

THE ROLE OF CONCEPTUAL STRUCTURE AND BACKGROUND
KNOWLEDGE IN CATEGORY LEARNING

by

MATTHEW CARL JOHNSON, B.S., M.S.

A DISSERTATION

IN

PSYCHOLOGY

Submitted to the Graduate Faculty
of Texas Tech University in
Partial Fulfillment of
the Requirements for
the Degree of

DOCTOR OF PHILOSOPHY

Approved

Accepted

December, 1998

ACKNOWLEDGEMENTS

I would like to thank Dr. Roman Taraban for his patience and also for his valuable guidance in the preparation of this dissertation. I would also like to thank Dr. Philip Marshall, Dr. William Maki, and Dr. Katherine Gannon for the time and effort they contributed and for their expertise. This dissertation could not have been possible without the support and love of my parents Jack and Margaret Linquist, and my Savior, Jesus Christ. Many thanks to Kim Rynearson for moral support, and Maria Maldia for help in data collection. This dissertation is dedicated to the memory of Dr. Jason E. Albrecht.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
ABSTRACT	v
LIST OF TABLES	vii
CHAPTER	
I. INTRODUCTION	1
Exemplar Models and the Role of Background Knowledge	3
Guided Learning From Examples: Integrating Background Knowledge and Exemplar Similarity	11
Rationale for Experiments	14
II. EXPERIMENT 1: A CONTROL STUDY	21
Method	24
Results	29
Discussion	34
III. EXPERIMENT 2: THE EFFECTS OF BACKGROUND KNOWLEDGE AND CONCEPTUAL STRUCTURE ON CATEGORY LEARNING	36
Method	41
Results	49
Model Fits	58
Discussion	65
REFERENCES	82
APPENDICES	
A. EXTENDED LITERATURE REVIEW	86
B. SHORT TEXT CONTAINING BACKGROUND INFORMATION USED IN EXPERIMENT 1	107

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
ABSTRACT	v
LIST OF TABLES	vii
CHAPTER	
I. INTRODUCTION	1
Exemplar Models and the Role of Background Knowledge	3
Guided Learning From Examples: Integrating Background Knowledge and Exemplar Similarity	11
Rationale for Experiments	14
II. EXPERIMENT 1: A CONTROL STUDY	21
Method	24
Results	29
Discussion	34
III. EXPERIMENT 2: THE EFFECTS OF BACKGROUND KNOWLEDGE AND CONCEPTUAL STRUCTURE ON CATEGORY LEARNING	36
Method	41
Results	49
Model Fits	58
Discussion	65
REFERENCES	82
APPENDICES	
A. EXTENDED LITERATURE REVIEW	86
B. SHORT TEXT CONTAINING BACKGROUND INFORMATION USED IN EXPERIMENT 1	107

C.	LONG TEXT CONTAINING BACKGROUND INFORMATION USED IN EXPERIMENTS 1 AND 2	109
D.	COMPREHENSION QUESTIONS AND CORRESPONDING ANSWERS (IN PARENTHESES) FOR SHORT AND LONG TEXTS	115

ABSTRACT

Two experiments were conducted in order to determine whether background information acquired by reading from text differentially influences category learning relative to when no background information is provided. Experiment 1 was a control study that compared short and long versions of text containing information that describe the characteristics of different plant features (e.g., roots, stems, leaves, and flowers) and how each is able to adapt to the characteristics of desert and mountain environments. Seventy-two participants sorted eight drawings of plants into two categories (desert and mountain plants) and read either a short list or a longer, more elaborate text describing the characteristics of plant features. Then after reading, they answered comprehension questions over the text until they mastered the information, and then they re-sorted the plants again. The results indicated that learners applied what they had read when re-sorting as evidenced by fewer errors relative to initial sort patterns. Experiment 2 compared the learning of linearly separable and non-linearly separable concepts for groups of participants that either received no background information (no text), read background information from a text to a high level of mastery (comprehensive text), or read at their own discretion prior to learning (available text). After participants completed one error free run through the eight training stimuli, they classified old training items and eight new transfer items. The results indicate that requiring participants to fully comprehend the text (the comprehensive text group) facilitated learning of the linearly separable concepts, but not non-linearly separable concepts. This finding is consistent with the assumption that learning is enhanced when the items to be

learned do not violate background knowledge, as was the case for the comprehensive text group who learned the linearly separable concepts. In addition, transfer performance for the available and comprehensive text groups was driven more on what they had read relative to how similar the items were to past examples. However, exemplar similarity was predominantly used by all groups after learning non-linearly separable concepts. These findings were interpreted as supporting a mixed representational model that accounts for both exemplar similarity and background knowledge.

LIST OF TABLES

1.	Sample classification problem.	6
2.	Feature notation, category assignments, and observed response proportions after the first sorting opportunity for stimuli used in Experiment 1.	30
3.	Mean number correct and standard deviations (in parentheses) for initial and final sorting as a function of text length and text set.	32
4.	Feature coding and estimated classification probabilities for the training and transfer items for the linearly separable (LS) category structure used in Experiment 2.	39
5.	Feature coding and estimated classification probabilities for the training and transfer items for the non-linearly separable (NLS) category structure used in Experiment 2.	40
6.	Mean number of trials to reach learning criterion as a function of category structure and text condition.	50
7.	Mean proportion correct based on text information for critical transfer items as a function of category structure and text condition.	52
8.	Mean ratings of features prior to training (after training in parentheses) as a function of question type, category structure, and text condition.	55
9.	Mean response latencies (in seconds) during transfer as a function of category structure and text condition.	57
10.	Mean transfer classification probabilities as a function of text condition for the linearly separable (LS) structure.	61
11.	Mean transfer classification probabilities as a function of text condition for the non-linearly separable (NLS) structure.	62
12.	Model fits to all test items and the residual errors for critical transfer items. . .	63
13.	Values of T parameters for the two mixed representational models.	64

CHAPTER I

INTRODUCTION

The past several decades of research on category learning and processing have produced several distinct approaches to accounting for category learning and transfer. One approach emphasized the weighting and summing of item features (Beach, 1964; Reed, 1972; Rosch & Mervis, 1975) or distance from a prototype (Posner & Keele, 1968). Another approach emphasized the storage of items in memory and subsequent classification based on the within- and between-category similarity of items to one another, with similarity itself characterized by an exponential metric based on feature overlap (Medin & Schaffer, 1978; Estes, 1986; Kruschke, 1992).

These approaches, and others, like the explanation-based approaches of Mooney (1993) and Wisniewski and Medin (1991) have generally been tested as mutually-exclusive possibilities for explaining categorization. Other evidence from studies that have investigated the effects of background knowledge, however (e.g., Allen & Brooks, 1991; Hayes & Taplin, 1995; Nakamura, 1985) have suggested that these models might not be mutually exclusive. In particular, these researchers examined whether providing background knowledge to participants (or using category labels that activated pre-existing knowledge) produced fewer errors during learning compared to groups who did not receive background information. In these studies, participants were either supplied with background knowledge in the form of explicit rules (Allen & Brooks, 1991), lists of feature characteristics (Nakamura, 1985), or by using stimuli with features that activated pre-existing knowledge when given meaningful category labels (Hayes & Taplin, 1995). The importance of examining the effects of background knowledge in these contexts sets the stage for contrasting the relative contributions of both similarity- and knowledge-based processes. The present research will further test the need for positing more than one

model (strategy) for category learning and transfer or whether a single model is sufficient to account for learning and transfer within several distinct learning contexts. It will also address several limitations of the findings reported by Nakamura (1985) and Hayes and Taplin (1995), particularly with respect to how background knowledge is implemented during categorization.

The present research will attempt to determine whether background information acquired from reading text differentially influences category learning relative to when no background information is provided. One purpose of this research is to examine the usefulness of text as a means of supplying learners with background knowledge when sorting instances (Experiment 1). The major goal was to examine the relative contributions of background knowledge in learning new concepts (Experiment 2). These will be accomplished by comparing learning rates between two sets of training stimuli, which are drawings of fictitious plants, and transfer performance for new stimuli and old training items. The motivation for comparing classification performance between these two sets is that they differ with regard to their consistency with background information provided to learners prior to learning. One training set includes items that are each consistent with the background information stated in the text and another set includes training items that violate background knowledge (i.e., some of the training items have features that match the background information pertaining to the contrasting category). Classification patterns observed for transfer items will uncover whether categorization decisions are driven more by exemplar similarity, or by applying background knowledge since both of these sources of information make opposite predictions regarding category membership for some of the transfer items.

A second goal is to compare two types of categorization models that differ with respect to how category information is used by comparing model predictions against the data obtained from human participants. One type of model is a purely exemplar-based

model and the other type of model is an additive features model. A third possibility is an extended exemplar-based model that combines background knowledge with exemplar similarity. Recent investigations have argued in favor of models that address both similarity-based processes and background knowledge (Hayes & Taplin, 1995; Nakamura, 1985). However, the way in which background knowledge and similarity are incorporated requires clarification. Thus, comparisons between these models will clarify whether classification behavior is best understood according to the principles of exemplar similarity, additive features, or a combination of exemplar similarity and additive features.

This chapter is organized as follows. First, an exemplar-based model of categorization (e.g., the context model of Medin and Schaffer, 1978) will be discussed with particular emphasis placed on the role that background knowledge plays during categorization. This will be followed by a discussion of how background knowledge has been investigated to test different theories of category learning. Finally, the rationale for the current experiments will be presented.

Exemplar Models and the Role of Background Knowledge

According to similarity-based views of category learning, categories are represented as either a summary representation (e.g., prototypes: Posner & Keele, 1968) or as a set of past examples (Estes, 1986; Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1986; Taraban & Palacios, 1993) that are stored in memory. Categorization stems from a matching process that operates on a feature by feature basis where the features of a probe item are compared against the features that make up category prototypes or features that make up items represented in long term memory. The similarity between two items increases exponentially as the number of common features between items increases. In the case of prototype models, items cohere to a category

when the features present in the item match those of the category prototypes (see also Smith & Medin, 1981).

One class of categorization models that has been mentioned already includes models that operate according to the principles of exemplar similarity. These models have gained favor over prototype models because they have been highly successful in explaining a wide range of classification behavior (Murphy & Medin, 1985). In particular, exemplar-based models do not assume that concepts must be linearly separable in order for them to be learnable. As indicated above, research by Medin and Schwanenflugel (1981) indicated that non-linearly separable concepts were actually easier to learn than linearly separable ones, a finding that is easily accommodated by exemplar models. In addition, exemplar models do not assume that over the course of learning, prototypes, or the “central tendencies” of concepts, become abstracted and serve as the basis for classifying new items during classification. Rather, over the course of learning, individual items are stored in memory along with their associated category labels.

The earliest exemplar model was developed by Medin and Schaffer (1978) and is referred to as the context model. Other derivatives of the context model include Nosofsky’s (1986) generalized context model, Estes’ (1986) array model, Kruschke’s (1992) ALCOVE model, and the exemplar-based back propagation model by Taraban and Palacios (1993). The underlying assumptions in all of these models is that categorization results from comparing presented items to a set of items stored in memory based on the features present in the probe item and those present in the items stored in memory. According to exemplar-based models, during the early stages of learning, each item is stored in memory along with its corresponding category label. Thereafter, when items are presented for classification, the learner computes the similarity of the presented exemplar to each of the category items stored in memory (which is described in more

detail below), sums its similarity to all members associated with each category, computes the probability of each category, and generates a response based on these probabilities.

To give a concrete example of this process, consider the following simple category structure listed in Table 1. It is adapted from Estes (1993) and consists of two exemplars in each category. The two members of each category vary on two binary valued stimulus dimensions (e.g., color: red or black; and shape: circle or square). Each exemplar is represented in memory according to the values of the features it possesses, where color values are represented on the first stimulus dimension (f1) and shape is represented as values on the second stimulus dimension (f2). The method for computing similarity between two exemplars begins with a feature by feature comparison using Equation 1 (Medin & Schaffer, 1978).

$$\text{Similarity}(A1, A2) = \prod s_i \tag{1}$$

In this example, Item A1 is compared against Item A2. When the features values of dimension i for Items A1 and A2 match, $s_i = 1$, but when they do not match, a value of s_i is entered, where $0 \leq s_i \leq 1$. The values of s from Equation 1 are then combined multiplicatively for each item. Thus, the similarity between two items drops off exponentially as the number of mismatching feature values increases. When exemplar A1 is compared against itself, it matches on both features and a value of 1 is entered (e.g., $1 \times 1 = 1$). When compared against the other item in Category A (A2), it mismatches on one of the features, and its similarity is computed as $1 \times s$, or s . To compute the similarity of item A1 to items stored in Category A, the products of each comparison between item A1 and all items in Category A are summed. The similarity between item A1 and those from Category B is computed in the same way. Item A1 matches on only one feature when compared to item B1 (e.g., $s \times 1 = s$), and it matches none of the

features in item B2 (e.g., $\underline{s} \times \underline{s} = \underline{s}^2$). Thus, the similarity to Category B is $\underline{s} + \underline{s}^2$. The probability of classifying item A1 as a member of Category A is determined by summing its similarity to all Category A items and dividing this product by its summed similarity to items from both categories using Equation 2 (Medin & Schaffer, 1978).

$$P(A | a1) = \frac{\sum_{a \in A} \text{Similarity}(a1, A)}{\sum_{a \in A} \text{Similarity}(a1, A) + \sum_{b \in B} \text{Similarity}(a1, B)} \quad (2)$$

When $\underline{s} = 0$, the probability of classifying item A1 as a member of Category A = 1.0, but when $\underline{s} = 1$, the probability of classifying the item as a member of Category A is at chance (0.5). Thus, the values of the similarity parameter also serves as an index for determining feature salience. In most situations, the values of \underline{s} are estimated from the data in order to minimize the error between the model predictions and observed outcomes.

Exemplar models at present have no mechanism to accommodate background knowledge, other than memory for past examples. Wisniewski and Medin (1994)

Table 1

Sample classification problem.

Item	Feature Values	f1	f2	Category	Similarity to all Category A items	Similarity to all Category B items
A1	red circle	0	0	A	$1 + s$	$s + s^2$
A2	red square	0	1	A	$1 + s$	$s + s^2$
B1	black circle	1	0	B	$s + s^2$	$1 + s$
B2	black square	1	1	B	$s + s^2$	$1 + s$

indicate that the effects of background knowledge could also be incorporated in exemplar models by adjusting feature salience prior to learning, or by limiting the set of features that are to be considered. This operates by adjusting the values of the similarity parameters accordingly, as if some degree of learning had already taken place. Similarly, background knowledge could also limit the number of features that are considered in the exemplar model by setting the similarity values for irrelevant features to one, which essentially eliminates irrelevant features in the model. Currently, the exemplar model considers all features that are present in the stimuli when computing similarity. Only over the course of learning will irrelevant features be eliminated, if it is determined that they play no role in determining category membership. At present, there is no mechanism in exemplar models that allows them to adjust feature salience prior to learning or to restrict the set of features that are considered.

Another difficulty with these possibilities is that it does not allow background knowledge to be modified in light of new evidence, or after some degree of learning over items has taken place. In order to accommodate the effects of background knowledge within exemplar models, a mechanism that allows background knowledge to change over the course of learning is needed. In addition, by incorporating a mechanism for background knowledge in exemplar models, it is possible to test whether background knowledge influences classification performance independently of exemplar similarity.

There have been several experiments that directly pitted exemplar similarity against background knowledge. One set of experiments that have pitted the contributions of exemplar similarity against background knowledge were reported by Allen and Brooks (1991). In the experiments reported by Allen and Brooks (1991), one group of participants was initially supplied with a perfectly predictive rule for determining category membership and another group was not given the predictive rule. Providing perfectly predictive rules is equated with situations in which participants are supplied

with appropriate background knowledge prior to learning. All participants then classified a set of drawings of fictitious creatures into categories that corresponded to “builders” and “diggers.” During a transfer phase, participants classified old and new items. Some of the items in the transfer phase were highly similar to one of the training items and in the same category as the similar old item based on the rule given before learning (i.e., positive match items). Other transfer items were highly similar to one of the old training items, but in the opposite category as the similar old item based on the rule given (i.e., negative match items). The results indicated that both experimental groups produced more errors on the negative match items. More errors were observed for the no-rule group on negative match items relative to those given the perfectly predictive rule. What this suggests is that exemplar similarity exerted influence on categorization decisions, even when application of the perfectly predictive rule would have led to perfect classification. This effect persisted even when participants were alerted to the negative match items, and when instructions emphasized accurate responding more than speed. In other words, the effects of exemplar similarity competes with background knowledge when they are in conflict, as evidenced by higher error rates for items that were highly similar to old training items, but inconsistent with the perfectly predictive rule. The implication of this finding for the present research is that exemplar similarity may exert greater influence than background knowledge, particularly when similarity relations between old and new items are strong (i.e., when the features of an old training item and a new test item differ on only one of the five feature dimensions). This issue will be revisited again in a later section.

In related work, Hayes and Taplin (1995) also examined the independent contributions of similarity-based influences (prototype and nearest old exemplar similarity) and background knowledge by comparing categorization performance when participants were given category labels that activated their background knowledge

compared to when they were given nonsense labels. In order to test the contributions of each of these factors, it was necessary to ensure that these factors were statistically independent. In their experiments, exemplar similarity was computed according to the number of mismatching features a probe item had with respect to its most similar (nearest) old training exemplar. If a probe item mismatched on only one feature with respect to its corresponding old training item, a value of 1 was entered as the predictor value. Prototype similarity was computed in a similar manner, except similarity was expressed according to the number of features that deviated from category prototypes. Thus, if a particular probe item possessed two feature values that were different from that of the prototype for that category, a prototype distance of 2 was given for that item. The rationale for computing exemplar and prototype distance in this way was based on the notion that individual cases of high similarity exert more influence in classification decisions than average similarity summed across all cases (Medin & Schaffer, 1978). This was also used because as the number of exemplars in a category is increased, both exemplar and prototype models produce similar predictions with regard to transfer classification performance. However, there is reason to suspect that formulating exemplar similarity in this way may be too simplistic in that it ignores the contribution of other exemplars. That is, for any given item, when exemplar similarity is computed according to only one item, it fails to capture the similarity relations that hold among other items from the same category. For this reason, any conclusions made regarding exemplar similarity may be misleading. The other predictor variable that was considered by Hayes and Taplin was participants' functional knowledge of the stimuli. For each item, an independent group of participants rated the usefulness of each item as a pounding or cutting tool. The average ratings were then used to examine the effects of background knowledge independently from prototype and exemplar distance.

The results obtained for transfer classification performance (Hayes & Taplin, 1995, Experiment 2) indicated that those given category labels that activated pre-existing knowledge (e.g., pounders or cutters) based categorizations decisions on functional knowledge to a greater extent than those given nonsense labels (e.g., kupod or davit). Functional knowledge accounted for more variance than prototype distance and nearest old exemplar distance for the those who were given meaningful category labels. However, the influence of nearest old exemplar distance accounted for less variance in classification accuracy than prototype distance for those given nonsense labels. Even though the way in which exemplar similarity was used by Hayes and Taplin should be used with caution, the results mirror previous findings in that categorization is strongly influenced by information that is not observable from the stimuli, but inferred from pre-existing knowledge.

At this stage, it is apparent that the effects of background knowledge have profound influences on category learning. In sum, background knowledge facilitates learning when the concept to be learned does not include members that violate background knowledge, and it impedes learning when the concept includes items that violate background knowledge. The studies put forward up to this point have provided preliminary evidence that both similarity-based information and background knowledge are used during learning, and that both of these sources tap different types of conceptual information. It remains to be determined is whether supplying background information in the form of text produces the same effects as those observed by Nakamura (1985), Allen and Brooks (1991), and Hayes and Taplin (1995).

Guided Learning from Examples: Integrating Background Knowledge and Exemplar Similarity

Guided learning from examples represents a naturalistic way in which one can acquire information about categories. As in direct instruction, it first requires that background information be provided to participants followed by a learning phase over several examples of the category (e.g., Nakamura, 1985). Intuitively, guided learning represents a more naturalistic learning paradigm than methods that are purely inductive or purely direct. This is because learning includes both uncovering relevant facts that disambiguate category membership and experience with the members that make up categories. Typically, a set of facts that link features to categories are provided before learning over several examples and the effects of these facts are contrasted against a group who does not receive such facts. The information is then used to selectively weight features that are relevant for classification, identify abstract features from concrete features, or to provide a justification for why certain members are members of one category and not another (Medin, 1989; Wisniewski & Medin, 1991, 1994).

Nakamura (1985) conducted one of the earlier studies that incorporated guided learning. Categorization performance using linearly or non-linearly separable category structures was compared between groups of participants that either received or did not receive background information. Linearly separable category structures are those that can be correctly partitioned according to an additive weighting of feature values whereas non-linearly separable category structures cannot be partitioned based on summed feature weights (Medin & Schwanenflugel, 1981). The rationale for comparing linearly and non-linearly separable category structures was to uncover whether learning operated by summing feature weights, which would produce a learning advantage for linearly separable categories over non-linearly separable ones, or whether learning operated by the multiplicative combining of feature weights, which would result in better

performance for non-linearly separable categories than linearly separable ones.

