

## INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

**The quality of this reproduction is dependent upon the quality of the copy submitted.** Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning  
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA

**UMI**<sup>®</sup>  
800-521-0600



**ON THE RULES-TO-EPIISODES TRANSITION IN CLASSIFICATION:  
GENERALIZATION OF SIMILARITY AND RULES WITH PRACTICE**

By

**TIMOTHY J. WOOD**

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfilment of the Requirements

for the degree

Doctor of Philosophy

McMaster University

(c) Copyright by Timothy Wood, March 1998

DOCTOR OF PHILOSOPHY (1998)  
(Psychology)

MCMASTER UNIVERSITY  
Hamilton, Ontario

**TITLE:** On the rules-to-episodes transition in classification:  
Generalization of similarity and rules with practice

**AUTHOR:** Timothy J. Wood, B. A. (University of Lethbridge)

**SUPERVISOR:** Professor Lee R. Brooks

**NUMBER OF PAGES:** ix, 142

## ABSTRACT

When classifying a novel object, people often rely on similarity to previously learned instances to help them identify the category of the object. People also rely on more analytic knowledge, like classification rules, to identify the category of a novel object. The coexistence of two different procedures that could be used to classify the same novel object raises the issue of coordination; that is, what is the relation between rule-based and similarity-based classification.

Learning a new category is much like acquiring any cognitive skill, therefore it is hypothesized that current theories of skill acquisition should provide a useful framework for studying the relation between rule-based and similarity-based classification. A popular view of skill acquisition, Logan's Instance Theory (1988), suggests that skilled performance can develop rapidly, possibly after only a few practice trials on specific instances. The underlying basis of this rapid skill development is attributed to a transition from reliance on slow analytic procedures to reliance on faster retrieval-based procedures. Given what is known about memory, this transition to a retrieval-based classification procedure makes a great deal of intuitive sense and should account for the faster, more efficient classification performance that is characteristic of highly practiced individuals. The experimental conditions were designed to test an extreme example of this transition from rule-based to retrieval-based categorization. If a transition to a fast, retrieval-based procedure is a reasonably automatic result given sufficient practice, then it should occur even when participants are given a classification rule that is simple, perfectly predictive, and easy to apply.

Although there was some evidence of participants relying on similarity to prior instances for classification, this reliance on similarity never approached the levels that would be expected if participants had abandoned the rule. In fact, a reliance on similarity

was limited to novel stimuli that were so similar to the training stimuli that they were actually falsely recognized as old. When people realized a novel stimulus was new, they relied on the classification rule. This “retrieve if believed old, rule if believed new” strategy held across a number of manipulations all designed to facilitate an increased reliance on similarity including extensive practice, training with multiple similar neighbors, speeded classification, and reducing the novelty cues present within the transfer stimuli. Continued reliance on rules to classify novel items believed to be new and false recognition of novel items believed to be old raises questions about the role of similarity and the assumption that a transition from rules to retrieval is a relatively automatic occurrence of skilled classification.

## TABLE OF CONTENTS

TITLE PAGE.....	i
DESCRIPTIVE NOTE.....	ii
ABSTRACT.....	iii
TABLE OF CONTENTS.....	v
LIST OF TABLES.....	vii
LIST OF FIGURES.....	viii
CHAPTER 1: Introduction.....	1
Thesis overview.....	3
CHAPTER 2: literature review.....	6
What is meant by rules and retrieval .....	6
Evidence for the use of classification rules.....	13
Evidence against the use of classification rules.....	15
Alternatives to rule-based classification.....	16
Evidence for the use of both rule-based and retrieval-based knowledge.....	18
The importance of instructions for the control of rule-based and retrieval-based knowledge.....	19
Other instructional manipulations that control rule-based and retrieval-based knowledge.....	25
Non-instructional control of rule-based and retrieval-based knowledge.....	27
A summary of the factors that control the rule-based and retrieval-based knowledge.....	30
Theoretical models of the control of rule-based and retrieval-based knowledge.....	31
Coordinating rules and retrieval with the development of skilled classification.....	36
Determining the use of rules and retrieval.....	43
Overview of the experiments in this thesis.....	46

CHAPTER 3	
Experiment 1: The effect of practice.....	48
CHAPTER 4	
Experiment 2: The effect of many similar training items.....	61
CHAPTER 5	
Experiment 3: The effect of a deadline.....	72
CHAPTER 6	
Experiment 4: The role of novelty detection.....	83
Experiment 5: Unconfounding novel cues from number of neighbors.....	94
CHAPTER 7	
Experiment 6: Making Old/New Explicit.....	104
CHAPTER 8	
General Discussion.....	117
An assessment of the relations between similarity and false recognition.....	122
A reassessment of the relations between rule-based and retrieval-based classification .....	126
CHAPTER 9	
References.....	131
APPENDIX A.....	140
APPENDIX B.....	141
APPENDIX C.....	142



## LIST OF TABLES

Table 1 : Examples of categories structures around lists of fixed features.....	9
Table 2: An example of the procedure used to determine the relative importance of rules and retrieval.....	44
Table 3 : Mean responses to high practiced and low practiced transfer stimuli (Experiment 1).....	58
Table 4 : Mean responses to transfer stimuli that are similar to multiple training instances (Experiment 2).....	69
Table 5: Mean responses to transfer stimuli when a response deadline is imposed at transfer (Experiment 3).....	78
Table 6 : Mean responses to transfer stimuli when a novelty cue is present versus when a novelty cue is reduced (Experiment 4).....	91
Table 7 : Mean responses to transfer stimuli when a novelty cue and number of neighbors is unconfounded from each other (Experiment 5).....	101
Table 8 : Proportion of old responses to transfer stimuli when the old/new discrimination was made explicit (Experiment 6).....	110
Table 9 : Proportion of rule-inconsistent responses and the probability of a rule-inconsistent response conditional on the old/new discrimination when that discrimination was made explicit (Experiment 6).....	112

## LIST OF FIGURES

Figure 1 : An example of an artificial grammar.....	13
Figure 2 : Examples of some of the stimuli used in the extra practice experiment (Experiment 1).....	49
Figure 3 : Mean response times, corresponding fitted power functions, and mean error rates for high practiced and low practiced training stimuli (Experiment 1).....	55
Figure 4 : Examples of two families of stimuli used in the multiple neighbor experiment (Experiment 2).....	64
Figure 5 : Mean response times, corresponding fitted power function, and mean proportion of error when there are multiple similar training neighbors (Experiment 2).....	68
Figure 6: Examples of some of the stimuli used in the deadline experiment (Experiment 3).....	74
Figure 7 : Mean response times, corresponding fitted power function, and mean proportion of errors when there is a response deadline (Experiment 3).....	76
Figure 8 : Examples of some of the stimuli used in the novelty reduced experiment (Experiment 4).....	84
Figure 9 : Mean response times, corresponding fitted power functions, and mean error rates for the novelty present and novelty reduced conditions. (Experiment 4).....	88
Figure 10 : Examples of some of the stimuli used in the experiment in which the number of similar neighbors and the degree of novelty is unconfounded from one another (Experiment 5).....	95

Figure 11 : Mean response times, and fitted power functions when the number of similar neighbors is unconfounded from the novelty alert (Experiment 5)..... 98

Figure 12 : Examples of some of the stimuli used in the experiment in which the old/new discrimination was made explicit (Experiment 6)..... 105

Figure 13 : Mean response times and fitted power functions when the old/new discrimination was made explicit (Experiment 6). ..... 108

## CHAPTER 1

### *Introduction*

When classifying an object that is novel, adults often rely on simple rules to decide whether that particular object is a member of a category. These rules usually consist of lists of features that are important for that category. For example, bird guide books contain information about the location or nature of certain defining colorations, art history books encourage readers to look for certain styles or mediums characteristic of a period or artist, and medical students are taught to look for certain symptoms characteristic of specific illnesses when determining a diagnosis. In addition to using this rule-based knowledge, adults are also able to remember other familiar objects that may be similar to the novel object. The category of the novel object can then be determined on the basis of its similarity to other known members of a category. Given that either type of knowledge can be used to classify a novel object, an issue of interest how is do we know when to use one type of knowledge or the other; that is, how do we coordinate our use of rule-based and retrieval-based processes when classifying novel objects.

In our everyday environment, we have considerable practice interacting with, and therefore classifying, a host of examples from specific categories. Given that practice is also the significant factor in acquiring a skill, it is possible that current theories of skill acquisition may provide a useful framework for studying how the use of rules and retrieval are coordinated when classifying objects. A common finding in the skill acquisition literature is that there is a transition in how we perform a task as we become more practiced. As novices, our performance is initially slow and deliberate and we often rely on simple rules and explicit examples to perform a task. As we become more practiced at that task we start to rely on other forms of knowledge and our performance

becomes faster, and less deliberate. A similar transition occurs in classification. There is a transition from relying on slow and explicit rule-based knowledge to relying on other less explicit, faster types of knowledge. Subjectively, this increase in expertise is often experienced as gradually being able to “just recognize” the category of an instance with the importance of checking with the original rule being less important.

There are two broad intuitions from the skill acquisition literature that could account for a transition from deliberate rule application to more automatic recognition of category membership. Both intuitions assume that we start with the slow, deliberate use of a rule-based procedure. According to the first intuition, with practice, we develop a more proceduralized version of the classification rule (Anderson, 1982, 1993). Although the process of applying the proceduralized rule is less explicit and is much faster than applying the original rule, it is still a rule-based process in the sense that the rule-relevant features are still the primary basis of categorization. According to the second intuition, we accumulate a stock of interpreted instances with practice, and these instances are then retrieved whenever a familiar item or a highly similar item is presented (Logan, 1988). Retrieval of prior instances is also less explicit and much faster than applying the original rule but unlike the proceduralized rule, it is not considered a rule-based process because the effective cues are those that facilitate retrieval of a prior instance rather than those that are relevant to the rule.

