Learning at Different Levels of Abstraction

Bradley C. Love
Department of Psychology
The University of Texas at Austin
Austin, TX 78712 USA
love@psy.utexas.edu

Abstract

Previous category learning research and the SUSTAIN (Supervised and Unsupervised STratified Adaptive Incremental Network) model of category learning suggest that preferred category level (in a hierarchy of categories) shifts towards lower-level (i.e., more specific) categories when stimuli are perceived to be more distinctive. This shift is in accord with work in expertise. In their domain of expertise, experts excel (relative to novices) at classifying stimuli at lower category levels, but their advantage is attenuated with higher-level categories. The work described here directly tests (within a single study) this predicted interaction between category level and stimulus distinctiveness using well controlled artificial stimuli. The results are consistent with prior work utilizing natural stimuli. The results are also informative for evaluating whether attention is dimension-wide (i.e., all items are represented in a common multi-dimensional space of the same extent) or cluster specific (i.e., different conceptual clusters can stress different stimulus dimensions so that different aspects of different stimuli are stressed). The results suggest that attention is not dimension-wide. Instead, attention can stress different aspects of different stimuli. The implications of these findings for models of category learning are discussed.

Introduction

Humans frequently utilize and acquire category knowledge at multiple levels of abstraction. For example, the same object can be classified as a vehicle, as a car, or as a 1978 Lincoln Continental. Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976) argue that objects are most easily classified at the intermediate category level which most effectively partitions the world into informative clusters. However, Tanaka and Taylor (1991) have found that different groups of people tend to prefer different levels of abstraction with experts preferring lower-level categories (i.e., narrower or finer grained categories) compared to novices who prefer higher-level categories (i.e., broader or more abstract categories).

One domain in which all adult humans are experts is the domain of face perception. Medin, Dewey, and Murphy (1983) found that people are faster to associate unique names to photographs of nine female faces than they are to categorize the photographs into two categories. The logical structure of the two categories is shown in Table 1. One possible explanation for the relative ease of identification learning is that the stimuli used in Medin et al. (1983) were rich and distinct, varying along many dimensions not listed in Table 1, such as the shape of the face, the type of nose, etc.. This *idiosyncratic* information makes each stimulus item more distinct. Experts may be more sensitive to idiosyncratic information than novices. In the absence of idiosyncratic information, common wisdom holds that identification learning should be harder than category learning. In other words, the ease of category acquisition interacts with the nature of the stimuli such that learning at lower levels of abstraction (with identification learning being the lowest level of abstraction) becomes easier relative to learning at higher levels of abstraction as stimuli become more distinct.

Table 1: The logical structure of the two categories used in the category learning condition. The four dimensions were hair color, smile type, hair length, and shirt color.

Category A	Category B
1112	1122
1212	2112
1211	2221
1121	2222
2111	

Results from the category learning literature support this conclusion. Shepard, Hovland, and Jenkins (1961) trained subjects on the six category learning problems listed in Table 2 and found that Type I was the easiest to master, followed by Type II, followed by Types III-V, followed by Type VI. The Type IV problem has a family resemblance structure that resembles the category structure used in Medin et al. (1983). In the Type IV problem, each category consists of an underlying prototype (111 for category “A” and 222 for category “B”) and any item that matches a prototype on two out of three dimensions is a member of the corresponding category. The more difficult to master Type VI problem, while not an identification learning problem, requires subjects to memorize all eight stimulus items because no regularities exist across any pair of dimensions. This data, along with Medin et al.’s (1983) data, suggest that preferred category level and stimulus distinctiveness interact with learning being facilitated more at lower category levels with distinctive stimuli.