Sometimes the two models make the same predictions, sometimes they make contrasting predictions, especially for specific items. However, if learners are provided with background information that is biased toward linearly separable categories, then these categories should be easier to learn than non-linearly separable categories.

In the experiment reported by Nakamura (1985), participants given background information (or theory instructions) were instructed that one kind of flower attracted birds whereas another type of flower attracted bees. They were also told that bees can see color, like to land, can smell, and are active at night (which correspond to flowers with the features bright in color, wide petals, fragrant odor, and open at night, respectively), and that birds cannot see color, like to hover, cannot smell, and are active during the day (which correspond to flowers with the features dull in color, narrow petals, odorless, and open during the day, respectively). Another group of participants (the standard learning group) was given no additional information. Nakamura reported that the linearly separable category was easier for those given theory instructions to learn than those given standard instructions, but the opposite pattern of results was obtained when non-linearly separable structures were given. For non-linearly separable categories, possession of background information disrupted learning to a greater extent than those given standard instructions.

Since the background information used by Nakamura (1985) was more closely matched with the structure of the linearly separable categories than the non-linearly separable ones, it is not surprising that an advantage in learning linearly separable categories over non-linearly separable ones was observed. This finding is interesting with respect to other studies which have shown no learning advantage for linearly separable concepts over non-linearly separable ones (e.g., Wattenmaker, Dewey, Murphy, & Medin, 1986). The reason put forward by Nakamura for the advantage of

linearly separable concepts is that for those given non-linearly separable structures to learn, exactly two of the training items (one from each category) fit the information pertaining to the contrasting category than to the correct category. As a result, whenever participants applied their background knowledge (e.g., that bees are attracted to flowers that are open during the night), and received corrective feedback (e.g., that a flower that is open during the day attracts bees that are active during the night), they were more likely to incorrectly classify the exception items, since they were inconsistent with what they knew about the habits of birds and bees. Although this finding suggests that background knowledge and category structure are closely linked, it did not address whether participants continued to apply their background knowledge even when it led to errors of classification. In addition, there was no mention of whether similarity, background knowledge, or some interaction of these operates during transfer. Inclusion of transfer items that would differentiate these possibilities was also ignored. These issues will be more fully addressed in a later section.

Subsequent research indicates that one possible reason for the disruptive effects of inconsistent background knowledge is that it continues to place more weight on irrelevant features even when feedback after classification suggests the opposite. For example, Livingston and Andrews (1995) observed that when background information supplied prior to learning initially made salient features that were not diagnostic of category membership, these features were not replaced by new, diagnostic features over the course of learning. Rather, learners continued to consider the irrelevant features in addition to newly uncovered diagnostic features that were learned through experience with training items. This initial bias to consider irrelevant features persists because occasionally, such irrelevant features prove to be observed in members of the correct category. This finding further suggests that learning impediments occur when one's prior knowledge emphasizes irrelevant features, but through experience with category examples,

participants gradually incorporated the relevant features into their theories. The implication of this finding is similar to that reported by Nakamura (1985), in that people's background knowledge represents a potential constraint on learning as do the items to be learned. These results also provide a preliminary indication as to the importance of background knowledge during categorization, and indicate that both exemplar-similarity and background knowledge need to be implemented in a unified model of categorization.

Rationale for Experiments

The present research has two objectives. The first objective is to examine categorization performance when learners have access to domain-specific background information, specifically when the information is acquired by reading from text. This is motivated by recent findings which suggest that utilization of background information interacts with a similarity-based component during learning (e.g., Hayes & Taplin, 1995; Wisniewski & Medin, 1994). However, when the effects of background knowledge are investigated, it is typically presumed that participants possess relevant background knowledge prior to experimentation, as evidenced whenever meaningful category labels that activate pre-existing knowledge are used to test the effects of background knowledge (Hayes & Taplin, 1995). By comparing learning and transfer performance between groups of participants that have different levels of access to background information, it is possible to test the influence of previously acquired knowledge during categorization. In addition, by incorporating text as a means for acquiring background information, the current experiments may also provide a baseline for establishing how certain types of background information are acquired and used during learning, an endeavor that has not received much empirical attention within the context of categorization.

The second objective of the present research is to contrast two broad theories of category learning, those that are based purely on exemplar similarity and those that include a mechanism for incorporating background knowledge and exemplar similarity. Specifically, the model predictions of a purely exemplar-based model, expanded exemplar-based models that accommodate background knowledge, and an additive features model will be fitted to human data in order to determine which model best captures human categorization performance. In addition, the data will also be fitted to the mixed representational model used by Nakamura (1985).

As indicated above, Nakamura (1985) was able to demonstrate a learning advantage for those given background information, but only for linearly separable categories. To model this effect, Nakamura used a mixed representational model that included an exemplar component and a prototype component, where the latter component corresponds to the background information given to participants prior to learning the items. The method by which the exemplar component was computed for fitting the classification probabilities was identical to that used by Medin and Schaffer (1978). In addition, the prototype component was calculated in a similar manner. That is, the model treated the prototype as if it were a focal exemplar (Estes, 1994). Thus, categorization probabilities were computed by weighting both similarity to past examples and similarity to a prototype, where prototype similarity was computed using a multiplicative similarity rule. The primary difficulty with this formulation is that the information provided to participants prior to the training phase emphasized additive combinations of features, which was not captured by the prototype component of the model. Thus, in its present formulation, the mixed representational model of Nakamura treats the background information given to participants as nothing more than an instance of a prototype. Close examination of the model fits were as predicted for the linearly separable categories. The influence of the prototype component was greater for the theory group than the standard

group. However, for the non-linearly separable data, the prototype component was also greater than zero for both instructional groups. Taken at face value, this indicates that participants used the background knowledge to a greater degree than exemplar similarity when learning the non-linearly separable concepts. But in reality, the mixed model tested by Nakamura represented the prototype twice. It was included in the exemplar component since the items that correspond to the prototypes were included as training items, and it was also included in the prototype component since it contained all the features that were provided as background information prior to learning. Thus, with these problems of the model in mind, it is less clear whether the mixed exemplar-prototype model tested by Nakamura accurately captured the nature of the background information that was given prior to learning. What is needed (at a minimum) is a model that more closely separates exemplar similarity from an additive feature solution, which was not captured by the model tested by Nakamura.

Experiment 1 represents a control study designed to uncover whether the texts that will be used in Experiment 2 to supply learners with background information pertaining to the characteristics of desert and mountain environments, and the adaptations of the different plant features with regard to the environments, are learnable. This will be assessed by having participants answer several questions over the material stated in the text. The experiment will also examine whether any of the training items to be used in Experiment 2 are biased toward any particular plant category based on the how the items are initially sorted into categories. In addition, the effectiveness of the text as a means of supplying learners with background information will also be examined by observing sorting patterns after participants are able to answer questions over the text to a high level of mastery. That is, sorting patterns observed after participants read the texts will be compared against sort patterns observed prior to learning the text. If participants are able to infer from the text that an item must possess two out of three critical features (which is

not directly stated in the text), then sorting patterns after reading the text should include a higher proportion of sorts based on the two out of three feature (additive) rule compared to initial sort patterns.

In order to more fully uncover the relative contributions of exemplar- and knowledge-based influences. Experiment 2 will compare categorization performance using two different category structures in which within- and between-category similarity is systematically varied across training items. The purpose is to examine whether possession of background information (acquired from text) attenuates the effects of exemplar similarity relative to those who are not supplied with background information prior to learning the items. It may also further identify the conditions under which background knowledge is applied during categorization. For example, when background knowledge indicates that certain features or feature combinations are specific to one category, this knowledge should be applied whenever it leads to correct classifications. However, when feedback after classification indicates that category-relevant features or feature combinations are not congruent with one's background knowledge, the background knowledge may be modified to account for the new information, or it may be discounted altogether and replaced by memorization of the items (Wisniewski & Medin, 1994). Examination of transfer classification patterns for critical transfer stimuli and participants' ratings of featural importance will test these possibilities. Since the present research looks closely at the influence of exemplar similarity and background knowledge on critical transfer items, it will address the shortcomings of the research reported by Nakamura (1985), and Hayes and Taplin (1995). In these studies, there were no diagnostic items that were tested during transfer. Thus, the current investigation may provide a more strict test of when transfer is driven by similarity to past examples, background knowledge, or both.

One of the category structures that will be used corresponds to a linearly separable (LS) structure and was constructed according to an additive rule similar to that used by Allen and Brooks (1991). Another way to consider this type of structure is in terms of disjunctions of conjunctions (e.g., IF an item has [f1 AND f2], OR [f1 AND f3], OR [f2 AND f3], THEN it is a DESERT PLANT, where f1, f2, and f3 correspond to specific feature values like tap roots, woody stems, and bladed leaves, respectively). This type of additive rule represents the simplest possible rule involving three of the four features used to construct the items. The second type of category structure used corresponds to a non-linearly separable structure and was also constructed according to the above additive rule, but one of the training items from each category represents an “exception” to the additive rule. That is, the “exception” items include features that are related to survival in one environment (based on the text), but the actual category label given to it places it in the opposite category (relative to the text). By incorporating this type of category structure, it is possible to determine whether learning is disrupted when some of the training items from one category possess features that are suited for survival in the opposite category. In addition, it will determine whether such inconsistencies force learners to abandon the background information that was provided to them. For example, when background information is acquired prior to category learning, this may initially cause learners to classify the exceptions in the contrasting category. Because of this, learners may interpret that the text does not apply to all category members, and as a result, the information from the text may be disregarded during learning and transfer. No recent studies have directly examined this possibility with the type of text materials that are examined in the present experiment. What evidence that is available indicates that learners are sensitive to previous examples of a category even when they are given a perfectly predictive rule beforehand (Allen & Brooks, 1991). Performance on critical transfer items will indicate whether categorization is driven more by exemplar similarity

or whether it is driven by applying background knowledge, or both. In sum, the experiments will attempt to answer the following questions:

1. To what degree does having access to background knowledge facilitate the learning of linearly separable categories, as measured by the number of training blocks to reach learning criterion for each of the different groups?
2. To what degree does having access to background knowledge disrupt the learning of non-linearly separable categories, as measured by the number of training blocks to reach learning criterion for each of the different groups?
3. Is transfer performance best explained by the principles of exemplar similarity, additive features, or by a combination of additive features and exemplar similarity, based on the classification patterns for diagnostic items and overall fits of quantitative models to all transfer items?
4. To what extent does background knowledge change over the course of learning as measured by ratings given prior to and after training?

Experiment 1 will represent a control study designed to identify potentially biased stimuli and will also compare the learning of text to be used in Experiment 2. Experiment 2 will provide an answer to the four questions listed above by examining the role of background knowledge, as it is acquired by reading from text, on learning of linearly separable and non-linearly separable concepts. It will also contrast purely exemplar based models of category learning, additive feature models, and models that include mechanisms for accommodating background knowledge (e.g., the mixed representational model used by Nakamura, and a mixed representational model that includes an additive feature component). This will be carried out, in part, by fitting the transfer classification performance observed in human participants to four quantitative models. If it is observed that a mixed representational model with additive features produces better overall fits to the human data, particularly with respect to the text conditions, this would lend credence

to the assumption that a purely exemplar based model fails to capture the effects of background knowledge. Experiment 2 will also determine whether participants modify their reliance on the background information they had received by comparing ratings over the importance of features, which will be taken prior to a learning phase and immediately after it. If participants estimate that the information that they had read prior to learning does not apply to all training items (as in the non-linearly separable conditions), they should produce lower ratings for those features that were rated as being highly important after the training phase. If this is the case, participants should also give higher ratings to features that they initially rated as being not important after the training phase.

CHAPTER II

EXPERIMENT 1: A CONTROL STUDY

The available research up to this point indicates that the application of relevant background knowledge during categorization produces increases in learning, as evidenced by some reduction on the number of trials to reach some learning criterion, relative to when no background information is provided (Heit, 1997). However, the conditions under which acquisition of background knowledge can be transferred to categorization performance after it is acquired by reading from text remains unknown. In most studies, background knowledge of a particular domain is presumed (e.g., Pazzani, 1991; Spalding & Murphy, 1996; Wisniewski & Medin, 1994), or for rule-based categories, the classification rule is made explicit prior to learning (Allen & Brooks, 1991). In other studies, background information is presented prior to learning as a simple list describing the characteristics of the features present in items to be learned (Nakamura, 1985). However, in many natural learning situations, learners are required to form concepts by integrating information they had read from a larger, more complex text. In addition, this text is rarely expressed as a list or as simple rules, but as information that specifies why a given property present in category members supports category membership. One example of this would be learning about the adaptations of plants from a textbook where not only the functions or characteristics of relevant features are presented, but are presented in conjunction with explanations why certain attributes support category membership.

The primary goal of the present experiment is to compare sorting patterns of plants when background information is acquired either by reading from a list or by reading from a larger, more integrated text. The primary assumption to be tested is that possession of relevant background information (acquired from reading both short and

long texts) will be used by participants to uncover the functional relationships between the features making up the stimuli and the appropriate categories. Application of the background information should therefore lead participants to selectively attend to the relevant features and ignore those that are not relevant. Likewise, understanding the relevant background information may also force participants to consider more than one feature at a time when making classification decisions, as is the case when sorting according to family resemblance (Ahn & Medin, 1992).

To get a clearer picture as to how the text relates to the categories, the text indicates that each of the plant features (with the exception of flowers) is able to accommodate two of the three demands of desert or mountain environments. For example, the text listed under Set A in Appendix B indicates that (a) extremely hot temperatures, (b) extended periods without rain, and (c) intense sunlight are the three properties of desert climates. If a plant is to survive in a desert climate it must be able to adapt to these three conditions. Regarding the information pertaining to one of the plant features (e.g., tap roots), the text states that tap roots help plants survive in extreme heat and extended periods without rain, a condition that satisfies two of the three environmental demands for desert plants. Thus, in order for a given plant to be classified as a desert plant, it must also possess a feature that also helps the plant tolerate or adapt to intense sunlight. This means that in order for an item to be a desert plant, it must also have a woody stem or bladed leaves (or both).

Inspection of the text reveals that the information pertaining to desert and mountain environments is not different from what these actual environments are like. However, the information pertaining to how each of the various plant features are designed to adapt to environments represents new information that participants do not know since these adaptations do not correspond to feature adaptations that occur in the real world. Rather, such information represents adaptations that are plausible for the

features. The information pertaining to the functions of each of the plant features are designed so that it is possible to link the features to the applicable environments. The exception to this applies to the feature flowers, which the text indicates it possesses adaptations that are not related to adaptation in either environment.

It is expected that sorting patterns after reading both short and long text will include fewer single dimension sorts compared to the first sorting opportunity. Thus, if participants are able to acquire the information from the texts, it is predicted that they will use this when sorting items. This prediction will be supported if it is frequently observed that sorting patterns after reading are made according to an additive rule than that observed prior to reading text, or if the number of single dimension sorts are reduced relative to the first sorting opportunity. Such a finding will also support recent findings which indicate that family resemblance sorting patterns are more likely to occur when relevant background information can be applied to the items relative to when no background knowledge can be applied, particularly when the features that make up the stimuli are integrated to a common theme (Spalding & Murphy, 1996).

A second goal of the experiment is more normative in nature. Specifically, it is designed to uncover stimuli that are potentially biased to one of the plant categories in addition to determining whether the background information provided by both short and long versions of the text is learnable. This latter issue will be clarified by comparing performance on comprehension questions over two versions of the text, whereas the former will be revealed by sorting patterns on the first sorting opportunity. If it is observed that the binomial probability of sorting a given item into one of the categories is less than .05, this would indicate that such an item is biased to one of the categories. If this occurs, then any biased items would have to be replaced by items that are not biased toward one of the categories if it is to be used in the subsequent experiment. The rationale for comparing short and long texts is to uncover whether sorting patterns during

the second sorting opportunity differ with regard to the type of text used. In addition, comparisons between the two texts will also reveal potential learning differences between short and long texts, as evidenced by the number of attempts to answer comprehension questions over the texts to criterion.

Method

Participants

There were 72 volunteers enrolled in introductory psychology courses at Texas Tech University that served as participants in the experiment. All participants received partial course credit for their participation.

Materials

Texts. The short and long texts that were used in the experiment are listed in Appendix B and C, respectively. The readability of each of version of the text (Sets A and B) for both short and long texts were computed according to the Flesch scale and the Flesch-Kincaid grade level. Readability according to the Flesch scale is computed according to the average number of syllables per word and the average number of words per sentence. Higher readability scores reflect text that is more understandable relative to texts that have low readability scores. Readability according to the Flesch-Kincaid grade level is computed the same as that for the Flesch scale, but scores indicate a grade school reading level. Both sets of the short text are 235 words in length and have a reading difficulty level of 66.2 according to the Flesch scale with a Flesch-Kincaid grade level of 7.0. For the long texts, Set A was 1,035 words in length and has a Flesch score of 63.3, and a grade level of 9.0, and Set B was 1,036 words in length with a Flesch score of 62.8 and a grade level of 9.2.

Both short and long texts were printed on cards and each included illustrations of each of the eight features that occurred in the actual stimuli centered above the text. Each of the eight features were drawn in their appropriate locations as they appeared in the actual stimuli, but were drawn on stencil drawings of generic plants so that the illustrations did not correspond to any of the target stimuli. The reason for this was to prevent potential priming effects that could be caused by including actual stimuli, while at the same time, providing participants with an opportunity to see what each of the features stated in the text refer to in the actual stimuli.

The text itself contains information that describes the characteristics of desert and mountain environments in addition to a description of the functions and capabilities of each of the plant features. The information pertaining to the functions of each of the plant features is designed so that it is possible to link the features to the applicable environments, with the exception of the feature flowers in which the text indicates that it possesses adaptations that are exclusive to neither environment. Inspection of both the short and long texts reveals that the information pertaining to desert and mountain environments is not different from what these actual environments are like, and this information is not different from what people already know about these environments. However, the information pertaining to how each of the various plant features are designed to adapt to environments represents new information that participants do not know. For the short version of the text, this information is simply listed next to the labels of each of the environments and plant features. No supplemental information or further explanations are given as to why each of the features are capable of adapting to various environmental demands (see Appendix B). In contrast, the long version of the text contains supplemental information for each of the features in addition to explanations as to why the features are able to adapt to their environments. This is designed to provide learners with a more rich set of information relative to the short version. For example,

the supplemental information about roots (from Set A) indicates that they anchor the plant in order to prevent it from being uprooted and that they absorb moisture and nutrients from the soil. This information, although supplemental, is applicable to both tap and fibrous roots but is not mentioned in the short versions of the text. In addition, more specific information pertaining to each of the different types of features provide an explanation as to why they are able to adapt to environments. For example, the information pertaining to bladed leaves in the long text indicates that they protect the plant from extreme heat by curling up, and since they have a small surface area, they require a large supply of sunlight in order to synthesize food for the plant. In the short version, the text merely states that this type of leaf help plants survive in extreme heat and intense sun.

Probe Questions. Eleven questions and were printed on sheets of paper and are listed in Appendix D. These questions are designed to uncover whether participants were able to learn the information that was stated in the text. Each question requires participants to identify either the characteristics of desert and mountain environments, or the specific adaptations of each of the different types of plant features. The purpose of which is to ensure that all participants are able to fully comprehend the information from the text before proceeding to the final sorting task. For the 11 questions, there were 26 possible correct responses, since there is more than one correct answer for each question. For example, the first question requires four responses, the second and third questions require three responses each, and the remaining eight questions require two responses each (see Appendix D for the correct responses to each of the questions).

Stimuli. The stimuli include eight black and white drawings of plants that vary on four binary valued stimulus dimensions printed on 6 x 6 in. cards. The plants are abstractly represented as 0s and 1s for each of the feature dimensions. The four respective feature dimensions of the stimuli correspond to the type of root (tap [0] or

fibrous [1]), stem (woody [0] or herbaceous [1]), leaves (bladed [0] or compound [1]), and flower (headed [0] or spiked [1]) present for each plant. Thus, a value of 0 for each of the feature dimensions identifies a plant that has a tap root, woody stem, bladed leaves, and a headed flower. Likewise, a plant that possesses a fibrous root, herbaceous stem, compound leaves, and a spiked flower is identified with values of 1 on each of the respective feature dimensions.

Design and Procedure

Participants were tested in groups of up to four and were randomly assigned to receive text from either short or long versions of Set A or Set B. Thus, a 2 (Text: short or long) X 2 (Set: A or B) between-groups design was incorporated. Each of the plant categories was counterbalanced so that any given item was a desert plant for half of the participants, conditional on whether participants were randomly assigned to receive text from Set A or Set B. The experiment proper was conducted in three phases: an initial sorting phase, a learning phase over the text, and a re-sorting phase.