Reflections about the normal use of memory retrieval in our everyday environment would support the idea that this transition to a faster, less effortful classification procedure is retrieval-based, and that this retrieval-based process also extends to novel but similar-to-old stimuli. For example, many everyday stimuli provide conditions that should encourage a transition from relying on a rule to relying on memory. Classes of natural objects, such as birds or art, and some medical disorders, such as those in dermatology or radiology, consist of individually distinctive instances. These distinctive

instances provide good encoding and retrieval cues that should aid the use of memory. In addition, we seldom interact with completely novel examples of a category that may require a rule. Most instances of a category that we see in our everyday environment are either already familiar, or are highly similar to other known instances. This high level of familiarity should also provide good retrieval cues and facilitate the use of memory. A third finding about the normal use of memory is that under most conditions in our everyday environment rules and retrieval give the same response, so the continued use of an effortful rule does not make sense considering that the alternative is a less effortful process. Again, this should be conducive to using retrieval-based processes as a basis of classification. Another observation about the normal use of memory is that while highly practiced individuals are able to list the rule-relevant features that should be used to classify members of a category, often the features they list do not seem to be the ones they use to classify objects (Riesbeck & Schank, 1989). This finding suggests that highly practiced individuals are no longer meticulously checking the features specified in a rule. Furthermore, even when a rule is remembered, it is often remembered with difficulty, again suggesting that practiced individuals are relying on non rule-based knowledge. Anyone who has had to teach someone else to drive a car can attest to how difficult it is to explicitly remember the simple rules we learned as beginners.

In short, if the instances being classified are sufficiently distinctive to support efficient retrieval of prior instances, if the proportion of old and similar items is sufficiently high to reward use of memory, and if the cost of an error is low, then a transition to a fast classification process should reflect an increased reliance on memory and this use of memory should extend to items that are old and novel but similar-to-old.

### ***Thesis Overview***

The subsequent chapters of this thesis document an attempt to support a classification framework in which it was assumed that a transition from a slower, rule-based process to

a faster, retrieval-based process occurs relatively automatically as a part of our normal everyday category learning. In Chapter 2, past evidence for the use of rule-based and retrieval-based classification will be reviewed. Several factors that have been found to influence the use of these types of knowledge will be reviewed as will four potential models that were designed to account for how we use rules and retrieval. The chapter ends with a discussion of how two current theories in skill acquisition could account for the transition to a faster, less effortful classification strategy and then will describe a novel hypothesis about how rules and retrieval are used to classify novel objects.

The experiments in Chapter 3, 4 and 5 of this thesis were designed to test an extreme example of this rules to episodes transition. If the transition to a faster retrieval-based classification procedure is a relatively automatic occurrence given the right retrieval conditions and given enough practice then it should occur even when participants are given a simple, perfectly predictive, and easy to apply rule. In addition, to giving participants a rule, other factors were manipulated that were thought would facilitate the use of a retrieval-based process. These factors included extensive practice, training with multiple similar training neighbors and speeded responding during the transfer task. If a transition to retrieval-based processes occurs rapidly, even under these extreme circumstances, then it implies that most of our everyday support for a category is instance-based. If the transition to retrieval-based process occurs slowly and only after considerable practice, then it may suggest that a transition to retrieval is more complicated than was assumed and its use may be limited only to highly familiar items or for new items that are very similar to highly familiar items. Either finding would be theoretically interesting because its occurrence would support one of two very different views of everyday categorization.

These first three experiments provided some evidence for a rapid transition to retrieval-based classification because participants tended to classify old familiar stimuli

faster than the novel stimuli. This transition to retrieval, however, did not extend to novel but similar-to-old stimuli. Although there was some evidence of an increased reliance on retrieval when classifying these stimuli, it was limited to stimuli participants believed were old. If participants realized an item was novel they relied on the rule.

The experiments in Chapters 6 and 7 were designed to confirm that participants do a recognition check before classifying the stimuli and adopt a “retrieve if believe old, rule if believed new” classification strategy. The factor that appears to control the use of this strategy is the presence or absence of cues that signal novelty. When the cues are reduced, participants tend to treat the novel items as old and produce responses consistent with similarity to prior instances. When the cues are present, participants treat the items as new and produce responses consistent with the rule. This result is interesting especially considering that many models of similarity-based classification assume that recognition and classification reflects two different uses of similarity and that good old/new discrimination reflects the use of memory whereas poor old/new recognition reflects the use of a rule (Nosofksy, 1988; Medin, 1986; also Cho & Mathews, 1996) .

The final chapter of this thesis presents a summary of the experiments and a discussion of the relations between similarity and false recognition, the relations between rule-based and retrieval-based classification, and concludes that a transition to a fast retrieval-based classification process is more complicated than originally assumed.



## CHAPTER 2

### *The coordination of rules and retrieval: a literature review*

#### *What is meant by rules and retrieval*

To illustrate what is meant by the terms rules, or rule-based knowledge, and retrieval, or retrieval-based knowledge, consider the following example. Imagine that you are trying to become more cultured and have decided to learn about art. Your friends often talk about impressionism versus expressionism so you decide to learn how to tell the two artistic styles apart. After examining a book, which contains reproductions of famous paintings, you notice that all the impressionist paintings appear to have short brush strokes with color boundaries between strokes while the expressionist paintings seem to consist of unrealistic looking objects that are broken up into pieces. This observation can then serve as a classification rule that will allow you to classify any further paintings as impressionist or expressionist. Anything with color swatches is an impressionist painting and anything with an unrealistic subject matter is an expressionist painting.

A retrieval-based classification strategy is also possible. Rather than looking for perceptual characteristics that define one style over another, you could try to remember information specific to each individual painting; information such as the name of the artist, the name of the painting, the content, and the artistic style. If you were later shown another painting that you had studied, you may recognize it, thereby remembering its category. Classifying a painting in this fashion is an example of a retrieval-based classification strategy. An additional classification strategy, related to retrieval, also exists. Imagine that you saw a painting that was not one that you had studied. You can still rely on retrieval to classify the painting by basing your classification on how similar it is to a painting that you remembered studying. This latter strategy is an example of a

similarity-based classification strategy.

The above example serves to illustrate how one might acquire and use a rule-based classification strategy versus how one might acquire and use a retrieval-based classification strategy. The example is a little simplistic however, in that there are a number of additional factors that can influence what strategy one uses. For example, one could start off trying to look for specific perceptual characteristics that could be used in a rule but not notice anything that distinguishes impressionism from expressionism and therefore decide to memorize the paintings. One could also start off trying to memorize the paintings and in the process notice that the paintings by Monet all seem to have bright color patches. After checking other impressionist paintings and noticing that this observation seems to be true of them as well, a useful classification rule is developed. It is also possible that one may begin by using a rule to classify the paintings but after seeing the same set of paintings repeatedly, start recognizing them as ones that had been studied and abandon the rule. In addition, doing a second task at the same time one is classifying a painting may make it hard to learn a cognitively demanding rule, while an emphasis on accuracy may override any desire to rely on similarity. In short, although the example serves to illustrate the two classification strategies, there are a number of other factors that can influence how these two types of categorical knowledge are used.

The conventional category learning experiment is often structured in a fashion similar to the art example above. The experiment first consists of a training phase, in which participants are presented with a series of stimuli and are asked to place the stimuli into one of the experimenter-defined categories. The training phase continues until a certain number of trials have occurred, or until the participant has reached a certain level of performance. The training phase is often followed by a second phase, or transfer task, in which participants classify a series of stimuli that consist of some they had studied and some that are novel. The purpose of the transfer task is to assess what was learned during

the training phase by measuring the degree to which the learned information generalizes to novel items.

To control the information that participants use during the training task, experimenters have tended to use artificial stimuli that allow easy manipulation of the structure of the categories when creating stimuli. There are two methods that have typically been used to create artificial stimuli. One method manipulates a list of fixed features, and the second method uses an artificial grammar.

The first method of creating stimuli assumes that our mental representation of a category consists of lists of features. A category, such as table, may consist of a list of features like number of legs, color, type of surface, and size of surface. Each example of a table possesses some combination of these features. When using this method to create stimuli, the experimenter first invents a number of possible features that they want their stimuli to have. These features may represent a variety of properties such as weight, color, shape; or sex, hobbies, type of vehicle. Additionally, there could be any number of potential features but every stimulus in a category consists of the same number and type of features. What differs between stimuli are the values that each feature may possess. Typically these values are binary, for example black versus white, round versus square, or male versus female.

This method of creating stimuli results in a category being represented as an array of ones and zeros. Table 1 contains a number of potential examples of arrays used to represent categories. In each case the column represents a feature and each row corresponds to a particular member of the category. When creating these arrays of ones and zeros the experimenter also needs some method of determining which stimuli belong to one category and which stimuli belong to another. A number of category structures exist including: simple rules involving one feature, complex rules involving more than one feature, features that are correlated with one another, similarity between the

---

Table 1.  
Examples of categories structures around lists of fixed features.

---

**Panel A:** A category structured around a single defining feature. All the members of Category A have a one on the first feature and all the members of Category B have a zero.

	<u>Category A</u>	<u>Category B</u>
Training stimuli	101010 110101 101011 110100	010101 001010 010100 001011
Transfer stimuli	111000	000111

**Panel B:** A category structured around a two out of three feature additive rule. All the members of Category A have a value of one on at least two of the first three features and the members of Category B have a value of zero on at least two of the first three features.

	<u>Category A</u>	<u>Category B</u>
Training stimuli	101000 110111 011010	010111 001000 100101
Transfer stimuli	110101	001010

**Panel C:** A category structured around a pair of correlated features. Whenever a Category A member has a value of one on the first feature it also has a value of one on the second feature.

	<u>Category A</u>	<u>Category B</u>
Training stimuli	110101 111010 111000 110111	010101 001010 001000 010111
Transfer stimuli	110000	011111

(Continued)

Table 1 (Continued)

**Panel D:** A category with a family resemblance, or prototype structure. The training items within each category have three of the five features in common with each other and four of the five features features in common with their prototype).

	<u>Category A</u>	<u>Category B</u>
Prototype	00000	11111
Training Stimuli	10000 01000 00100 00010 00001	01111 10111 11011 11101 11110
Transfer Stimuli	11000	00111

category members and a prototype and similarity between individual items.

If a category is structured around a rule that consists of one defining feature then all members of the category would have the identical value on that feature. For example, Panel A of Table 1 illustrates a structure in which all the members of one category have a value of one on the first feature and all members of the second category have a value of zero on the first feature. The value of zero or one on the first feature then serves as the rule for determining category membership. If a new stimulus, like 111100, is then presented for classification, the rule can be used to classify the item as belonging to Category A. Categories may also be structured around a rule that consists of more than one feature. Panel B of Table 1 illustrates two categories structured around a two out of three feature additive rule. In this case, all members of Category A have a value of 1 on at least 2 of the first three features while all members of Category B have a value of zero on at least two of the first three features. The presence of a one or a zero on at least two of the first three features serves as the classification rule and the application of this rule would lead to a new item like 111100 being classified into Category A. Panel C of Table 1 illustrates a category structured around a pair of correlated features. When a Category A exemplar has a value of one on the first feature it also has a value of one on the second

feature. Thus a novel exemplar like 110000 belongs to Category A but an exemplar like 011111 belongs to Category B.