One category learning model can capture this interaction. SUSTAIN (Supervised and Unsupervised STratified Adaptive Incremental Network) has successfully fit Shepard et al.’s (1961) and Medin et al.’s (1983) data using the same set of

Table 2: The logical structure of the six classification problems tested in Shepard et al. (1961). The three binary valued dimensions correspond to size (small or large), shape (triangle or square), and color (light or dark)

Input	I	II	III	IV	V	VI
111	A	A	B	A	B	B
112	A	A	B	A	B	A
121	A	B	B	A	B	A
122	A	B	A	B	A	B
211	B	B	A	A	A	A
212	B	B	B	B	A	B
221	B	A	A	B	A	B
222	B	A	A	B	B	A

parameters (Love & Medin, 1998a; Love & Medin, 1998b).[1] SUSTAIN is a connectionist model that clusters similar items together. When items are clustered together inappropriately (i.e., similar items from incompatible categories are placed in the same cluster), SUSTAIN adds a new cluster in memory to encode the misclassified item. For example, if SUSTAIN is applied to stimulus items and classifies them as members of the category mammals or the category birds it will develop one or more clusters (i.e., prototypes) for the bird category and one or more clusters for the mammal category. When SUSTAIN classifies a bat for the first time, the bat item will strongly activate a bird cluster because bats are similar to birds (both bats and birds are small, have wings, and fly). After incorrectly classifying the bat as a bird, SUSTAIN will create a new cluster to encode the misclassified bat item. The next time SUSTAIN classifies a bat, this new cluster will compete with the other clusters and will be the most strongly activated cluster (i.e., it will be more similar to the current stimulus than any other cluster), leading SUSTAIN to correctly classify the novel bat as a mammal and not as a bird. The new cluster would then become a bat prototype (a subcategory of mammal). The primary difference between SUSTAIN and exemplar models is that SUSTAIN can cluster examples together in memory (like a prototype model). Unlike a prototype model, SUSTAIN can form multiple clusters (i.e., prototypes) per category.

When applied to Shepard et al.'s data, SUSTAIN recruits fewer clusters for the simpler problems. For the simplest problem, the Type I problem, SUSTAIN only recruits two clusters (one for each category). For the most difficult problem, the Type VI problem, SUSTAIN resorts to recruiting eight clusters (one for each item; each stimulus is memorized). When applied to Medin et al.'s (1983) data, SUSTAIN recruits more clusters (nine clusters; one for each stimulus item) in identification learning condition than in the category learning condition (the modal solution involves seven clusters). It is important to note that abstraction does not occur in the identification learning condition (i.e., each cluster responds to only one item), but does occur in the category learning condition. What is interesting about these data fits is that in one case memorizing more items (acquiring more fine grained clusters or subcategories) led to more efficient learning, while in the other case it led to less efficient learning. The critical difference between these two data sets is the distinctiveness of the stimuli.

Two factors conspire to cause SUSTAIN's performance to interact with the nature of the stimuli. As the stimuli become more distinctive, clusters that respond to multiple items (i.e., prototypes) are not as strongly activated. In other words, the benefit of abstraction is diminished with distinctive stimuli. This occurs because distinctive items sharing a cluster are not very similar to each other (i.e., within cluster similarity is low). Notice that the diminished benefit of abstraction negatively impacts performance in the Medin et al.'s (1983) category learning condition, but does not affect identification learning. In identification learning, each item forms its own cluster (within cluster similarity is maximal). When SUSTAIN is altered so that it does not form abstractions in either condition, but instead recruits a subcategory unit for each item, SUSTAIN fails to predict the interaction or the identification learning advantage, suggesting that abstraction is critical for capturing this effect. Without abstraction, the inferred category structures (i.e., the clusters recruited) are identical for both conditions.

The second factor that leads SUSTAIN to predict that distinctiveness and category level should interact is that the effects of cluster competition are attenuated with distinctive stimuli. As items become more distinctive, the clusters that are recruited tend to be further separated in representational space (i.e., the clusters match on fewer dimensions and mismatch on more dimensions). In other words, the clusters become more orthogonal to one another. The more distinctive the clusters are, the less they will tend to compete with one another. For instance, when a distinctive stimulus is presented to SUSTAIN, it will tend to strongly activate the appropriate cluster and will only weakly activate the competing clusters. Reduced cluster competition with distinctive stimuli favors both identification and category learning, but differentially benefits identification learning (or more broadly, learning at lower levels of abstraction) because there are generally more clusters present (i.e., potential competitors) in identification learning. Simulations support this analysis. When SUSTAIN is modified so that clusters do not compete, SUSTAIN reaches criterion more often and overall accuracy is much higher in the category learning condition.