Prior to meeting with the participants, all eight items were randomized. Initially, participants were informed that they would be sorting the eight drawings into two equal sized categories that correspond to desert and mountain plants. Each participant was told that the drawings represent plants from desert and mountain environments. They were also told that the sorting task may seem difficult, but to nevertheless sort each plant into one of the categories in a manner that seems sensible to them. Specifically, all participants were given the following instructions:

This experiment is about how people construct categories. In front of you, there are eight cards and each card has a drawing of a plant on it. I would like you to sort the plants into two categories that correspond to desert and mountain plants. Carefully examine each plant and sort them in a way that is sensible to you. Four

of the plants should be desert plants and four should be mountain plants. After you sort the plants, put them in their appropriately labeled envelope and give them back to me. After this, I will give you a card with some information about environments and plants on it. Please read this information carefully because you will be required to answer questions about it to a high level of mastery before proceeding to the next part of the experiment.

Each participant was then given two empty envelopes and the eight randomly ordered drawings of plants. All participants were given as much time as needed to sort the items.

After doing this, participants returned the envelopes to the experimenter, who then scored the items based on the number that fit the three feature additive rule as determined by the set they received (see below for specific details on how the items were scored). After participants had returned the envelopes, they were then given a card containing either the short or long text from one of the sets, depending on what text and set they had been randomly assigned to receive. Participants were also given as much time as needed to read the text. After they had read the text, participants returned the cards and were then required to answer all probe questions. They wrote their responses to the 11 questions in the 26 spaces that were provided on a sheet of paper containing the questions. They were also allowed as much time as needed to answer questions. Once the participants completed the set of questions, they returned them to the experimenter who then tabulated the number of correct responses to all the probe questions. If the participants did not provide 24 out of 26 possible correct responses, they were given the text they had previously read and were asked to review it. This procedure of reading and testing was repeated for each participant until criterion was met. This was done in order to ensure that all of the information stated in the text was understood before participants were allowed to proceed to the final sorting task.

After reaching criterion, participants were then given the eight drawings of plants once again and sorted them again into one of the two plant categories using the same procedure as that was used during initial sorting. However, before sorting the plants again, they were given the following instructions.

Now, I would like you to sort the plants again into desert and mountain plant categories. As before, you should put four of the plants into the envelope that is labeled desert plants and four in the envelope marked mountain plants. When you are sorting this time, you should use the information you just read to help determine what environment each plant is best suited for. When you have finished sorting, please return the envelopes to me.

The entire procedure took less than one hour to complete.

Results

A 2 (Set: A or B) X 2 (Text: short or long) factorial analysis of variance was conducted using difference scores between the first and last sorting opportunity as the dependent variable. Note that difference scores reflect the number correct during the second sorting opportunity subtracted from the first sorting opportunity. An item was counted as correct only if it was sorted into the category that corresponds to the correct additive rule for that category. This means that for half of the participants, Items 1 through 4 in Table 2 correspond to desert plants and Items 5 through 8 correspond to mountain plants (for those receiving text from Set A). There was a maximum of 8 possible correct for each testing occasion and a minimum of 0 correct. Regarding difference scores, a score of 8 would be obtained only when participants failed to place any of the items in the correct category during the first phase and subsequently placed all of the items in their correct category during the final sorting phase. A score of -8 would be obtained only when all items are initially classified correctly and when all are

Table 2

Feature notation, category assignments, and observed response proportions after the first sorting opportunity for stimuli used in Experiment 1.

Item	Category	Feature dimensions				Response Probability
		f1	f2	f3	f4	
1	A	0	0	1	0	.59
2	A	0	1	0	0	.44
3	A	0	1	0	1	.43
4	A	1	0	0	1	.47
5	B	1	1	0	0	.41
6	B	1	0	1	0	.55
7	B	0	1	1	1	.57
8	B	1	1	1	1	.53

Note. Values of 0 for feature dimensions f1, f2, f3, and f4 correspond to tap roots, woody stems, bladed leaves, and headed flowers, respectively. Values of 1 on for these respective feature dimensions are fibrous roots, herbaceous stems, compound leaves, and spiked flowers. Category assignments of stimuli correspond to Set A text. Numbers beneath Response Probability refer to the observed proportions that each plant was sorted as a desert plant during the first sorting opportunity.

incorrectly classified during final sorting phase. There were no main effects for Set or Text, nor was there a significant interaction between Set and Text (all $F_s < 1$, $p > .05$), a finding which indicates equivalence between sets and texts in terms of improvement in sorting accuracy.

Regarding comprehension of text material, a (Set: A or B) X 2 (Text: short or long) factorial analysis of variance was conducted using number of correct responses to probe questions during the first testing opportunity (i.e., after reading the text once) as the dependent variable. The number of correct responses to probe questions was scored according to whether participants were able to provide responses similar to those in

Appendix D. Thus, verbatim recall of facts exactly as they were stated in the text was not required for correct responses. As in the above analysis, there were no main effects for Set or Text, nor a significant Set X Text interaction (all $F_s < 1.5$, $p > .05$), a finding which also indicates equivalence between sets and texts in terms of question answering accuracy. A similar pattern of results emerged when question answering accuracy was measured during the second testing opportunity (i.e., after reviewing the text once). Note that participants were required to reach a criterion of 24 out of 26 correct responses to questions before proceeding to the final sorting task. Analyses conducted on the second question answering phase also failed to reveal any significant main effects or interactions (all $F_s < 2.1$, $p > .05$). This finding further indicates equivalence between the two sets and between the short and long texts. When the number of attempts to reach criterion were used as the dependent variable, only a significant main effect of Text was observed [$F(1, 68) = 5.36$, $MSE = 0.31$, $p < .05$]. What this finding indicates is that participants required significantly more attempts to reach criterion performance when given the long text ($M = 2.16$, $SD = .50$) relative to the short text ($M = 1.85$, $SD = .61$).

Since it was possible for items to be sorted into both categories during initial sorting, an analysis of binomial response probabilities was conducted in order to uncover those stimuli that are biased to one of the categories. The observed proportions of desert responses are listed in Table 1. None of the items were biased to one of the categories. All binomial probabilities were less than .05 for each item. Other analyses were conducted on the frequency of sorting strategies observed during the first sorting opportunity. These were divided into single-dimension sorting patterns (i.e., sorting based on the values of one of the four feature dimensions) and sorting patterns that did

Table 3

Mean number correct and standard deviations (in parentheses) for initial and final sorting as a function of text length and text set.

Condition	Initial Sort		Final Sort	
Short Text				
Set A	4.00	(1.73)	6.82	(1.02)
Set B	4.12	(1.80)	6.71	(1.21)
Long Text				
Set A	4.00	(2.00)	6.32	(1.38)
Set B	3.68	(1.76)	6.63	(1.17)
Total				
Set A	4.00	(1.86)	6.57	(1.20)
Set B	3.90	(1.78)	6.67	(1.19)

not conform to a single-dimension. A chi-square was performed on the initial sorting patterns and revealed that participants displayed a tendency to consider the feature roots as salient [$\chi^2(5, N = 72) = 91.17, p < .05$]. It should be noted that all participants initially sorted according to a single feature dimension. No participant considered more than one feature during the first sorting opportunity. This pattern changed after reading the texts. A chi-square was also performed on the frequency of sorting patterns after reading the text and revealed that participants were more likely to consider more than one feature when sorting [$\chi^2(5, N = 72) = 34.83, p < .05$], as evidenced by more frequent sorting strategies that were based on additive features than any single dimension. In addition, participants never sorted unidimensionally using flowers as the relevant feature, which is to be expected if participants understand that the background information

pertaining to flowers is specific to neither category. Taken together, these findings indicate that participants were able to apply the information that they had acquired from the texts during the final sorting opportunity.

Further evidence that supports this finding was obtained by contrasting the number of items correctly sorted during the first opportunity with the number of items correctly sorted during the second opportunity. Sorting patterns during both sorting opportunities were scored as follows. First, sorting accuracy was defined according to the number of items sorted into desert and mountain plant categories that matched the additive feature combinations of the target categories. For initial and final sorting accuracy, the number correct was dependent on the set each participant was randomly assigned to receive. Thus, if a participant was assigned to receive text from Set A, items placed in the desert plant category were counted as correct if they possessed at least two of the features tap roots, woody stems, and bladed leaves, and items placed in the mountain plant category were counted as correct if they had at least two of the features fibrous roots, herbaceous stems, and compound leaves. For those assigned to receive text from Set B, items were counted as correct as above, but were given opposite category labels relative to those who were given Set A. For those given the short version of the text, sorting accuracy significantly improved after reading, [$t(33) = 7.37, p < .05$], with a mean difference of 2.71 ($SD = 2.14$) observed between the number of items correctly sorted during the second sorting opportunity subtracted from the number correct after the first sorting. Comparable improvement was also observed for those given the long version of the text, [$t(37) = 7.59, p < .05$], with a mean difference of 2.63 ($SD = 2.14$) observed between the number of items correctly sorted during the second sorting

opportunity subtracted from the number correct after the first sorting opportunity. The mean number correct for each sorting opportunity is listed in Table 3.

Discussion

Overall, the results indicate that both long and short texts are learnable. All participants were able to reach criterion within four attempts, but most required only two attempts. Since the long version of the text produced the same pattern of results as the short text, only the long version of the text will be retained for use in Experiment 2. Incorporating the short text in this experiment was simply designed to uncover any knowledge effects resulting from an annotated version of the long text, if such effects were to exist.

The results also indicate that none of the stimuli (which correspond to training items that were used in Experiment 2) were biased to any of the categories during the initial sorting opportunity. Although participants did consider roots to be a salient feature during initial sorting, they did not consistently place any specific type of root into one of the categories. Given the high salience of this feature, it is reasonable to assume that participants would regard this type of feature as important when differentiating the categories during the initial sorting opportunity. To illustrate a plausible reason for this finding, it is possible that in the initial sorting task participants drew on pre-existing knowledge. For example, one of the characteristics that differentiates desert and mountain environments is the amount of moisture each receives. Because most people know that desert climates tend to be much drier than mountain environments, and because most participants know that the main function of roots is to absorb moisture, it is

likely that participants drew upon this knowledge during initial sorting. As a result, participants were more likely to sort unidimensionally using roots as the salient feature than the other features.

Perhaps the most significant finding of the current experiment is the sorting patterns that were observed during the second sorting opportunity. Previous research has often failed to observe sorting patterns that conform to a family resemblance structure (Ahn & Medin, 1992), where family resemblance represents the type of structure that occurs in most natural categories (Rosch & Mervis, 1975). In the present experiment, participants often considered more than one feature during the second sorting opportunity over any single feature, a finding which indicates that the text was more likely to produce sorting patterns that match the structure of natural categories. One important implication of these findings is how the sorting patterns were influenced by background knowledge. Recent findings indicate that family resemblance sorting patterns are more likely to be observed when the features activate the learner's background knowledge and when they share a common underlying theme that is relevant for the category (Spalding & Murphy, 1996). The reasoning behind this is that background knowledge provides a mechanism for establishing coherence in categories (Murphy & Medin, 1985). In the present experiment, the background knowledge was supplied to participants prior to final sorting, and the effects of the text produced sorting patterns that mirror sorting patterns that have been observed when pre-existing background knowledge was assumed. Because of this, it can be argued that learning from text is capable of building a knowledge base comparable to that in which pre-existing knowledge is assumed.

CHAPTER III

EXPERIMENT 2: THE EFFECTS OF BACKGROUND KNOWLEDGE AND CONCEPTUAL STRUCTURE ON CATEGORY LEARNING

One main goal of Experiment 2 was to determine whether the background knowledge acquired from reading text would produce the same effects that have been observed under conditions where established general knowledge is assumed. It extended Experiment 1 by looking in detail at learning rates to a high level of mastery of the training items. Transfer items were introduced in order to examine generalization patterns to new stimuli. Examination of learning rates was necessary in order to determine whether possession of background information produces fewer trials to reach criterion relative to when no background information was available, particularly for the linearly separable categories. The purpose of which is to replicate earlier findings that have demonstrated a facilitation effect of background knowledge (Hayes & Taplin, 1995; Nakamura, 1985). The experiment also determined whether classification accuracy is disrupted when some training items of a particular category violate background knowledge by comparing learning rates and transfer performance between three experimental groups and over two category structures. One group was required to read text and answer questions over it to a high level of mastery (labeled as the Comprehensive Text group). Another group was simply given the option to review the text but they were not required to answer comprehension questions over the material (labeled as the Available Text group). Finally, a third group (the No Text group) were not given text but simply started with the learning phase over the training items.

Regarding the two different category structures, one category structure included training items that could be correctly partitioned according to an additive rule and the other category structure included items that were exceptions to the additive rule. The

former category structure is labeled as linearly separable (LS) since all training items from each category can be separated by a linear boundary. That is, LS items can be separated by a weighted sum of feature values such that all desert items fall on one side of a linear boundary and all mountain items fall on the other side of a linear boundary. In contrast, the latter category structure is labeled as non-linearly separable (NLS) since there is no linear boundary that separates all desert training items from all mountain items.

Regarding learning performance for the LS categories, one possible outcome was that the background information acquired from the text would be applied by learners during learning, and as a result, it should produce fewer trials to reach criterion relative to participants who did not receive background information. This finding would provide support for Nakamura's (1985) facilitation effects that were observed when background information was provided to participants that biased them to learn LS categories. However, unlike Nakamura, this additive weighting of features was incorporated in a mixed representational model. Recall that Nakamura examined categorization based on similarity to past examples and to prototypes. Unlike other formulations, the prototype component in Nakamura's model was treated as if it were a focal exemplar, and it retained the multiplicative similarity rule for both exemplar and prototype components. For this reason, it treated background knowledge as if it were nothing more than a highly representative exemplar. However, the way in which background information was supplied to the theory group suggests that a more correct approach would be to treat the background knowledge as a special case of additive features.

Regarding learning of NLS categories, it is unknown whether background information acquired from text will be disregarded by participants once it is determined that it does not apply to all items. Previous research has been relatively silent on this issue. In the event that this should occur early in learning, learning rates for training

items should not differ between those who were given text and those who were not given text. A second possibility is that application of background knowledge will continue to mislead participants over the course of several blocks of trials into incorrectly classifying the items that are exceptions to the additive rule. This possibility would be supported if the number of trials to reach criterion during learning are greater than the number of learning trials for those who did not have access to the text. Those who were not given access to the text should make classification decisions based on the surface features of the items, and as a result, they should require fewer learning trials to reach mastery than those who received text.

A transfer phase was also incorporated for the purposes of identifying whether new items were classified using exemplar similarity alone or whether individuals also incorporated background knowledge extracted from text. The transfer stimuli included items that were highly diagnostic for this purpose. Items T4 and T6 in Table 4 and items T4 and T6 in Table 5 should be categorized in the opposite category relative to their similarity to training items if the information acquired from the text (Set A) was correctly applied. Thus, if classification is guided more by application of one's background knowledge, these items should be classified as desert plants because they possess two features that are adapted for survival in desert environments. However, if classification is guided more by exemplar similarity, these items should be classified as mountain plants since they are highly similar (i.e., they differ on only one feature) to more training items from the mountain category than the desert category.

Also presented in Tables 4 and 5 are the predicted category values based on a linear regression where each of the individual feature values for each training item are used as predictor variables. For the LS category structure, the transfer items in Table 4 are separated into the correct categories. Using 0.5 as a boundary, all training items from the desert category have a predicted category value that is less than 0.5 and all mountain

Table 4

Feature coding and estimated classification probabilities for the training and transfer items for the linearly separable (LS) category structure used in Experiment 2.

Item	f1	f2	f3	f4	Similarity		regression	Text
					(s = .10)	(s = .30)		
Training Items								
A1	0	0	1	0	.90	.72	.548	D
A2	0	1	0	0	.91	.76	.619	D
A3	0	1	0	1	.90	.74	.857	D
A4	1	0	0	1	.97	.79	.786	D
B1	1	1	0	0	.89	.70	.167	M
B2	0	1	1	1	.90	.73	.405	M
B3	1	0	1	0	.90	.74	.095	M
B4	1	1	1	1	.98	.86	.048	M
Transfer Items								
T1	0	0	0	0	.91	.78	1.000	D
T2	0	0	0	1	.94	.82	1.238	D
T3	0	0	1	1	.50	.51	.786	D
T4	1	0	0	0	.38	.47	.548	D
T5	1	1	0	1	.50	.49	.404	M
T6	0	1	1	0	.62	.53	.167	M
T7	1	1	1	0	.07	.25	.287	M
T8	1	0	1	1	.34	.39	.332	M

Note. The values for each feature dimension are the same as in Table 2. Numbers under Similarity refer to estimated classification probabilities when $s = .10$, and when $s = .30$. Numbers under regression refer to predicted categories (desert = 1; mountain = 0) based on a regression analysis using weighted sums of feature values as predictor variables. Letters under Text identify items consistent with information about desert (D) or mountain (M) climates, based on the information stated in the text (Set A). For transfer items, all numbers refer to the probability of classifying the items as desert plants.

Table 5

Feature coding and estimated classification probabilities for the training and transfer items for the non-linearly separable (NLS) category structure used in Experiment 2.

Item	f1	f2	f3	f4	Similarity		regression	Text
					(s = .10)	(s = .30)		
Training Items								
A1	0	0	1	0	.99	.89	1.024	D
A2	0	1	0	0	.91	.76	.810	D
A3	0	1	0	1	.89	.70	.429	D
A4	1	0	1	0	.97	.82	.548	M
B1	1	1	0	0	.89	.70	.333	M
B2	0	1	1	1	.90	.73	.453	M
B3	1	0	0	1	.98	.84	.143	D
B4	1	1	1	1	.98	.86	.024	M
Transfer Items								
T1	0	0	0	0	.91	.78	1.000	D
T2	0	0	0	1	.52	.53	.619	D
T3	0	0	1	1	.50	.51	.643	D
T4	1	0	0	0	.38	.44	.524	D
T5	1	1	0	1	.26	.30	.047	M
T6	0	1	1	0	.65	.62	.834	M
T7	1	1	1	0	.36	.41	.358	M
T8	1	0	1	1	.34	.37	.167	M

Note. The values for each feature dimension are the same as in Table 2. Numbers under Similarity refer to estimated classification probabilities when $s = .10$, and when $s = .30$. Numbers under regression refer to predicted categories (desert = 1; mountain = 0) based on a regression analysis using weighted sums of feature values as predictor variables. Letters under Text identify items consistent with information about desert (D) or mountain (M) climates, based on the information stated in the text (Set A). For transfer items, all numbers refer to the probability of classifying the items as desert plants.

items have predicted category values that are greater than 0.5. Thus, for the LS category structure, the predicted categories resulting from a weighted sum of feature values does

not differ with respect to the information from the text that links features to their respective categories. However, both a weighted sum solution and appealing to the information stated in the text produce different predictions relative to exemplar similarity for transfer items T4 and T6. For these items, as was stated above, the classification probabilities based on exemplar similarity should produce classification probabilities that are opposite to the predictions resulting from a weighted sum of feature values or by applying background knowledge acquired from the text.

Unlike the LS structure, the NLS structure included training items that fit the contrasting category based on a weighted sum of feature values. For the transfer items, item T4 (and also T6 for the LS structure) represent items that are highly diagnostic for separating classification patterns that were based on exemplar similarity from those that were based on applying the information from the text. As with the LS category structure, item T4 is highly similar to more training items from the mountain plant category relative to training items from the desert plant category. However, it contains two of the three critical features that identify it as a desert plant, based on the information stated in the text. Thus, if participants considered the information from the text, they should label this item as a desert plant. In contrast, if participants disregarded the information from the text over the course of learning, as is predicted for the NLS items, they should base classification decisions according to exemplar similarity and label it as a mountain plant.

Method

Participants

Exactly 144 volunteers enrolled in introductory psychology courses at Texas Tech University served as participants in the experiment. All participants were given partial course credit for their participation.

