentner, D. & Ratterman, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds), *Perspectives on Thought and Language: Interrelations in Development*, pp. 225–277, London: Cambridge University Press.
- Gibson, B. M. & Wasserman, E. A. (2004). Time-course of control by specific stimulus features and relational cues during same-different discrimination training. *Learning and Behavior*, *32*, 183–189.
- Gleitman, L. R., Gleitman, H., Miller, C. & Ostrin, R. (1996). Similar, and similar concepts. *Cognition*, *58*, 321–376.
- Gluck, M. A. & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 225–244.
- Gureckis, T. & Love, B. C. (2004). Common mechanisms in infant and adult category learning. *Infancy*, *5*, 173–198.
- Hahn, U. & Ramscar, M. (2001). Conclusion: Mere similarity? In U. Hahn & M. Ramscar (Eds), *Similarity and Categorization*, pp. 257–272, New York: Oxford University Press.
- Hampton, J. A. (2001). The role of similarity in natural categorization. In U. Hahn & M. Ramscar (Eds), *Similarity and Categorization*, pp. 13–28, New York: Oxford University Press.
- Johansen, M. K. & Palmeri, T. J. (2002). Are there representational shifts during category learning? *Cognitive Psychology*, *45*, 482–553.
- Jones, M. & Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive Psychology*, *55*, 196–231.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.

- Kruschke, J. K. (1996). Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *22*, 3–26.
- Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science*, *12*, 171–175.
- Lacroix, G. L., Giguere, G. & Larochelle, S. (2005). The origin of exemplar effects in rule-driven categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *31*, 272–288.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press.
- Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review*, *9*, 829–835.
- Love, B. C. & Gureckis, T. M. (2007). Models in search of a brain. *Cognitive, Affective, & Behavioral Neuroscience*, *7*, 90–108.
- Love, B. C. & Jones, M. (2006). The emergence of multiple learning systems. *Proceedings of the 28th Cognitive Science Society*, pp. 507–512, Mahwah, NJ: Erlbaum.
- Love, B. C., Medin, D. L. & Gureckis, T. (2004). SUSTAIN: A network model of human category learning. *Psychological Review*, *111*, 309–332.
- Love, B. C., Tomlinson, M. & Gureckis, T. (in press). The concrete substrates of abstract rule use. *The Psychology of Learning and Motivation: Advances in Research and Theory*.
- Maddox, W. T. & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, *53*, 49–70.
- Marcus, G. F., Vijayan, S., Bandi Rao, S. & Vishton, V. (1999). Rule learning by seven-month-old infants. *Science*, *283*(5298), 77–80.
- Markman, A. B. & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, *23*, 431–467.
- Markman, A. B. & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592–613.
- Markman, A. B. & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, *13*, 329–358.
- Marr, D. (1982). *Vision* San Francisco: W. H. Freeman.
- Medin, D. L. & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.
- Medin, D. L. & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning & Memory*, *7*, 355–368.
- Murphy, G. L. (2002). *The Big Book of Concepts* Cambridge, MA: MIT Press.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.
- Nosofsky, R. M. & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*, 266–300.
- Nosofsky, R. M., Palmeri, T. J. & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *101*(1), 53–79.
- Nosofsky, R. M. & Zaki, S. F. (1998). Dissociations between categorization and recognition in amnesic and normal individuals. *Psychological Science*, *9*, 247–255.
- Nosofsky, R. M. & Zaki, S. F. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *28*, 924–940.

AU: Please update this reference if the book has already been published.

- Palmeri, T. J. & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *21*, 548–568.
- Posner, M. I. & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 241–248.
- Pothos, E. M. & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science*, *26*, 303–343.
- Reber, P. J., Gitelman, D. R., Parrish, T. B. & Mesulam, M. M. (2003). Dissociating explicit and implicit category knowledge with fMRI. *Journal of Cognitive Neuroscience*, *15*(4), 574–583.
- Reed, S. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*, 382–407.
- Rehder, B. & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*, 1–41.
- Rodrigues, P. M. & Murre, J. M. J. (2007). Rules-plus-exception tasks: A problem for exemplar models? *Psychonomic Bulletin & Review*, *14*, 640–646.
- Rosch, E. & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.
- Sakamoto, Y. & Love, B. C. (2004). Schematic influences on category learning and recognition memory. *Journal of Experimental Psychology: General*, *33*, 534–553.
- Sakamoto, Y. & Love, B. C. (2006). Vancouver, Toronto, Montreal, Austin: Enhanced oddball memory through differentiation, not isolation. *Psychonomic Bulletin & Review*, *13*, 474–479.
- Sakamoto, Y., Matuska, T. & Love, B.C. (2004). Proceedings of the international conference of cognitive modeling. In M. Lovett, C. Schunn, C. Lebiere, & P. Munro (Eds), *Dimensionwide vs. Exemplar-Specific Attention in Category Learning and Recognition* Vol. 27, pp. 261–266, Mahwah, NJ: Lawrence Erlbaum.
- Shepard, R. N., Hovland, C. L. & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75* (13, Whole No. 517).
- Smith, J. (2002). Exemplar theory's predicted typicality gradient can be tested and disconfirmed. *Psychological Science*, *13*, 437–442.
- Smith, J. D. & Minda, J. P. (2001). Journey to the center of the category: The dissociation in amnesia between categorization and recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *27*, 984–1002.
- Smith, J. D., Minda, J. P. & Washburn, D. A. (2004). Category learning in rhesus monkeys: A study of the Shepard, Hovland, and Jenkins tasks. *Journal of Experimental Psychology: General*, *133*, 398–414.
- Thompson, R. K. R. & Oden, D. L. (2000). Categorical perception and conceptual judgements by nonhuman primates: The paleological monkey and the analogical ape. *Cognitive Science*, *24*(3), 363–396.
- Tomlinson, M. & Love, B. C. (2006). From pigeons to humans: Grounding relational learning in concrete examples. In Y. Gil & R. J. Mooney (Eds), *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-2006)*, pp. 199–204, AAAI Press.
- Tomlinson, M. & Love, B. C. (2008). Monkey see, monkey do: Learning relations through concrete examples. *Behavioral and Brain Sciences*, *31*, 150–151.
- Vampaemel, W. & Storms, G. (1008). In search of abstraction: The varying abstraction model of categorization. *Psychonomic Bulletin & Review*, *10*, 29–43.
- Wittenstein, L. (1953). *Philosophical Investigations* Oxford, England: Blackwell. (G. E. M. Anscombe, trans.)

- Yang, L. X. & Lewandowsky, S. (2004). Context-gated knowledge partitioning in categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 1045–1064.
- Young, M. E., Ellefson, M. R. & Wasserman, E. A. (2003). Toward a theory of variability discrimination: Finding differences. *Behavioral Processes*, **62**, 145–155.
- 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.