Experiment

SUSTAIN's ability to fit Medin et al.'s studies on item and category learning is notable because other models cannot predict the advantage for identification learning or the interaction between learning task and stimulus distinctiveness. More importantly, SUSTAIN offers a framework for understanding the results. At the same time time, it seems impor-

[1] The data actually fit was from Nosofsky et al.'s (1994) replication of Shepard et al. (1961).

tant to place SUSTAIN's account of these findings on firmer ground. To begin with, one should be cautious about accepting SUSTAIN's characterization of Medin et al.'s (1983) results. SUSTAIN's successful fit of Medin et al.'s (1983) studies depended on the choice of input representation. The idiosyncratic information in each photograph was represented by adding a number of input dimensions. Each item had a unique value on each added dimension. This manipulation had the effect of making all the items less similar to each other and making between and within category similarity virtually the same in the category learning condition. This input representation led SUSTAIN to predict that identification learning should precede category learning with distinctive stimuli.

The general intuition that guided my choice of input representation seems justified. Unlike artificial stimuli, the photographs do vary along a number of dimensions. Still, replicating the results from Medin et al. (1983) under more controlled circumstances with artificial stimuli would bolster my claims. Also, it is possible that there may be something "special" about faces (c.f., Farah, 1992).

The stimuli used in the Experiment were schematic cars that varied on a few dimensions. Like the Medin et al.'s (1983) study, subjects were assigned to either an identification or a category learning task. To manipulate the distinctiveness of the stimuli, some subjects viewed stimuli that were uniquely colored, while other subjects viewed stimuli that shared a common color. In essence, this experiment augments Medin et al.'s (1983) design with two non-distinctive stimuli conditions. The key prediction SUSTAIN makes is that category level and distinctiveness should interact such that identification learning performance should improve more than category learning performance as the stimuli become more distinctive. The choice of stimuli in this experiment directly tests SUSTAIN's characterization of the Medin et al. (1983) results.

Unfortunately, whether or not the interaction crosses over (i.e., whether or not identification learning with distinctive stimuli turns out to be faster than category learning with distinctive stimuli) cannot be predicted by SUSTAIN because the size of the effect depends on the saliency of the color dimension (the distinctive dimension). A crossover interaction can be accommodated by SUSTAIN, but is problematic for many other models.

The Experiment's design also allows for a secondary prediction to be tested related to cluster encoding and attention. When clusters only respond strongly to one item (i.e., "exception" clusters), does the cluster focus on the distinctive item information? Conversely, when a cluster encodes a number of items (i.e., an "abstraction" or "rule" cluster), does the cluster focus on what is general to the items and suppress distinctive item information? If the answer to both of these questions is "yes", the results would strongly suggest that attention is not dimension-wide (e.g., Nosofsky, 1986; Kruschke, 1992), but is cluster specific (c.f., Aha & Goldstone, 1992). In other words, different clusters can attend to different stimulus dimensions.

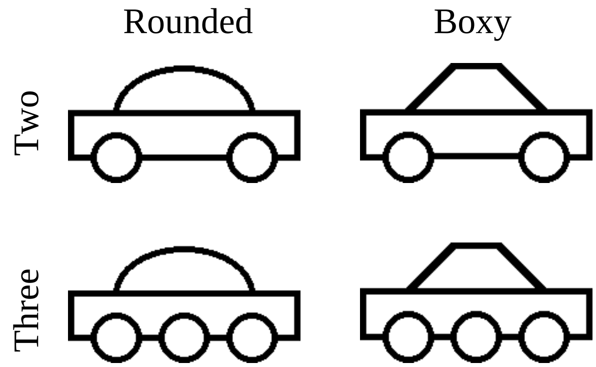

Figure 1: The stimuli varied in size (small or large), the number of wheels (two or three) and the shape of the cockpit (rounded or boxy shaped). Only one car size is shown in the figure. In the Experiment, eight different items were used.

