Categorization by Elimination: A Fast and Frugal Approach to Categorization

Patricia M. Berretty (berretty@mpipf-muenchen.mpg.de)
Peter M. Todd (ptodd@mpipf-muenchen.mpg.de)
Philip W. Blythe (blythe@mpipf-muenchen.mpg.de)
Center for Adaptive Behavior and Cognition
Max Planck Institute for Psychological Research
Leopoldstrasse 24, 80802 Munich Germany

Abstract

People and other animals are very adept at categorizing stimuli even when many features cannot be perceived. Many psychological models of categorization, on the other hand, assume that an entire set of features is known. We present a new model of categorization, called Categorization by Elimination, that uses as few features as possible to make an accurate category assignment. This algorithm demonstrates that it is possible to have a categorization process that is fast and frugal--using fewer features than other categorization methods--yet still highly accurate in its judgments. We show that Categorization by Elimination does as well as human subjects on a multi-feature categorization task, judging intention from animate motion, and that it does as well as other categorization algorithms on data sets from machine learning. Specific predictions of the Categorization by Elimination algorithm, such as the order of cue use during categorization and the time-course of these decisions, still need to be tested against human performance.

1. Introduction

Hiking through the Bavarian Alps, you come upon a large bird gliding over a meadow. You pull out your European bird guidebook to identify it. From the shape of its body, you assume that this is a bird of prey, so you turn to the section on raptors in the guide. To determine the exact species, you next use size to narrow down your search to a few kinds of hawks; then you use color to eliminate a couple more species; and finally with one last cue--tail length--you can make a unique classification. Using only four cues (or features), you correctly identify this bird as a sparrow hawk. You could take out your binoculars and check more cues to support this identification, but for a rapid decision these few cues are enough.

How would this categorization process proceed if a rabbit rather than a human were watching the bird? The rabbit would not be interested in knowing the exact species of bird flying overhead, but rather would want to categorize it as predator or not, as quickly as possible--the Rabbit's Guide to Birds has only two short sections. While the rabbit could also use several cues to make its category assignment, as soon as it finds enough cues to decide "predator"--for instance, that this bird is gliding--it will not bother gathering any more information, and instead will head for shelter. Obviously, in the case of the rabbit, speed is of the essence when categorizing birds as predators or nonpredators. Humans face similar circumstances where rapid categorization is called for, making use of only whatever information is immediately available. Being able to categorize rapidly the intention of another approaching person as either hostile or courting, for instance, will enable the proper reactions to ensure the most desirable outcome.

In this paper, we consider the case for a "fast and frugal" (à la Gigerenzer & Goldstein, 1996) model of categorization, akin to the lexicographic process of bird identification described in the first paragraph. This model, which we call *Categorization by Elimination* (CBE), uses only as many of the available cues or features as are necessary to first make a specific categorization. As a consequence, it often uses far fewer cues in categorizing a given stimulus than do the standard cue-combination models, yielding its fast frugality. This information-processing advantage can be crucial in a variety of categorization contexts where speed is called for, as in identifying threats. On the other hand, the accuracy of this approach typically rivals that of more computationally extensive algorithms, as we will show. We therefore propose Categorization by Elimination as a parsimonious psychological model, as well as a potentially useful candidate for applied machine-learning categorization tasks.

Categorization by Elimination is closely related to Tversky's Elimination by Aspects (EBA) model of choice (Tversky, 1972). After describing competing psychological and machine-learning models of categorization in the next section, we discuss the background of elimination models in section 3. We present the Categorization by Elimination model in section 4. Most other recent models of human categorization focus on the use of two or three cues, situations in which CBE can show little advantage. Therefore, we have experimentally investigated a multiple-cue categorization task in which we can compare our model with others in accounting for human performance with seven cue dimensions. We describe this study, which involves categorizing animate motion trajectories into different behavioral intentions, in section 5. CBE does as well as linear categorization methods, and does not overfit the data as neural networks seem to. Next, in section 6 we look at how well our algorithm does alongside some of the multiple-attribute categorization methods developed in psychology and machine learning on standard data sets from the latter field. This comparison shows that Categorization by Elimination can often compete in accuracy with more complex methods. Further, if minimizing the number of cues used is sought to maximize computational speed, CBE usually emerges as the clear winner. Finally in section 7 we consider some of the challenges still ahead, including how to test CBE against human learning data.