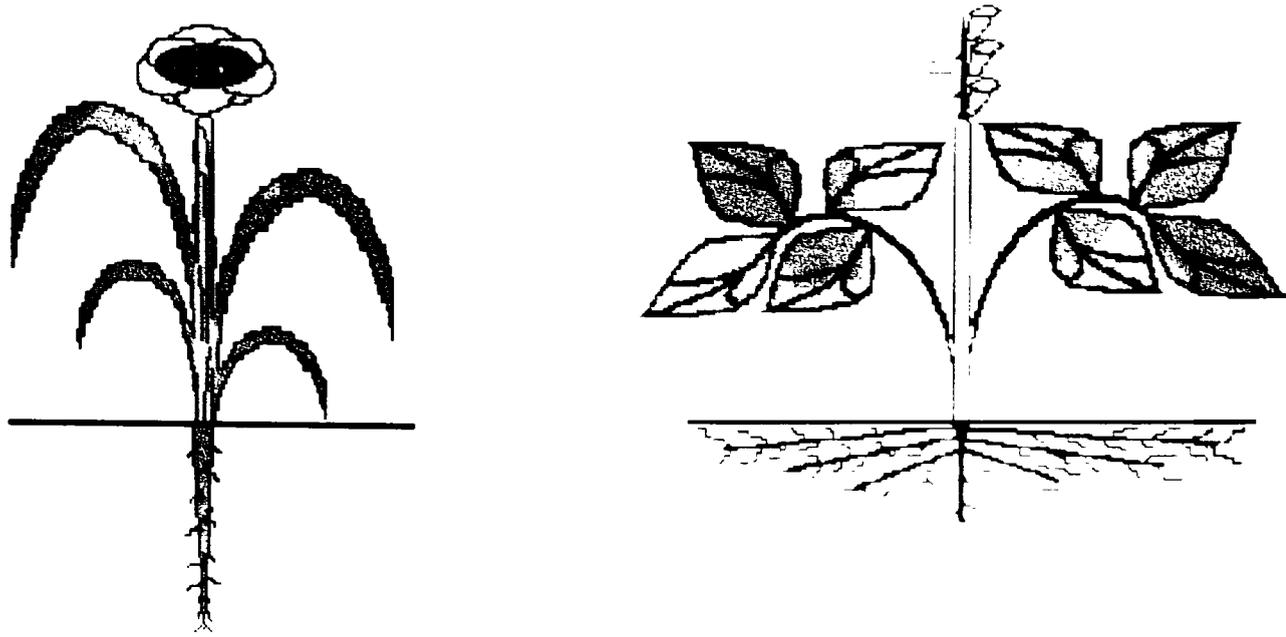


Figure 1

Sample drawings of plants used in Experiment 2.

Materials

Texts. The long version of the text containing background information, and the comprehension questions covering the information stated in the text were the same as that used in Experiment 1. As in Experiment 1, there were two versions of the text (Sets A and B), which differed from each other with respect to the type of information associated with each feature. For example, the information concerning the adaptations of tap roots indicates that they are capable of reaching water located far underground and that they are capable of protecting the plant from intense heat (Set A). In contrast, the information about tap roots listed in Set B text indicates that they absorb water near the surface and protect the plant from cold temperatures. Both sets of text are listed in Appendix C.

Plants. The plants that were used during classification included 16 black and white drawings of plants, which were created using a computer graphics software package. The plants were constructed by exhausting all possible combinations of the different feature types with the constraint that each plant could only have one type of root, stem, leaf, and flower, and these represented the four stimulus dimensions present in each plant. In addition, each of the feature dimensions was binary valued, which means that there were two different types of roots, stems, leaves, and flowers. Examples of two stimuli are listed in Figure 1. The plant on the left side of Figure 1 has a tap root, woody stem, bladed leaves and a headed flower, and is abstractly represented with the binary feature values 0000, which correspond to the different features for each of the four respective feature dimensions. In contrast, the plant on the right possesses a fibrous root system, herbaceous stem, compound leaves, and a spiked flower. It is abstractly represented as 1111 for the four respective feature dimensions. Since each plant was abstractly represented using binary feature values, it was possible to describe all of the stimuli according to the different features they possessed.

Eight of the 16 plants corresponded to training items and eight were new transfer items. Four of the eight training plants were designated as desert plants and four were designated as mountain plants. These labels were counterbalanced for half of the participants such that a given training item corresponded to a desert plant for half of the participants and as a mountain plant for the other half.

Category Structures. For the two types of category structures that were tested, one of the structures was linearly separable (LS) and the other structure was non-linearly separable (NLS). The LS structure was composed of training items that could be perfectly classified based on an additive combination of feature components, specifically over the first three feature dimensions. To get a clearer picture of this, all desert plants had at least two of the features tap roots, woody stems, and bladed leaves, whereas all

mountain plants had at least two of the features fibrous roots, herbaceous stems, and compound leaves. Thus, summing these three critical feature dimensions produced a linear boundary that separated all desert plants from all mountain plants. For the LS category structure, all training items could be correctly classified based on the information that members of one category must possess features that are adapted to survival in one of the environments. This means that learners must understand that tap roots, woody stem, and bladed leaves, were adapted for survival in desert environments and that fibrous roots, herbaceous stems, and compound leaves were adapted for survival in mountain environments (see Set A in Appendix C). For the NLS category structure, not all of the training items could be classified according to application of the information stated in the text. Rather, one item from each category matched the information from the contrasting category. These items (items A4 and B3 in Table 5) represented exception items. An additive features model based on a regression analysis indicated that item B3 is a mountain plant (correct) and A4 is a desert plant (correct). Regression also indicated that A3 is a mountain plant (incorrect). Therefore, whether the additive model is based on a regression or rule-based analysis, there is at least one exception in the NLS learning set.

Procedure

All stimuli and instructions were presented on an IBM compatible PC using Authorware Academic software Version 3.5 with the exception of the background information and probe questions. The text and the probe questions were presented on separate cards and on sheets of paper as in Experiment 1. All participants were tested individually or in pairs by the experimenter and were randomly assigned to one of the six experimental conditions. The experiment included a training phase followed by a transfer phase for all participants. Of the 144 participants, the last 48 who participated

performed a rating task prior to and immediately after completion of the learning phase. In addition, these 48 participants also provided ratings as to the confidence in the accuracy of their classification decisions during the transfer phase of the experiment.

During the training phase, participants were instructed that they would be shown eight drawings of plants several times. For each plant they were asked to indicate whether the plant was a desert plant or a mountain plant by pointing to and clicking the appropriately labeled response option using the mouse controller, or by typing the letters D or M when making their respective category selections. Participants were also told that the task may initially seem difficult, but it will become easier because the computer would provide corrective feedback after each category selection. The comprehensive text group was given access to text containing background information prior to the learning phase. These participants were instructed that specific information pertaining to environments, roots, stems, leaves, and flowers, could be obtained by reading from the text that was placed in front of them. They were also told to read the information carefully because they would be required to answer questions over the material to a high level of mastery before proceeding. As in Experiment 1, these participants were required to answer all 11 probe questions listed in Appendix D, and they were not allowed to proceed to the learning phase until they provided 24 out of 26 correct responses to the 11 questions. In the event that these participants did not answer questions to criterion, they were also informed that they would be required to review the text and answer questions over it again until the criterion was met. Immediately after participants in the comprehensive text group reached criterion on the probe questions, they advanced to the learning phase of the experiment. Those in the available text group were given similar instructions, but unlike the comprehensive text group, they were not required to answer questions over the material stated in the text before proceeding to the learning phase of the experiment. Rather, they were simply be told that they had the opportunity to read

information about desert and mountain environments and how various plant structures such as roots, stems, leaves, and flowers were adapted to survival in various climates. These participants thus had the opportunity to be more selective with regard to the background information. As a result, they had the option to completely disregard the background information if they choose to do so and could have proceeded immediately to the training phase of the experiment. However, they were encouraged to read the information at least one time before going on to the learning phase. Both text groups were also allowed to review the text during the learning phase of the experiment if they choose to do so. Those participants in the no text group were not given the text, but simply started with the learning phase.

The 48 participants who were required to make ratings (eight per condition) were instructed to rate the importance of different plant parts (e.g., tap roots) with respect to how well these parts helped plants survive in desert and mountain environments using a six-point Likert type scale. These participants made ratings prior to and immediately after the learning phase. The purpose of the ratings was to track changes in participants' estimates regarding the importance of the features with respect to how they helped plants survive in desert and mountain environments. Each question was of the form "How important are X for survival in Y environments?" where the values of X correspond to the different plant parts and the value of Y was represented by desert or mountain labels. One example question was: How important are tap roots for survival in desert environments? For each question, participants made ratings by typing a number from one to six indicating their selection. A rating of one indicated that a given plant structure (e.g., tap roots) was not at all important for survival in a given environment, whereas a rating of six indicated that a given plant structure was very important for survival. Participants provided ratings for those features present in the actual stimuli (i.e., tap and fibrous roots, bladed and compound leaves, woody and herbaceous stems, and headed

and spiked flowers) with respect to their how important they were to survival in both desert and mountain environments. In addition, they also made ratings as to the importance of other types of plant parts that were not present in the actual stimuli (i.e., sponge and bulbous roots, needled and velvet leaves, trunked and thorny stems, and thistle and bell flowers) for use as distractor items. Thus, there were a total of 32 questions that were given ratings.

During the learning phase of the experiment, the computer randomly displayed the first training item from one of the category structures. It appeared on the computer screen until the participant made a category selection. After doing this, the computer displayed corrective feedback below the drawing (e.g., Correct or Incorrect), and the plant remained on the screen for three additional seconds before the computer randomly selected another drawing from the training set. In addition, there was a one second interstimulus interval between trials in which the screen remained blank.

The eight drawings of plants were presented in blocks of eight trials with no obvious separation between blocks. One block corresponds to one run through each of the eight training stimuli. No more than 24 training blocks were allowed during the training phase. After 24 blocks, or when the participants classified all training items in one block with no errors, the computer proceeded to the transfer portion of the experiment. Those participants who made ratings at the beginning of the experiment were required to provide ratings for the 32 rating questions again after completion of the training phase. Throughout the experiment, the computer recorded participants' classification decisions for each item and for each run immediately after participants made a classification decision.

During the transfer phase, participants were told that the computer would once again present drawings of plants. They were also informed that some of the plants were old items and some corresponded to new items. For each plant, they were asked to

indicate whether the plant was a desert plant or a mountain plant using the same procedure as during the training phase. However, participants were not given feedback as to the accuracy of their classifications nor were they given access to the text. They were also asked to make classification decisions as accurately as possible and to use what they had learned from the training phase to help them classify the plants. After this, the computer randomly displayed the eight training items and the eight new transfer items one time. All participants were given as much time as needed to make their selections. Those who were required to make ratings prior to and after the learning phase were also required to rate their confidence in the accuracy of their classification decisions. Immediately after these participants selected a category, they were asked to type a number from one to six using a Likert-type scale. A confidence rating of one indicated that participants were not at all confident in their selections and a rating of six indicated that they were very confident in the accuracy of their selections. The entire procedure took 25 to 45 minutes to complete.

Design

For the training phase, a 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or comprehensive text) X 2 (Ratings: ratings or no ratings) X 2 (Set: A text or B text) factorial design was incorporated. All of the four factors were manipulated as between-subjects independent variables. The dependent variable was the number of learning blocks completed until a learning criterion of one error-free run through each of the eight training stimuli was met. Each block was defined as one run through each of the eight training stimuli. For the transfer phase, the above design was also incorporated. However, the dependent variable was the proportion of times each item was classified as a desert plant during one run through each of the 16 plants. Transfer items included the 8

training items plus an additional 8 new items (i.e., those items not presented during the training phase).

Results

Initial Learning

For the training phase, a 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or comprehensive text) X 2 (Rating: ratings or no ratings) X 2 (Set: A or B) factorial analysis of variance was performed using the number of blocks to reach learning criterion as the dependent variable. All factors were between-groups factors. The Ratings factor was included in order to examine potential bias effects on the dependent variable that may have resulted from the inclusion of the rating task. Likewise, the factor Set, which corresponds to the version of text that was received, and hence the category labels of items, was also included in order to identify whether the counterbalancing of stimuli to category labels was biased to either desert or mountain plant categories. A four factor analysis of variance revealed significant main effects for Category Structure [$F(1, 120) = 11.96, p < .001$], Text [$F(2, 120) = 3.08, p < .05$], and a significant Category Structure X Text interaction [$F(2, 120) = 3.07, p < .05$]. In addition, there were no main effects for Ratings and Set (all $F_s < 1, p > .05$), nor did these factors contribute to any significant interaction effects. Post hoc t-tests using a Bonferroni adjustment to control familywise error were also performed on the Category Structure X Text interaction components. These analyses indicated that participants in the comprehensive text group took significantly fewer trials to reach criterion in the LS condition relative to the NLS condition [$t(46) = -4.09, p < .001$]. In addition, there were no significant differences between the number of trials to reach criterion for the available text and the no text groups when performance in the LS condition was contrasted to performance in the NLS condition (see Table 6). For those learning linearly separable

Table 6

Mean number of trials to reach learning criterion as a function of category structure and text condition.

Category Structure	Text Condition		
	No Text	Comprehensive	Available
Linearly Separable	14.13 (5.87)	7.58 (4.64)	12.25 (6.37)
Non-Linearly Separable	15.17 (6.73)	14.33 (6.62)	14.13 (6.34)

Note. Standard deviations are in parentheses.

items, the comprehensive text group required fewer learning blocks relative to the no text group [$t(46) = -4.28, p < .001$], and the available text group [$t(46) = -2.90, p < .01$].

There were no reliable differences between groups when learning the non-linearly separable categories. Thus, merely providing access to the background information (as was the case for the available text group) failed to produce a significant decrease in the number of training blocks to reach criterion relative to the no text group for both LS and NLS category structures. However, supplying learners with background knowledge facilitated learning only when participants were required to read the text and answer questions over it (as was the case for the comprehensive text group), and this facilitation was only observed when learning LS categories. The implication of this finding will be addressed in the Discussion.

Transfer Performance

Analyses of transfer performance for the critical transfer items were conducted using a 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or

comprehensive text) X 2 (Rating: ratings or no ratings) X 2 (Set: A or B) factorial analysis of variance. Only the data from those participants who reached a learning criterion of one error free block within 24 blocks of learning trials were retained for analysis of transfer phase classification. Thus, six participants in the no text condition failed to reach criterion (4 in the LS and 2 in the NLS condition), and four in the available text condition (2 learning LS and 2 learning NLS categories). As a result, the transfer data from these participants was not included in the following analyses. In contrast, none of the participants in the comprehensive text group failed to reach a learning criterion of one error free training block. Regarding the dependent variable, it was calculated for the critical transfer items such that they were coded with a value of 1 when participants made classification decisions using the three feature additive rule. When they were classified based on similarity to past examples, they were coded with a value of 0 (see Table 7). All factors were between-groups factors as in the previous analysis. A four factor analysis of variance revealed significant main effects for Category Structure [$F(1, 110) = 7.71, p < .001$] and Text [$F(2, 110) = 9.41, p < .001$]. In addition, a significant interaction between Category Structure and Text was also obtained [$F(2, 110) = 3.17, p < .05$]. There were no main effects for Set and Ratings, nor were there any reliable interactions associated with these factors. Post hoc t-tests for the critical transfer items were also conducted in order to identify whether classification is driven by background knowledge or by exemplar similarity. Of critical importance are the text conditions contrasted with the no text condition, particularly for the LS category structure. Post hoc t-tests using a Bonferroni adjustment to control familywise error revealed that participants in the comprehensive text group based classification decisions more on a summing of feature values for the critical transfer items relative to the no text group [$t(44) = 4.63, p < .001$], but not when compared to the available text group [$t(45) = 1.29, p > .05$]. In addition, the comprehensive text group relied more on an additive

solution when learning LS concepts relative to NLS concepts [$t(46) = 3.71, p < .01$], as did the available text group [$t(43) = 3.17, p < .01$]. These results are summarized in Table 7. This pattern of results supports the findings of earlier studies (e.g., Allen & Brooks, 1991; Nakamura, 1985). The results suggest that individuals use background knowledge when the category structure is linearly separable, but less so or not at all when the category structure is non-linearly separable. In the latter case, individuals may rely more on exemplar similarity than the former case.

Other analyses were conducted that tested predictions from the regression analysis based on additive features from those based on similarity. Means approaching 1 indicate that participants made classification decisions based on weighting feature values, whereas means approaching 0 indicate that decisions were based on exemplar similarity. A 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or comprehensive text) X 2 (Rating: ratings or no ratings) X 2 (Set: A or B) factorial analysis of variance was performed on the critical items as in the previous analysis. However, since

Table 7

Mean proportion correct based on text information for critical transfer items as a function of category structure and text condition.

Category Structure	Text Condition		
	No Text	Comprehensive	Available
Linearly Separable	.25 (.33)	.71 (.36)	.56 (.43)
Non-Linearly Separable	.23 (.29)	.33 (.35)	.46 (.39)
Non-Linearly Separable (T4)	.42 (.43)	.45 (.41)	.58 (.45)

Note. Standard deviations are in parentheses. Means approaching 1 reflect consistency with three-feature additive solution. Means approaching 0 reflect reliance on exemplar similarity.

predictions of the regression analysis are different from the predictions of exemplar similarity only for item T4 (NLS), classification patterns for was used as the dependent variable for the non-linearly separable condition. A four factor analysis of variance revealed only a main effect for Text Condition [$F(2, 110) = 4.73, p < .01$]. There were no other main effects or interactions that were significant at the .05 level. However, a Category Structure X Text interaction approached reliability [$F(2, 110) = 2.64, p = .075$]. Examination of the main effect of Text revealed that participants in the comprehensive text ($M = .58, SD = .45$), and the available text conditions ($M = .57, SD = .46$) made classification decisions on the basis of weighting feature values to a greater extent than participants in the no text condition ($M = .33, SD = .43$). These results are also summarized in Table 7.

Ratings

Participants' ratings as to the importance of the features with regard to adaptation in desert and mountain environments were conducted on initial ratings and on difference scores between initial and final rating opportunities. In these analyses, the actual questions were divided into four categories: text-consistent, text-inconsistent, irrelevant, and distractors. An example of a text-consistent question is one in which participants read from the text that tap roots are adapted for survival in desert conditions and are asked the question: How important are tap roots to survival in desert environments? An example of a text-inconsistent question is one in which participants read from the text that tap roots are adapted for survival in desert conditions and are asked the question: How important are tap roots to survival in mountain environments? In contrast, the irrelevant questions are those that refer to flowers. Finally, the distractor questions are over the remaining features that are not present in any of the stimuli. A 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or comprehensive text) X 2

(Question Type: text-irrelevant, text-consistent, irrelevant, and distractor) mixed factor analysis of variance was performed on the initial ratings, and also on the difference scores. Of more theoretical importance are the difference scores between initial and final ratings. These difference scores were computed by subtracting final ratings from initial ratings for the purpose of tracking changes in participants' estimates prior to and after observing several items regarding the importance of the features with respect to how they help plants survive in desert and mountain environments. A three factor analysis of variance failed to detect any reliable changes that occurred between initial and subsequent rating opportunities. That is, there were no significant main effects or interactions. However, analysis of initial ratings produced a significant main effect for Question Type [$F(3, 312) = 17.35, p < .001$] and a significant Text X Question Type interaction [$F(6, 312) = 3.15, p < .01$]. Inspection of Table 8 indicates that those participants in the no text group tended to make ratings for all questions near the middle of the scale, whereas those in the comprehensive and available text groups tended to give high importance ratings for text-consistent items and low ratings for text-inconsistent items. The prediction was that the no text group should have shifted their ratings after exposure to the training items and that the NLS-text participants would shift after training, but not the LS-text participants. Thus, this rating task showed less of an effect of text than the classification and transfer tasks did.

In order to examine transfer confidence, an analysis was also performed on participants confidence ratings that were made after each category selection. A 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or comprehensive text) X 2 (Set: A or B) factorial analysis of variance was performed using confidence ratings as the dependent variable. Confidence ratings approaching 6 indicated that participants were very confident in their category selections. Low ratings indicate that participants were not at all confident. There were no significant main effects or

Table 8

Mean ratings of features prior to training (after training in parentheses) as a function of question type, category structure, and text condition.

Probe type	Text Condition		
	No Text	Comprehensive	Available
	<u>Linearly Separable</u>		
Text-Consistent	4.00 (4.67)	5.27 (5.35)	5.17 (5.36)
Text-Inconsistent	3.02 (3.22)	1.72 (1.65)	2.48 (2.73)
Irrelevant	3.75 (3.46)	3.40 (3.28)	4.12 (4.06)
Distractor	3.55 (3.63)	3.32 (3.11)	3.19 (3.57)
	<u>Non-Linearly Separable</u>		
Text-Consistent	3.89 (3.94)	5.27 (5.14)	5.16 (4.96)
Text-Inconsistent	4.02 (3.83)	2.02 (2.21)	2.77 (2.44)
Irrelevant	3.59 (4.03)	3.90 (3.50)	3.66 (4.06)
Distractor	3.88 (3.69)	3.39 (3.30)	3.42 (3.38)

Note. Numbers are mean ratings on a six-point Likert scale where low values reflect low importance and high values indicate high importance.

interactions, but participants in the comprehensive text group ($M = 5.10$, $SD = .56$) tended to provide slightly higher ratings than the no text group ($M = 4.64$, $SD = .65$). Participants in the available text group made ratings that were in between the other groups ($M = 4.88$, $SD = .78$).

