13

Modeling Learning Under the Influence of Culture

Bradley C. Love and Todd M. Gureckis

Category learning is an incredibly broad topic. Researchers with heterogeneous goals and methods from various traditions are working to understand the nature of human categorization. Diversity is desirable if findings can eventually be placed in a common theoretical framework. However, one danger is that the study of categorization could fracture into isolated communities that will not benefit from insights outside their cadre.

One current tension is between researchers engaged in mathematical modeling of laboratory studies and those working in domains that explore humans' real-world knowledge, such as researchers exploring the influence of culture on category formation. Both communities have good reasons to be weary of each other. Researchers working in more naturalistic and meaningful domains fault laboratory researchers and modelers for focusing on tasks that do not approximate the richness of real-world categorization. Conversely, some modelers might be hard-pressed to see how research in less well-controlled settings will eventually lead to a mechanistic understanding of categorization.

For example, Lynch, Coley, and Medin (2000) explore how different types of tree experts have different conceptualizations of the same trees (as evidenced by their typicality ratings). Although their results are illuminating and rich, the groups they consider differ in a number of respects, making it hard to develop a causal story that ends in a mechanistic account of becoming an expert. On the other end of the naturalistic–tightly controlled spectrum, Nosofsky, Gluck, Palmeri, McKinley, and Glauthier (1994) revisit Shepard, Hovland, and Jenkins's (1961) classic learning problems and fit models to the results. Although data collected in well-controlled laboratory studies that use random assignment and counterbalanced stimuli promote model development, how such models could speak to Lynch et al.'s results remains unclear.

In this chapter, we attempt to ease this tension. We apply a model of category learning developed through consideration of data from laboratory tasks to understanding how culture affects conceptual organization. In particular, we applied the SUSTAIN model of category learning (Love & Medin, 1998b; Love, Medin, Gureckis, 2004) to cross-cultural data on experts' conceptual organization of biological kinds (Medin et al., in press). Experts from both cultures have equivalent knowledge in many ways. Nevertheless, the two groups of experts organize their knowledge in different ways. One challenge is to develop a mechanistic understanding of how culture affects the organization of these experts' knowledge.

Basic Challenges

Developing a mechanistic understanding of how culture affects conceptual organization requires addressing a number of issues that are not usually a concern when modeling data drawn from laboratory tasks. Some of these challenges are discussed below.

Knowledge Representation

The stimuli in laboratory experiments are usually constrained by the experimenter to admit one clear interpretation. For example, a stimulus set consisting of triangles and squares that vary in color and size can be represented in a straightforward manner.[1] In such constrained cases, the stimulus dimensions are clear. When one considers less constrained domains, specifying the appropriate representations is more challenging. For example, what are the dimensions that represent an expert's knowledge of ecology? Fully addressing this representational problem is beyond the grasp of the field. A more realistic goal is to approximate experts' knowledge. If it is based on reasonable assumptions, such an approximation would enable formal modeling that could bolster our understanding of the psychological processes underlying expert performance.

Task Formalization

The most common induction task used in laboratory studies is classification learning. In classification learning, the learner is presented with a stimulus, assigns it to one category in a set of mutually exclusive categories, and then receives corrective feedback. The dependent measure is based on classification accuracy. In such experiments, the learning task and key data are clear. A learning trial for a human subject closely corresponds to a learning trial for a model.

[1]However, see Love and Markman (2003) for results suggesting that mental representations of even this very constrained stimulus set are more elaborate than previously believed. In particular, when spatial working memory is available, color and size act as predicates with shape serving as the argument.

Identification of the learning problem is not as clear in the case of the development of expertise within a cultural context. Models require specific training regimens. The modeler must decide what the learning problem is, how long to train the model, and what the feedback pattern should be. Mapping a model simulation to the learning history of an individual member of a culture is a much more challenging task than fitting a model to the participant's data in a laboratory study. In the former case, training will have to be somewhat idealized and the mapping qualitative.

Simulation Evaluation

Like task formalization, the evaluation of simulations of laboratory classification-learning studies is straightforward. For these studies, a model simulation is successful to the extent that the error pattern generated by the model corresponds to that of human subjects. When the goal of a simulation is for the model's conceptual organization to correspond to that of a human from a specific cultural group, the evaluation procedure is greatly complicated. Both model and human conceptual organization must be operationalized within a sensible metric.

SUSTAIN

In this section, the Supervised and Unsupervised STratified Adaptive Incremental Network (SUSTAIN) model is introduced. SUSTAIN has accounted for an array of challenging data sets spanning a variety of category-learning paradigms, including classification learning (Love & Medin, 1998b), learning at different levels of abstraction (Love & Medin, 1998a), inference learning (Love, Markman, & Yamauchi, 2000), development trends in learning (Gureckis & Love, 2004), and unsupervised learning (Gureckis & Love, 2002, 2003). SUSTAIN's formal description will not be discussed here. Instead, the focus will be on SUSTAIN's general operation and underlying principles. Those interested in knowing more about the equations that actualize SUSTAIN's principles can read Love, Medin, and Gureckis (in press).

SUSTAIN represents categories by one or more clusters. Clusters can be seen as category subtypes or covert categories (see Berlin, 1974; Brown, 1974). As shown in the right panel of Figure 13.1, SUSTAIN's clusters mediate the relationship between inputs (e.g., stimulus presentation) and output (e.g., category assignment). SUSTAIN begins with one cluster centered on the first training item. Additional clusters are recruited in response to surprising events. In unsupervised learning, a surprising event is exposure to a sufficiently unfamiliar or novel stimulus. In supervised learning, a surprising event is a classification error (e.g., incorrectly predicting a bat is a bird). When a surprising event does not occur, the current stimulus is assigned to the dominant cluster (i.e., the cluster most activated or similar to the current item), and this dominant cluster moves toward the current stimulus so that the cluster converges to the centroid or prototype of its members.