A different type of structure, based on the similarity among the category members rather than a rule, is also used to structure categories. Similarity is usually defined in terms of the number of features that overlap between items. By this logic, items that have five out six features in common are more similar to one another than items that have one out six features in common. Although similarity is defined in terms of the overlap in features, how it is used varies. Panel D of Table 1 illustrates a use of similarity that requires a family resemblance structure. In this case, there is no one feature that has the identical value across all the members of a category. Therefore, a simple rule to define category membership does not exist. However, the members of Category A tend to have more features in common with each other than they do with the members of Category B. Thus, all the Category A exemplars have three out of five features in common as do all the Category B items. In addition, each category has a special member called a prototype. The prototype is that category member that is the most similar to every member of the category, or alternatively, the most average category member. For example, every exemplar in Category A has four out of five features in common with the prototype exemplar (00000), making the prototype the item that is most similar to every category member. Once a prototype has been learned, all stimuli can then be classified on the basis of their similarity to the prototype. For example, a novel item like 11000 belongs in Category A because it has three features in common with the Category A prototype but only two features in common with the Category B prototype.

There is a second use of similarity that does not require people to learn any special categorical representations like a rule or a prototype. This is called a similarity-based retrieval account of classification and according to this view stimuli are classified on the basis of their similarity to one another. For example, exemplars that are similar to one

another tend to be grouped together in a category, much like the family resemblance training stimuli in Panel D of Table 1. The major difference between this use of similarity and the use of similarity in a prototype model is that novel items are classified on the basis of their similarity to other specific training exemplars. Participants do not have to learn a prototype before they can successfully classify novel stimuli. Therefore, a transfer stimulus like 11000 could be classified as a member of Category A because it has four features in common with the training exemplars 10000 and 01000.

The second method of creating artificial stimuli for a classification study assumes that categories are structured around a complex set of underlying rules. Stimuli, created using this method, typically consist of sets of letter strings that were derived from an artificial grammar. Figure 1 contains an example of an artificial grammar that was used in Vokey and Brooks (1992). By starting at the entry node labeled with the number "1" and then following the arrows until the grammar is exited, stimuli can be created simply by keeping track of which letters are picked up as an arrow is traversed. The resulting letter strings form a category and are all structured around the same set of underlying rules. Stimuli are considered to be "grammatical", or belong to the same category, if the order of letters within a string is allowable according to the rules of the grammar. Stimuli are "nongrammatical" or do not belong to the category if the order of letters within the string are not allowable according to the rules of the grammar. For example, the novel string VXVR is possible under the rules of the grammar illustrated in Figure 1, and is therefore grammatical while the letter string MMT is not possible under the rules of the grammar and is therefore nongrammatical.

Stimuli created using an artificial grammar are also similar to one another. Like the categories structured around lists of features, similarity is determined by the number of letters that overlap between items. The more letters that are identical within the string,



Training Stimuli      VXVT  
                              VMT  
                              MVXRM  
                              MVXRMVR

Transfer Stimuli      VXVR - grammatical  
                              MMX - nongrammatical

Figure 1. An example of an artificial grammar used in Vokey and Brooks (1992) and several possible letter strings created using this grammar. In this example, grammatical means a letter string follows the rules of the artificial grammar and nongrammatical means a letter string does not follow the rules.

the more similar two items are. According to a similarity-based retrieval account, a novel letter string like "VXVR" may be classified as a member of Category A because it has three letters in common with the training string "VXVT" rather than because it does or does not follow the rules of the grammar.

***Evidence for the use of classification rules.***

The early focus of category learning studies was to investigate how the underlying regularities in the structure of a category were abstracted. That is, category learning



consisted of determining what feature or features were necessary to classify all members of the category and, once a successful rule had been found, classification was based on the presence or absence of these features (See Panel A and Panel B of Table 1). There were three key assumptions to this rule-based approach. First, once a rule was determined, the features used in that rule were defining. All members of the category had to have the particular feature or combination of features that were used in the rule. Second, all members of a category were equal members of the category with no one exemplar being better than another. This occurred because an exemplar either met the requirements of the rule and was in the category or did not meet the requirements and was not in the category. The third assumption was that the features, inherent to a category, were discrete. That is, the features clearly had one value or another, like black and white, and there was no uncertainty as to the value of a feature.

Within these early studies, two approaches were used to explain how the underlying regularities were learned, one an associationist account, the other a hypothesis testing account. Hull (1920) argued for an associationist account of rule-based classification. He presented participants with several lists of 12 pseudo-chinese characters and asked participants to learn how each character was pronounced. Although the same 12 pronunciations were used across the lists, the characters associated with a specific pronunciation always differed across the lists. Embedded within each character, however, was a specific feature that was consistently paired with a particular pronunciation. Hull found that as participants learned the lists, their accuracy at using the correct pronunciation for a character increased. He concluded that participants were abstracting the underlying feature that was paired with each pronunciation.

Although learning in the associationist view was fairly passive, the hypothesis generation approach assumed that people were actively forming and testing potential hypotheses about the structure of the category (Bruner, Goodnow, & Austin, 1956;

Haygood & Bourne, 1965). A series of studies reported by Bruner et al. (1956) represent the paradigmatic approach to this rule-based view of categorization. In their studies, the stimuli were either presented in an array or were presented one at a time. The stimuli themselves varied in four dimensions with each dimension having three possible values: color (red, green, black), shape (circle, square, cross), number of objects on a card (one, two, three), and the number of borders on a card (one, two, three). Usually, one stimulus was designated as an exemplar of a concept and the participants' task was to determine, with feedback, what the underlying rule was that determined membership. Using this procedure, Bruner et al. described several strategies that people use to help them learn the underlying classification rule and the relative ease of learning different types of classification rules (see also Haygood & Bourne, 1965).

#### *Evidence against the use of classification rules*

In the mid 1970's evidence had accumulated that was difficult to account for with a rule-based view of classification. For example, several researchers (Rosch, 1975; Mervis & Rosch, 1981; Rips, Shoben & Smith, 1973; McCloskey & Glucksberg, 1978) found that not all members of a category were considered to be equally good examples of the category. Rosch (1975) found that members of a category tend to have a graded structure; that is, some members of a category were rated as better members of a category than other members. Similarly, Rips et al. (1973), using a semantic verification task, found that some exemplars were verified as category members faster than other exemplars. McCloskey and Glucksberg (1978) also found that there was considerable disagreement between participants, and within the same participants after a one month delay, as to whether an exemplar belonged to a category. This disagreement was more pronounced for exemplars that were of medium typicality compared to exemplars that were highly typical or atypical of their categories.

Other researchers found that category members did not appear to have a set of

defining features. For example, Rosch and Mervis (1975) had participants list the features that belonged to specific members of a category and then looked to see how many of these features were shared with other members of the same category. They found that most of the features that were listed were not true of every member in the category. Similarly, Bellezza (1984) asked participants to list the features for a series of nouns and then return a week later and repeat the task. She found considerable variation in the features mentioned both within and between participants, particularly for less concrete nouns.

There is also evidence that participants do not actively test hypotheses when trying to learn categories and may not even have to be consciously aware of the existence of rules to be able to use them (Reber 1989; cf. Brooks, 1978). Reber (1976, 1967) presented participants with lists of letter strings that were created so that they conformed to an artificial grammar. Participants studied these letter strings and then, immediately before a transfer task, they were told that the letter strings were derived from a complex set of rules and that they were to classify a new set of letter strings as to whether they followed the same rules or did not follow the same rules. Although the underlying rules were extremely difficult to induce, Reber found that participants were able to perform the task at above chance levels even though they had problems stating what the rules were. Reber suggested that participants were able to learn the underlying rules of the grammar to some degree and that this rule learning occurred unconsciously. In addition, when participants explicitly tried to figure out what the underlying rules were, they actually performed worse than when they simply memorized the strings.

### ***Alternatives to rule-based classification***

Findings that contradicted the notion that classification rules underlie categorization led researchers to devise alternative models of classification that did not require the use of rules. One alternative was the family resemblance view of classification (Rosch, 1975;

Rosch & Mervis, 1975; Mervis & Rosch, 1981). As shown in Panel D of Table 1, this view assumes that stimuli consist of a set of features that are characteristic but not necessary for the category. Classification of an item is based on whether it has a sufficient number of these features to distinguish membership in one category from another. That is, a category is determined by how many features a stimulus shares with the members of one category compared to another. A subsequent model of classification that derived from the family resemblance view was the prototype view of classification. According to this view, through exposure to possible members of a category, some form of average or prototype is created that consists of all the characteristic features (Posner & Keele, 1968; Homa, Sterling & Trepel, 1981; Reed, 1972). Stimuli are then classified on the basis of their similarity to the prototypes of various categories that have been learned.

A second alternative to rule-based models was one in which classification did not require the abstraction of a rule or a prototype but was based on information specific to individual items (Brooks, 1978, 1987; Medin and Schaffer, 1978; Hintzman, 1986; Nosofsky, 1986). According to this exemplar-based view of classification, categories consist of separate examples that people have experienced. These exemplars are stored in memory and are grouped together into a category because they are similar to one another in some way. Classification of an exemplar involves retrieving similar stored exemplars from memory and if the retrieved exemplars are in a particular category then the presented exemplar will likely be in the same category.

Although formal exemplar-based models all share the assumption that classification is a nonlinear function of the similarities between exemplars, these models have been instantiated in a number of ways. For example, the Context Model of Medin and Schaffer (1978), and the Generalized Context Model of Nosofsky (1986, 1988a, 1988b) compares the features of a to-be-classified item to other items within a category. The comparison returns a value that represents a ratio between the similarity of the

to-be-classified item to the members of the category and the similarity of the to-be-classified item to all items in memory. The higher the value the more likely the to-be-classified item is in that category. Hintzman's (1986) Minerva II model compares the features of a to-be-classified item to a pool of stored memory traces. The resulting comparison returns an echo from the pool of traces that contains information about the category of the to-be-classified item and also the intensity of the similarity comparison.