2. Existing Categorization Models

Many different models of categorization have been proposed in both the psychological and machine learning literature. Psychologists are primarily concerned with developing a model that best describes human categorization performance, while in machine learning the goal is to develop an optimally-performing model--that is, one with the highest accuracy of categorization. These two goals are not necessarily mutually exclusive; indeed, one of the main findings so far in the field of human categorization is that people are often able to achieve near optimal performance (that is, categorize a stimulus set with minimal errors--see Ashby & Maddox, 1992). As a consequence, some models, including neural networks and SCA (Miller & Laird, 1996) are often aimed at filling both roles.

However, the majority of psychological studies of categorization have used simple stimuli that vary on only a few (2-4) dimensions, unlike the typical high-dimensional machine learning applications. It remains to be seen whether humans can also be optimal at categorizing multi-dimensional objects. In addition, the predominant psychological models of categorization have not addressed the issue of constraints, such as limited time and information. What might the categorization process be when there are both time and information constraints, either because there is an overwhelming number of possible cues to use or only a subset of cues available? Here we briefly review some of the currently popular categorization models for human categorization and machine learning with these questions in mind. Throughout the remainder of the paper we use the terms *cues, aspects, dimensions*, and *features,* as appropriate, to all mean roughly the same thing.

The predominant theories of categorization in the psychology literature include exemplar models (Nosofsky, 1986), decision bound models (Ashby & Gott, 1988), and neural network models (e.g. ALCOVE--see Kruschke, 1992). Each of these categorization models assumes that the stimuli may be represented as points in a multidimensional space. Furthermore, these models all assume that humans integrate features--that is, combine multiple cues to come to a final judgment--and that we usually use all of the cues that are present--that is, do not discard any available information.

Exemplar models (Brooks, 1978; Estes, 1986; Medin & Schaffer, 1978; Nosofsky, 1986) assume that when presented with a novel object, humans compute the similarity between that object and all the possible categories in which the novel object could be placed. Similarity is a function of the sum of the distances between the object and all the exemplars in the particular category. The object is placed into the category with which it is most similar.

Nosofsky's (1986) generalized context model (GCM) allows for variation in the amount of attention given to different features during categorization (see also Medin & Schaffer, 1978). Therefore, it is possible that different cues will be used in different tasks. However, this attention weight remains the same for the entire stimulus set for each particular categorization task, rather than varying across different objects belonging to the same category (in contrast to our new method, as we will see).

Decision Bound Theory (or DBT--see Ashby & Gott, 1988) assumes that there is a multidimensional region associated with each category, and therefore that categories are separated by bounds. An object is categorized according to the region of perceptual space in which it lies. Similarly, neural network models (e.g., Kruschke, 1992) learn hyperplane boundaries between categories, capturing this knowledge in their trainable weights. In both cases, all of the cues available in a particular stimulus are used to determine the region of multidimensional space, and hence the associated category, in which that stimulus falls.

These psychological models all categorize by integrating cues and using all the cues available (except in GCM if a cue has an attention weight of 0). In addition, training these models to learn new categories is a relatively simple process. But the memory requirements assumed by these models do differ: for example, GCM assumes that all exemplars ever encountered are stored and used when categorizing a novel object, while DBT does not need to store any exemplars. In comparison, our CBE algorithm does not integrate all available cues, is similarly easy to train, and typically requires little memory.