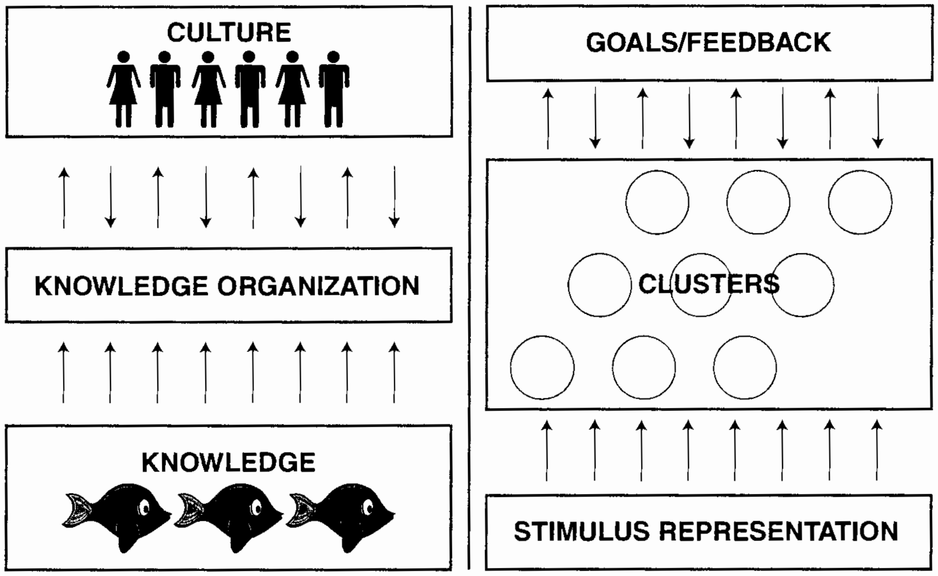

Figure 13.1. SUSTAIN's mapping to constructs from the Medin et al. (in press) studies is shown. An expert's knowledge of fish corresponds to SUSTAIN's stimulus representations. An expert's knowledge organization corresponds to SUSTAIN's clustering of the stimuli. The effect of cultural orientation on an expert's knowledge organization corresponds to SUSTAIN's training signal (i.e, its goals and feedback).

The Key Principles of SUSTAIN

With this general understanding of the operation of the model in mind, we now examine the six key principles of SUSTAIN.

PRINCIPLE 1: SUSTAIN IS DIRECTED TOWARD SIMPLE SOLUTIONS. At the start of learning, SUSTAIN has only one cluster that is centered on the first input item. It then adds clusters (i.e., complexity) only as needed to accurately capture the category structure of the learning task. Its selective attention mechanism further serves to bias SUSTAIN toward simple solutions by focusing the model on the stimulus dimensions that provide consistent information.

PRINCIPLE 2: SIMILAR STIMULUS ITEMS TEND TO CLUSTER TOGETHER. SUSTAIN clusters similar items together. For example, different instances of a bird subtype (e.g., sparrows) could cluster together and form a sparrow cluster instead of leaving separate traces in memory for each instance. Clustering is an unsupervised process because cluster assignment is done on the basis of similarity, not feedback. SUSTAIN's clusters are shown in Figure 13.1.

PRINCIPLE 3: SUSTAIN RELIES ON BOTH UNSUPERVISED AND SUPERVISED LEARNING PROCESSES. As discussed above, SUSTAIN can cluster on the basis of similarity (an unsupervised process). SUSTAIN's operation is also affected by supervision when available. Consider the example of SUSTAIN learning to classify stimuli as members of the category *mammals or birds*. Let's assume that a cluster representing four-legged land creatures has already been acquired by the model, as well as another cluster representing small, winged creatures that fly. The first time SUSTAIN is asked to classify a bat, the model will predict that a bat is a bird because the bat stimulus will be more similar to the existing bird cluster than to the existing mammal cluster. After receiving corrective feedback (supervision), SUSTAIN will note its error and create a new cluster to store the anomalous bat stimulus. Now, when this bat or one similar to it is presented to SUSTAIN, it will correctly predict that the bat is a mammal. This example also illustrates how SUSTAIN can entertain more complex solutions when necessary through cluster recruitment (see Principle 1).

Principle 4: Clusters Are Recruited in Response to Surprising Events. As the previous example illustrates, surprising events lead to new clusters being recruited. In unsupervised learning, a surprising event is simply exposure to a stimulus that is not sufficiently similar to any existing cluster (i.e., a very novel stimulus).

Principle 5: The Pattern of Feedback Matters. As the bird–mammal example in the preceding section illustrates, feedback affects the inferred category structure. Prediction failures result in a cluster being recruited; thus different patterns of feedback can lead to different representations being acquired. This principle allows SUSTAIN to predict different acquisition patterns for different learning modes (e.g., inference vs. classification learning) that are informationally equivalent but differ in their pattern of feedback. Likewise, the order of item presentation in unsupervised learning can affect how items cluster together.

Principle 6: Clusters Compete. Clusters can be seen as competing explanations of the input. The strength of the response from the winning cluster (the cluster to which the current stimulus is most similar) is attenuated in the presence of other clusters that are somewhat similar to the current stimulus (compare with Sloman's (1997) account of competing explanations in reasoning).

Summary of SUSTAIN

SUSTAIN represents categories by one or more clusters. SUSTAIN starts simple and adds clusters in response to surprising events. The cluster solutions that SUSTAIN uncovers are driven by both the structure of the environment and the learner's goals (i.e., what task the learner is trying to master). This property will prove critical in enabling SUSTAIN to account for different conceptualizations of biological kinds by members of different cultures.

Cross-Cultural Studies of Wisconsin Fish Experts

We now describe the general findings to which SUSTAIN was applied. The data are drawn from Medin et al. (in press). Medin et al. collected data from two groups of expert fishermen, majority-culture sports fishermen and Native Americans of the Menominee tribe living on a reservation in Wisconsin. The reservation contains numerous lakes, streams, rivers, and ponds. Cultural beliefs and tribal practices within the Menominee community emphasize a nature-centered biology of fish. The majority-culture fishermen live in communities proximate to the Menominee reservation. Outdoor recreation, including fishing clubs, play an important role in the lives of the majority-culture members. Many majority-culture fishermen participate in fishing contests with considerable rewards. Fishing for food is a relatively more important activity for the Menominee, whereas fishing for sport is relatively more important for majority-culture fishermen.