Transfer Latency

Additional analyses were conducted on categorization response latencies during the transfer phase of Experiment 2. Response latencies were calculated for each transfer

item such that the timer started simultaneously with the presentation of a given plant and stopped immediately after participants made a category selection. A 2 (Category Structure: LS or NLS) X 3 (Text: no text, available text, or comprehensive text) X 2 (Rating: ratings or no ratings) X 2 (Set: A or B) factorial analysis of variance was performed on the response latencies averaged across all transfer items. This analysis revealed a significant main effect of Text [$F(2, 110) = 4.87, p < .01$], and interaction between Category Structure and Text [$F(2, 110) = 3.14, p < .05$]. In addition, the Text X Category Structure X Set interaction approached statistical reliability [$F(2, 110) = 2.28, .05 > p > .10$]. However, there were no other significant main effects or interactions. Post hoc analyses using Tukey's HSD were conducted on the main effect of Text. These analyses revealed no reliable differences in response latencies between those who did not receive text ($M = 4.22, SD = 1.18$) and those in the available text group ($M = 4.61, SD = 1.45$), or between those in the available text group and those in the comprehensive text group ($M = 5.28, SD = 2.14$). The only reliable difference between groups occurred between the no text group and the comprehensive text group, with the latter requiring more time to make classification decisions than the former group. A series of t-tests were conducted on the significant Text X Category Structure interaction. These analyses indicate that the source of the interaction was due to significantly higher reaction times for those participants who received comprehensive text instructions and learned LS items compared to the other groups. More specifically, participants in the LS-comprehensive text group took more time to make classification decisions than those in the LS-available text group [$t(44) = 2.40, p < .05$], and the LS-no text group [$t(43) = 3.66, p < .01$]. The LS-comprehensive text group also took more time to classify than the NLS-comprehensive text and NLS-available text groups [$t(46) = 2.10, p < .05$], and the NLS-no text group [$t(45) = 3.03, p < .01$]. These results are summarized in Table 9.

Table 9

Mean response latencies (in seconds) during transfer as a function of category structure and text condition.

Category Structure	Text Condition		
	No Text	Comprehensive	Available
Linearly Separable	4.18 (0.79)	5.90 (2.17)	4.58 (1.62)
Non-Linearly Separable	4.28 (1.48)	4.65 (1.96)	4.65 (1.28)

Note. Standard deviations are in parentheses.

The most relevant findings were obtained from the Text X Category Structure interaction. It is apparent that participants in the LS-comprehensive text group took more time to make categorization decisions than the other groups. One reason for this increase in reaction time for the LS-comprehensive text group was that they were classifying the transfer items according to a rule (i.e., analytically) to a greater extent than the other groups. In contrast, participants from the other groups predominantly classified the transfer items according to their similarity to previously learned exemplars (i.e., analogy) than those in the LS-comprehensive text group (Allen & Brooks, 1991; Wattenmaker, McQuaid, & Schwartz, 1995). To clarify this distinction, analytic processes are equated with the use of background knowledge during classification and are contrasted with analogical processes, which include exemplar-based categorization strategies (Wattenmaker et al., 1995). Or stated differently, the use of background knowledge during categorization requires more deliberate or analytical processes whereas responding on the basis of exemplar similarity is more perceptually driven and automatic (Smith & Sloman, 1994). For example, in an experiment reported by Allen and Brooks

(1991), participants were slower to make categorization decisions when given a perfectly predictive classification rule prior to learning relative to participants who did not receive an informative rule. In particular, reaction times were slowest when exemplar similarity conflicted with the rule that was given to participants. The implication that is relevant to the current findings is that analytical processing consumes more processing time relative to analogical processes. As a result, whenever participants make classification decisions by applying background knowledge, they required more time to make classification decisions than when responding on the basis of exemplar similarity.

Model Fits

Classification probabilities for transfer items were fitted to three models. One model represents a purely exemplar based model, and is formally identical to the context model of Medin and Schaffer (1978). The other model represents an expanded version of the context model, except it includes a prototype component as tested by Nakamura (1985). It has an exemplar component and a prototype component with five free parameters. Four parameters represent similarity parameters for each of the four stimulus dimensions, and these similarity parameters are shared by both the exemplar and prototype components. Nakamura's prototype component model not only computes the similarity of presented items to previously stored instances, but it also computes the similarity of a probe item to category prototypes in the same manner. For this reason, it is not different from a purely exemplar based model since it treats prototypes as focal exemplars (Estes, 1994). Thus, there is reason to suspect that this model fails to accurately account for the effects of background knowledge since it treats background knowledge as similarity to a prototype. The fifth parameter is a bias parameter that weights the degree to which classification is driven by the prototype component.

The third model is a mixed model that has an exemplar component in addition to an additive feature component for accommodating the contribution of text-based background knowledge associated with each of the features. It differs from the mixed representational model tested by Nakamura (1995) in that it treats background knowledge as a weighted sum features values, not as an instance of a prototype. In the mixed model, the probability of classifying a given item in Category A was determined according to the following formula (Equation 3):

$$P(A | a1) = (1 - T) * \frac{\sum_{a \in A} Sim(a1, A)}{\sum_{a \in A} Sim(a1, A) + \sum_{b \in B} Sim(a1, B)} + T * [(f1 * k_1) + (f1' * k_2) + (f2 * k_3) + (f2' * k_4) + (f3 * k_5) + (f3' * k_6) + (f4 * k_7) + (f4' * k_8)] \quad (3)$$

Where $f1$, $f2$, $f3$, $f4$, represent the values of the features present in a presented item and $f1'$, $f2'$, $f3'$, and $f4'$ represent the complements of the features present in a given item.

The parameters $k1$ through $k8$ represent parameters that represent background knowledge parameters based on the features present in a given item and their complements. The parameter T represents a bias parameter that weights the contribution of the additive features component in the mixed model. It can have a value between 0 and 1. Values approaching 1 indicate complete reliance on the additive sum of feature weights whereas values approaching 0 reflect reliance on exemplar similarity. The exemplar component of the mixed model is identical to that of Medin and Schaffer's (1978) context model. It is also identical to the exemplar component in the model tested by Nakamura (1985).

With this mixed representational model, there are eight parameters for the additive

features component, four for the exemplar component, and one bias parameter. In addition, a purely additive features model (the additive features component of the mixed model) was also fitted to the transfer data. As indicated above, the additive features model is based on a linear regression where each feature present in a given stimulus and their complementary values are used as predictors of category membership.

These four models were fitted to the transfer data separately for each Category Structure X Text condition (see Tables 10 and 11 for LS and NLS transfer data, respectively). The parameter values were estimated from the data for each model in order to determine the parameter values that minimized the error between model predictions and the observed transfer classification probabilities. The overall fits of each model are presented in Table 12. From this, it is clear that each model is very good at accounting for the transfer results, with the exception of the additive features model. There is a slight advantage for the mixed model over the other models, but this advantage may be due to the disparity in the number of free parameters between the models.

Close examination of the variance that is accounted for by the models indicates that the models with exemplar components fared equally well. In particular, the exemplar model of Medin and Schaffer (1978) yielded fits that were virtually identical to that of Nakamura's (1985) model. The mixed model that includes an additive features component also yielded closely matched that of the other exemplar models. However, the fit of this model was not due entirely to exemplar similarity. In particular, the additive features component of the model contributed to categorization almost equally as the exemplar component (for the comprehensive text group learning LS categories) based on the values of the \underline{T} parameter. Note that the \underline{T} parameter weights the degree to which

Table 10

Mean transfer classification probabilities as a function of text condition for the linearly separable (LS) structure.

Item	f1	f2	f3	f4	Text condition		
					No text	Comprehensive	Available
Training items							
A1	0	0	1	0	.76	.75	.86
A2	0	1	0	0	1.0	1.0	.91
A3	0	1	0	1	.95	.87	.86
A4	1	0	0	1	.71	.79	.86
B5	1	1	0	0	.43	.29	.18
B6	0	1	1	1	.10	.33	.18
B7	1	0	1	0	.00	.08	.00
B8	1	1	1	1	.05	.00	.14
Transfer items							
T1	0	0	0	0	.95	1.0	.95
T2	0	0	0	1	.95	.96	.82
T3	0	0	1	1	.33	.71	.50
T4	1	0	0	0	.33	.79	.59
T5	1	1	0	1	.76	.29	.32
T6	0	1	1	0	.76	.38	.45
T7	1	1	1	0	.10	.00	.09
T8	1	0	1	1	.05	.00	.18

Note. For transfer items, all numbers refer to the probability of classifying the items as desert plants.

Table 11

Mean transfer classification probabilities as a function of text condition for the non-linearly separable (NLS) structure.

Item	f1	f2	f3	f4	Text Condition		
					No text	Comprehensive	Available
Training items							
A1	0	0	1	0	.95	.96	.68
A2	0	1	0	0	.95	.88	.91
A3	0	1	0	1	.71	.92	.82
A4	1	0	1	0	.76	.71	.86
B5	1	1	0	0	.38	.25	.41
B6	0	1	1	1	.19	.42	.36
B7	1	0	0	1	.10	.17	.32
B8	1	1	1	1	.24	.04	.09
Transfer items							
T1	0	0	0	0	.90	1.0	.95
T2	0	0	0	1	.62	.83	.77
T3	0	0	1	1	.43	.63	.45
T4	1	0	0	0	.43	.46	.59
T5	1	1	0	1	.00	.04	.18
T6	0	1	1	0	.90	.79	.59
T7	1	1	1	0	.76	.41	.36
T8	1	0	1	1	.10	.13	.14

Note. For transfer items, all numbers refer to the probability of classifying the items as desert plants.

Table 12

Model fits to all test items and the residual errors for critical transfer items.

Model	Text Condition								
	No text			Comprehensive			Available		
	R^2	T4	T6	R^2	T4	T6	R^2	T4	T6
	<u>Linearly-Separable</u>								
Mixed	.979	-.020	.039	.992	.014	.027	.989	.026	-.015
Exemplar	.969	.020	.076	.984	.083	.074	.969	.055	-.020
Nakamura	.959	.022	.075	.984	.081	.079	.945	.057	-.020
Additive	.788	-.241	.253	.876	.136	-.125	.791	-.091	.080
	<u>Non-Linearly Separable</u>								
Mixed	.968	.046	.000	.977	.043	-.030	.946	-.032	-.060
Exemplar	.948	.063	-.015	.951	.069	-.009	.807	.126	-.131
Nakamura	.963	.072	-.033	.954	.094	-.023	.887	.094	-.035
Additive	.813	-.139	-.039	.870	-.060	-.055	.770	-.071	-.088

Note. There were 13, 4, 5, and 8 free parameters in the Mixed, Exemplar, Nakamura, and Additive feature models, respectively.

classification is driven by an additive features solution relative to exemplar similarity. When the value of the \underline{T} parameter approaches 1, this indicates that classification was driven primarily by an additive features solution. In contrast, the \underline{T} parameter in Nakamura's model weights the degree to which categorization is driven by similarity to prototypes. As can be seen in Table 13, the values of the \underline{T} parameter in the model tested by Nakamura are near 0.0, which indicates that categorization was driven almost exclusively according to exemplar similarity. In contrast, the values of the \underline{T} parameter

for the mixed model indicate that participants in the comprehensive text group who learned the LS categories based their responses to a greater extent on additive features than in the other conditions. At least qualitatively, this indicates that the mixed model is better equipped to capture the effects of background knowledge than Nakamura's model. That is, when learners are required to fully learn the adaptations of plant features, this may have biased learners to make categorization decisions of the basis of additively weighting features, particularly when learning LS categories. The model tested by Nakamura does not make this distinction.

Table 13

Values of T parameters for the two mixed representational models.

Model	Text Condition		
	No text	Comprehensive	Available
	<u>Linearly Separable</u>		
Mixed	.01 (.06)	.56 (.09)	.01 (.08)
Nakamura	.01 (.05)	.03 (.04)	.00 (.05)
	<u>Non Linearly Separable</u>		
Mixed	.01 (.04)	.03 (.06)	.06 (.05)
Nakamura	.13 (.06)	.05 (.05)	.08 (.07)

Note. Numbers in parentheses refer to standard errors of the parameters.

Inspection of Table 13 also indicates that for conditions other than the comprehensive text – LS condition, participants based classification decisions to a greater extent on the exemplar component of the model than the additive features component, based on the values of the bias (\underline{T}) parameter. This finding also matches the results obtained with the critical transfer stimuli. Note that participants learning LS categories in the no text group classified on the basis of exemplar similarity to a greater extent than participants in the comprehensive text condition, a finding which is not captured by the mixed prototype-exemplar model which was tested by Nakamura (see Table 7). In addition, transfer performance for those who learned NLS concepts was more consistent with responding on the basis of similarity to previous items than by additively weighting features, a finding which is also captured by the mixed model. Thus, there is at least preliminary evidence in favor of the mixed model over the exemplar-prototype model of Nakamura.

Discussion

The goal of this experiment was to answer four research questions. Specifically, the experiment examined whether learning about the adaptations of different plant features by reading from text prior to the training phase facilitated learning of linearly separable categories relative to when no text was provided, as measured by the number of training blocks to reach learning criterion. It also examined whether having knowledge about the adaptations of the plant features disrupted learning of non-linearly separable concepts compared to when no information about their adaptations was given prior to learning the plant categories. The motivation for examining these two questions was to

see if the text manipulations used in the current experiments would replicate earlier findings of knowledge effects. The experiment also examined classification performance for critical transfer items in which classification patterns for these items identified whether participants' classification decisions were based primarily on exemplar similarity, a summing strategy over the feature values present in the items, or a combination of both. This was motivated by a lack of research examining knowledge effects for transfer performance. Much of the research examining the role of background knowledge has focused exclusively on the speed at which concepts are attained (e.g., Heit, 1997; Pazzani, 1991), or the accuracy in which concepts are sorted into categories (Ahn & Medin, 1992; Murphy & Allopena, 1994; Spalding & Murphy, 1996). Very few studies have additionally examined the effects of background knowledge on transfer performance. With the exception of the study by Allen and Brooks (1991), those that have examined transfer have not incorporated transfer items that are highly diagnostic for separating the effects of background knowledge and similarity (e.g., Hayes & Taplin, 1995; Nakamura, 1985). Lastly, the experiment also examined the extent to which participants revised their estimates concerning the importance of the different plant features with respect to how critical they are for survival in desert and mountain environments, as measured by participants' ratings given prior to and immediately after the learning phase. These questions are addressed in turn in the following paragraphs.

Regarding learning of the linearly separable category structures, the results indicate that when participants in the comprehensive text group learned about the adaptations of the different plant features prior to the training phase, they took significantly fewer blocks of trials to reach learning criterion relative to participants in

the available text and no text conditions. This finding suggests that learning is increased when participants are able to apply what they had learned from the text to the set of training items, particularly when all training items are labeled in a manner that does not violate the facts stated in the text. This finding also replicates the knowledge effects that were reported by Nakamura (1985). In the Nakamura study, a group of participants was given background information (i.e., theory instructions) describing the habits of birds and bees, and another group of participants was not given this information (i.e., standard instructions). Those who were given background information were required to memorize it before they were allowed to proceed to the learning phase of the experiment, which is similar to the requirements of the comprehensive text group in the current experiment. Nakamura reported that the theory instructions produced significantly fewer classification errors during the training phase relative to those who received standard instructions, which was replicated in the current experiment with comprehensive and no text groups. The reason for this effect was interpreted as being due to the storage of different levels of category information. For example, the background information that was given to participants prior to training emphasized the importance of additively combining feature values. When this information was applied during learning, it benefited the learning of linearly separable category structures relative to the non-linearly separable structures since none of the linearly separable training items violated the additive rule.

The current research also extends Nakamura's (1985) findings to the learning contexts that were adopted in the present experiment. That is, the knowledge effects that occurred as a result of being provided with a background theory were replicated with larger, more integrated texts. However, the present research goes further in

demonstrating that the improvement in learning for the comprehensive text group compared to the available text group was due to the nature of the learning requirements for the text. Participants in the comprehensive text group were required to answer comprehension questions over the facts in the text whereas those participants in the available text condition were not required to do this. Thus, these results argue that mere access to the background information is insufficient to produce significant increases in learning, since participants in the available text group required significantly more training blocks to reach learning criterion relative to those in the comprehensive text group. In addition, learning performance for the available text group did not significantly differ from the no text group, which suggests that these groups based classification decisions to a greater extent on exemplar similarity than on background knowledge. To get a clearer picture of this process, if it is assumed participants in the no text group had no prior knowledge regarding how the plant features were adapted to survival in either desert or mountain environments, then they should make classification decisions according to the principles of exemplar similarity (Medin & Schaffer, 1978; Nosofsky, 1986). In contrast, if participants in the available text group applied the information they had read from the text prior to the learning phase, then their performance on the training items should more closely mirror that of the comprehensive text group than the no text group.

The results over training performance for the non-linearly separable category structures failed to detect any reliable differences between the three text conditions. This finding does not lend support for the notion that having background knowledge concerning the adaptations of plant features disrupts learning of non-linearly separable concepts relative to those who do not have such background knowledge, as was reported

by Nakamura (1985). Greater disruption should have occurred for the comprehensive text group relative to the no text group since the background information had been learned prior to training would have misled participants on some of the training items. More specifically, participants the comprehensive text condition learned that in order for a given item to be a desert plant, it must have two of the features tap roots, woody stems, and bladed leaves. When they observed the training items, one of the plants with two of these three critical features was labeled as a mountain plant, which is contrary to the information they had learned prior to training. As a result, these participants should have been misled by the information they had received prior to training. That is, they should have made more classification errors for the items that did not fit the background information compared to the other text groups, as was reported by Nakamura (1985). In contrast, participants in the no text group were less likely to be biased regarding how the features present in the items are associated with the category labels since they were not given background information prior to the learning phase. Because of this, the participants in the no text condition should have based categorization decisions on exemplar similarity, which should have led to fewer errors relative to those in the comprehensive text group.

Why then was there not a reliable difference in learning for these text conditions? One possibility is that the participants in the comprehensive text condition eventually adopted a similar learning strategy as those in the no text condition. That is, when participants in the comprehensive text condition were presented exemplars that violated the information that they had learned prior to training, they stored the individual exemplars in memory and made classification decisions based primarily on exemplar

similarity. There is very little empirical evidence that addresses this question directly. However, an experiment reported by Allen and Brooks (1991, Experiment 4) revealed that when participants are given a classification rule prior to learning that produces errors when it is applied, “memory for prior instances takes on a special status.” (p. 12). Similarly, when this background knowledge is in error, classifying items on the basis of similarity to previously encountered examples represents a classification strategy that may be worth considering. Support for this hypothesis was reported by Allen and Brooks who observed that classification was driven more by similarity to previous examples and not according to the rule that was given to participants prior to learning. Evidence for responding on the basis of exemplar similarity for the NLS categories is corroborated by transfer performance on the critical transfer items, which will be addressed next.

The third main finding of this experiment concern categorization performance for the critical transfer items (items T4 and T6 in Tables 4 and 5). These items are highly diagnostic for separating classification behavior that is based primarily on exemplar similarity from that which is based on applying the information from the text. For both category structures, item T4 is highly similar to more training items from the mountain plant category relative to training items from the desert plant category. However, it contains two of the three critical features that identify it as a desert plant, based on the information stated in the text. Similarly, item T6 is highly similar to previously learned desert plants than old mountain plants, even though it has two of the three critical features that are present in desert plants. Thus, if participants consider the information from the text, they should label item T4 as a mountain plant and item T6 as a desert plant. In contrast, if participants disregard the information from the text over the course of

learning, they should base classification decisions according to exemplar similarity and label item T4 as desert plant and item T6 as a mountain plant.

The results revealed that participants in the comprehensive text and available text conditions classified the critical transfer items according to the information from the text to a greater extent than those in the no text condition, but only when learning linearly separable concepts. In contrast, participants in the no text condition classified primarily on the basis of exemplar similarity. When participants learned non-linearly separable concepts, classification patterns for the critical transfer items did not significantly differ across the text groups. Participants in the three text conditions based classification decisions primarily according to exemplar similarity and not according to the information from the text.

The importance of this finding is twofold. First, it indicates that learners in the comprehensive text condition continued to apply the information they had learned from the text into the transfer phase. This is to be expected if the background information matches the category labels assigned to all training items, as is the case for the linearly separable categories, but not for the non-linearly separable structures. Likewise, if participants classified the training items on the basis of exemplar similarity, it is expected that they would continue to do so for the transfer items. There is no a priori reason why participants would adopt a different categorization strategy during transfer as they had done during the initial training phase. However, this interpretation can be contrasted with the findings of Allen and Brooks (1991), who reported that similarity to prior instances influenced categorization decisions to a greater extent than application of a perfectly predictive rule. In their experiments, participants were given a sufficient and perfectly

predictive additive rule for determining category membership. Participants classified drawings of fictitious creatures (e.g., builders and diggers) where builders were defined according to an additive rule. Specifically, half of the participants were told that builders must have at least two of the features long legs, spots, and an angular body. During transfer, participants classified new items that either fit the rule and were highly similar to an old item from the same category (positive match items), or items that fit the rule and were highly similar to an old item the contrasting category (negative match items). If participants classified the items on the basis of similarity, then more errors should occur for the negative match items than the positive match items. Indeed, that is what Allen and Brooks found. Namely that similarity to previous items influenced categorization decisions to a greater extent than applying a previously learned rule. In the current research, the influence of previous items (similarity) was less pronounced for the comprehensive text group that learned LS categories. A likely reason for the disparity between Allen and Brooks' findings, and the current results is that the effects of similarity were stronger in the Allen and Brooks study. That is, in the Allen and Brooks study, negative match items differed on only one feature when compared to its most similar training item, and this item was in the opposite category based on the rule given prior to learning. In addition, the stimuli contained five binary valued features. Thus, negative match items shared four of the five features as its most similar training item, and this training item was in the opposite category as that defined by the rule. The present study only used stimuli that had four binary valued features. Thus, there was greater featural overlap in the Allen and Brooks study than in the present study, and this could have led participants to be more sensitive to similarity than in the present study.