A more recent alternative to these exemplar-based models are connectionist-based models (Kruschke, 1992; McClelland & Rumelhart, 1985). These models consist of a network of interconnected units that send and receive excitatory and inhibitory signals to and from other units, in the process forming patterns of activation that correspond to memory traces. These traces are superimposed onto other similar traces during learning and eventually form stable representations that correspond to a category. When a novel stimulus is presented to a network, its pattern of activation spreads to other similar learned patterns which, in turn, turn off dissimilar patterns. Eventually the activity of the network settles onto one representation, categorizing the novel stimulus on the basis of its similarity to previous exemplars.

***Evidence for the use of both rule-based and retrieval-based knowledge***

Despite the existence of evidence that is difficult to account for with a rule-based view of classification and the development of alternative models that emphasize retrieval processes, there remains considerable evidence that we rely on rule-based knowledge for classifying objects. For example, Malt (1990) found that for natural kind categories, people believe there are defining features that separate category members from category nonmembers. Keil (1989) found evidence that suggests that categories consist of a conceptual core, or an essence (Medin & Ortony, 1989), which is made up of a set of defining features. Medin and colleagues (Medin, Wattenmaker & Hampson, 1987; Ahn & Medin, 1992; see also Regehr & Brooks, 1995) found that when participants sorted an

array of perceptual stimuli into two categories, they had a strong tendency to sort the stimuli on the basis of single defining features. Finally, Martin and Carramazo (1980) asked participants to classify a set of schematic faces into two categories, but rather than collapsing over participants' results they analyzed each individual participant separately and found that participants were making their decisions on the basis of a simple rule that involved different combinations of features.

The presence of evidence that supports the use rules and the existence of models that argue for the importance of retrieval-based classification has led several researchers to suggest that we are able to use both rules and retrieval as a basis of classification (Elio & Anderson, 1981; Brooks, 1978, 1987; Reber, 1989; Kemler-Nelson, 1989). The possible use of two types of knowledge for classification raises the question as to how the acquisition and use of rules and retrieval are coordinated. In answering this question, a number researchers have investigated factors that selectively influence the acquisition of one type of knowledge over another. A number of these factors have been identified ranging from instructional manipulations to varying the memorability of the stimuli.

***The importance of instructions for the control of rule-based and retrieval-based knowledge***

Perhaps the most common factor manipulated in classification studies has been to vary the study instructions, thus creating experimental conditions in which participants selectively process stimuli in one fashion or another. The basic premise of this approach is that if participants are instructed to selectively attend to particular features then they will acquire a rule, but if they are instructed to attend to individual exemplars then they will rely on memory. The importance of this type of instructional manipulation was demonstrated by Medin, Dewey, and Murphy (1983). They presented participants with photographs of people that belonged to two categories. Some of the participants were given instructions to learn a first name for each photograph. Medin et al. hypothesized

that this should encourage participants to learn about the individual exemplars of the categories and promote a retrieval-based classification strategy. Other participants were told that the faces fit into two categories, based on the last names, and were to pay attention to specific features, like hair color, hair length, smile, and shirt color. By focusing on specific features within the stimuli, Medin et al. hypothesized that participants should learn information consistent with a rule. After training was complete, participants were asked to classify a new set of photographs that consisted of the study photographs and also some new photographs.

Medin et al. found that participants who learned the first names during the study task made fewer errors at study and learned the categories faster than the participants who learned the last names at study. This learning benefit did not transfer to novel stimuli, however, because participants who learned the last names at study made fewer errors when classifying novel stimuli than the participants who learned the first names at study. Medin et al. interpreted this pattern of results to suggest that, although memory for individual items aids category learning, transfer of that learning requires participants to have some knowledge about the features that characterize the stimuli in the category.

Other researchers have explicitly given participants a classification rule in an attempt to see how rules and retrieval are coordinated. For example, Nosofsky, Clark and Shin (1989) had participants classify stimuli into two categories. One group of participants learned the categories without the benefit of a classification rule, whereas a second group learned the categories with the benefit of a rule. After training, participants were first given a recognition test and then completed a transfer classification task. Nosofsky et al. reasoned that if participants had learned about individual exemplars during the study task, and were using that knowledge during classification, then their ability to discriminate old from new exemplars in the recognition task should be good. If, on the other hand, participants had focused their attention on the rule during the study phase, then they may

not learn about the individual exemplars and their ability to discriminate old from new exemplars in the recognition task would be poor. Nosofsky et al. found that when no rule was given during the study task, participants were able to discriminate old items from new items in the recognition task and their performance in the classification task was fit well by the Generalized Context model. These findings suggested that participants were relying on a retrieval-based classification strategy. When a rule was given during the training task, Nosofsky et al. found that participants' recognition performance was at chance while a rule-based model provided a better fit of the classification results than the Generalized Context Model. These findings suggested that participants were relying on a classification rule. Nosofsky et al. concluded that classification could occur by either rules or retrieval, depending on the strategy adopted during the study task. Similarly, Allen and Brooks (1991) and Regehr and Brooks (1993) found that when participants were not given a classification rule at training, their basis of transfer to novel but similar-to-old stimuli tended to reflect the use of similarity-based retrieval. In contrast to Nosofsky et al. (1989), however, when participants were explicitly given a classification rule at training the basis of transfer reflected a mixture of both rule-based and similarity-based knowledge, rather than exclusive use of a rule.

In short, explicit instructions to use a rule, or to attend to specific features, leads to the acquisition of rule-based knowledge, and instructions to learn a category without a rule, or to learn a category by attending to the individual items, leads to the use of retrieval-based knowledge. Other researchers have found, however, that deliberate attempts to learn a rule actually leads to worse rule-based performance if the categories are learned implicitly. For example, in a category learning task, in which the stimuli were derived using an artificial grammar, Reber (1976; Reber, Kassin, Lewis, Cantor, 1980) asked one group of participants to look for the underlying rules that determined category membership while studying a set of training items. A second group of participants



simply observed the training stimuli without any mention of the rules, or that there would be a later classification test. After training, this observation group was told about the existence of the underlying rule structure and then both groups classified a set of novel exemplars into those that were grammatical, or followed the same rules as the training stimuli, and those that were nongrammatical, or did not follow the same rules as the training stimuli. Reber found that participants who were told of the existence of the rule at training took longer to learn the training items. In addition, their ability to pick out the rule consistent letter strings was worse than participants who only observed the stimuli.

A potential argument against Reber's interpretation of his findings could be that the task was not really a classification task because participants did not have to learn whether the items belonged to one category or another. If the experiment had been designed with the conventional category learning procedure, then participants who were told to learn a rule would have provided better rule-based performance. Brooks (1978) used a more conventional task and found results similar to those of Reber's. Participants were trained to categorize sets of letter strings that were derived from two underlying grammars. In one condition, participants were told that the categories followed a complex set of rules and that they were to classify the stimuli into two categories. In a second condition, the strings from the two grammars were paired with animal names and city names and the participants' task was to memorize the pairs of stimuli. What participants did not know was that cities and animals from the new world were all paired with one type of grammar and the cities and animals from the old world were paired with the other grammar. At transfer, participants were asked to classify novel letter strings into old world or new world categories. Like Reber, Brooks found that participants who were informed of the existence of an underlying rule structure at training did worse at classifying the novel letter strings than participants who had done the paired associate learning and had not been given explicit instructions about the underlying rules.

There has been some debate as to how to interpret this finding that implicit category learning instructions lead to an increased learning of rule-based knowledge while explicit instructions to learn a rule leads to decreased learning of rule-based knowledge. Reber (1976, 1989; Reber & Allen, 1978) interpreted this finding in terms of the degree to which rule learning has occurred. According to Reber, most rule learning occurs unconsciously and therefore the study task is important for determining whether the rules will actually be learned. When participants actively search for the underlying rules at study, they tend to generate false rules that interfere with learning the underlying rule structure. Participants then perform poorly when classifying novel stimuli that require the rules for determining the category. When a paired associate study task is used at training, like that used by Brooks (1978), participants are encouraged to rely on memory to perform the task and learn little about the underlying rule structure of the stimuli. When novel stimuli are then presented for classification, participants have no rule-based knowledge and so are forced into relying on memory as a basis of classification. It is only when participants passively observe the letter strings at study, without explicitly trying to learn a category, that they are able to unconsciously learn the underlying structure. This knowledge can then be used if they are later asked to classify novel stimuli.

Brooks (1978) has also suggested that the study task is important for determining whether participants rely on rules or retrieval. However, he suggested that the study task influences participants' memory for specific items rather than the unconscious learning of the underlying rules. Brooks hypothesized that the basis of transfer for classifying novel stimuli is often by analogy to studied items. When a paired associate or observation study task is used, then the study conditions are conducive to encoding information about individual items. Memory for the items can then be used at transfer as part of a similarity-based classification process. When participants are explicitly trying to discover the underlying rules at study, they learn little about the individual training items which

then impairs their ability to rely on memory for classifying novel stimuli.

Evidence consistent with a retrieval-based interpretation of implicit learning has been found in a number of other studies that have not used stimuli derived from artificial grammars. For example, Wattenmaker (1991) presented participants with fictitious names and descriptions of people in which either two or three features were correlated with the categories they were learning. One group of participants, the intentional category learning group, was told that two categories existed and they were to try to discover ways of classifying the descriptions. A second group of participants were not told of the existence of the categories but either memorized the descriptions or, in one experiment, thought of someone with the same name and then rated how well the description fit that person. At transfer, all the participants were told that the two categories existed and were then presented with pairs of correlated features from one of the stimuli that they had studied. Their task was to indicate whether the correlated features co-occurred more often in one category or the other. Wattenmaker hypothesized that if participants had learned a rule at study, either implicitly or explicitly, then their knowledge of how often other features co-varied across the categories should be poor. If, however, participants had learned about individual exemplars, then they should be able to retrieve those exemplars from memory when making their judgements about the correlated features. He found that the participants who intentionally learned the category were less likely to correctly classify the correlated features compared to participants who had received the incidental instructions. Wattenmaker concluded that intentional category learning instructions had forced participants to look for rules that could delineate the category while implicitly learning the categories led to the encoding of individual stimuli.