One possibility is that the cultural orientations of the two groups play a role in shaping conceptual organization of fish species. Medin and his colleagues' (in press) results support this possibility. Majority-culture and Menominee fish experts sorted cards depicting 44 native fish species as they saw fit and then justified why they constructed the groupings as they did. These sorts and justifications indicated that the species are primary organized along fishing-related goals (e.g., a bait fish, a prestigious catch) for majority-culture members, whereas the Menominee data suggested that the Menominee's conceptual organization unfolds more along ecological lines (e.g., the habitat of the fish). In general, the Menominee were less extreme in their sorts. Although ecology played an important role in their sorts, fishing-related goals were also evident to a moderate extent. This moderation was not seen with majority-culture fishermen for whom fishing-related goals predominated and ecological explanations were rare. Another feature of the data was that there was less within-group agreement for the Menominee than for the majority culture. Table 13.1 summarizes the basic findings from Medin et al.

Differences in conceptualization between the cultures are evident. It is important to note that nothing in the data indicates higher expertise in one group as opposed to the other. Both groups have extensive firsthand experience with native fish species. When tested without time pressure, both groups display the same knowledge of the 44 species (e.g., identification, habitat, fish interactions). Whereas previous work has shown that experts who have different kinds of expertise in a common domain can show differences in free sorting (Medin, Lynch, Coley, & Atran, 1997; Proffitt, Coley, & Medin, 2000), the interactions of the Menominee and majority-culture fishermen with local fish species are highly similar. One explanation is that cultural differences between the two groups lead to differences in knowledge organization.

Table 13.1. The Basic Findings From Medin et al. (in press)

	Majority	Menominee
Ecological organization	Low	High
Prestige (fishing) organization	Very high	Medium
Within-culture agreement	High	Low

Simulations

SUSTAIN captures conceptual organization in its clusters. For SUSTAIN to generate the pattern of findings in Medin et al. (1997), SUSTAIN simulations of the Menominee must result in qualitatively different clusterings than do simulations of the majority culture. In particular, the cluster structure for Menominee simulations should be organized along both fishing-related and ecological lines, with an emphasis on ecology, whereas the majority simulations should yield clustering dominated by fishing-related goals. Also, agreement of cluster solutions across simulations should be greater for majority simulations than for Menominee simulations.

In the introductory section of this chapter, basic challenges in modeling the effects of culture on conceptual organization were discussed. Here, these three basic challenges, knowledge representation, task formalization, and simulation evaluation, are considered in the context of Medin et al. (in press).

Knowledge Representation

Following Medin et al. (in press), we assumed both groups had the same knowledge of the fish. The inputs to SUSTAIN represented this knowledge, whereas clustering represented knowledge organization. Thus the input representation for both groups was identical. We wanted as objective and neutral a representation as possible of the fish that captured or correlated with aspects of the experts' knowledge. The chosen solution was to create a spatial representation of the taxonomic relations among the 44 indigenous species considered in the Medin et al. (in press) study.

A 44×44 matrix was constructed that represented taxonomic distance. Each entry in the matrix represented the distance in the scientific taxonomy between two species. For example, two species sharing the same genus had a distance of 1, whereas two species that matched at the family level had a distance of 2. This dissimilarity matrix was subjected to Sammon's (1969) nonmetric multidimensional scaling. Five clear dimensions emerged. The values along these five dimensions were used to represent each fish. Following SUSTAIN's standard operation, the clusters were also located in this same five dimensional space.

Task Formalization

Medin et al. hypothesized that differences between the two groups of experts were driven by "habits of mind" or characteristic ways of thinking that were culturally prescribed. One simple hypothesis is that cultural pressures for majority-culture fishermen lead to an organization of fish species around fishing-related goals, whereas the Menominee's culture emphasizes ecological properties.

To model this influence on knowledge organization, majority-culture simulations were trained to predict the prestige of fish species (e.g., a northern pike is a prized catch, an American eel is not), whereas Menominee simulations were trained to predict the habitat of the species. According to SUSTAIN's fifth principle, the pattern of feedback guides cluster organization. In the present case, the pattern of feedback stresses ecological organization for the Menominee and fishing-related organization for the majority fishermen.

To implement this scheme, each of the 44 fish was scored for habitat and prestige. For ecology, fish were scored along the following scale: 0 indicates that the fish lives exclusively in lakes, 0.5 indicates that the fish lives in both rivers and lakes, and 1 indicates that the fish lives exclusively in rivers. For example, a bluegill was scored as a 0, a walleye as a 0.5, and a brown trout as a 1. The same three point scale was used for low-, medium-, and high-prestige fish.

SUSTAIN was trained on these two category systems using parameters from Love et al. (in press). Fish were assigned to the cluster to which they were most similar. If feedback indicated an incorrect cluster assignment, a new cluster was recruited to encode the fish. An assignment was considered incorrect when the cluster's position on the critical dimension (i.e., habitat for Menominee simulations, prestige for majority-culture simulations) was more than 0.501 units away from the current fish's value on this dimension.

Fish were presented to SUSTAIN in blocks. A block is the presentation of the whole set of fish in a random order. Each simulation consisted of 50 blocks of training. Because each simulation represents an individual member of the culture, the average results were pooled over thousands of simulations to obtain an accurate estimation of the group mean. Following training, clusters not responding to any of the 44 fish were removed. Through this training regimen, the final cluster organization reflected both the goals of the learner (either predicting ecology or prestige) and the correlational structure of the stimulus set.

One interesting question is how supervised training affects conceptual organization. To offer a "cultureless" baseline, a third set of simulations were also included that involved unsupervised learning. In these simulations, SUSTAIN was not trained on habitat or prestige. Instead, new clusters were formed when no existing cluster was sufficiently similar to the current stimulus. The threshold for recruiting a new cluster was set to 0.31 to roughly equate the number of clusters recruited with the Menominee and majority-culture simulations. These unsupervised simulations can be thought of as SUSTAIN's attempt to carve nature at its joints (Hempel, 1965; Rosch,1978).

Simulation Evaluation

The most challenging aspect of this project was developing a method for assessing the agreement between our modeling results and the results from Medin et al. (in press). In this section, we detail methods for deciding whether SUSTAIN simulations follow the same pattern as the populations studied by Medin et al.

EVALUATING ECOLOGICAL AND FISHING-RELATED ORGANIZATION. The result of each SUSTAIN simulation is represented by a matrix F that details which fish share a cluster. We define the matrix F such that the entry f_{ij} is 1 if fish i and fish j are in the same cluster at the conclusion of training and 0 otherwise. Only entries from cells in the upper diagonal (i.e., the shaded region of F in Figure 13.2) are included in the analyses because these cells carry all the relevant information. Entries along the main diagonal are necessarily 1 (each fish is in the same cluster as itself), and the lower diagonal mirrors the upper diagonal (i.e., $f_{ij} = f_{ji}$).