Another difference between Allen and Brooks (1991) and this study was that participants in the present research were never explicitly given a classification rule prior to learning. The text groups could have inferred an additive feature rule from the text, but they were not required to do so prior to the learning phase. Nevertheless, the participants in the comprehensive text responded as if they were using a rule during transfer for the critical items to a greater extent than the other groups, but only when learning the LS categories. When learning NLS categories, they responded as if they were using exemplar similarity. Thus, the key distinction between the Allen and Brooks study and the present one may lie in the procedures that were implemented to provide learners with background knowledge. In the present study, the information that linked the features to specific adaptations may have made them more memorable since they drew on conceptual information. That is, participants may have used the information from the text to form a rule that was more salient than in the study by Allen and Brooks. Recall that in the Allen and Brooks study, the rule simply listed the features of the rule. Participants were not given explanations as to why a particular feature (e.g. spots) was associated with a particular category (e.g., builders). Because of this, participants may have been less likely to apply the rule during transfer than in the present study. As a result, participants adopted classification strategies that were more consistent with similarity than a rule.

A second important finding of the present study is that the transfer results suggest that participants were more likely to disregard the information from the text when they discovered that it did not apply to all of the training items. Specifically, categorization patterns for those learning non-linearly separable concepts indicated that transfer was

driven more by exemplar similarity than by applying the information from the text. If participants in the text conditions classified on the basis of the text, their performance on the critical transfer items for the non-linearly separable concepts should mirror those that were observed for the linearly separable concepts. However, this was not the case in the present experiment. Thus, there is now preliminary evidence indicating that both exemplar similarity and background knowledge are crucial for transfer performance. Any account of categorization must address these in order to get a more complete picture of categorization.

Two of the models that were fit to the transfer data in the present study include a mixed exemplar-prototype model tested by Nakamura (1985), and a mixed model that includes an exemplar component and an additive features component. The model tested by Nakamura, as indicated in the introduction, fails to accurately capture the effects of background knowledge. In particular, Nakamura's model did not include an additive features component. Rather, background knowledge was expressed in terms of category prototypes, and they were treated as if they were actual items in the model. That is, the similarity of a probe item to the prototype component was calculated using a multiplicative similarity rule as in the exemplar component. The difficulty with this formulation is that it treats prototypes and exemplars the same. Thus, the model tested by Nakamura could be considered as a special type of exemplar model. Indeed, there were no qualitative differences between the performance of the Nakamura model and a pure exemplar model when they were fitted to the transfer data.

A growing body of data up to this point indicates that similarity to past examples or prototypes and background knowledge are in many ways qualitatively different

(Wisniewski & Medin, 1994), and models that are based exclusively on similarity (either to previous items or prototypes) are incomplete at best (Murphy & Medin, 1985). The mixed model tested in this study was designed to address this shortcoming of the Nakamura model by including an additive feature component. For the most part, the mixed model produced slightly better overall fits to the transfer data than the other models. However, both a pure exemplar model and Nakamura's model provided excellent fits to the data. Thus, interpretive caution is warranted. However, qualitatively, the mixed model seems to be capturing the effects of background knowledge to a greater extent than the other models as evidenced by an increased importance of the additive features component for the comprehensive text group who learned LS categories relative to the other models. The overall fits to the transfer data also rule out models that are exclusively based on additive features, as evidenced by relatively lower overall fits to the data. Thus, the current results highlight the importance of both exemplar representations and background knowledge of features in categorization.

Recent findings in other category learning domains have also reached similar conclusions (e.g., Hayes & Taplin, 1995; Lin & Murphy, 1997; Spalding & Murphy, 1996). In particular, Hayes and Taplin (1995) found that functional knowledge of features, when activated by meaningful category labels, influenced the ease in which concepts were learned. This knowledge accounted for more variance than prototype distance and exemplar similarity, but only when participants were given meaningful category labels. Research by Lin and Murphy (1997) also indicates that background knowledge affects the speed and accuracy of learning. In their experiments, participants were faster at detecting missing features that were considered important by background

knowledge. Other research points to the relations among features as evidence for both exemplar and knowledge based influences (Spalding & Murphy, 1996). That is, when any two pairs categories are constructed according to the same logical structure.

Participants will construct categories that fit a family resemblance structure to a greater degree when the features are meaningfully integrated (e.g., when they fit an underlying theme) relative to when the features do not fit an underlying theme. The important aspect of this is that participants generally do not construct (sort) according to family resemblance in most experimental situations. Although this was not a major concern in the present study, the results from Experiment 1 indicate that after reading the texts, participants were more likely to sort drawings of plants according to an additive rule (which also conforms to family resemblance since there were no single features that defined category membership). Thus, during initial sorting, the features were probably not interpreted as being consistent with an underlying theme (e.g., adaptation and survival), but after reading, they were.

The results over participants' ratings of the importance of plant features with respect to their ability to adapt to desert or mountain environments failed to reveal any significant differences between the ratings made prior to and immediately following the learning phase. What this suggests is that at the end of the learning phase, participants rated the plant features as being no more or no less important to survival in either desert or mountain environments compared to their initial ratings. This finding is similar to that reported by Livingston and Andrews (1995) who observed that when participants were given background information prior to learning that emphasized the importance of features that are actually irrelevant for categorization, they continued to consider the

irrelevant features as being important for determining category membership over the course of learning. In addition, they also gradually began to consider features that were initially considered not important when it became apparent that such features were diagnostic of category membership. Thus, rather than abandoning features that are irrelevant, they become augmented with other features that are more diagnostic of category membership. The implication of this is that participants' knowledge of features and the categories that they are primarily associated with does not undergo total revision over the course of learning, but continue to be considered as being important for category membership. The reason for this is that irrelevant features prove to be correct on occasion. In the current experiment, there were no individual features present in any of the stimuli that were exclusive to either desert or mountain plant categories. Thus, features that were rated as being important (or unimportant) prior to the learning phase continued to be rated as being important (or unimportant) at the end of learning.

There were differences between the three text groups with regard to initial ratings. For example, the results indicated that those who received text provided high ratings for features that were consistent with the information from the text, and gave low ratings for inconsistent features. For example, participants who read that tap roots are adapted for survival in desert environments gave high ratings when asked how important tap roots are for survival in desert environments (text-consistent), but gave low ratings when asked how important tap roots are for survival in mountain environments (text-inconsistent). This pattern of results was expected for the both text conditions and served as a manipulation check as to how well participants were able to infer the relationship between the features mentioned in the text and the plant categories. In addition,

participants in the comprehensive and available text groups also consistently provided ratings toward the midpoint of the six-point scale when asked questions emphasizing irrelevant features (e.g., flowers), and distractor features (e.g., bulbed roots). In contrast, participants in the no text condition provided ratings around the midpoint of the scale for all question types, as was expected considering that they had not received any background information prior to making ratings. Taken together, these findings indicate that participants did not revise their estimates concerning the importance of the different plant features with respect to how critical they are for survival in desert and mountain environments, as measured by ratings given prior to and immediately after the learning phase. They also indicate that participants were aware of what they had read from the text both at the beginning of training and at the end of the training phase. Thus, any interpretation of the results that implies that participants forgot the text materials, and hence classified using some other strategy, are less tenable.

To summarize, the primary goals of this experiment were to determine whether background information supplied by reading from text prior to the training phase facilitated learning of linearly separable categories, and whether the text disrupted learning of non-linearly separable concepts, compared to when no information about the feature adaptations was given prior to learning the plant categories. The experiment also identified whether participants' classification decisions for critical transfer items were based primarily on exemplar similarity or according to a summing strategy over the feature values present in the items, as a function of the type of category structure they learned and the text condition they received. Lastly, the experiment examined participants' ratings given prior to and immediately after the learning, and focused on

their estimates concerning the importance of the different plant features with respect to how critical they are for survival in desert and mountain environments.

For the most part, these goals have been met. Specifically, the experiment replicated previous findings that were observed in learning linearly separable concepts (e.g., Nakamura, 1985), but failed to replicate the superiority of the no text condition over the comprehensive text condition for non-linearly separable concepts. Performance on critical transfer items identified a potential source for this effect. In particular, the transfer performance indicates that participants in the comprehensive and available text conditions primarily based categorization decisions according to exemplar similarity after learning non-linearly separable training items, which also suggests that they placed less weight on application of the background knowledge they received by reading from texts during training. In contrast, those who did receive background information concerning the adaptations of plant features applied this knowledge during training and again during transfer. However, examination of feature ratings indicated that participants did not revise their estimates of the importance of the features at the conclusion of the learning phase.

The findings from this study set the stage for future work examining conceptual learning from text and how it interacts with learning over examples. Many of the learning effects that were observed resulted only after participants were required to answer comprehension questions over the text to a high level of mastery before they were allowed to proceed to the learning phase. The purpose of including this constraint was to guarantee that participants fully understood the text before proceeding to the learning phase. It also had the goal of establishing one condition under which background

knowledge of a concept becomes established prior to learning. Future work may be aimed at other conditions (besides learning over examples) in which conceptual knowledge is acquired. A starting point would be to more closely manipulate other text-based factors (such as complexity) or instructions that emphasize other coding strategies, such as relational coding. Other avenues of continued study could incorporate a greater number of more naturalistic stimuli besides those that were incorporated in the current experiment. It is rarely the case that sets of plants in the real world are composed of four binary valued features. Such an endeavor will provide a clearer picture as to the interaction between background knowledge and conceptual structure.

The present findings also provided additional evidence for the possibility that additive feature models and exemplar similarity models are not mutually exclusive explanations for category learning and transfer. It also addressed the qualitative weaknesses of other models that have uncovered knowledge effects (e.g., Nakamura, 1985). Further research is necessary to refine empirical tests so that they are more indicative of whether one model is sufficient or whether individuals draw on multiple models (strategies) for category learning and transfer. In addition, future work examining similarity and knowledge effects over the course of learning may provide a stronger test of these possibilities. One recent model advanced by Erickson and Kruschke (1998) represents a connectionist learning model that represents both exemplars and categorization rules in independent modules. The important component of the model is representational attention that weights either the rule module or the exemplar module accordingly depending on the learning context. That is, as evidenced in the present study, some learning contexts that are best explained according to exemplar similarity,

and other learning contexts are better explained by appealing to the background knowledge the learner uses during classification. Future research must provide more comprehensive account of when background knowledge (either expressed as rules of feature weights) and exemplar representations interact over the course of learning and during transfer. The notion of incorporating representational attention in models of category learning is a step in the right direction to achieving this goal.

REFERENCES

- Ahn, W., & Medin, D. L. (1992). A two-stage model of category construction. Cognitive Science, 16, 81-121.
- Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. Journal of Experimental Psychology: General, 120, 3-19.
- Beach, L. (1964). Cue probabilism and inference behavior. Psychological Monographs, 78, 21-37.
- Bruner, J., Goodnow, J., & Austin, G. (1956). A study of thinking. New York: John Wiley.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. Journal of Experimental Psychology: General, 107, 107-140.
- Estes, W. K. (1986). Array models for category learning. Cognitive Psychology, 18, 500-549.
- Estes, W. K. (1993). Models of categorization and category learning. In G. V. Nakamura, R. M. Taraban & D. L. Medin (Eds.), Categorization by humans and machines: Vol. 29. The psychology of learning and motivation (pp. 15-56). San Diego, CA: Academic Press.
- Estes, W. K. (1994). Classification and cognition. New York: Oxford University Press.
- Estes, W. K. (1995). A general model of classification and memory applied to discourse processing. In C. A. Weaver III, S. Mannes, & C. R. Fletcher (Eds.), Discourse comprehension: Essays in honor of Walter Kintsch (pp. 35-48). Hillsdale, NJ: Lawrence Erlbaum.
- Estes, W. K., Campbell, J. A., Hatsopoulos, N., & Hurwitz, J. B. (1989). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 556-576.
- Gluck, M. A., & Bower, G. H. (1988). Evaluating an adaptive network model of human learning. Journal of Memory and Language, 27, 166-195.
- Hayes, B. K., & Taplin, J. E. (1995). Similarity-based and knowledge-based processes in category learning. European Journal of Cognitive Psychology, 7, 383-410.

- Heit, E. (1994). Models of the effects of prior knowledge on category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 1264-1282.
- Heit, E. (1997). Knowledge and concept learning. In K. Lamberts & D. Shanks (Eds.), Knowledge, concepts, and categories (pp. 7-41). Cambridge, MA: MIT Press.
- Homa, D., & Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. Journal of Experimental Psychology: Human Learning and Memory, 2, 322-330.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. Psychological Review, 99, 22-44.
- Kruschke, J. K. (1993a). Three principles for models of category learning. In G. V. Nakamura, R. M. Taraban & D. L. Medin (Eds.), Categorization by humans and machines: Vol. 29. The psychology of learning and motivation (pp. 283-326). San Diego, CA: Academic Press.
- Kruschke, J. K. (1993b). Human category learning: Implications for back propagation models. Connection Science, 5, 3-36.
- Lin, E. L., & Murphy, G. L. (1997). Effects of background knowledge on object categorization and part detection. Journal of Experimental Psychology: Human Perception and Performance, 23, 1153-1169.
- Livingston, K. R., & Andrews, J. K. (1995). On the interaction of prior knowledge and stimulus structure in category learning. The Quarterly Journal of Experimental Psychology, 48A, 208-236.
- Medin, D. L., & Schaffer, M. M. (1978). A context theory of classification learning. Psychological Review, 85, 207-238.
- Medin, D. L. (1989). Concepts and conceptual structure. American Psychologist, 44, 1469-1481.
- Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. Journal of Experimental Psychology: Human Learning and Memory, 7, 355-368.
- Mooney, R. J. (1993). Integrating theory and data in category learning. In G. V. Nakamura, R. M. Taraban & D. L. Medin (Eds.), Categorization by humans and machines: Vol. 29. The psychology of learning and motivation (pp. 189-218). San Diego, CA: Academic Press.

- Mumma, G. H. (1993). Categorization and rule induction in clinical diagnosis and assessment. In G. V. Nakamura, R. M. Taraban & D. L. Medin (Eds.), Categorization by humans and machines: Vol. 29. The psychology of learning and motivation (pp. 283-326). San Diego, CA: Academic Press.
- Murphy, G. L., & Allopena, P. (1994). The locus of knowledge effects in concept learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, *20*, 904-919.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. Psychological Review, *92*, 289-316.
- Nakamura, G. V. (1985). Knowledge-based classification of ill-defined categories. Memory and Cognition, *13*, 377-384.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, *10*, 104-114.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-classification relationship. Journal of Experimental Psychology: General, *115*, 39-57.
- Nosofsky, R. M. (1991). Typicality in logically defined categories: Exemplar-similarity versus rule instantiation. Memory and Cognition, *19*, 131-150.
- Nosofsky, R. M., Clark, S. E., & Shin, H. J. (1989). Rules and exemplars in categorization, identification and recognition. Journal of Experimental Psychology: Learning, Memory, and Cognition, *15*, 282-304.
- 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.
- Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. Journal of Experimental Psychology: Learning, Memory and Cognition, *18*, 211-233.
- Pazzani, M. J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. Journal of Experimental Psychology: Learning, Memory, and Cognition, *17*, 416-432.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. Journal of Experimental Psychology, *77*, 353-363.

- Reed, S. K. (1972). Pattern recognition and categorization. Cognitive Psychology, 3, 382-407.
- Rosch, E. & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. Cognitive Psychology, 7, 573-605.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In J. L. McClelland & D. E. Rumelhart (Eds.), Parallel distributed processing: Vol. 1. Foundations (pp. 318-362). Cambridge, MA: MIT Press.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. Psychological Monographs, 75, (13, Whole No. 517).
- Smith, E. E., & Medin, D. L. (1981). Categories and concepts. Cambridge, MA: Harvard University Press.
- Smith, E. E., & Sloman, S. A. (1994). Similarity- versus rule-based categorization. Memory and Cognition, 22, 377-386.
- Spalding, T. L., & Murphy, G. L. (1996). Effects of background knowledge on category construction. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 525-538.
- Taraban, R. M., & Palacios, J. M. (1993). Exemplar models and weighted cue models in category learning. In G. V. Nakamura, R. M. Taraban & D. L. Medin (Eds.), Categorization by humans and machines: Vol. 29. The psychology of learning and motivation (pp. 91-127). San Diego, CA: Academic Press.
- Wattenmaker, W. D., Dewey, G. I., Murphy, T. D., & Medin, D. L. (1986). Linear separability and concept learning: Context, relational properties, and concept naturalness. Cognitive Psychology, 18, 158-194.
- Wattenmaker, W. D., McQuaid, H. L., & Schwertz, S. J. (1995). Analogical versus rule-based classification. Memory and Cognition, 4, 495-509.
- Wisniewski, E. J., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. Cognitive Science, 18, 221-281.
- Wisniewski, E. J., & Medin, D. L. (1991). Harpoons and long sticks: The interaction of theory and similarity in rule induction. In D. H. Fisher, M. J. Pazzani, & P. Langley (Eds.), Concept formation: Knowledge and experience in unsupervised learning, pp. 237-278. San Mateo, CA: Morgan Kaufman.

APPENDIX A

EXTENDED LITERATURE REVIEW

The act of categorization is equated with treating different objects as being in some way equivalent to one another. From this, people are able to draw conclusions and make inferences about specific objects when the appropriate category labels are known. This serves to economize cognitive resources while at the same time maximizing informativeness. In addition, the ability to categorize objects in the world is useful to make accurate predictions about particular objects when given limited information, and it also serves an adaptive function in that it allows people to avoid potentially dangerous outcomes; for instance, when determining whether a particular plant is edible or poisonous. In such a case, one must be able to understand that edible and poisonous plants differ from one another in some important way. Generally speaking, in order to correctly make classification decisions, one must examine the properties that a particular object possesses -- including those features that may or may not be directly observable -- and compare such properties either to summary category representations, or to individual stored examples from that category. In order to better understand how various sources of information contribute to classification decisions, it is necessary to identify how categorical information is cognitively represented and utilized. From a theoretical standpoint, this issue is not trivial. Models of categorization must not only capture the underlying principles that explain how category information is mentally represented and processed, they must also be able to make accurate predictions regarding classification behavior that is exhibited in humans under a variety of experimental contexts. Thus, they must adopt particular assumptions regarding how categories (like edible and poisonous plants) are cognitively represented, and the mechanisms responsible for determining how a particular item is alike or different from other items.

Since the ability to accurately segregate items into different classes represents a fundamental component of human cognition, it follows, then, that in order to gain a clearer picture of human cognition in general, it is essential to identify the factors that affect how people learn new categories. One factor that affects how new concepts are learned is the context in which learning takes place. There are at least three ways in which people learn concepts: learning over several examples, direct instruction, and through guided learning over examples. Learning over several examples represents an inductive method for learning categories and has been used to investigate exemplar-based models of category learning (Estes, 1994, 1993, 1986; Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1986, 1984) among others. It differs from direct instruction, which is equated with developing conceptual knowledge from external sources such as lectures and text, and places less emphasis on learning specific examples of the concept in order for learners to arrive at an understanding of the concept. Finally, guided learning over several examples represents a combination of the two former methods in that direct instruction is given prior to learning over several examples. This method has been used to test models of category learning that emphasize interactive relationships between background knowledge and exemplar similarity (Allen & Brooks, 1991; Nakamura, 1985). Since each of these learning paradigms foster learning in different ways and make different assumptions with regard to how concepts are mentally represented, each will be discussed in turn before addressing how they are used to test models of category learning.

Learning Over Examples

The most widely investigated method for learning new concepts is by learning over several examples, which is a bottom-up, inductive method for learning new concepts (Medin, 1989). It is an inductive method because it requires learners to uncover commonalities and differences among category items, or rules that define category

membership based on the features present in the stimuli. As a result, it places greater emphasis on encoding surface features of items during learning than other sources of category information, particularly when neutral category labels are used (e.g., Category A versus Category B), and when the features that make up the stimuli do not activate pre-existing knowledge (Hayes & Taplin, 1995; Pazzani, 1991). Under this learning paradigm, participants are shown several items in succession and corrective feedback is given after each classification attempt when necessary. This is repeated until learners are able to categorize the set of training items with a high degree of accuracy. Then, after some learning criterion is met, participants are given the set of old items intermixed with new items. This paradigm has been used mostly to compare models of category learning that base classification decisions on similarity to category prototypes (Posner & Keele, 1968), or to individual items stored in memory (Estes, 1994, 1993, 1986; Kruschke, 1992; Nosofsky, 1986, 1984). Other uses of this method have been employed in studies of rule induction (Nosofsky, Clark, & Shin, 1989; Nosofsky, Gluck, Palmeri, McKinley, & Glauthier, 1994; Shepard, Hovland, & Jenkins, 1961), where the learning goal is to uncover rules that identify positive examples of a target concept, while at the same time, excluding negative examples.