Methods

Subjects Two hundred eighty-eight Northwestern University undergraduate students participated in the experiment for course credit or pay.

Stimuli Example stimuli are shown in Figure 1. In two of the four experimental conditions, each car was a different color. The eight colors were yellow, light blue, black, red, navy blue, pink, green, and grey. In the two other conditions, subjects viewed cars that were all the same color (either yellow, light blue, black, red, navy blue, pink, green, or grey).

Design and Overview The two variables (category level and distinctiveness) were crossed for a 2 X 2 between subjects factorial design. Subjects were randomly assigned to one of these four conditions.

The category level variable had two levels: identification learning and category learning. Subjects performing the identification learning task partitioned the eight items into eight categories (i.e., each stimulus formed its own category). Subjects performing the category learning task partitioned the eight items into two categories that had the same logical structure as Shepard et al.'s (1961) Type IV problem (see Table 2).

The distinctiveness variable had two levels: distinctive and non-distinctive. In the distinctive conditions, each item was a unique color. In the non-distinctive conditions, each item shared a common color. The learning phase ended when subjects completed consecutive error-free blocks of trials or after the completion of thirty-second block (each stimulus was presented in a random order once per block).

After completing the learning phase, sixty-six subjects completed two transfer blocks. Transfer trials were identical to learning trials with the exceptions that feedback was not provided and that each item was colored orange (a color not used during the learning phase).

Procedure

Text was displayed in black on a white background. Trials began with a message displayed in the upper left corner of the screen alerting the subject to prepare for the next trial. After 1500 ms, this message was removed and the stimulus was displayed along with a message below it indicating that the subject should respond. Subjects were instructed to push the spacebar as soon as they decided on a response. After pressing the spacebar, subjects were prompted for their response. Subjects pressed either the "A" or "B" key in the category learning conditions. In the identification learning conditions, subjects used keys "A" through "H" to indicate their response. After responding, subjects received feedback. When subjects were correct, the message "Correct!" was displayed at the bottom of the screen. When subjects were incorrect, a message alerted the subject to the error and the correct response was displayed at the bottom of the screen. Following the subject's response, the stimulus and all messages were displayed for 1500 ms. After another 1500 ms, the next trial began.

On transfer trials, text was displayed in black on a white background. Trials began with a message displayed in the upper left corner of the screen alerting the subject to prepare for the next trial. After 1500 ms, this message was removed and the stimulus was displayed along with a message below it indicating that the subject should respond. Subjects were instructed to push the spacebar as soon as they decided on a response. After pressing the spacebar, subjects were prompted for their response. Subjects pressed either the "A" or "B" key in the category learning conditions. In the identification learning conditions, subjects used keys "A" through "H" to indicate their response. Subjects did not receive feedback. Whether or not the subject responded correctly, the message "Thank You" was displayed. Following the subject's response, the stimulus and all messages were displayed for 1500 ms. After 1500 ms, the next trial began. All possible factors were counterbalanced or randomly varied.

Results

Criterion The mean of the number of blocks required by subjects in each condition is shown in Table 3. Table 3 also shows the mean of the reciprocal of the number of blocks required (this measure is less sensitive to outliers). A 2 X 2 (category level by distinctiveness) ANOVA was performed on both the untransformed scores and the transformed reciprocal scores. The transformed scores' distributions were more similar to the normal distribution than were the distributions of the untransformed scores. Means are given only for the untransformed scores. Subjects required more blocks (14.5 vs. 12.9 blocks) in the category learning conditions than in the identification learning conditions (untransformed: $F(1,218)=3.65$, MSe=140.5, $p = .06$; transformed: $F(1,218)=4.40$, MSe=.00608, $p < .05$). Subjects required more blocks (17.0 vs. 10.5 blocks) in the non-distinctive conditions than in the distinctive conditions (untransformed: $F(1,218)=60.95$, MSe=2348, $p < .001$; transformed: $F(1,218)=78.4$, MSe=.108, $p < .001$). The key prediction SUSTAIN makes is that category level and distinctiveness should interact such that identification learning performance should improve more than category learning performance as the stimuli become more distinctive. As predicted, distinctiveness and category level interacted such that distinctiveness sped up learning more (8.4 vs. 4.6 blocks) in the identification learning conditions than in the category learning conditions (untransformed: $F(1,218)=5.08$, MSe=195.8, $p < .05$; transformed: $F(1,218)=15.59$, MSe=.00216, $p < .001$).