Once the matrix F is constructed, it is converted to the vector $\vec{f}$ by concatenating the column entries forming the upper diagonal as shown in Figure 13.2. The vector $\vec{f}$ represents SUSTAIN's conceptual organization for a single simulation. Two other matrices and vectors must be constructed to evaluate this conceptualization. These matrices and vectors are used to evaluate the extent to which a SUSTAIN simulation is organized along ecological and fishing-related lines.

The first matrix, E, is used to evaluate ecological organization. Like F, this matrix consists of 44×44 entries. Each entry in the matrix indicates the ecological

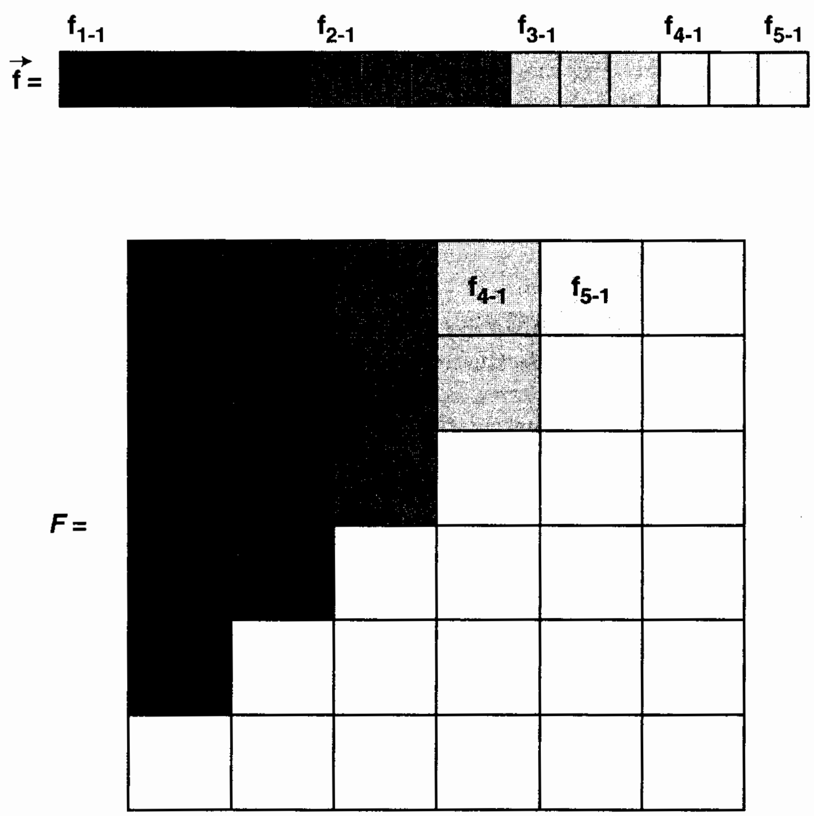

Figure 13.2. The matrix F details how fish cluster together. The matrix F used to evaluate the actual simulations was larger, with 44×44 cells. The vector $\vec{f}$ is constructed from the upper diagonal of the matrix F as shown in the figure.

compatibility (based on habitat) of two species. As in training, the three habitats were lake only, lake and river, and river only. Fish living in the same habitat had a compatibility of 1, fish overlapping in habitats (e.g., lake and river vs. river only) had a compatibility of 0, and fish in nonoverlapping habitats (lake only vs. river only) had a compatibility of –1. For example, the entry at the intersection of blue gill (a fish that only lives in lakes) and brown trout (a fish that only lives in rivers) is –1. This matrix E was converted into the vector $\vec{f}$ by the same method used to construct the vector $\vec{f}$ from the matrix F. One additional step was performed to create the vector $\vec{e}$—the vector $\vec{e}$ was normalized to have mean zero.

The same steps were used to build the matrix P, which was used to evaluate organization along fishing-related goals. The entries in matrix P were the fishing-goal compatibilities (based on prestige) of the two species. As in training, the three levels of prestige were *low, medium,* and *high.* As in the ecology evaluation, highly compatible fish had a value of 1 (e.g., low vs. low prestige), somewhat compatible fish had a value of 0 (e.g., high vs. medium prestige), and incompatible fish had a value of –1 (low vs. high prestige). The vector $\vec{p}$ was constructed from P as in the above cases and was normalized to have mean zero like vector $\vec{e}$.

With these three vectors, $\vec{f}$, $\vec{e}$, and $\vec{p}$, the conceptual organization of a SUSTAIN simulation can be evaluated. The organization of a simulation is determined by calculating the vector cosine of $\vec{f}$ with either $\vec{e}$ or $\vec{p}$ depending on whether organization is being evaluated along ecology or prestige. The vector cosine of two vectors $\vec{a}$ and $\vec{b}$ is calculated as follows:

$$\cos(q) = \frac{\vec{a}\vec{b}}{\overrightarrow{\|a\|}\,\overrightarrow{\|b\|}} \tag{1}$$

The vector cosine of two vectors measures their compatibility. Two vectors in the same direction (i.e., perfectly compatible) will have a vector cosine of 1, whereas two vectors in opposite direction (i.e., perfectly incompatible) will have a vector cosine of –1. Vectors that are orthogonal (i.e., unrelated or neutral) will have a vector cosine of 0. Figure 13.3 shows three pairs of vectors. The pair on the left are somewhat incompatible, the middle pair is neutral, and the pair on the right are somewhat compatible.

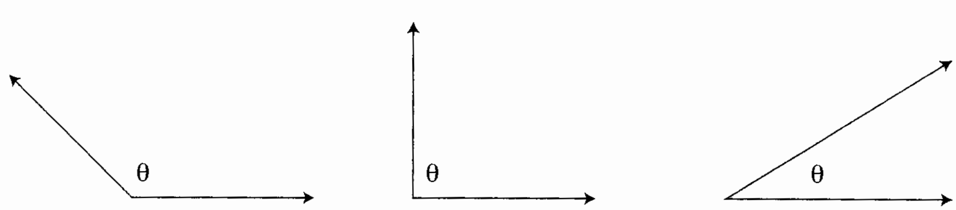

Figure 13.3. Incompatible (negative vector cosine), neutral (zero vector cosine), and compatible (positive vector cosine) vector pairs.