Since learning over examples represents a primarily bottom-up method for learning concepts, participants may display perfect classification accuracy even though they may not know why a given item is a member of one category and not another (Wisniewski & Medin, 1994). As a result, the only information that can be brought to bear on the learning task comes directly from the perceptual affordances in the stimuli coupled with the category labels that are provided as feedback after classification. One difficulty with this is that in real learning situations, information other than that pertaining to the perceptual characteristics of features is also used by learners when making classification decisions, as evidenced by faster learning rates for categories that

draw on pre-existing knowledge (Hayes & Taplin, 1995; Nakamura, 1985). This issue will be revisited in a later section.

Direct Instruction

A second method for learning concepts is through direct instruction. Unlike purely inductive methods, direct instruction allows learners to arrive at an understanding of target concepts directly from external sources such as lectures or text. It is equated with providing background information that identifies the functional or causal relationships between the features that occur in the stimuli and category labels, but can also include information pertaining to specific items (Estes, 1995). At present, there is very little research that contrasts this type of learning method with inductive methods within the context of learning new concepts. The research that is available indicates that acquiring information directly from external sources may provide explicit rules for determining category membership (Allen & Brooks, 1991), or it may provide a basis for developing intuitive theories that learners use to provide explanations for why certain features are linked to categories (Murphy & Medin, 1985; Wisniewski & Medin, 1994).

Clinical diagnosis represents a special case in which various psychiatric disorder categories are capable of being understood through direct instruction (Mumma, 1993). When making a diagnosis, the constellation of behaviors (features) observed in a patient can be classified based on whether they fit the criteria for a psychiatric disorder, and these criteria are typically expressed in terms of conjunctive-disjunctive rules. For example, a diagnosis of Major Depressive Episode requires that five out of nine symptoms be present during the same two-week period with at least one of two core features (e.g., depressed mood or diminished interest) being present as well. From this, it is possible to use the information acquired from external sources, such as diagnostic manuals, to determine whether a given example satisfies the conditions for a particular

disorder. This leads to speculation as to how common conjunctive-disjunctive rules are for other types of categories in addition to whether they represent a type of rule that can be easily learned from text. Allen and Brooks (1991) argue that such rules (which they describe as additive rules) are useful for conveying category information to novices. They argue that additive rules are capable of “generating family resemblance structures that have been held useful for verbally characterizing ill-defined, natural-kinds categories” (p. 5). This suggests that learning ill-defined category structures benefit from application of additive rules, and that additive rules can be easily learned from external sources. What remains to be determined is whether additive rules can be inferred from text and whether such rules can be applied to a set of training items.

What follows is a brief discussion of past and current theories of categorization, beginning with similarity-based views and leading up to more recent developments. For each theory, the specific assumptions regarding how category information is mentally represented and implemented during learning will be highlighted in addition to relevant empirical findings that have been used to generate and develop extant categorization models. Towards the latter portion, it will be argued that categorization models based exclusively on the principles of exemplar similarity must incorporate a mechanism to account for the effects of background knowledge in order to fully capture the essence of human categorization.

Early Theories: The Classical Approach

The classical approach represented an early attempt to identify the structural basis of categories. Under the classical view, categories are assumed to be represented as a summary description of an entire class of objects bound by singly necessary and jointly sufficient features. Features that are singly necessary for category membership require that they be included in every example of the category, and in order for a set of features

to be jointly sufficient, this implies that every exemplar that possesses a set of necessary features must be a member of the category. In essence, categories were believed to be composed of defining features, where all members were equally representative of the category. Under the classical view, learning is regarded as a function of induction of features or rules that define category membership. One of the first empirical explorations of categorization was reported by Bruner, Goodnow, and Austin (1956), whose influential investigations regarded category learning as a problem solving task. This research was more concerned with uncovering strategies people use when learning categories and was less concerned with providing support for the classical view, even though the majority of the category structures they investigated were structured according to defining features. They reported that people tend to use systematic strategies in their effort to discover some unknown concept. It was assumed that category representations were not formed through passive reception of repeated occurrences with category exemplars. Rather, concepts were formed through an active, strategic, hypothesis-testing process. For example, a particular item (e.g., a single red square with two borders) was designated as a positive instance of an unknown category prior to learning. The task of the participants was to discover the unknown category by formulating a hypothesis based on the single example. Subsequently, the participants were allowed to freely select any other instance from an array (or universe) of items, whereupon the experimenter provided feedback regarding whether the selected instance was a positive or negative example of the category. Additionally, after receiving feedback, participants were asked to provide a revised hypothesis regarding category membership. This was repeated until the participants correctly extracted the rule that bound all members of the category, and excluded all non-members. From this, participants took more trials to learn categories that were structured according to disjunctive rules relative to those that were constructed according to conjunctive rules; a finding which argues that the structure of categories

affects the probability of subsequent learning (see also Shepard, Hovland, & Jenkins, 1961). Thus, after category learning had been completed, it was assumed that the category representation consisted of a summary representation of whatever logical rule defined category membership (e.g., IF [color = red] ^ [shape = square] THEN object is a member of Category A).

One problem with the classical view is that most categories are rarely represented by defining features. Although some categories have singly necessary and jointly sufficient features (e.g., geometric shapes like triangles and squares), most do not. Rather, exemplars are graded with regard to how well they represent members of a category. For example, even well-defined categories like triangles may not be composed of members that are equally representative of the category. It is possible that people may consider some triangles to be more representative of the concept triangle than other members. For natural categories, this possibility is more likely. That is, robins are considered to be more representative or typical members of the category birds than penguins (e.g., Rosch & Mervis, 1975). Thus, any theory of categorization must be able to account for typicality effects. This limitation of the classical view is addressed by probabilistic theories.

A Move Away From Defining Features: The Probabilistic Approach

The probabilistic approach represents the first shift away from the classical view of categorization (Smith & Medin, 1981). Unlike the classical view, probabilistic approaches assume that categories are composed of members that can vary in typicality or category representativeness (Rosch & Mervis, 1975). By this same token, the probabilistic approach also rejects the notion that most categories are describable according to necessary and sufficient features. Thus, members of a category can possess attributes that are not shared with all other members of the category. Under this view,

highly representative exemplars of a category tend to have more features in common with other members of the category, but few or no features in common with members of contrast categories. In contrast, less representative members of a category tend to possess fewer features that are shared with other items from the same category, or possess many features that overlap with members of contrasting categories. Thus, the degree with which a particular item fits into a category is based on a weighted sum of the number of common and distinctive features that the item possess with respect to contrast categories (Reed, 1972). Rosch and Mervis (1975) used the notion of family resemblance as a way of referring to the structure of natural categories. Family resemblance is computed by summing the number of features in a particular item that are shared with other members of the category. Rosch and Mervis reported that this measure was highly correlated with learning rates, speed of classification, and ratings of typicality for both natural and artifact categories. However, when family resemblance and the degree of featural overlap with contrasting categories were in conflict, such items were less likely to be rated as being prototypical relative to when there was less overlap with contrasting categories. Thus, one key aspect of family resemblance is that it serves as a structural basis for categories, and it does so without resorting to the notion that categories are structured according to defining features.

The probabilistic approach does retain one aspect of the classical view by assuming that categories are represented by summary descriptions, which are expressed in terms of abstracted central tendencies (e.g., prototypes) or weighted feature lists (Smith & Medin, 1981). One key aspect of abstraction is that prototypes need not correspond to any specific item that had been previously encountered. However, specific items are often judged as being prototypical when they share more features in common with members of the same category and few features with members of contrasting categories (Rosch & Mervis, 1975). An early investigation using prototypes as the basis

for classification was reported by Posner and Keele (1968). In prototype models of categorization, it is assumed that over the course of learning, participants tend to form prototypes through abstraction of the average or modal feature values over all members of the category that are presented. Classification decisions are then based on a comparing presented items to the abstracted prototypes. Classification accuracy is increased when presented patterns are highly similar to abstracted prototypes, with similarity a decreasing function of the distance in multidimensional space between presented items and prototypes (Reed, 1972) or as the number of mismatching attributes increases (Smith & Medin, 1981). In the experiments reported by Posner and Keele (1968), participants were shown configurations of seemingly random dot patterns. Actually, the dot patterns were constructed from distortions from one of four prototypes such that they possessed high and low levels of variability relative to the prototypes. The prototypes themselves were never shown to participants during learning, but they consisted of a triangle, the letters M and E, and a random dot pattern. During the learning phase, participants indicated which of these categories a particular pattern was a member of, and received feedback regarding the accuracy of their selections. The rate at which the categories were learned depended on the variability of patterns to prototypes. Patterns that had a high degree of variability took longer to learn than patterns that had a low level of variability. However, during transfer, participants who were trained on patterns with high variability classified prototypical patterns more accurately than old and new patterns. Similar results were observed with participants who were trained on patterns with low variability, but their classifications tended to be significantly less accurate relative to those who were trained on high variability patterns, a finding which suggests that variability is an important factor in learning. Other research also supports this finding (Homa & Vosburgh, 1976). For example, Homa and Vosburgh reported that categories learned with mixed levels of distortion benefitted transfer performance for new patterns and prototypes when the

number of training examples was increased, but when training patterns were of uniformly low distortion, category size had negligible effects on transfer. In addition, this pattern of results was preserved after a delay of ten weeks, suggesting that category prototypes are resistant to forgetting.

Although this research indicates that classification performance for prototypical items is consistent with the notion that people form prototypes during learning, it does so at the expense of excluding other relevant information, like specific information represented by previous examples. Indeed, this has been regarded as one of the weaknesses of appealing to summary representations in general and to prototype models in specific (Smith & Medin, 1981). Other problems specific to prototype models is that they consider similarity to prototypes as an additive function of common and distinctive features. Because of this, prototype models predict that linearly separable categories are easier to learn than categories that are not linearly separable because they can be partitioned on the basis of an additive combination of weighted features. The problem with this distinction is that linearly separable categories are generally not easier to learn than non linearly separable categories. Actually, it is often the case that non linearly separable categories are easier to learn than linearly separable categories (Medin & Schwanenflugel, 1981), except when background information is biased towards learning linearly separable categories (Nakamura, 1985; Wattenmaker, Dewey, Murphy, & Medin, 1986). Perhaps the most notable problem with prototype models is that they require abstraction as a key component, or at least some mixture of abstraction and exemplar storage (Posner & Keele, 1969). If it can be demonstrated that abstraction is not a key component of categorization, this would seriously undermine the utility of prototype models in favor of alternative models of categorization that rely purely on exemplar storage.

Exemplar-Based Views of Categorization

The exemplar-based approach represents a third shift in theories of category learning (Smith & Medin, 1981). Exemplar models share some characteristics with prototype models in that they both assume that categorization is driven by statistical regularities between probes and stored representations, and both reject the notion that categories are explainable solely in terms of defining features. However, they depart from one another in terms of how similarity is computed and how category information is represented in memory. Unlike prototype or family resemblance models, the exemplar-based view asserts that categories are represented by individual items that are stored in memory and categorization of new material is driven by computation of its similarity to these stored examples using a multiplicative similarity rule (e.g., Medin & Schaffer, 1978). Specifically, categorization stems from a matching process that operates on a feature by feature basis where the features of a probe item are compared against the features that make up the items represented in long term memory. The similarity between two items increases when the number of common features between items increases, and when the number of distinctive features between items decreases. The implication is that items from the same category tend to have more features in common with one another relative to items from other categories.

Regarding learning, the exemplar-based approach first assumes that exemplars can be described according to a pre-specified and unambiguous feature space. That is to say that all items are composed of features that are easily perceptible to the learner, and the features that are entered into quantitative models are the same as those perceived by learners. Second, learning operates by selecting features from the feature space. Third, selection of features is driven by statistical properties of features such that common and distinctive features become discriminated, and is equated with attending to diagnostic features and ignoring features that are not diagnostic of category membership. Fourth,

classification is based on a matching function to determine which features of an item match and mismatch those in the category (Estes, 1993; Medin & Schaffer, 1978; Nosofsky, 1984).

This version of the exemplar-based model listed above is referred to as the context model by Medin and Schaffer (1978). Another version of this model was proposed by Nosofsky (1986) who adapted the model to fit continuous dimension stimuli using multidimensional scaling procedures. In Nosofsky's generalized context model (GCM), exemplars are represented as points in multidimensional space with similarity computed according to distance in this space. Thus, when the distance between any two exemplars in this space increases, the similarity between them decreases exponentially. One feature of these models is that they are able to account for prototypicality effects without the need for prototype abstraction. For example, items that correspond to prototypes presented during transfer tend to be classified more accurately than other items because they are more similar to stored examples within the category relative to items from contrasting categories. In addition, Medin and Schaffer (1978) demonstrated that in some situations, many previously learned training patterns were classified more accurately than prototypical patterns, a finding which is contrary to that predicted by prototype models.

The context model and versions derived from it have been excellent at accounting for a wide range of classification phenomena, and have been able to overcome many of the problems associated with earlier theories. However, empirical tests of the context model have typically focused entirely on transfer performance, and in doing so, they have often failed to examine classification performance over the course of learning. In order to account for sequential learning, it was argued that learning in exemplar-based models was based on a gradual accumulation of exemplars in memory such that the strength of any given exemplar is increased by a constant amount each time it is presented (Estes,

Campbell, Hatsopoulos, & Hurwitz, 1989). Other versions of the context model were augmented with mechanisms that were sensitive to recency effects and background noise (e.g., Nosofsky, Kruschke, & McKinley, 1992). However, simply adding additional free parameters to provide more accurate predictions may fail to provide an account of the dynamics of learning. This limitation in exemplar models was initially observed by Gluck and Bower (1988) who argue that learning is error-driven (cf. Estes et al., 1989; Kruschke, 1993; Nosofsky, Kruschke, & McKinley, 1992). For example, in the experiments reported by Gluck and Bower (1988), participants were required to classify exemplars containing four medical symptoms into two disease categories. However, the probabilities of each of four symptoms given a particular disease category was varied as was the frequency with which exemplars occurred in the two disease categories. Specifically, one of the disease categories was appropriate for 25% of the exemplars (i.e., the rare disease) and the other disease category was appropriate for 75% of the items (i.e., the common disease). In addition, one of the symptoms (Symptom 1) occurred more frequently in the rare category than other symptoms and less frequently than other symptoms in the common category. The actual probability of Symptom 1 given the rare category was .6 and .2 given the common category. However, since the base rate of the categories were varied during learning, the normative probability Symptom 1 was equal for both rare and common disease categories. That is, Symptom 1 occurred with equal frequency in both disease categories. The most notable finding reported by Gluck and Bower, however, was that participants failed to apply category base rates during transfer. Specifically, participants estimated a higher probability of a rare disease when given only Symptom 1 relative to the common disease category. This finding has also been replicated by Estes et al. (1989) and Nosofsky et al. (1992), but is difficult for purely exemplar-based models to account for. However, such a finding is readily captured by a

network model that operates on error-driven learning principles, like the adaptive network models of Gluck and Bower (1988).

The Adaptive Network Model

Two versions of the adaptive network model were tested by Gluck and Bower (1988). They include the component- and the configural-cue model, and both differ from one another with regard to the type of architecture that is employed. The component-cue model is implemented as a connectionist network with two layers of processing nodes: an input layer which codes featural values, and an output layer that corresponds to category labels. The configural-cue model is also a two-layered connectionist network, but it differs from the component-cue model only in terms of how feature values are represented on the input layer. Rather than only allowing stimulus dimensions to be coded in terms of single features, the configural-cue model also allows stimulus dimensions to be coded in terms of pairs, triples, and n-tuples of features on the input layer. Depending on the nature of the input patterns, both versions can be equipped to handle feature dimension values that are either present or absent (additive features), or when one of two positively existing feature values are present on each trial (substitutive features). Focusing specifically on the component-cue model, when additive features are present, each feature is coded by a single input node that detects the presence or absence of a particular feature, but when substitutive features are present, each feature dimension is coded by a pair of input nodes. All nodes on the input layer are fully connected to each of the category nodes represented on the output layer where the activations of all input nodes are multiplied by the current connection weights and are summed to form outputs. The connection weights are adjusted after each trial using an interactive, error-driven learning rule (i.e., the delta rule) such that the discrepancy between predicted and observed response probabilities over the course of learning become reduced. The benefit

of this type of learning rule is that it gradually changes the connection weights between features and category labels, unlike the augmented version of the context model that incrementally increases the strengths of exemplars by a constant amount each time they are presented during learning (Nosofsky et al., 1992). As a result, the adaptive network model is often better able to predict learning performance than purely exemplar-based approaches because the error-driven learning component allows the connection weights to be adjusted less when there is little errors of classification, as would be the case during later stages of learning relative to earlier stages. In addition, the interactive nature of learning in connectionist networks allows the individual cues to compete with one another to become associated with the alternative categories. Feature values that are relatively better predictors become more highly associated with the category such that larger connection weights develop between predictive cues and category labels over the course of learning. In terms of how the adaptive network model is able to predict base rate neglect, it learns that Symptom 1 is a relatively poor predictor of the common disease category relative to other features. Through the use of the interactive learning mechanism, the connection weight between Symptom 1 and the rare disease category becomes larger over the course of learning than the connection weight between Symptom 1 and the common disease category. As a result, the model was able to correctly predict base rate neglect phenomena more accurately than an exemplar-based learning model that did not incorporate error-driven learning (Nosofsky et al., 1992).

One problem with the component-cue model tested by Gluck and Bower (1988) is that is formally identical to a multiplicative-similarity prototype model in that activation of nodes on the input layer are multiplied by connection weights that are linked to the output layer where they are then summed to determine categorization probabilities (Nosofsky, 1991). Thus, the implication for the component-cue model is that its utility may be limited only to category structures that are linearly separable. However, even for

linearly separable category structures, the component-cue model is unable to provide accurate predictions regarding learning and transfer performance relative to models that include exemplar-based category representations (Nosofsky et al., 1992).

Regarding transfer performance, the component-cue model predicts that items corresponding to modal prototypes would be classified with equal or greater accuracy relative to previously presented training items. In contrast, exemplar-based models allow for old training items to be classified with greater accuracy relative to prototypical items, depending on the similarity relations between previously presented items. Taken together, the data reported by Nosofsky et al. (1992) and Estes et al. (1989) indicate that both exemplar-based category representations and connectionist learning rules need to be combined within a unified model. The goal of such an endeavor is designed to more closely capture both principles of error-driven learning and exemplar storage.

Exemplar-based Connectionist Models

Two models of category learning have incorporated exemplar representations combined with connectionist learning rules are ALCOVE (Kruschke, 1992), and the Exemplar-Based Back Propagation Model (EBP; Taraban & Palacios, 1993). ALCOVE is an extension of the context model (Medin & Schaffer, 1978; Nosofsky, 1984), but is formalized within a connectionist framework that selectively attends to the relevant stimulus dimensions and ignores those stimulus dimensions that are irrelevant. The Exemplar-based Back Propagation model (Taraban & Palacios, 1993) also incorporates three layers of processing units, but it diverges from ALCOVE in that it possess no special mechanism for selective attention. Rather, selective attention is incorporated by limiting the number of connections between units so that units on the input layer are connected only to nodes on the hidden layer that possess the same features. In both models, categories are represented by individual exemplars stored in memory, with

associations being formed between activated exemplars and the category labels (represented on the output layer) during learning.

Regarding the architecture of ALCOVE, it consists of three interconnected layers: an input, hidden, and output layer. The input layer consists of nodes that code individual stimulus dimensions of presented items, and each node is gated by dimensional attention weights such that activation of each node is determined by the psychological scale value of each item on each stimulus dimension. The hidden layer includes a set of nodes that correspond to individual exemplars, which are localized as specific points in multidimensional space, but can be randomly distributed across the stimulus space. These nodes become activated according to their similarity to presented patterns with similarity expressed as an exponentially decreasing function of distance in multidimensional space when distance is increased. Over the course of several learning trials, the dimensional attention strengths in ALCOVE adjust themselves so that exemplars from different categories become less similar, and exemplars within categories become more similar. As a result, ALCOVE learns to increase the attention strength on dimensions that are relevant to classification, and to decrease the attention strength on the irrelevant dimensions. This has the effect of stretching or shrinking the receptive fields for each exemplar unit in stimulus space along dimensional axes. The output layer consists of nodes that correspond to category labels and where associations are formed between individual exemplars and category labels over the course of learning. Activation of output nodes is determined by a linearly combination of activations from exemplar nodes. Response probabilities are determined by the magnitude of a particular category node's activation value relative to the sum of all category activation values as in the adaptive network models tested by Gluck and Bower (1988).