Another approach to psychological modeling is captured in the discrete symbol-processing framework of Miller and Laird's (1996) Symbolic Concept Acquisition (SCA) model. Here rules are built up incrementally for classifying stimuli according to specific features, beginning with very general rules that test a single feature and progressing to more detailed rules that must match the stimuli on many features. While there are similarities between this approach and CBE (in particular, the order in which features are processed can be related to our cue validity measure), one major difference is that new stimuli are first checked against rules using all available cues, and only when this fails are fewer cues tested against the more general rules. In contrast, CBE begins with a single cue, and only adds new ones if necessary, thereby minimizing computation. The earlier EPAM symbolic discrimination-net model (Feigenbaum & Simon, 1984) tests rules in the efficient general-to-specific method we advocate, but the rest of our approach is distinct.

In machine learning, predominant categorization theories include neural networks, Classification and Regression Trees (or CART--see Breiman, Friedman, Olshen, & Stone, 1984), and decision trees (e.g., ID3--see Quinlan, 1993). The goal of these machine learning models is usually to maximize categorization accuracy on a given useful data set. Algorithm complexity and speed are not typically the most important factors in developing machine learning models, so that many are not psychologically plausible.

One model that does attempt psychological plausibility by applying selective attention to unsupervised concept formation is Gennari's (1991) CLASSIT. This system classifies objects initially using a subset of the available cues determined by their attentional salience. However, all cues must still be considered before a final decision is reached, due to a "worst case" stopping rule.

Thus, even though many of the machine learning models (e.g., CART and CLASSIT) use only a few cues during a given categorization, the process of setting up the algorithm's decision mechanisms beforehand, including determining which cues to use, can be very complex. In contrast, our CBE algorithm has a simple learning phase, and still maintains comparable accuracy using few cues.

3. Elimination Models

Motivated by the concerns raised in section 1, we wanted to develop a fast and frugal categorization method that combines the best aspects of both the psychological and machine learning models. From the psychological models we used the concepts of simple training and decision processes and a small memory load. From the machine learning models we took the notion of categorizing stimuli without using all available cues. This combination led us to look into elimination models.

Classical elimination models were conceived of for choice tasks (Restle, 1961; Tversky, 1972) In a sequential elimination choice model, an object is chosen by repeatedly eliminating subsets of objects from further consideration, thereby whittling down the set of remaining possibilities. First a particular subset of the original set is chosen with some probability, using a particular feature to determine the subset members. Subsequent subsets are chosen in the same manner, with successive features, until only one object remains.

The most widely known elimination model in psychology is Tversky's (1972) Elimination by Aspects (EBA) model of probabilistic choice. One of the motivating factors in developing EBA as a normative model of choice was that there are often many relevant cues that may be used in choosing among complex alternatives (Tversky, 1972). Therefore, part of any reasonable psychological model of choice should be a procedure to select and order the cues to use from among many alternatives. In EBA, the cues, or aspects, to use are selected according to their utility for some decision (for instance, to choose a restaurant from those nearby, what they serve and how much they charge might be the most important aspects). Possible remaining choices that do not possess the current aspect being used for evaluation (for instance, restaurants that do not serve seafood) are eliminated from the choice set. Furthermore, only aspects that are present in the most recent choice set are considered (for instance, if all nearby seafood restaurants are cheap, then expense will not be used as an aspect to distinguish further among them). Additional aspects are used only until a single choice can be made, which is different from the categorization models described above that use all cues.

4. Categorization by Elimination

Our new Categorization by Elimination algorithm is a noncompensatory lexicographic model of categorization, in that it uses cues in a particular order, and categorization decisions made by earlier cues cannot be altered (or compensated for) by later cues. In CBE, cues are ordered and used according to their validity. For our present purposes we define validity as a measure of how accurately a single cue categorizes some set of stimuli (i.e., percent correct). This is calculated by running CBE only using the single cue in question, and seeing how many correct categorizations the algorithm is able to make. (If using the single cue results in CBE being unable to decide between multiple categories for a particular stimulus, as will often be the case, the algorithm chooses one of those categories at random--in this case, cue validity will be related to a cue's discriminatory power.) Thus if size alone is more accurate in categorizing birds (or more successful at narrowing down the possible categories) than shape alone, size would have a higher cue validity than shape. (There are other ways that cues can be ordered besides by validity, such as randomly or in order of salience, which we are currently exploring.)