Thus for evaluation of the ecological organization of a SUSTAIN simulation, the vector cosine of $\vec{f}$ and $\vec{e}$ is calculated, whereas fishing-goal–related organization is evaluated by calculating the vector cosine of $\vec{f}$ and $\vec{p}$. This scoring procedure has a number of advantages. One favorable property is that positive and negative values have a clear interpretation. Any positive value indicates consistent organization, whereas any negative value indicates inconsistent organization. Furthermore, the normalization of $\vec{e}$ and $\vec{p}$ ensures that simulations that store all fish in a common cluster and simulations that store each fish in a separate cluster result in a vector cosine of zero. Also, random clustering of fish will on average result in a zero vector cosine.

EVALUATING CONSENSUS. To evaluate consensus, the $\vec{f}$ vectors from two different simulations are selected and percent agreement is calculated. Corresponding entries (i.e., 0 and 0 or 1 and 1) count as an agreement, whereas mismatches count as a disagreement. For example, if bluegill and brown trout shared a cluster for the first simulation but did not for the second simulation, this would result in a disagreement.

AVERAGING RESULTS. The presentation order of the fish in training was randomized for each simulation. Different presentation orders can result in different clusterings. Both an evaluation of conceptual organization and agreement involve averaging over 5,000 simulations to ensure that the reported means are stable. Averaging over 5,000 simulations results in estimates of the mean stable enough that error bars are not needed (i.e., the results will replicate at the level of precision reported).

Results and Discussion

The pattern of results from the SUSTAIN simulations corresponded to that of the Medin et al. (in press) studies (see Table 13.1 for a summary of their findings). The evaluation of SUSTAIN's cluster organization for the Menominee and majority-culture simulations is shown in Figure 13.4. As predicted, the majority-culture simulations had a cluster organization highly compatible with fish-related goals that was not highly aligned with ecology. The Menominee simulations were more moderate and had a strong ecological organization that was also somewhat aligned with fishing related goals.

The differences between these two groups of simulations were driven by their different training regimens. That cluster organization is most consistent with prestige when simulations are trained to predict prestige is not surprising. However, it is somewhat unexpected that correlational structure of the stimuli and feedback leads to a cluster organization consistent with multiple goals in simulations trained to predict ecology but not in simulations trained to predict prestige. As Medin et al. (in press) found with the Menominee, Menominee simulations satisfy multiple goals.

Two SUSTAIN simulations for the same population can differ in their clustering solutions because of differences in the order in which the items are presented. Analysis of "individual" differences data for SUSTAIN simulations

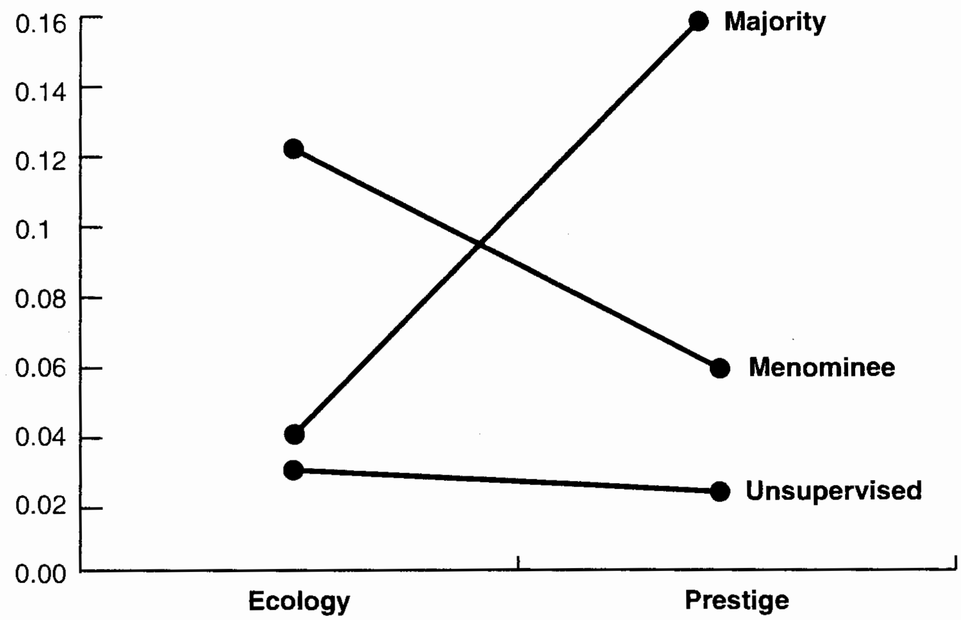

Figure 13.4. Mean measures for ecological and fishing-related organization are shown for the three sets of simulations.

showed distinct patterns of variations for the two groups. In majority-culture simulations, simulations high on one measure tended to be high on the other ($r = +.43$), whereas the correlation was weaker and negative in Menominee simulations ($r = -.23$). Medin et al. (in press) was kind enough to examine the human data to see whether this pattern held for human experts. Consistent with the SUSTAIN simulations, the number of fishing-related and ecological sorting justifications provided by majority-culture fishermen correlated positive ($r = +.56$), whereas the correlation was negative and weak ($r = -.08$) for the Menominee.

Both the Menominee and majority-culture simulations involved supervised learning. This supervision or goal-directedness played a strong role in shaping conceptual organization. One interesting comparison with these simulations is the unsupervised simulations in which clustering was guided by similarity. In these simulations, the structure of the environment (i.e., the distribution of fish features) was the primary factor in guiding cluster formation. As can be seen in Figure 13.4, the unsupervised simulations resulted in weakly consistent organizations for both measures.

Table 13.2 provides additional information about how many clusters were recruited in each simulation to represent the 44 fish. Although this was not predicted, majority-culture simulations resulted in more clusters than Menominee simulations (21.3 vs. 17.4). This difference mirrors the sorting results of Medin et al. (in press). Majority-culture fishermen created 9.5 piles in which to sort the 44 fish on average, compared with 8.5 piles for the Menominee fishermen.