Kemler-Nelson (1984) also compared intentional and implicit category learning conditions. Participants were presented with schematic faces that belonged to two categories. In one condition, participants were explicitly told to learn the categories, and

in the other condition participants were asked to judge whether they had seen the item before. Transfer stimuli, which could be classified by either a single defining feature or by their family resemblance, were created. The stimuli were designed, however, so that reliance on rules and reliance on family resemblance would result in opposite categorical responses. Kemler-Nelson hypothesized that if participants had learned a rule in the implicit learning condition then their knowledge of the family resemblance structure should be poor and they should use the rule when classifying the transfer stimuli. She found that participants in the intentional condition tended to use a defining feature for classification whereas participants in the incidental group tended to base their classifications on the family resemblance structure of the faces.

In short, encouraging participants to learn a category by using a rule appears to lead to an increased reliance on rule-based knowledge. However, there is some debate as to whether learning a category under implicit conditions leads to a greater reliance on a rule, or to reliance on retrieval.

***Other instructional manipulations that control rule-based and retrieval-based knowledge***

A variety of other instructional manipulations, which do not explicitly emphasize the use of rules or retrieval, have also been found to have an effect on the degree that either type of knowledge is used. For example, encouraging participants to respond quickly facilitates the use of similarity-based responding. Evidence consistent with this claim was found by Smith and Kemler-Nelson (1984). They used a speeded classification task, in which the categories were structured around the perceptual features of size and brightness. These perceptual features are considered to be separable dimensions (Garner, 1974); that is, they are features that do not readily combine together to form a unitary whole. When classifying stimuli that consist of separable dimensions, participants have a tendency to treat one of the dimensions as defining. Smith and Kemler-Nelson found that

when speed was not emphasized, participants based their classification on the presence or absence of one of the perceptual features. When speed was emphasized, participants based their classification on overall similarity. A similar finding occurred when participants were encouraged to respond with their first impression. Smith and Kemler-Nelson (1984) told one group of participants to respond with their first impression while a second group were encouraged to be slow and accurate. First impression participants were faster and relied more on similarity when classifying the stimuli than the group who were encouraged to be slow and accurate.

Ward, Foley, and Cole (1986) investigated these findings further. They used a procedure similar to that used by Smith and Kemler-Nelson, but with two types of stimuli. One type of stimuli varied along integral dimensions; that is the categories were structured around perceptual features that readily combine into unitary wholes. For example, perceptual features like hue and saturation are considered to be integral dimensions because stimuli that vary along these features tend to be classified on the basis of overall visual similarity. The other type of stimuli, used in the experiment, varied along separable dimensions. Ward et al. found that when classifying stimuli that consisted of integral dimensions, participants relied primarily on similarity, and instructions to respond quickly had little effect. When classifying stimuli that consisted of separable dimensions, participants relied primarily on a single feature. However, the number of single feature classifications was reduced, and the number of responses based on similarity increased, as the speed of responding was increased. A similar finding occurred when Ward et al. contrasted participants who were told to respond with their first impression and participants who were told to be accurate. They found that when participants classified stimuli that consisted of integral dimensions, they relied primarily on similarity. When classifying stimuli that consisted of separable dimensions, participants relied primarily on a single feature to classify the stimuli. However, the

number of responses based on a single feature diminished, and the number of responses based on similarity increased, if participants responded with their first impression.

Performing a second task concurrently with a classification task also encourages participants to rely on similarity. In a third experiment, Smith and Kemler-Nelson (1984) instructed one group of participants to count backwards from a number while classifying stimuli that consisted of separable dimensions. A second group of participants classified the same stimuli but without the instruction to count backwards. The authors found that without the concurrent task most of the responses were based on a single feature but with the concurrent task the majority of the responses were similarity-based.

#### ***Non-instructional control of rule-based and retrieval-based knowledge***

Not all studies that have investigated factors that influence the use of rules and retrieval have concentrated on the role of instructions. Some studies have shown that the materials used in an experiment can also contribute towards determining the use of rules or retrieval. For example, the ease with which a rule can be learned is one of these factors.

Reber et al. (1980), used stimuli derived from an artificial grammar to demonstrate the importance of making a classification rule salient. In one condition, training stimuli were mounted on a board and the items were arranged in columns, with each column representing an aspect of the underlying rule structure. The purpose of this training procedure was to make the underlying rules more salient and easy to learn. Participants in the other training condition were presented with the training stimuli in a random order so that the underlying rules were not salient, and therefore harder to learn. In addition, within each training condition, one group of participants was given the conventional implicit learning instructions to just observe the items and a second group of participants was instructed to figure out the underlying rules. During a subsequent transfer task, participants classified a new set of stimuli that either followed or did not follow the same

rules as the training stimuli. Reber found that making the rule structure salient during training did not influence participants performance if they also observed the stimuli at study. Making the rule structure salient did influence participants' performance if they were explicitly told, during training, to figure out the underlying rules. These findings suggest that making a rule salient only helps if participants are actually looking for one.

Cho and Mathews (1996) also found that when a classification rule is easy to learn participants tend to rely on rule-based knowledge for classification. They presented participants with fictitious diseases to classify. These diseases were represented either with sequences of eight binary letters or with bar charts. Cho and Mathews hypothesized that the binary characters would make it easy for participants to discover a classification rule whereas the bar graphs would make it harder to discover a classification rule. In addition, in one experiment, the stimuli were created so that a single defining feature determined category membership, again making rule discovery easy. In a second experiment the rule involved correlated features making the rule difficult to discover.

To measure the degree to which participants were using rules or retrieval, Cho and Mathews used a classification task followed by a recognition task as transfer measures. They argued that good performance in a classification task requires the use of a classification rule and good performance in a recognition task requires the use of retrieval. If the knowledge participants learn during the study task matches the knowledge required for the transfer tasks, then participants should be able to perform well. If the knowledge learned during the training task does not match what was required for the transfer tasks then performance should be impaired. The authors found that when a single defining feature was used, participants classified the letter strings better than the bar graphs but recognition performance was at chance. They interpreted this finding to suggest that an easy to learn rule facilitated its use but prevented participants from learning about the individual exemplars. When the classification rule consisted of

correlated features, there was no difference when classifying the letters versus the bar charts but participants could reliably tell the old from the new diseases in the recognition task, Cho and Mathews interpreted this finding to suggest that when a rule was difficult to learn it facilitated the learning of individual cases.

Allen and Brooks (1991) also found that when stimuli were structured so that a classification rule was easy to use, performance in a later transfer task reflected the use of that rule. During the training task, participants were given a perfectly predictive classification rule and practiced using that rule to classify a set of imaginary animals. The animals either consisted of a list of features or were drawn so that the visual appearance of each feature corresponded to the value of the feature. Allen and Brooks hypothesized that participants would find it easier to use the rule when classifying the lists of features than the pictures and that this rule-based strategy would transfer to novel but similar-to-old stimuli. They found that participants were in fact more likely to apply the rule when classifying novel stimuli that consisted of the lists of features than they were when classifying pictures of the animals.

Memorability of the stimuli is another factor that influences participants' use of rules and retrieval. Regehr and Brooks (1993) investigated this issue in a series of experiments similar to those reported by Allen and Brooks (1991). Two types of stimuli were used in the experiments. In one case, the stimuli were imaginary animals constructed from unimorphic features, that is, features that were identical in appearance from animal to animal. For example, if any two animals had a short tail then the tails were identical. In the other case, stimuli were imaginary animals constructed from individuated features, that is, features that were unique in appearance from animal to animal, even when the animals were in the same category. For example, even if two animals had a short tail, the tails appeared different from one another. Regehr and Brooks noted that animals created with unimorphic features all tended to look like one another, while animals created with



individuated features all tended to look distinctive from one another. Regehr and Brooks hypothesized that this distinctiveness would make the animals more memorable by providing better encoding and retrieval cues than the unimorphic animals. This increase in memorability should lead to a corresponding increase in participants' reliance on retrieval processes when classifying stimuli, despite training with a perfectly predictive classification rule. The authors found that participants relied primarily on a classification rule when classifying novel stimuli that were similar to the unimorphic stimuli but increased their reliance on retrieval when classifying novel stimuli that were similar to the individuated animals. A similar finding was reported in an earlier study by Brooks (1976). He had participants study words constructed with an artificial alphabet. All the words were constructed from the same set of discrete characters but in one condition the words consisted of a horizontal listing of the characters. In a second condition the characters were listed vertically. The characters were designed so that when they were listed horizontally all the words tended to resemble one another, making each word nondistinctive from every other word. When the characters were listed vertically they resembled a glyph making each word distinct from every other word. Brooks found that glyphic words were consistently pronounced faster than the nonmemorable words.

***A summary of the factors that control the use of rule-based and retrieval-based knowledge***

In summary, the focus of a number of recent classification studies has been to investigate factors that preferentially bias the selective acquisition and use of rule-based knowledge and retrieval-based knowledge. A number of factors have been identified. When instructed to use a classification rule, participants will do so if that rule is given to them or if it can be easily discovered. If it is not easily discovered, then instructions to use a rule will actually impair rule-based performance. When participants are required to learn a category without the benefit of a rule, they tend to rely on retrieval of exemplars

under intentional category learning instructions. When learning a category under implicit learning conditions, there is some debate as to whether participants rely on retrieval (Brooks, 1978) or on an unconsciously learned rule (Reber, 1989). Responding quickly, responding with a first impression, performing a concurrent task simultaneously, and using stimuli that are memorable are other factors that facilitate the use of retrieval-based classification.

***Theoretical models of the control of rule-based and retrieval-based knowledge***

Although a large proportion of the research in category learning has focused on identifying factors that influence the use of rule-based or retrieval-based knowledge, other research has tried to develop theoretical models to account for the role that these factors play in coordinating rules and retrieval.

One model designed to account for the coordination of rules and retrieval assumes that simple rules do not exist and that similarity between items is sufficient to account for how we classify objects. As suggested by Medin (1986), the factor that determines whether participants' performance will appear to be rule like, or will appear to be guided by retrieval, is the strategy that participants adopt during the study phase. If participants adopt an analytic learning strategy, one which involves selective attention to specific features, then they will produce results that appear to be determined by the use of a classification rule. If participants adopt a nonanalytic learning strategy, one in which the whole exemplar is encoded rather than a subset of features, then they will produce results more consistent with the retrieval of other similar exemplars. However Medin argues that in both cases, similarity-based comparisons drive the classification process.