Table 3: The mean number of blocks required for each condition. In parenthesis, the mean of the reciprocals of the number of blocks required is shown.

	Identification Learning	Category Learning
Non-Distinctive	17.1 (.0661)	16.8 (.0753)
Distinctive	8.7 (.130)	12.2 (.100)

Table 4: The proportion correct for learning trials. In parenthesis, the proportion of subjects reaching the learning criterion is shown.

	Identification Learning	Category Learning
Non-Distinctive	.76 (.98)	.85 (.88)
Distinctive	.90 (1.00)	.90 (.96)

A series of t-tests were conducted to probe individual cell differences. All differences were statistically significant at the .01 level except for the comparison (17.1 vs. 16.8 blocks) between the identification learning/non-distinctive condition and the category learning/non-distinctive condition (untransformed: $t < 1$; transformed: $t(109)=1.54$, $p = .13$).

Table 4 shows the proportion subjects that reached the learning criterion (the completion of consecutive error-free blocks) for each condition. Subjects reached the learning criterion more often in the identification learning conditions than in the category learning conditions ($p < .05$ by a binomial test). Subjects also reached the learning criterion more often in the distinctive conditions than in the non-distinctive conditions ($p < .05$ by a binomial test). Individual cell differences were not probed because of ceiling effects. Cell differences and interactions are explored in other analyses.

Learning Trial Accuracy In addition to analyzing the number of required learning blocks, the accuracy data were analyzed. We assume that after reaching the learning criterion subjects would respond correctly on the remaining trials if they maintained their motivational level. We scored the remaining post criterion blocks accordingly. The proportion correct for each condition is shown in Table 4.

A 2 X 2 (category level by distinctiveness) ANOVA was performed with the subjects' accuracy rates serving as the dependent variable. Subjects were more accurate (.88 vs. .83) in the category learning conditions than in the identification learning conditions ($F(1,218)=16.45$, MSe=.132, $p < .001$). Subjects were more accurate (.90 vs. .81) in the

Table 5: The proportion correct for transfer trials.

	Identification Learning	Category Learning
Non-Distinctive	.97	.93
Distinctive	.47	.81

distinctive conditions than in the non-distinctive conditions (F(1,218)=60.65, MSe=.485, $p < .001$). The key prediction SUSTAIN makes is that category level and distinctiveness should interact such that identification learning performance should improve more than category learning performance as the stimuli become more distinctive. As predicted, distinctiveness and category level interacted such that distinctiveness led to a larger improvement in accuracy (.14 vs. .05) in the identification learning conditions than in the category learning conditions (F(1,218)=12.75, MSe=.102, $p < .001$).

A series of t-tests were conducted to probe individual cell differences. All differences were statistically significant at the .01 level except for the comparison between the identification learning/distinctive condition and the category learning/distinctive condition ($t < 1$).

Transfer Trial Accuracy Sixty-six subjects engaged in transfer trials after finishing the learning phase. The proportion correct for transfer trials for each condition is shown in Table 5. A 2 X 2 (category level by distinctiveness) ANOVA was performed with the subjects' accuracy rates serving as the dependent variable. Subjects were more accurate (.87 vs. .72) in the category learning conditions than in the identification conditions F(1,62)=15.62, MSe=.2838, $p < .001$). Subjects were more accurate (.95 vs. .64) in the distinctive conditions than in the non-distinctive conditions (F(1,62)=82.92, MSe=1.506, $p < .001$). Distinctiveness and category level interacted such that identification learning led to a larger decrement in accuracy (.342 vs. -.0478) in the distinctive conditions than in the non-distinctive conditions (F(1,62)=34.34, MSe=.6238, $p < .001$).