CBE assumes that cue values are divided up into bins (either nominal or continuous) which correspond to certain categories. These bins form the knowledge base that CBE uses to map cue values onto the possible corresponding categories. As an example, the *size* cue dimension for birds could be divided into three bins: a *large size* bin (which could be specified numerically, e.g. "over 50 cm") corresponding to the categories of eagles, geese, and swans; a *medium size* bin corresponding to crows, jays, and hawks; and a *small size* bin corresponding to sparrows and finches.

To build up the appropriate bin structures, the relevant cue dimensions to use must be determined ahead of time. At present we construct a complete bin structure before testing CBE's categorization performance, but learning and testing could also be done incrementally. In either case, bins can be constructed in a variety of ways from the training examples--in the next two sections, we present two possibilities.

A flowchart of CBE is shown in Figure 1. Given a particular stimulus to categorize, an initial set of possible categories is assumed, along with the ordered list of cue dimensions to be used. The categorization process begins by using the cue dimension C with the highest validity. Next a subset S of the possible categories is created containing just those categories that correspond to the first cue C's value for the current stimulus object (this subset is determined through the binning map described earlier). If only one category corresponds to that cue value, the categorization process ends with this single category. If more than one category corresponds to the current cue value, that set of possible categories S is kept, and the cue with the next highest validity, C^*, is checked. The set of categories S corresponding to the previous cue C's value is intersected with the set, S^*, of categories corresponding to the present cue C^*'s value. This is CBE's elimination step.

If only one category remains in the new set intersection, the algorithm terminates at this point with that one category. If more than one category remains, this intersection becomes the new set S of remaining possibilities, the next cue is checked, and the process is repeated. If the intersection is empty, then the present cue is ignored, the prior set S of categories is retained, and the next cue is evaluated. This process of checking cues and using them to reduce the remaining set of possible categories continues until a single category remains, or until all the cues have been checked, in which case a category is chosen from the remaining set at random.

This algorithm has several interesting features. It is frugal in information, using only those cues necessary to reach a decision. It is non-compensatory, with earlier cues eliminating category-choice possibilities that can never be replaced by later cues. The binning functions used to associate possible categories with particular cue values can be as simple or detailed as desired, from one-parameter median cuts to multiple-cutoff mappings. And the exact order of cues used does not appear to be critical: in preliminary tests, different random cue orderings vary the algorithm's categorization accuracy by only a few percentage points (but, interestingly, the number of cues used with different orderings *does* vary more widely).

CBE is clearly similar to EBA in several aspects, though there are some important differences. First, EBA is a probabilistic model of choice while CBE is (in its current form) a deterministic model of categorization. Second, in CBE cues are ordered before categorizing so that the same cue order is used to evaluate each object. In EBA, aspects are selected probabilistically according to their weight. Therefore, the order of aspects is not necessarily the same for each object. Third, as mentioned previously, EBA only