Table 13.2. The Mean and Modal Number of Clusters Recruited for Each Set of Simulations

Condition	Mean	Mode
Majority culture	21.3	22
Menominee	17.4	17
Unsupervised	20.1	17

One additional prediction was that consensus should be less for the Menominee than for the majority-culture simulations. This prediction held. Agreement was 92% and 90% for the Menominee and majority-culture simulations, respectively. The agreement is very high in both cases because of the sparseness of the F matrix (i.e., most fish do not cluster with one another). To put these numbers in perspective, we calculated the agreement between the Menominee and majority-culture simulations. The agreement level was 90%. One factor that led to higher agreement for the majority-culture simulations was that these simulations resulted in more clusters being recruited on average (see Table 13.2). If each fish was assigned to its own cluster, agreement would be 100%. Another explanation is that training on prestige more highly constrains cluster organization than does training on ecology. Support for this view is the extreme values for both measures for the majority-culture simulations and the moderate values for the Menominee simulations.

Future Directions

Although the present work does make an important bridge between work inside and outside the laboratory, more can be done to narrow the gap. In this section, we consider possible ways to make models more effective tools for understanding and exploring the role that culture plays in influencing conceptual organization. This section is organized by the basic modeling challenges previously discussed. The greatest focus is on the first challenge, knowledge representation. In terms of the three levels illustrated in Figure 13.1, knowledge representation involves the bottom level. Readers not interested in our proposed solution for this level may skip this lengthy section.

Knowledge Representation

Representing experts' knowledge of fish for the simulations by scaling the taxonomic distance matrix was expedient and fairly effective. However, expert knowledge in general is more nuanced. In this section, we consider methods for improving the quality of expert knowledge representations without resorting to hand-coded representations. Because most current learning models are bound to process spatial representations, we will confine the discussion to spatial representations. Of course, a complete account of representation would also need to take into account symbolic relations (Gentner, 1983).

Ideally, stimulus representations would be derived from measures taken from the experts, as opposed to surrogate measures such as taxonomic distance. Also, measures should be open ended. One very open-ended measure is to have the expert speak freely about the subject matter of interest. For example, an expert could verbalize everything he or she knows about brown trout. In this section, we will discuss a method that takes such dialogues and automatically translates them into dimensional representations over which learning models can operate. Some initial results will be discussed.

The method, Corpus Scaling Analysis (CSA), works by taking a distance matrix generated by transforming dialogues collected from human experts (replacing the taxonomic distance matrix used in the previous simulations) and submitting this distance matrix to multidimensional scaling (MDS) (see Burgess & Conley, 1998, for a similar approach). A corpus approach is used to create spatial representations of the experts' dialogues. In terms of the schematic shown in Figure 13.1, CSA translates from the knowledge box in the left panel to the stimulus representation box shown in the right panel.

Corpus approaches construct term representations by exploiting information contained in word co-occurrence patterns (Burgess & Lund, 1997; Landauer & Dumais, 1997). According to corpus approaches, the meaning of a term is determined by the contexts in which the term appears. For example, the concepts denoted by the terms *circle* and *square* are somewhat similar to each other because *circle* and *square* appear in similar contexts, such as "The child drew a red circle" and "The child drew a red square." Of course, *circle* and *square* do not have the same exact meaning because they can appear in different contexts, such as "You are such a square" or "Circle the runway, and then land." It is unclear whether this contextual information is driving the acquisition of word meaning in humans or is simply an observable correlate of the underlying semantics. However, work in semantic priming suggests that feature correlation patterns (e.g., has wings, can fly, has feathers) play a role in semantic organization (McRae, de Sa, & Seidenberg, 1997). Nevertheless, whether these effects are being driven by extensional or intentional forces is not entirely clear. What is clear is that word usage patterns are a rich source of information.

Here, Landauer and Dumais's (1997) Latent Semantic Analysis (LSA) method is adopted because of its popularity, broad application, and use of context in determining word meaning. The similarity of two terms in LSA is calculated by taking the vector cosine of the vectors representing the terms. For all the analyses presented here, the LSA group's default corpus was used. The Touchstone Applied Science Associates, Inc. corpus consists of readings in academic subjects for students from third grade through the first year of college. The corpus consists of 37,651 concatenated texts containing 92,409 unique terms. The key parameter in LSA is the number of factors included in the singular valve decomposition reconstruction. The LSA group's recommended solution involving 300 factors was adopted. Thus all the defaults for LSA were used, and no attempt was made to optimize performance for the particular analyses considered here.

Of course, actual dialogues from experts contain numerous words. One simple method for representing dialogues or documents is simply to sum the vectors for the document's constituent words. This simple linear operation was used to construct dialogue representations in the analyses presented here. With these representations of the dialogues, the similarity of two dialogues is their vector cosine. These vector cosines are then arranged in a similarity matrix and are subjected to MDS for the final solution. For example, in the Medin et al. (in press) study discussed in a previous section, the similarity matrix would contain 44×44 entries because 44 fish were included in the study.

Most learning models operate over a fixed input space or dimensionality. Therefore if a new item is to be included as a transfer item (e.g., a new fish species is introduced to a fish expert), then a method is needed to map it into the space CSA yields. To accomplish this goal, the pairwise similarities to the original set of documents–dialogues are calculated and these similarities are transformed to distances. The position of the novel document in the final space is determined by minimizing a stress measure. The following function is minimized in order to determine the position of the novel document in the final space:

$$\sum_{i=1}^{n} (f(x_i) - d_i \tag{2}$$

where n is the number of original documents, x_i is the euclidean distance in final space between the new document's position and original document i, d_i is the dissimilarity (i.e., transformed similarity) between the new document and original document i, and f is a transformation function. In the case of metrical MDS, f is simply a linear transformation used for scaling purposes. Once a novel document is mapped into the final space, the learning model can make predictions about the document and train on the document if corrective feedback is available. In terms of the SUSTAIN simulations discussed previously, the final space corresponds to the Stimulus Representation box shown in Figure 13.1.

Demonstrations of CSA. Before CSA can be applied to expert populations, its properties and performance must be better understood. This section takes an initial step toward evaluating CSA. Evaluating whether CSA has correctly identified the appropriate set of dimensions assumes the existence of a correct answer that can be known. In light of this verification issue, CSA was applied to a stimulus set with known and agreed-upon dimensions. Success on such a task is a prerequisite for exploring more abstract domains, such as experts' representation of biological kinds. The stimulus set considered consists of eight items that are described by three dimensions: color (blue or red), shape (square or triangle), and size (small or large). This stimulus set is frequently used in category-learning research (Medin & Schaffer, 1978; Nosofsky et al., 1994; Yamauchi, Love, & Markman, 2002).