The method of using similarity to account for rule-like performance has been formalized in the Context Model (Medin & Schaffer, 1978; Medin & Smith, 1981) and in other variants of the model such as Nosofsky's (1986) Generalized Context Model and Kruschke's (1992) ALCOVE model. The key assumption to the Context Model is that

each encounter with an exemplar is stored in memory and whenever a new exemplar is presented to be classified it retrieves other similar stored exemplars from memory. The probability that a particular exemplar will be classified into a particular category is determined by summing the similarities of that exemplar to all the members of a particular category and then dividing by the sum of the similarities of that exemplar to all other exemplars stored in memory. The similarity comparison is determined by measuring how similar each feature of an exemplar is to the features of the stored exemplars and then assigning a value to that comparison. The values are then combined multiplicatively so that the more similar two items are the more likely they will be in the same category.

The Context Model has some additional assumptions that allow it to account for both rule-based and retrieval-based results. First, not all items in memory need to enter a similarity comparison so that, at an extreme, classification can be based on the retrieval of one exemplar from memory. Additionally, selectively attending to a feature, or a subset of features, will increase the weight of those features in the similarity comparison and decrease the weight given to unattended features. At an extreme, the status of one particular feature can drive the classification procedure producing results analogous to the use of a single defining feature.

Evidence supporting the Context Model account of coordinating rules and similarity was described in Medin and Smith (1981). Participants were instructed to learn two categories of schematic faces, but under three different encoding conditions. A standard group learned the categories by relying only on feedback as to whether they were right or wrong. A second group, a rule plus exception group, was told to pay attention to a specific feature (nose length) but to remember two faces that were exceptions to the rule. A third group of participants, the prototype group, were told to learn what the most average face in each category would look like. Medin and Smith hypothesized that the

different study instructions should encourage participants to attend to different aspects of the stimuli depending on the instructions they received. If so, then what participants learn about the categories should differ and change the way the stimuli are classified. For example, the prototype group should attend to all of the features of the stimuli until they learn which features are important for the prototype. The rule plus exception group should only attend to how long the nose is on each stimulus, so that this feature should become important.

The transfer phase consisted of a recognition test, a transfer test with novel but similar-to-old stimuli, and a speeded classification test using only training stimuli. The three different encoding conditions all produced differences in how well the participants learned the items, how well they recognized old versus new items, their accuracy at transfer, and their speed of responding. Based on their assumptions about what information was important for each encoding task, Medin and Smith changed the weights associated with each feature before the model made its similarity comparisons. The authors found that the Context model provided a good fit of the data. They concluded that the different learning instructions had not induced participants to use different classification strategies; Rather, they were all using similarity and just changing what parts of the stimuli they were encoding.

A second attempt to model the coordination of rule-based and retrieval-based knowledge assumes that rules and similarity are separate sources of information but are in a compensatory relationship with each other. The more we use one type of knowledge the less we use the other. For example, according to Reber (1989; Reber & Allen, 1978), we rely on memory for individual exemplars when the training task requires us to attend to individual items. This use of memory partially interferes with the learning of rule-based information so that the more we learn about individual items during the learning task the less likely we will learn information that will aid the use of a rule. Under

conditions that do not require attention to individual items, as when simply observing the stimuli, or when a rule is very easy to learn, people will attend to the underlying structure of the stimuli and learn information useful for a rule. Reber also adds that this learning of an underlying rule is our favoured mode of learning, as it occurs relatively unconsciously. In a variant of this method of coordination, Cho and Mathews (1996) suggest that category learning actually begins with the unconscious learning of individual exemplars. During this learning process, however, we often unconsciously notice a feature or combination of features that seem predictive for the categories that we are learning. Once this occurs, the predictiveness of the features that were noticed becomes conscious and forms the basis of a classification rule.

Brooks (1987; 1978) and colleagues (Whittlesea, 1987; Whittlesea, Brooks, Westcott, 1994; Vokey & Brooks, 1992; Whittlesea and Dworkin, 1993; Jacoby & Brooks, 1984) have also developed a model that assumes that rules and retrieval are separate forms of knowledge. These studies contrast two learning styles. One style is an analytic mode of learning, in which participants break apart the stimuli into its parts. By comparing these parts across stimuli, participants are able to abstract out general knowledge, like rules or prototypes. The second learning style is a nonanalytic mode of learning, which places less emphasis on comparing parts of the stimuli and more emphasis on remembering individual items. The degree that either learning style is used is a function of the processing demands inherent in the tasks and the materials that are used in experiments. In addition, these two learning styles are independent of each other; that is, each can operate simultaneously without interfering with the other (Vokey & Brooks 1992; Allen & Brooks, 1991; Regehr & Brooks, 1993).

Evidence supporting this method of coordinating rules and similarity was reported in Vokey and Brooks (1992). All of the stimuli were letter strings derived from an artificial grammar. Participants studied the letter strings during a training task and were tested with

a subsequent recognition test followed by a classification test. The stimuli for the transfer tasks consisted of novel letter strings that varied in terms of whether they followed, or did not follow, the rules of the grammar, and whether they were similar, or were not similar, to a training item. Vokey and Brooks found separate effects of similarity use and rule use in the recognition test. Participants had more false alarms to stimuli that were consistent with the rules than stimuli that were not consistent with the rules. Participants also made more false alarms on transfer stimuli that were similar to a training stimulus than they did to transfer stimuli that were not similar to a training stimulus. In the subsequent classification task, Vokey and Brooks again found separate effects of similarity use and rule use. Participants were more likely to consider a novel item as grammatical if it followed the rules than if it did not follow the rules and they were more likely to consider a novel letter string as grammatical if it was similar to a training item than if it was not similar to a training item.

A fourth method of coordinating rule-based and similarity-based knowledge is the rules plus exception model (RULEX) formalized recently by Nosofsky and colleagues (Nosofsky, Palmeri, & McKinley, 1994; Palmeri & Nosofsky, 1995; also Ward & Scott, 1987). According to RULEX, category learning consists of participants actively searching for classification rules, using hypothesis testing procedures similar to those described by the early rule-based models (Bruner et al. 1956). Once a participant finds a rule that is sufficient, then that rule is used for classifying all stimuli except those that may belong to a category but are exceptions to the rule. These exceptions are then stored directly in memory so that classification relies on a combination of retrieval-based and rule-based knowledge.

To test the RULEX model, Nosofsky et al. (1994) asked participants to classify, into two categories, a set of stimuli that varied on four binary dimensions. During a transfer phase, participants classified the training stimuli and a set of novel but similar-to-old

stimuli. Predictions based on the RULEX model and the Context model were then fitted to the observed transfer results. Nosofsky et al. found that the RULEX model provided just as good a fit of the transfer results as the Context model, despite the learning procedure being different in the two models.

In summary, there are a number of possible methods of coordinating the use of rules and retrieval when classifying objects. According to the Context model (Medin & Schaffer, 1978; Medin, 1986), and its variants (Nosofsky, 1986; Krushcke, 1992), learning strategies influence what aspects of a stimulus participants rely on when making similarity comparisons, and these similarity comparisons can produce results that are consistent with either retrieval or rule application. According to a model derived from work done by Reber and colleagues (Reber 1989; Reber & Allen, 1978; Cho & Mathews, 1996) rule learning and retrieval are separate forms of knowledge but are in a compensatory relation with each other. The relative importance of either is determined by the study task. A model based on work done by Brooks and colleagues (Brooks, 1978,1987; Vokey & Brooks, 1992; Whittlesea et al, 1994) also suggests that rules and similarity are separate types of knowledge but their acquisition and use occur simultaneously and are independent of each other. The relative importance of either is a by-product of the processing demands inherent in the study tasks, the transfer tasks, and the materials that are used in an experiment. Finally, according to the RULEX model (Nosofsky et al. 1994), participants use explicit rules for classification, with retrieval only being used to store any items that are exceptions to the rule.

***Coordinating rules and retrieval with the development of skilled classification***

Much of our interaction with categories in the real world occurs in conjunction with considerable practice at classifying instances of a category. For the most part, however, category learning research has ignored the role that practice might play in coordinating the use of rule-based and retrieval-based knowledge. Occasionally, studies have used

extensive training procedures as part of their tasks, but these studies have also tended to use experimental procedures that encourage the use of just one type of knowledge. For example, most classification studies have used stimuli that are simply lists of abstract features or have used perceptual stimuli that consist of unimorphic features. As argued by Regher and Brooks (1993; Allen & Brooks, 1991), these types of stimuli selectively encourage the use of analytic classification strategies. Therefore, it is difficult to determine how rules and retrieval are coordinated with practice if the experiment is designed to encourage the development of only one type of categorical knowledge.

In addition to the role it may play in classification, practice is significant for another reason. It is crucial for the development of skilled performance in cognitive tasks. The role that practice plays, however, depends on which of two views of skilled performance one adopts. According to one view, skilled performance develops slowly over time and only as a result of slow deliberate practice (Ericsson & Lehmann, 1996; Anderson 1982). As a result of this slow development of skilled performance, most learning occurs in a controlled and cognitively effortful manner and deliberate practice is necessary to make the underlying procedures more efficient. The contrasting view is that skilled performance can develop rapidly, possibly after only a few practice trials on specific items (Logan, 1988). The role of practice, according to this view, is to support a type of learning that is fast and is not cognitively demanding. The experiments in this thesis were designed to set up conditions that would support the latter view but use a category learning task to do so.

How might rule-based and retrieval-based knowledge be coordinated as one becomes skilled at categorization? An intuitive possibility is that when initially learning a category we rely on simple explicit rules for classification. Categorization based on these explicit rules tends to be slow and cognitively effortful. As we become more practiced, a transition occurs. The strategy that we use seems to change from the slow, effortful use of



explicit rules to a strategy that is faster, less effortful, and feels more like we immediately recognize the category of an object. Given that one view of skill acquisition has suggested that this transition can occur relatively quickly, it is important to understand how the use of rules and retrieval coordinated as this transition occurs.

There are two broad intuitions that account for a transition from the slow effortful classification performance of novices to the faster less effortful classification performance of more skilled individuals. One intuition is based on Anderson's ACT-R model of skill acquisition (Anderson, 1993; also Anderson, 1982; Singley & Anderson, 1989). According to this model, performance is initially guided by the deliberate use of examples and general rules that the novice uses to perform the task. The use of this information tends to be slow and cognitively demanding. With practice, novices gradually learn how to skip some of the steps required by the general rule and are able to reorganize and combine the remaining steps so that the application of the rule becomes faster and less demanding. This reorganized rule is called a production rule (Blessing & Anderson, 1996; Anderson, 1982). An account of category learning based on the ACT-R model would predict that a novice would start off deliberately using a classification rule to decide if an exemplar is a member of a category. With practice, however, the novice learns shortcuts in the application of the rule and eventually the use of the rule will become faster and less demanding.