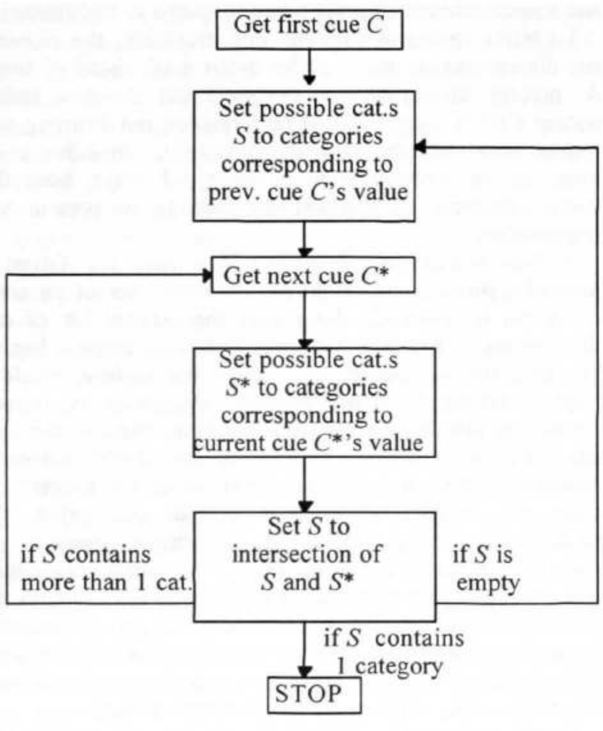

Figure 1: Flow Diagram of CBE

chooses aspects that are present in the *current* set of remaining possible choices, and therefore the process never terminates with the empty set. However, to select such an aspect, all candidates must be examined to determine which aspects are still possible to use. CBE does no such checking ahead of time for appropriate cues to use, but rather takes this circumstance into account in its behavior when the intersection of current and previous possible category sets comes up empty.

5. CBE and Human Data

Under what conditions might CBE be a plausible description of human categorization? We expect the most evidence for CBE to come from situations in which categorization may be affected by time and cue availability constraints. As mentioned in the introduction, one specific domain where time and number of available cues are limited is in inferring intention from motion. Blythe, Miller, and Todd (1996) conducted an experiment in which the subjects' task was to infer intention from motion of two animated bugs shown moving about on a computer screen. The movement patterns had all been previously generated by other subjects instructed to engage in various types of interaction by each controlling the motion of a single on-screen bug. The six possible categories of interactive motion were: pursuit, evasion, fighting, courting, being courted, and playing. For example, one subject's task would be to have their bug pursue the other bug, while the other subject would move their bug to evade their pursuing opponent. Next, new subjects viewed the recorded bug interactions and through forced choice, categorized the interactions as a specific type of intentional motion.

Seven salient cues of motion were calculated for each of the recorded motion patterns (see Blythe, Miller, & Todd, 1996, for details). These cues were used to compare different categorization models with each other and against human performance. (While we cannot be sure that these are the exact cues used by the human subjects, it is a reasonable set to start with.)

The four models tested were CBE and three traditional cue-integrating compensatory algorithms: unit tallying (counting up the total number of cues that indicate one category versus another, using the same bin mapping as CBE), weighted tallying (adding up weighted votes from all the cues that indicate one category versus another, again using CBE's binning along with weights determined by correlation), and a three-layer feed-forward neural network model trained by backpropagation learning (see Gigerenzer & Goldstein, 1996, for more details on the first two).

The bin structure used for CBE and the tallying algorithms were determined by considering the distribution of cue values for each category and placing the bin boundaries at points of minimum overlap between categories. As a result, some cue values could be mapped onto too few possible categories (e.g. if pursuit was usually fast and courtship usually slow, the fast velocity bin would only map to pursuit, and thus would miss all those instances of rapid, excited courtship motion). Thus this bin mapping made perfect categorization impossible for CBE in this domain, and yet it did surprisingly well. Table 1 lists

Table 1: Categorization accuracies and average number of cues used for subjects and models (chance = 16.7%).

Method	Cat Acc	Avg Cues
Subject	49.33%	?
CBE	65.33%	3.77
Wtally	64%	7
Utally	63.33%	7
Nnet	88.33%	7

the average categorization accuracies of the human subjects, the categorization accuracies for the four models, and the average number of cues used (this value is unknown for the human subjects). Since there were six possible categories, chance performance is 16.7%.