The key question is whether CSA can transition from a term space consisting of 92,409 dimensions to a final space consisting of the appropriate three dimensions. To answer this question, CSA was applied to short descriptions of the eight stimuli (e.g., "small, red, triangle"). Each document consisted of three words describing the stimulus's features.

The results of applying CSA to these eight documents are shown in Table 13.3 and illustrated in Figure 13.5. Three clear dimensions of color, shape, and size emerge. The eigenvalues for these three dimensions are 0.292, 0.080, and 0.065. The eigenvalue for the fourth dimension was less than 0.001, indicating a three-dimensional solution. Notice that the values of each dimension are clearly represented. For example, negative numbers on the first dimension

Table 13.3. CSA Recovers the Three Dimensions of Color, Shape, and Size

Text	Dim. 1	Dim. 2	Dim. 3
Blue, square, small	–1.91	–1.21	–0.63
Blue, square, large	–1.81	–0.75	1.01
Blue, triangle, small	–1.97	0.93	–1.23
Blue, triangle, large	–1.86	1.36	0.82
Red, square, small	1.52	–1.15	–0.51
Red, square, large	1.49	–0.78	0.89
Red, triangle, small	2.30	0.58	–1.13
Red, triangle, large	2.25	1.01	0.77

Note. All coordinates are multiplied by 10.

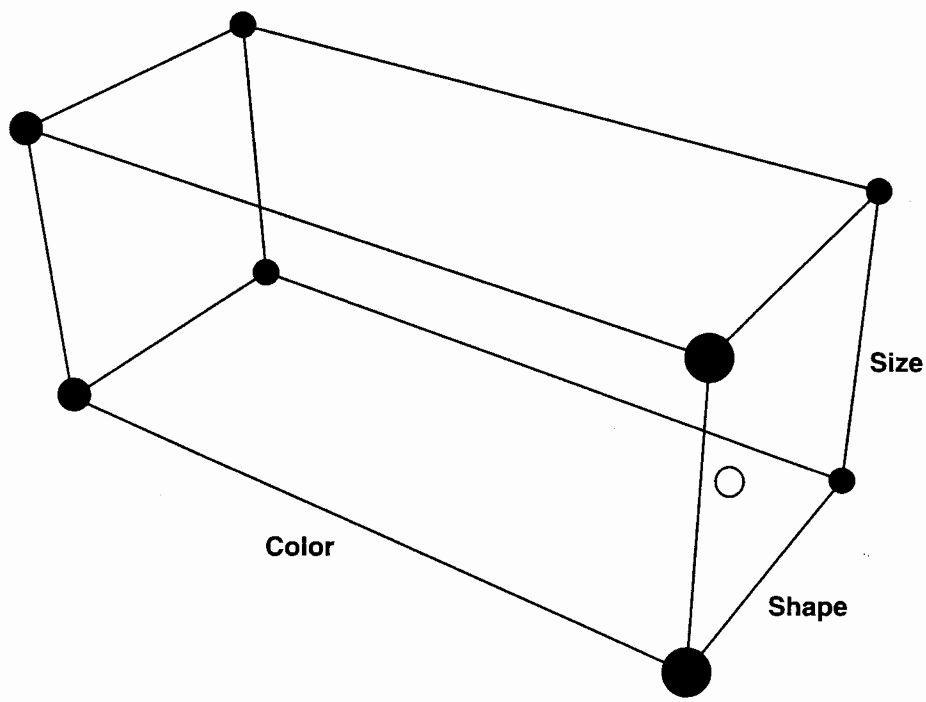

Figure 13.5. The eight coordinates from Table 13.3 are shown as circles connected by line segments. The two largest circles are nearest to the viewer in three-dimensional space. The positions of the eight points approximate a rectangular box. The hollow, small circle is the result of mapping "red, square, medium" into the space.

indicate a blue stimulus, whereas positive values indicate a red stimulus. The values along dimensions clearly cluster. For example, the blue stimuli are all closer to one another on the first dimension than they are to any red stimulus.

One question is whether CSA can successfully map new documents into its space. To demonstrate this ability, the novel document "red, square, medium" was mapped into the space shown in Table 13.3 according to Equation 1. One would expect the result to be close to "red, square, small" at position (1.52, –1.15, –0.51) and "red, square, large" at position (1.49, –0.78, 0.89), splitting the difference on the third dimension, which represents size. Indeed, the position (see Figure 13.5) of "red, square, medium" maps to (1.62, –0.84, 0.02).

One interesting question is whether the the proper three-dimensional space could be recovered by including only LSA's three most salient dimensions (300 dimensions were used in the previous demonstration). When CSA is run with LSA selecting out three dimensions, the scaling result does not produce the three correct dimensions. The first dimension loosely corresponds to size (although the values are not clearly separated), and the other two dimensions are uninterpretable. LSA cannot "know" ahead of time which three dimensions are relevant to the current problem.

Although the previous results support CSA, the demonstrations did not involve descriptions generated by human subjects, which are likely to be more irregular than the experimenter-generated documents. Twenty-five human subjects were shown the stimuli (as geometric figures, not text) sequentially on a computer screen and provided text descriptions on a worksheet. These worksheets were transcribed verbatim. All descriptions for a given stimulus were concatenated into a document, and the resulting eight documents were fed to CSA. The recovered space mirrors the previous result and is shown in Table 13.4. Only the first three eigenvalues were nonzero (0.027, 0.019, and 0.016).

The emergence of three clear dimensions corresponding to color, shape, and size is impressive given the nature of the data. The subjects' task was unconstrained, and many of their descriptions amounted to noise. For instance, some descriptions referred to the previously displayed stimulus (e.g., "The same color as the previous shape"), which is noise in the analysis given that the stimulus presentation order was randomized for each subject. The size descriptions tended to show a lot of variation (e.g., "the size of four sugar cubes," "1/2 cm," "one half cms," "1 inch," "1.2 inches," along with spelling errors). Somewhere in this morass, CSA found the signal.

Much more work must be done to determine whether CSA is a feasible method for transforming experts' verbal descriptions into knowledge representations. The initial results presented here are encouraging. CSA was able to recover the three relevant physical dimensions with little human intervention.