The transition to faster, less effortful performance in a task may also be explained by Logan's instance theory (Logan 1988, 1992). Novices begin a task using a general algorithm that is slow, effortful, and requires conscious awareness. With each exposure of an instance, novices also encode information about that specific instance that can later be retrieved if it is ever encountered again. Encoding and retrieval of these instances are assumed to be obligatory, in that information is stored and retrieved every time an instance is encountered. It is also assumed that retrieval requires little effort or conscious

awareness, and therefore tends to be faster than the algorithm. In addition, use of the general algorithm and retrieval operate simultaneously, with both processes racing against each other to form the basis of a response. As increasingly more instances are stored, the likelihood that a retrieved instance will win the race increases. Eventually, given enough instances in memory, retrieval will win the race on virtually every trial and the importance of the algorithm will diminish. An account of category learning based on Logan's Instance theory would predict that novices start off deliberately using a classification rule, but with repeated exposure to category members, specific instances are recognized as a member of a category, and the importance of the classification rule fades away.

A transition from slow effortful performance to performance that is faster and less effortful has been found when learning a variety of skills including solving alpha-arithmetic problems (Logan, 1988; Blessing & Anderson, 1996; Anderson, Fincham & Douglass, 1977), solving logic proofs (Neves & Anderson, 1981 as discussed in Anderson, 1982), solving arithmetic problems (Charness & Campbell, 1988; Carlson & Lundy, 1992), computer programming (Singley & Anderson, 1989), lexical decision (Logan, 1988, 1990), memory search (Logan and Stadler, 1991) and visual numerosity judgements (Lassaline & Logan, 1993; Palmeri, 1997). Although the underlying mechanisms differ, both the ACT-R model and the Instance theory have been put forth by various researchers to account for the transition found in these studies. However, a transition to faster performance with practice is only one result of learning a skill. A second result of learning a skill is the ability to generalize that skill to novel situations that are similar to those encountered during training. Factors that influence the acquisition of a skill often have very different effects when generalization of that skill is required (Schmidt & Bjork, 1992). Furthermore, the ACT-R model and Instance theory make very different predictions as to what aspects of a skill will transfer to novel

situations.

Transfer of a skill to novel situations is a strength of the ACT-R model. Although production rules are considered to be specific to the domain in which they were acquired, as long as the underlying procedures required to perform a task remain consistent, a production rule will generalize to novel situations. For example, Anderson and Fincham (1994; also Anderson, Fincham & Douglass, 1997) gave participants rules to memorize that, in one experiment, allowed them to solve a series of alpha-arithmetic problems and, in another experiment, a series of word problems. In both experiments participants were given part of the information and they had to solve the problem and produce an output (e.g.,  $42 / 7 = ?$ ). At transfer, participants solved problems that were new but had the same form as the training problems (e.g.  $45 / 5 = ?$ ) or were similar to the training problems except that participants were given what had been the output and had to solve the problems in reverse (e.g.,  $42 / ? = 6$ ). With practice, Anderson and Fincham found that participants solved the training problems faster and were able to transfer that skill to novel problems that were of the same form as the training problems. When the form of the transfer problem was changed, however, Anderson and Fincham found that participants' performance suffered. The authors concluded that a production rule will generalize to novel items, but only when the procedures required by the rule do not change.

If the development of a skill reflects the use of a production rule, and if this rule can be transferred to novel situations in which it still applies, then it should be relatively difficult to find conditions in which similarity-based responding increases with the development of a skill. This does not appear to be the case, however, as there are a number of examples of similarity-based transfer given skilled performance. For example, in dermatology, doctors are initially trained by being told to look for certain symptoms characteristic of a disease. Given enough medical experience, skilled performance

should reflect the fast use of these rules. Brooks, Norman and Allen (1991) presented slides of dermatological cases to general practitioners and medical residents for diagnosis and then presented new slides that were of the same diseases but varied in their similarity to the original slides. The new slides that were similar to the original slides were diagnosed more accurately than the new slides that were not similar to the original slides, and this benefit of similarity occurred despite the same underlying rule-based procedures being applicable in all the novel slides. Palmeri (1997) trained participants on a visual numerosity task until their performance became automatized, and then presented them with novel shapes that were distortions of the training shapes. Palmeri found that novel shapes that were similar to the training shapes were responded to faster than novel shapes that were not similar to the training shapes. Further, moderately similar shapes were responded to faster than shapes that had a low degree of similarity. If the skilled performance of participants reflected the use of a production rule for counting the dots, then it should have generalized to all novel stimuli.

Although the ACT-R model can account for transfer in some tasks, these findings suggest that it does not capture all aspects of skilled performance. Relatively complex tasks such as solving math problems, logic problems, or computer programming, may require the execution of distinct sets of steps, a notion that is consistent with learning production rules. On the other hand, some perceptual stimuli may be relatively easy to remember (Paivio, 1971) and may not require the serial execution of a set of steps to process. Thus, tasks like those used by Brooks et al. (1991) and Palmeri (1997) may lend themselves to more retrieval-based processing than rule-based processing. The importance of a classification rule for skilled classification of novel perceptual stimuli will be one of the issues examined in this thesis.

Although transfer of a skill is a strength of the ACT-R model, it is a weakness of the Instance theory. As instantiated by Logan (1988), the Instance theory was designed to

account for the facilitated processing of old familiar items. For transfer to occur, the range of similarity that is allowed is extremely narrow. In most cases, if an item is new, even if it is similar to the training items, it does not benefit from that similarity and induces the use of algorithmic processes. The highly specific nature of this memory retrieval process is problematic because similarity-based comparisons to stored exemplars are a central component to many models of classification. For example, according to a prototype view (Posner & Keele, 1968; Reed, 1972; Homa, Sterling & Trepel, 1981), classification of all exemplars depends on their similarity to a category prototype. The more similar an item is to the prototype of a category the more typical it is and the more likely it will be placed in that category. With exemplar-based models (Brooks, 1978; Medin & Schaffer, 1978; Hintzman, 1986; Nosofsky, 1986; Kruschke, 1992) classification of novel items is determined by their similarity to previously learned examples. The probability of classifying a novel item as a member of a category is a function of its similarity either to a single prior item from within a category or to a pool of items from within a category. Similarity comparisons to stored exemplars is an important construct in a number of other areas of cognition as well, including problem solving, (Ross, 1987, 1989; Gentner, Ratterman, & Forbus, 1993), recognition memory (Murdock, 1982; Gillund & Shiffrin, 1984; Jacoby & Brooks, 1984), word identification (Rueckl, 1992; Andrews; 1989) and connectionist modelling (Hinton, McClelland & Rumelhart, 1986).

The vast set of findings that show the importance of similarity in memory retrieval, and the models that were derived from this set of findings, suggest that similarity is a critical component to memory. Therefore, if a transition to a faster, less effortful classification strategy involves an increased reliance on retrieval, then there should also be evidence of an increased reliance on similarity-based responding when classifying novel items that are similar to stored instances. The problem with this conclusion is that

few classification studies have created conditions that would facilitate a transition to a faster, less effortful similarity-based strategy. Even when experiments were designed to facilitate the use of similarity-based classification, the stimuli that were used tended to lead to analytic classification strategies (Regehr and Brooks, 1993). This emphasis on analytic classification strategies may limit the role that memory is allowed to play in a classification task. The overall objective of the research reported here was to create experimental conditions that would be conducive to using memory as a basis of classification and then examine whether a model of classification based on Logan's Instance theory could account for the pattern of generalization effects obtained with skilled classifiers.

#### *Determining the use of rules and retrieval*

To determine the degree to which participants use rules or similarity to classify members of a category, an opposition procedure was used that allows an estimate of the importance of each. As described in studies by Allen and Brooks (1991) and Regehr and Brooks (1993 ; also Neal, Hesketh & Andrews, 1995), participants practiced classifying a set of training stimuli and were then given a subsequent transfer test in which they classified a new set of stimuli. This new set consisted of the training items and also some new items that were similar to the training items. Some of the new items were similar to a specific training item and were in the same category according to an underlying rule (good transfer items, or GT items) and some of the new items were similar to a training item but in the opposite category according to an underlying rule (bad transfer items, or BT items).

The logic of using this opposition procedure is that GT and BT stimuli should be classified differently depending on whether participants rely on the rule or on retrieving other similar instances as a basis of classification. Consider the example displayed in Table 2. According to the classification rule, the GT stimuli in the top panel of the table

Table 2.

An example of the procedure used to determine the relative importance of rules and retrieval. If the stimuli are classified according to the rule, then GT and BT stimuli will consistently be placed in opposite categories, despite both stimuli being one feature away from, and thus similar to, a training stimulus. If the stimuli are classified on the basis of their similarity to a studied item then GT and BT animals will consistently be placed in the same category as the training stimulus, despite the BT stimuli not following the rule for that category.

**Rule:** Anything with the value of “1” on the first feature is in Category A and anything with the value of “0” on the first feature is in Category B.

	Training stimuli	GT stimuli	BT stimuli
	101010 110101	111010 100101	001010 010101
rule-based response	“A”	“A”	“B”
similarity-based response	“A”	“A”	“A”
	010101 001010	000101 011010	110101 101010
rule-based response	“B”	“B”	“A”
similarity-based response	“B”	“B”	“B”

are in Category A because they have a value of “1” on the first feature. The BT stimuli in the top panel are in Category B because they have a value of “0” on the first feature. The GT stimuli in the bottom panel of Table 2 are in Category B because they have a value of “0” on the first feature, while the BT stimuli in the bottom panel have a value of “1” on the first feature and should be in Category A. Similarity to the training stimuli produces a different pattern of responses. Notice that both the GT and BT stimuli are one feature away from, and thus are equally similar to, a specific training stimulus. Based on this similarity, the GT and BT stimuli in top panel are in Category A and the GT and BT stimuli in bottom panel are in Category B. Because the category of the BT items changes

depending on whether one is using the rule or similarity to a prior instance, the relative importance of either type of knowledge is determined by contrasting performance across the GT and BT stimuli. If participants are relying solely on a rule, then they should tend to put the GT and BT stimuli in opposite categories, and therefore there should be relatively few responses that are inconsistent with the classification rule. If participants are relying solely on similarity to prior instances, then they should tend to put the GT and BT stimuli in the same category and therefore the BT stimuli should tend to have responses that are inconsistent with the classification rule.