As can be seen in Table 1, human subjects performed well above chance in this task, and the four categorization algorithms performed better still. The neural network did suspiciously far better than the human subjects, indicating that it has possibly been overtrained on this data. (When tested on generalization ability on a further untrained set of motion stimuli, the network's performance drops to 68%, while the other three algorithms hover around 56%.) The tallying algorithms and CBE are all much closer to human performance, but CBE achieves its accuracy while using only about half of the cues of the others.

The difference in accuracy between subjects and the algorithms can be explained in part by the fact that the algorithms are "trained" on all the stimuli, either through the binning process or neural network learning (300 motion patterns in this case). In contrast, subjects must make their categorizations without previous exposure to these stimuli (under the assumption that they would already know the cue structures of these categories through their experiences outside the lab). To make a more fine-grained assessment of how well each categorization algorithm matches the human data, we are performing analyses of the case-by-case categorizations made by subjects and algorithms. But even without this detailed analysis, CBE emerges as a parsimonious contender among categorization algorithms in this multi-cue domain, and the clear winner when time and information-availability constraints are taken into account.

6. CBE and Machine Learning Algorithms

It is difficult to compare CBE to existing categorization models on multiple-cue human data beyond the domain just presented, because few other experiments have been performed with more than three or four cues. Instead, as an alternative test of CBE's general accuracy potential, we examined how well CBE categorized various multi-dimensional objects using data from the UCI Machine Learning Repository (Merz & Murphy, 1996). We compared the performance of CBE to a standard exemplar model and a three-layer feed-forward neural network trained with backpropagation. Results are shown in Table 2 for categorization performance when trained on the full data sets and generalization performance when trained on half of each data set and tested on the other half. In addition, we include the best reported categorization performance we have found for each data set in the machine learning literature.

For the following comparisons, CBE used "perfect" binnings for the cue values. This means that the cue-value bins always map to the entire set of possible categories associated with those particular cue values (unlike the bins in the motion categorization example, where only the most prevalent categories in each bin were returned). With perfect binning, the same categorization accuracy is always obtained regardless of the order in which the cues are used. However, when the cues are ordered by validity, categorizations can be accomplished using the fewest number of cues.

Table 2 shows the results of these comparisons for three data sets. The first is the famous iris flower database in which there are 150 instances classified into three categories (different iris species) using four continuous-valued features (lengths and widths of flower parts). The next comparison used wine recognition data, in which 13 chemical-content cues are used to classify 178 wines as one of three particular Italian vintages. The third data set analyzed contains two mushroom categories, poisonous and edible, with 22 nominally valued dimensions, and 8124 total instances.

Overall, CBE does very well on these three sets of multi-feature natural objects, while using only a small proportion of the available cues. CBE even performs well in comparison with models specifically designed for the best possible performance on these particular data sets. We were not expecting CBE to outperform these specialized algorithms; merely being in the same ballpark is a powerful testament to this approach's potential accuracy across varied domains. Furthermore, these algorithms all employ the full set of cues, making the contrast with CBE's high accuracy through limited information all the more striking.

7. Future Work

The results we have presented here indicate that a fast and frugal approach to categorization is a viable alternative to cue-integrating compensatory models. By only using those cues necessary to first make a categorical decision, CBE can categorize stimuli under time pressure and information constraints. Moreover, if certain cues are missing (i.e. some feature values are unknown or cannot be perceived), CBE can still use the other available cues to come up with a category judgment (we are in process of collecting data on this type of generalization ability across different categorization algorithms). Yet CBE still performs very accurately, despite its limited use of knowledge, rivaling the abilities of much more complex and sophisticated algorithms (not to mention human subjects!).

The following issues still need to be explored. First, how should bin structures be created? Incremental learning can build the cue-value bins gradually as more and more stimuli are seen. But how far should this learning process go, and in what way should it proceed? We have presented two alternatives here, and there are many others possible. One important issue to explore further is the performance tradeoff between accuracy and the amount of knowledge captured in the bin structure (CBE's memory requirements).