Task Formalization

A second challenge in modeling expert performance is task formalization. Our task formalization of the Medin et al. (in press) studies was fairly simplistic. SUSTAIN was trained to predict habitat for the Menominee simulations and

Table 13.4. CSA Recovers the Three Dimensions of Color, Shape, and Size

Text	Dim. 1	Dim. 2	Dim. 3
Blue, square, small	−7.10	−5.65	−2.35
Blue, square, large	−4.13	−3.85	6.77
Blue, triangle, small	−7.18	4.99	−3.84
Blue, triangle, large	−3.60	5.72	4.94
Red, square, small	3.22	−4.49	−5.41
Red, square, large	7.33	−5.26	3.02
Red, triangle, small	3.76	3.11	−5.31
Red, triangle, large	7.69	5.43	2.18

Note. All coordinates are multiplied by 100.

prestige for the majority-culture simulations. This difference in training was intended to capture the different orientations of the two cultures and to explain how differences in conceptual organization arise. Ideally, future work would bear a closer resemblance to the learning situations that humans actually face. More detailed information about the type and frequency of activities from members of the two cultures across development would also allow for predictions to be made about how one becomes an expert.

Simulation Evaluation

For the Medin et al. (in press) studies, the simulation evaluation was effective in bridging the simulations and the results on conceptual organization. Of course, new sources of data will require new methods for evaluating models. For example, expert data from a picture-naming task (e.g., Tanaka & Taylor, 1991) would require new model evaluation methods to link human and model performance.

Conclusions

In this chapter, a model developed from consideration of data from the laboratory was applied to understanding cross-cultural data. SUSTAIN offers a mechanistic explanation of how the goals emphasized by a culture can influence conceptual organization. SUSTAIN is simultaneously sensitive to the structure of the environment and the goals of the learner. These two drives allowed SUSTAIN to capture the commonalities and differences between the Menominee and majority-culture fishermen. Although both groups have similar experiences and interactions with biological kinds, their respective cultures appear to influence how they organize this information.

The primary challenges faced in applying SUSTAIN to the Medin et al. (in press) data were specifying knowledge representation, task formalization, and simulation evaluation. The choices we made proved successful and did not involve free parameters. Nevertheless, more progress can be made on all three fronts. Although the relationship between SUSTAIN simulations and categorization within the two cultures is specified, the mapping is somewhat coarse. A number of facets of real-world learning are not present in the current simulations. As more constraining data become available, SUSTAIN can be further specified.

We would be remiss if we did not acknowledge the role Doug Medin played in this work. He is no stranger to developing psychological models (e.g., Medin & Schaffer, 1978) and has been involved in the development of SUSTAIN. It was Medin's suggestion to apply SUSTAIN to the data considered here. As always, his suggestions were very helpful in framing the project.

Finally, we would like to suggest that studies in the field and in the laboratory are not as unrelated as some might believe. The same principles allow SUSTAIN to account for data (with the same parameter settings) from both communities. Theories generated in the laboratory may be more relevant to those working outside the laboratory than is generally assumed. Similarly, work outside the laboratory offers exciting challenges and perspectives on categorization for researchers working in the laboratory.

References

Berlin, B. (1974). Further notes on covert categories and folk taxonomies. *American Anthropologist, 76,* 327–331.

Brown, C. H. (1974). Unique beginners and covert categories and folk biological taxonomies. *American Anthropologist, 76,* 325–327.

Burgess, C., & Conley, P. (1998). Developing semantic representations for proper names. In *Proceedings of the Annual Conference of the Cognitive Science Society* (pp. 185–190). Hillsdale, NJ: Erlbaum.

Burgess, C., & Lund, K. (1997). Modeling parsing constraints with high-dimensional context space. *Language and Cognitive Processes, 12,* 177–210.

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

Gureckis, T., & Love, B. C. (2002). Who says models can only do what you tell them? Unsupervised category learning data, fits, and predictions. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 399–404). Hillsdale, NJ: Erlbaum.

Gureckis, T., & Love, B. C. (2003). Towards a unified account of supervised and unsupervised learning. *Journal of Experimental and Theoretical Artificial Intelligence, 15,* 1–14.

Gureckis, T., & Love, B. C. (2004). Common mechanisms in infant and adult category learning. *Infancy, 5,* 173–198.

Hempel, C. G. (1965). *Aspects of scientific explanation.* New York: The Free Press.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104,* 211–240.

Love, B. C., & Markman, A. B. (2003). The nonindependence of stimulus properties in human category learning. *Memory and Cognition, 31,* 790–799.

Love, B. C., Markman, A. B., & Yamauchi, T. (2000). Modeling classification and inference learning. *Proceedings of the Seventeenth National Conference on Artificial Intelligence* (pp. 136–141). Menlo Park, CA: AAAI Press.

Love, B. C., & Medin, D. L. (1998a). Modeling item and category learning. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 639–644). Mahwah, NJ: Erlbaum.

Love, B. C., & Medin, D. L. (1998b). SUSTAIN: A model of human category learning. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence* (pp. 671–676). Cambridge, MA: MIT Press.

Love, B. C., Medin, D. L., & Gureckis, T. (2004). SUSTAIN: A network model of human category learning. *Psychological Review, 111,* 309–322.

Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory and Cognition, 28,* 41–50.

McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General, 126,* 99–130.

Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology, 32,* 49–96.

Medin, D. L., Ross, N., Atran, S., Cox, D., Waukau, H. J., Coley, K., et al. (in press). The role of culture in the folkbiology of freshwater fish. *Cognition.*

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

Nosofsky, R. M., Gluck, M. A., Palmeri, T. J., McKinley, S. C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory and Cognition, 22,* 352–369.

Proffitt, J. B., Coley, J. D., & Medin, D. L. (2000). Expertise and category-based induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 811–828.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.

Sammon, J. W. (1969). A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers, C-18,* 401–409.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs, 75*(13, Whole No. 517).

Sloman, S. A. (1997). Explanatory coherence and the induction of properties. *Thinking and Reasoning, 3,* 81–110.

Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology, 23,* 457–482.

Yamauchi, T., Love, B. C., & Markman, A. B. (2002). Learning nonlinearly separable categories by inference and classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 585–593.