Analyses of the Attentional Clustering Hypothesis The tests presented here explore the possible existence of imperfect rule or abstraction clusters (focused on the non-distinctive dimensions) and exception clusters (focused on the distinctive dimension). Only the category learning data are relevant to the analyses presented here.

The Type IV problem used in the category learning conditions has essentially two types of items in each category: the prototypes (111 and 222) and all the other items. Within a category, each "other" item matches the prototype on two out of the three stimulus dimensions and mismatches on the third. Drawing this distinction between prototypes and other items proves useful in evaluating the attentional clustering hypothesis.

Although it is very difficult to analyze a subject's data and identify which items were exceptions and which items clustered with other items, SUSTAIN offers some direction. According to SUSTAIN, it is much more likely for an "other" item to be an exception than it is for a prototype item to be an exception (prototypes are much more likely to cluster with other items). Drawing on this knowledge, 2 X 2 (item type by distinctiveness) ANOVAs can be performed to evaluate the hypothesis. If the hypothesis is correct, the advantage of the distinctive category learning condition over the non-distinctive category learning condition should largely be attributable to the greater ease in memorizing exceptions (i.e., non-prototype items) made possible by focusing on distinctive information. Thus, in the learning data, we should observe an interaction in which item type and distinctiveness interact such that non-prototype items benefit more from distinctiveness than prototype items. This prediction was confirmed (analysis below).

Table 6: The proportion correct for Learning trials (to the left of the slash) and for transfer trials (to the right of the slash) by item type.

	Others	Prototypes
Non-Distinctive	.83 / .92	.92 / .96
Distinctive	.89 / .77	.94 / .93

Conversely, for the transfer trials (where the distinctive information is obscured), the advantage of the non-distinctive condition should largely be attributable to better performance on the non-prototype items. In other words, the two factors should interact such that the distinctiveness hurts performance more for the non-prototype items than for the prototype items. This predictions was confirmed (analysis below).

Accuracy Data A 2 X 2 (item type by distinctiveness) ANOVA was performed on the learning data from the category learning conditions with subjects' accuracy rates serving as the dependent variable (see Table 6). Subjects more accurately (.93 vs. .86) classified the prototypes than than the other items (F(1,110)=139.18, MSe=.264, $p < .001$). Subjects were more accurate (.92 vs. .87) in the distinctive conditions than in the non-distinctive conditions (F(1,110)=7.39, MSe=.104, $p < .01$). One prediction of my proposed account of clustering and attentional focus is that item type and distinctiveness should interact such that learners in the distinctive conditions should see greater facilitation with non-prototype items than with prototype items in comparison to the non-distinctive conditions. As predicted, distinctiveness and item type interacted such that distinctiveness led to a larger improvement (.06 vs. .02) in the classification of non-prototype items in the distinctive conditions in comparison to the non-distinctive conditions (F(1,110)=7.33, MSe=.0139, $p < .01$).

Transfer Data A 2 X 2 (item type by distinctiveness) ANOVA was performed on the transfer data from the category learning conditions with subjects' accuracy rates serving as the dependent variable (see Table 6). Subjects more accurately (.94 vs. .85) classified the prototypes than than the other items (F(1,32)=15.78, MSe=.155, $p < .001$). Subjects were more accurate (.94 vs. .85) in the non-distinctive conditions than in the distinctive conditions and this result was marginally significant ($F(1,32)=2.91$, MSe=.125, $p < .10$). One prediction of my proposed account of clustering and attentional focus is that item type and distinctiveness should interact such that transfer performance in the distinctive conditions should decline more for non-prototype items than for prototypes in comparison to item performance in the non-distinctive conditions. As predicted, distinctiveness and item type interacted such that distinctiveness led to a larger decline in performance (.15 vs. .03) in the classification of non-prototype items in the distinctive conditions in comparison to performance in the non-distinctive conditions ($F(1,32)=5.49$, MSe=.054, $p < .05$).