Second, more data from human performance on categorizing multi-dimensional objects needs to be collected

Table 2: Categorization accuracies and average number of cues used for various models on three data sets

Model	Train/Test Set Size	Iris		Wine		Mushroom	
		Cat Acc	Avg Cues	Cat Acc	Avg Cues	Cat Acc	Avg Cues
CBE	Full	91.33 %	40.00 %	96.63 %	20.74 %	91.71 %	26.11 %
	Half	92.40 %	26.24 %	90.37 %	15.83 %	91.66 %	26.16 %
Nnet	Full	97.67 %	100 %	100 %	100 %	86.21 %	100 %
	Half	97.07 %	100 %	95.95 %	100 %		
Best Reported		98 % (James,	1985)	100 % (Aeberhard	et al., 1992)	95.00 % (Schlimmer,	1987)

and analyzed to compare CBE with other categorization models. We are particularly interested in investigating the patterns of misclassifications, learning curves, and predicted time-courses associated with CBE and human performance. The intriguing finding in our intention from motion data that categorization accuracy varied little with changes in cue order can also be studied experimentally.

Third, category base-rates and payoffs for right and wrong classifications should be incorporated into the model. For example, with the mushroom categories described in the previous section, if a mushroom remains uncategorized as poisonous or safe even after all the cues have been used, it seems reasonable to err on the side of caution and guess that the mushroom is poisonous.

With these further explorations and extensions to CBE, we will come to understand the algorithm's behavior better, and be able to make it a better model of human behavior in turn. For now, though, we have shown evidence for the view that the mind need not amass and combine all the available cues when telling a hawk from a dove, or a threat from a flirt--fast and frugal does the trick.

References

Aeberhard, S., Coomans, D., & de Vel, O. (1992). *Comparison of classifiers in high dimensional settings* (Tech. Rep. 92-02). James Cook University of North Queensland, Dept. of Computer Science and Dept. of Mathematics and Statistics.

Ashby, F.G., & Gott, R. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 33-53.

Ashby, F.G. & Maddox, W.T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 50-71.

Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees.* New York, NY: Chapman & Hall, Inc.

Blythe, P. W., Miller, G. F., & Todd, P. M. (1996). Human simulation of adaptive behavior: Interactive studies of pursuit, evasion, courtship, fighting, and play. *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 13-22). Cambridge, MA: MIT Press/Bradford Books.

Brooks, L. R. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B.B. Lloyd (Eds.), *Cognition and Categorization*, (pp. 169-211) Hillsdale, NJ: Erlbaum.

Estes, W. K. (1986). Array models for category learning. *Cognitive Psychology*, 18, 500-549.

Feigenbaum, E. A., & Simon, H.A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, **8**, 305-336.

Gennari, J. H. (1991). Concept formation and attention. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Society* (pp. 724-728). Hillsdale, NJ: Erlbaum.

Gigerenzer, G. & Goldstein, D. G., (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650-669.

James, M. (1985). *Classification algorithms*. New York, NY: John Wiley.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22-44.

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

Merz, C. J., & Murphy, P. M. (1996). UCI Repository of machine learning databases [http://www.ics.uci.edu/~mlearn/MLRepository.html]. Irvine, CA: University of California, Department of Information and Computer Science.

Miller, C. S. & Laird, J. E. (1996). Accounting for graded performance within a discrete search framework. *Cognitive Science*, 20, 499-537.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39-57.

Quinlan, J. R. (1993). *C4.5: Programs for machine learning*. Los Altos: Morgan Kaufmann.

Restle, F. (1961). *Psychology of judgment and choice*. New York: Wiley.

Schlimmer, J.S. (1987). *Concept acquistion through representational adjustment* (Tech. Rep. 87-19). University of California, Irvine. Doctoral dissertation, Department of Information and Computer Science.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4), 281-299.