During learning, when a probe item is presented, activation propagates from the input nodes to the output nodes. At the output layer, the obtained activations are then

compared against desired output activation levels. For a positive example, the model produces feedback by comparing observed outputs against teacher values which are coded as +1 when the item is a member of the category, and -1 when the item is not a member of the category. If the activation of a particular output node exceeds the teacher value (e.g., if it is greater than +1 for a positive example, or conversely, if it is less than -1 for a negative example) the difference is not counted as an error by the model. ALCOVE then attempts to reduce the discrepancy between obtained output values and desired output values by modifying the association and attention weights using gradient descent correction on error. The association and attention weights are then adjusted proportionally to the error produced by the model using back propagation (Rumelhart, Hinton, & Williams, 1986). When ALCOVE is fitted to learning data, it has four free parameters which represent learning rates for association weights between output nodes and hidden nodes, learning rates for attention weights between hidden and input nodes, a response scaling parameter, and a specificity parameter which determines the width of the receptive fields for exemplar units.

ALCOVE represents one of the more powerful models in that it is able to account for a wide range of categorization phenomena. This is due in part by its ability to capture some of the underlying principles that guide learning. One principle is that categories are represented by exemplars stored in memory (consistent with exemplar-based theories). Another principle is that learning is guided by selectively attending to features that are relevant for classification and ignoring those that are irrelevant, and a third principle is that learning is error driven (Kruschke, 1993a, 1993b). Evidence of the importance of selective attention to relevant features is typified in filtration versus condensation learning tasks. In the case of categories that are comprised of two feature dimensions, a filtration task requires one to attend only to one relevant feature dimension (e.g., size) while ignoring the irrelevant feature dimension (e.g., color) whereas condensation

requires attending to both relevant feature dimensions (e.g., size and color). When these two tasks are contrasted, filtration tasks are easier to learn than condensation tasks because they require that the learner only attend to one of the features. ALCOVE is capable of reliably uncovering the filtration advantage because it is able to selectively attend to the relevant feature, but when attention is divided between two features (as in condensation), learning is more difficult in ALCOVE. Standard backpropagation models (e.g., Rumelhart, Hinton, & Williams, 1986) typically fail to produce filtration advantages because they are incapable of selective attention to features (Kruschke, 1993a, 1993b).

Interactive Models

The above approaches to category learning all share one common component. Namely, they view category learning as a bottom-up, empirically-driven process. However, many recent findings indicate that such models neglect the role that background knowledge plays during category learning. In order to account for the effects of background knowledge, it is necessary to address top-down, or theory-based processes in addition to empirically-driven processes. Many current models have attempted to do just that, but each make different assumptions with regard to how background knowledge interacts with empirically-driven processes during learning. Models that include mechanisms for representing background knowledge include PostHoc (Pazzani, 1991), IOU (Mooney, 1993), the integration model (Heit, 1994), and the tightly coupled model of Wisniewski and Medin (1994). These models can be divided into those that include a theory-driven component that is separate from an empirical learning module (e.g., Post Hoc: Pazzani, 1991) from those that allow these two sources of information to interact during learning (e.g., IOU: Mooney, 1993).

Wisniewski and Medin (1994) indicate that a more correct approach to category learning must account for the interaction between exemplar similarity and one's background knowledge. They suggest that people's prior beliefs, goals, expectations, and theories restrict the set of features that are to be considered when making classification decisions. This pre-existing knowledge also is used to infer additional, often abstract features from observable ones. This approach does not imply that exemplar similarity should be abandoned. Rather, it suggests that the current view of similarity is insufficient on the grounds that it is too unconstrained to account for coherence in categories, and that categories are not reducible solely on the basis of the sum of the attributes they possess. For example, Murphy and Medin (1985) indicate that when complex stimuli are used, any two items can match (and mismatch) on an infinite number of attributes if there are no limits as to what is to be considered as a feature. In the case of perceptual categories, such as those used in many exemplar-based models, there is rarely a discrepancy as to what counts as a feature (e.g., color, shape, size, orientation, etc.). However, for natural categories (and many artificial categories), there are numerous possibilities as to what could be considered as a feature. For example, in the experiments reported by Wisniewski and Medin (1994), when participants were required to classify pictures that were labeled as being drawn by "creative" versus "non-creative" children, they searched for features that supported their intuitions regarding what constitutes creativity. From this search, participants inferred that drawings done by creative children were more likely to show "detail" or "action" relative to drawings that were done by non-creative children. Inherent in this claim is that features need not be represented as concrete features, as is often the case in perceptual categories. In order to determine whether participants based category judgments on concrete, or abstract features, those who were given theory-activating labels for categories were more likely to mention abstract features when probed. When concrete features were mentioned, they were often used to support the

abstract features. However, when participants were not given category labels that activated pre-existing knowledge (i.e., theory activating labels), they used concrete features more frequently than abstract features when making classification decisions.

It is not well understood how theories of a particular domain become established when they are constructed out of the categories they are supposed to explain. Thus, in order to understand how one applies their prior knowledge or beliefs during classification, it is important to determine how background information is acquired. In a manner consistent with the interactive view, Heit (1994) reported that the use of theory activating category labels cues the retrieval of previously stored category examples, which then exert their influence in conjunction with training items. Other research indicates that when attempting to classify novel objects (e.g., pounders or cutters), people draw on their previous experience with items that are used for pounding or cutting (like hammers and knives) in addition to their functional knowledge regarding features that a particular tool must have in order to pound or cut (Hayes & Taplin, 1995). Although the conditions under which prior theories become established remain unclear, the interactive view maintains that the theories held by the learner can be modified when given corrective feedback. Participants rarely modify their theories when given positive feedback. However, when given negative feedback, the features may be reinterpreted or the theory may be modified in light of new evidence (Wisniewski & Medin, 1994). In addition, Livingston and Andrews (1995) indicate that when prior knowledge initially selects features that are not diagnostic of category membership, these features are not always replaced by new, diagnostic features. Rather, learners continue to consider the irrelevant features in addition to supplemental features that are more diagnostic. Learning is impeded when one's prior knowledge emphasizes irrelevant features, but through experience with category examples, participants gradually incorporate the relevant features into their theories.

APPENDIX B
SHORT TEXT CONTAINING BACKGROUND
INFORMATION USED IN EXPERIMENT 1

Set A

All plants have the following parts: Roots, Stems, Leaves, and Flowers. These are used by plants in different ways to adapt to desert and mountain environments.

DESERT environments are characterized by: (a) extremely hot temperatures, (b) extended periods without rain, and (c) intense sunlight.

MOUNTAIN environments are characterized by: (a) extremely cold temperatures, (b) extended periods of rain, and (c) moderate sunlight.

There are two kinds of roots: TAP ROOTS and FIBROUS ROOTS.

Tap roots (see Figure 1) help plants survive in extreme heat and extended periods without rain. Fibrous roots (see Figure 2) help plants survive in extreme cold and extended periods of rain.

There are two kinds of stems: WOODY STEMS and HERBACEOUS STEMS.

Woody stems (see Figure 3) help plants survive in extended periods without rain and intense sun. Herbaceous stems (see Figure 4) help plants survive in extended periods of rain and moderate sun.

There are two kinds of leaves: BLADED LEAVES and COMPOUND LEAVES.

Bladed leaves (see Figure 5) help plants survive in intense heat and intense sun.

Compound leaves (see Figure 6) help plants survive in extreme cold and moderate sun.

There are two kinds of flowers: HEADED FLOWERS and SPIKED FLOWERS.

Headed flowers (see Figure 7) help plants survive by attracting large birds.

Spiked flowers (see Figure 8) help plants survive by attracting small insects.

Set B

All plants have the following parts: Roots, Stems, Leaves, and Flowers. These are used by plants in different ways to adapt to desert and mountain environments.

DESERT environments are characterized by: (a) extremely hot temperatures, (b) extended periods without rain, and (c) intense sunlight.

MOUNTAIN environments are characterized by: (d) extremely cold temperatures, (e) extended periods of rain, and (f) moderate sunlight.

There are two kinds of roots: TAP ROOTS and FIBROUS ROOTS.

Tap roots (see Figure 1) help plants survive in extreme cold and extended periods of rain. Fibrous roots (see Figure 2) help plants survive in extreme heat and extended periods without rain.

There are two kinds of stems: WOODY STEMS and HERBACEOUS STEMS.

Woody stems (see Figure 3) help plants survive in extended periods of rain and moderate sun. Herbaceous stems (see Figure 4) help plants survive in extended periods without and intense sun.

There are two kinds of leaves: BLADED LEAVES and COMPOUND LEAVES.

Bladed leaves (see Figure 5) help plants survive in extreme cold and moderate sun.

Compound leaves (see Figure 6) help plants survive in intense heat and intense sun.

There are two kinds of flowers: HEADED FLOWERS and SPIKED FLOWERS.

Headed flowers (see Figure 7) help plants survive by attracting small insects.

Spiked flowers (see Figure 8) help plants survive by attracting large birds.

APPENDIX C
LONG TEXT CONTAINING BACKGROUND
INFORMATION USED IN EXPERIMENTS 1 AND 2

Set A

ENVIRONMENTS: Of the many environments that are part of the Earth's ecosystem, two that are of particular importance include Desert and Mountain environments.

DESERT environments tend to be hot and dry with much sunlight and mild winters. However, few people realize that during the winter months, occasional rains deposit moisture deep below the surface of the soil. the moisture is only available for short periods lasting into the growing season. Few people also realize that extremely scorching temperatures are frequent during the growing season.

MOUNTAIN environments tend to be cool and rainy with relatively little sunlight and very cold winters. However, few people know that during the summer months, constant rains deposit moisture near the surface of the soil. This moisture remains available into the growing season. Few people also realize that extremely frigid temperatures are common during the growing season.

ROOTS: the roots of a plant are crucial in order for it to survive and adapt to its environment. roots anchor the plant in order to prevent it from being uprooted. Roots also absorb moisture and nutrients from the soil. Two common types include Tap and Fibrous root systems.

TAP ROOTS are long and thin and extend several feet below the surface of the soil (see Figure 1). However, most people do not realize that tap roots are capable of reaching water that is only found far underground. This type of root is also capable of protecting the plant when temperatures rise above 100 degrees.

FIBROUS ROOTS contain many branching fibers that extend several feet just underneath the surface of the soil (see Figure 2). Few realize that fibrous roots are most efficient when there is ample moisture located near the surface of the soil. In addition, this type of root is able to protect the plant when temperatures drop below 32 degrees.

STEMS: The stems provide support for plants. Stems are also capable of transporting water and nutrients and can synthesize food for the plant. there are many kinds of stems, but two that are very common are Woody and Herbaceous stems.

WOODY STEMS are very rigid and have a thick layer of bark on the outer surface (see Figure 3). Few realize that woody stems are very good at preventing moisture from escaping because they have an outer coating that water cannot easily pass through. Woody stems have few food-producing cells, which means that they can only produce food for the plant when there is much sunlight.

HERBACEOUS STEMS are very flexible and have a layer of tender, green tissue on the outer surface (see Figure 4). Few people know that herbaceous stems are unable to prevent moisture from escaping because they have an outer coating that water easily passes through. Herbaceous stems have many food-producing cells, which means that they can produce plenty of food for the plant even when there is little sunlight.

LEAVES: The leaves of a plant are used primarily to synthesize food for the plant. They convert carbon dioxide and sunlight into glucose and oxygen through the process of photosynthesis. Although there are many types, two that are very common include Bladed and Compound leaves.

BLADED LEAVES are a type of leaf that are thin and elongated (see Figure 5). Few people realize that bladed leaves are able to protect the plant when

temperatures rise above 100 degrees by curling up. Since these leaves tend to have a small surface area, they require a large supply of sunlight during the growing season in order to synthesize an adequate supply of food for the plant.

COMPOUND LEAVES are a type of leaf that are round and wide with many leaves attached to each leaf stalk (see Figure 6). Few people know that compound leaves are able to protect the plant when temperatures drop below 32 degrees by releasing a protective wax. Since these leaves have a large surface area, they require a small supply of sunlight during the growing season in order to synthesize an adequate supply of food for the plant.

FLOWERS: The flowers of a plant are used for reproduction. Their color, shape, and fragrance are designed to attract animals that pollinate the flowers. Of the many different varieties, two that are very common include Headed and Spiked flowers.

HEADED FLOWERS tend to be large and round with many petals radiating from the center (see Figure 7). Few people realize that headed flowers attract small birds that pollinate this kind of flower. It is also widely unknown that headed flowers produce a sweet fruit that is eaten by animals which then distribute the seeds.

SPIKED FLOWERS tend to be small and grouped together on a single stalk (see Figure 8). Few people realize that spiked flowers attract large insects that pollinate this kind of flower. It is also widely unknown that spiked flowers produce a bitter fruit that is eaten by animals which then distribute the seeds.

Set B

ENVIRONMENTS: Of the many environments that are part of the Earth's ecosystem, two that are of particular importance include Desert and Mountain environments.

DESERT environments tend to be hot and dry with much sunlight and mild winters. However, few people realize that during the winter months, occasional rains deposit moisture deep below the surface of the soil. The moisture is only available for short periods lasting into the growing season. Few people also realize that extremely scorching temperatures are frequent during the growing season.

MOUNTAIN environments tend to be cool and rainy with relatively little sunlight and very cold winters. However, few people know that during the summer months, constant rains deposit moisture near the surface of the soil. This moisture remains available into the growing season. Few people also realize that extremely frigid temperatures are common during the growing season.

ROOTS: the roots of a plant are crucial in order for it to survive and adapt to its environment. roots anchor the plant in order to prevent it from being uprooted. Roots also absorb moisture and nutrients from the soil. Two common types include Tap and Fibrous root systems.

TAP ROOTS are long and thin and extend several feet below the surface of the soil (see Figure 1). However, most people do not realize that tap roots are most efficient when there is ample moisture located near the surface of the soil. In addition, this type of root is able to protect the plant when temperatures drop below 32 degrees.

FIBROUS ROOTS contain many branching fibers that extend several feet just underneath the surface of the soil (see Figure 2). Few realize that fibrous roots are capable of reaching water that is only found far underground. This type of root is also capable of protecting the plant when temperatures rise above 100 degrees.

STEMS: The stems provide support for plants. Stems are also capable of transporting water and nutrients and can synthesize food for the plant. there are many kinds of stems, but two that are very common are Woody and Herbaceous stems.

WOODY STEMS are very rigid and have a thick layer of bark on the outer surface (see Figure 3). Few people know that woody stems are unable to prevent moisture from escaping because they have an outer coating that water easily passes through. Woody stems have many food-producing cells, which means that they can produce food for the plant even when there is little sunlight.

HERBACEOUS STEMS are very flexible and have a layer of tender, green tissue on the outer surface (see Figure 4). Few realize that herbaceous stems are very good at preventing moisture from escaping because they have an outer coating that water cannot easily pass through. Herbaceous stems have few food-producing cells, which means that they can only produce food for the plant when there is much sunlight.

LEAVES: The leaves of a plant are used primarily to synthesize food for the plant. The convert carbon dioxide and sunlight into glucose and oxygen through the process of photosynthesis. Although there are many types, two that are very common include Bladed and Compound leaves.

BLADED LEAVES are a type of leaf that are thin and elongated (see Figure 5). Few people know that bladed leaves are able to protect the plant when temperatures drop below 32 degrees by releasing a protective wax. Since these leaves have a large surface area, they require a small supply of sunlight during the growing season in order to synthesize an adequate supply of food for the plant.

COMPOUND LEAVES are a type of leaf that are round and wide with many leaves attached to each leaf stalk (see Figure 6). Few people realize that compound leaves are able to protect the plant when temperatures rise above 100

degrees by curling up. Since these leaves tend to have a small surface area, they require a large supply of sunlight during the growing season in order to synthesize an adequate supply of food for the plant.

FLOWERS: The flowers of a plant are used for reproduction. Their color, shape, and fragrance are designed to attract animals that pollinate the flowers. Of the many different varieties, two that are very common include Headed and Spiked flowers.

HEADED FLOWERS tend to be large and round with many petals radiating from the center (see Figure 7). Few people realize that headed flowers attract large insects that pollenate this kind of flower. It is also widely unknown that headed flowers produce a bitter fruit that is eaten by animals which then distribute the seeds.

SPIKED FLOWERS tend to be small and grouped together on a single stalk (see Figure 8). Few people realize that spiked flowers attract small birds that pollenate this kind of flower. It is also widely unknown that spiked flowers produce a sweet fruit that is eaten by animals which then distribute the seeds.

APPENDIX D
COMPREHENSION QUESTIONS AND CORRESPONDING
ANSWERS (IN BRACKETS) FOR SHORT AND LONG TEXTS

Short Text

- 1) What parts do all plants have?
[Roots] [Stems] [Leaves] [Flowers]
- 2) What are the characteristics of DESERT environments?
[intense heat] [extended periods of no rain] [intense sunlight]
- 3) What are the characteristics of MOUNTAIN environments?
[extreme cold] [extended periods of rain] [moderate sunlight]
- 4) What do TAP ROOTS do to help plants survive?
Set A: [they protect the plant when there is no rain]
[they protect the plant from intense heat]
Set B: [they protect the plant when there is much rain]
[they protect the plant from extreme cold]
- 5) What do FIBROUS ROOTS do to help plants survive?
Set A: [they protect the plant when there is much rain]
[they protect the plant from extreme cold]
Set B: [they protect the plant when there is no rain]
[they protect the plant from intense heat]
- 6) What do WOODY STEMS do to help plants survive?
Set A: [they protect the plant when there is no rain]
[they protect the plant from intense sunlight]
Set B: [they protect the plant when there is much rain]
[they protect the plant when there is moderate sunlight]

- 7) What do HERBACEOUS STEMS do to help plants survive?
- Set A: [they protect the plant when there is much rain]
 [they protect the plant when there is moderate sunlight]
- Set B: [they protect the plant when there is no rain]
 [they protect the plant from intense sunlight]
- 8) What do BLADED LEAVES do to help plants survive?
- Set A: [they protect the plant from intense heat]
 [they protect the plant when there is moderate sunlight]
- Set B: [they protect the plant from extreme cold]
 [they protect the plant from intense sunlight]
- 9) What do COMPOUND LEAVES do to help plants survive?
- Set A: [they protect the plant extreme cold]
 [they produce food when there is very little sunlight]
- Set B: [they protect the plant from intense heat]
 [they produce food when there is much sunlight]
- 10) What do HEADED FLOWERS do to help plants survive?
- Set A: [they attract small birds] [they produce a sweet fruit]
- Set B: [they attracts large insects] [they produce a bitter fruit]
- 11) What do SPIKED FLOWERS do to help plants survive?
- Set A: [they attracts large insects] [they produce a bitter fruit]
- Set B: [they attract small birds] [they produce a sweet fruit]

Long Text

- 1) What parts do all plants have?
- [Roots] [Stems] [Leaves] [Flowers]

- 2) What are the characteristics of DESERT environments?
 [intense heat] [lack of surface moisture] [intense sunlight]
- 3) What are the characteristics of MOUNTAIN environments?
 [extreme cold] [large supply of surface moisture] [moderate sunlight]
- 4) What do TAP ROOTS do to help plants survive?
 Set A: [they absorb moisture from deep underground]
 [they protect the plant when temperatures rise above 100 degrees]
 Set B: [they absorb moisture near the surface]
 [they protect the plant when temperatures drop below 32 degrees]
- 5) What do FIBROUS ROOTS do to help plants survive?
 Set A: [they absorb moisture near the surface]
 [they protect the plant when temperatures drop below 32 degrees]
 Set B: [they absorb moisture from deep underground]
 [they protect the plant when temperatures rise above 100 degrees]
- 6) What do WOODY STEMS do to help plants survive?
 Set A: [they prevent moisture from escaping]
 [they produce food only when there is much sunlight]
 Set B: [they allow excess moisture to escape]
 [they produce food when there is very little sunlight]
- 7) What do HERBACEOUS STEMS do to help plants survive?
 Set A: [they allow excess moisture to escape]
 [they produce food when there is very little sunlight]
 Set B: [they prevent moisture from escaping]
 [they produce food only when there is much sunlight]
- 8) What do BLADED LEAVES do to help plants survive?
 Set A: [they protect plants from hot temperatures by curling up]

[they produce food when there is much sunlight]

Set B: [they protect plants from cold temperatures by releasing wax]

[they produce food when there is very little sunlight]

9) What do COMPOUND LEAVES do to help plants survive?

Set A: [they protect plants from cold temperatures by releasing wax]

[they produce food when there is very little sunlight]

Set B: [they protect plants from hot temperatures by curling up]

[they produce food when there is much sunlight]

10) What do HEADED FLOWERS do to help plants survive?

Set A: [they attract small birds] [they produce a sweet fruit]

Set B: [they attract large insects] [they produce a bitter fruit]

11) What do SPIKED FLOWERS do to help plants survive?

Set A: [they attract large insects] [they produce a bitter fruit]

Set B: [they attract small birds] [they produce a sweet fruit]