Discussion

SUSTAIN predicts that category level and distinctiveness should interact such that identification learning performance should improve more than category learning performance as the stimuli become more distinctive. SUSTAIN also predicts that identification learning can become easier than category learning as the stimuli become more distinctive. These predictions were borne out in an analysis of the number of learning blocks required by subjects and in another analysis of accuracy. This robust effects occurred by simply manipulating the color of the stimuli.

The experiment makes a bridge to the Medin et al.'s (1983) study in which subjects performed identification and category learning tasks with rich stimuli (photographs of human faces). SUSTAIN suggests that the identification advantage in the Medin et al. studies arises from the distinctiveness of the stimuli. The experiment tested this prediction directly and completed Medin et al.'s (1983) design with two non-distinctive conditions to complement the two distinctive conditions, allowing for natural comparisons to be made between the various conditions.

Another interesting aspect of the data was that subjects apparently emphasized (or attended) to different stimulus dimensions depending on which stimulus was being classified. In particular, when a stimulus was encoded by its own cluster, the distinctive aspects of the stimuli were emphasized. In contrast, when items shared a cluster, the cluster emphasized the non-distinctive aspects of the cluster members.[2] These results suggest that attention is not dimension-wide and is instead cluster specific. While SUSTAIN (in its current form) exhibits dimension-wide attention, it can be modified so that attention is cluster specific. The interpretability of SUSTAIN, which arises from its internal cluster representation of categories, suggested the analyses related to item type and attentional focus. Although the results of these analyses actually run counter to the specifics of the model, they speak favorably of SUSTAIN's general framework and its utility for directing empirical investigations.

[2]The idea that divergent examples are encoded in terms of how they differ from a default type is in the spirit of Kruschke's (1996) ADIT model of inverse base rate phenomena.

Overall, the results of the Experiment are troublesome for current models of category learning. Existing models have difficulty accounting for the interaction between category level and stimulus distinctiveness, the advantage of identification learning with distinctive stimuli, and attention not being dimension-wide. The first two findings are predicted by SUSTAIN and the third can be accommodated. SUSTAIN's ability to cluster together similar items from the same category (an ability that exemplar models lack) allows it to account for the results from the Experiment.

Acknowledgments

I would like to thank Doug Medin for collaborating with me on this project. I would also like to thank Greg Ashby, John Kruschke, and James Tanaka for helpful comments on a related paper.

References

[1] AHA, D. W., AND GOLDSTONE, R. L. Concept learning and flexible weighting. In *n Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (1992), pp. 534–539.

[2] FARAH, M. Is an object an object an object and object? Cogntive and neuropsychological investigations of domain-specificity in visual object recogntion. *Current Directions in Psychological Science 164-169* (1992), 1.

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

[4] KRUSCHKE, J. K. Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition 22* (1996), 3–26.

[5] LOVE, B. C., AND MEDIN, D. L. Modeling item and category learning. In *Proceedings of the 20th Annual Conference of the Cogntive Science Society* (Mahwah, NJ, 1998), Lawrence Erlbaum Associates, pp. 639–644.

[6] LOVE, B. C., AND MEDIN, D. L. SUSTAIN: A model of human category learning. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence* (Cambridge, MA, 1998), MIT Press, pp. 671–676.

[7] MEDIN, D. L., DEWEY, G. I., AND MURPHY, T. D. Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, & Cognition 9* (1983), 607–625.

[8] NOSOFSKY, R. M. Attention, similairty, and the identification-categorization relationship. *Journal of Experimental Psychology: General 115* (1986), 39–57.

[9] NOSOFSKY, R. M., GLUCK, M. A., PALMERI, T. J., MCKINLEY, S. C., AND GLAUTHIER, P. Comparing models of rule based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition 22* (1994), 352–369.

[10] ROSCH, E., MERVIS, C. B., GRAY, W. D., JOHNSON, D. M., AND BOYES-BRAEM, P. Basic objects in natural categories. *Cognitive Psychology 8*, 3 (1976), 382–439.

[11] TANAKA, J. W., AND TAYLOR, M. Object categories and expertise: Is the basic level in the eye of the beholder. *Cognitive Psychology 23*, 2 (1991), 457–482.