This pattern of responding will occur, however, only when participants are relying solely on the classification rule or solely on similarity to prior instances when classifying the stimuli. An additional possibility is that participants will be at neither extreme and rely on some combination of rules and similarity for classification. The opposition procedure is sensitive to this possibility. Because rule-based and similarity-based processes lead to opposite responses on BT stimuli, they may interfere with each other if participants try to use both simultaneously. This interference may produce slow response times when classifying BT stimuli because participants will need time to resolve the conflict. Rules and similarity lead to the same response on GT items and therefore the interference between rules and similarity may not occur for these items. Irrespective of the category participants choose, whether or not interference occurred can be determined by comparing the response times between BT and GT items. A large difference between the two may indicate the presence of interference between a retrieval-based process and a rule-based process. No difference between the two may indicate that participants are relying on the same type of knowledge to classify the transfer stimuli.

Using this opposition procedure, Allen and Brooks (1991) found that when participants were not given a classification rule at the start of training, approximately 85% of the BT items were given a response inconsistent with the rule and there was little



difference in response times between GT and BT items. Similarly, Regehr & Brooks (1993) found that 77% of the BT items were given rule-inconsistent responses when participants did not train with a rule. These authors reasoned that the high proportion of incorrect responses occurred because, without an explicit rule, participants relied primarily on similarity to specific prior instances to help them categorize the novel transfer stimuli. When participants were given an explicit and perfectly predictive classification rule at the start of training, Allen and Brooks and Regehr and Brooks both found that BT items produced considerably slower response times compared to GT items and there was only a small difference in the number of rule inconsistent responses between BT and GT items. These authors reasoned that the differences in rule inconsistent responses reflected the use of prior similar instances as a basis of classification, but because the difference in accuracy was so small, the classification rule was also being used. Thus, learning a rule helps to maintain accuracy, but it does not prevent facilitation or interference from prior item-specific knowledge.

Analogous experimental conditions were used in this thesis for determining the role of rules and similarity as participants become skilled at classifying the stimuli. If a transition to a faster, less effortful classification strategy reflects an increased reliance on retrieval, and a decreased reliance on the rule, then the difference in the proportion of rule inconsistent responses between GT and BT items should be quite large. This difference should approach levels similar to what was obtained in the no rule training condition reported in Allen and Brooks (1991). If the transition to a faster, less effortful classification strategy reflects the use of a production rule, then the difference between GT and BT items should remain relatively small, reflecting the importance of the rule as a basis of responding for both types of transfer items.

### ***Overview of the experiments in this thesis***

The background view that is addressed by this thesis is that there are common

conditions in everyday category learning in which there is a gradual transfer of control from processes that are rule-based to processes that are retrieval-based. Given a large number of old and similar-to-old items, a category structure in which items that are similar usually belong to the same category, and sufficient practice, this transition may be relatively automatic. This transition makes particular sense when one considers that under many conditions in everyday category learning similarity and rule-based procedures give the same answer, the cost of an error is low, and there is no pressure to explicitly justify answers. All of these conditions are true of the experiments in this thesis and therefore should provide favorable conditions for a transition to retrieval to occur. The thesis experiments themselves test an extreme view of this prediction, in that participants were also given a classification rule and practiced using that rule during a training phase. If a transition to a retrieval-based classification process is a common occurrence in everyday category learning then it should still occur despite the presence of a rule that was simple, predictive, and easy to apply.

## CHAPTER 3

### *Experiment 1: The effect of practice*

The intent of this first experiment was to provide an initial demonstration that a transition to a faster, retrieval-based classification procedure occurs relatively automatically given sufficient practice. As with previous experiments in this series (Regehr & Brooks, 1993; Allen & Brooks, 1991) stimuli for this study were designed to meet the constraints necessary for measuring participants' use of rules and retrieval. As shown in Figure 2, stimuli were drawings of imaginary animals that consisted of five binary features (body shape, neck length, body markings, number of legs and length of legs). GT and BT items were created by changing the value of one of the three relevant features. GT items were those in which this change of a relevant feature left the item in the same category as the most similar training item; BT items were those in which the relevant feature change put it in the opposite category to the most similar training item. To promote the acquisition and use of rule-based knowledge, participants practiced using a perfectly predictive two out of three feature additive rule and were given feedback if they made a classification not in accordance with the rule. To promote the acquisition and use of retrieval-based knowledge, each training stimulus was constructed from individuated features. This use of individuated features results in training stimuli that are visually distinctive from each other and, as argued by Regehr and Brooks (1993), should provide good encoding and retrieval cues that are specific to each individual animal. In addition, each transfer stimulus had four of the five features in common with a specific training stimulus, and thus there was a high degree of similarity between the training and transfer stimuli.

**Rule :** Builders have any combination of two of the three following features, an angular body, spots, and six legs, otherwise they are diggers.

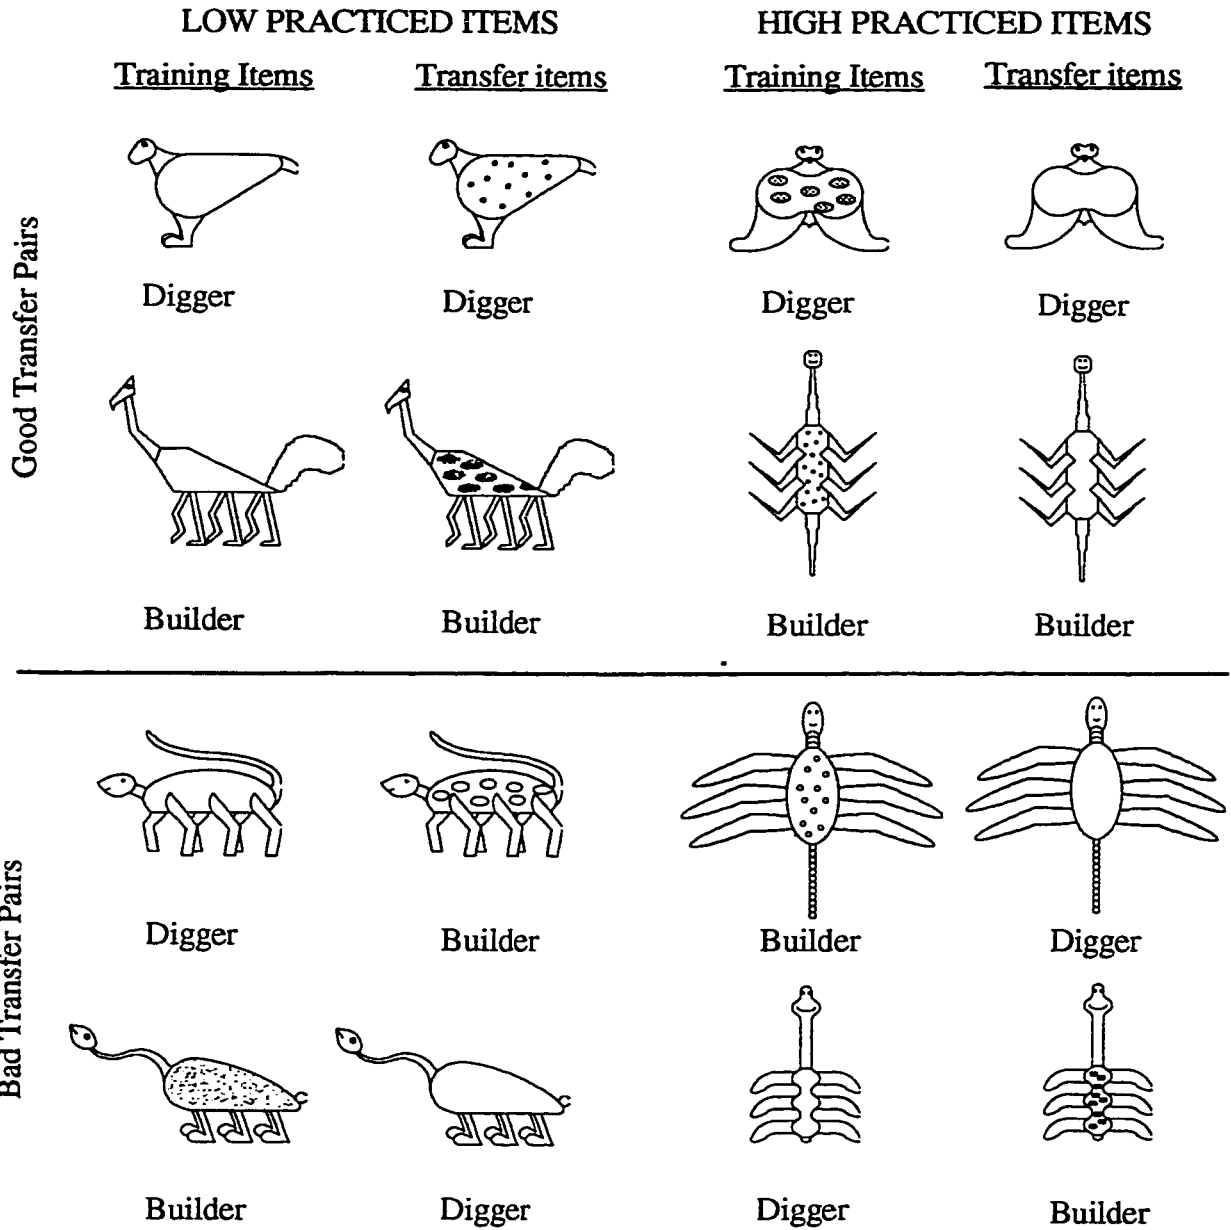


Figure 2. Examples of some of the stimuli used in the extra practice experiment (Experiment 1) and the corresponding rule-consistent response for each stimulus. Low practiced and high practiced items alternated across participants.

Finally, for every ten trials in the transfer phase, eight consisted of Old items. This high proportion of repeated stimuli should help retrieval because the presence of a large number of study items within a transfer list has been shown to help reinstate the study context and facilitate the use of memory retrieval (Jacoby, 1983; Allen and Jacoby, 1990).

The experimental procedure was similar to that used by Regehr and Brooks (1993) and Allen and Brooks (1991). Participants practiced using a rule to classify a set of training stimuli and then classified a set of transfer items consisting of Old, GT and BT items. The experiments differed from the prior studies, however, in the amount of practice each participant received. In Allen and Brooks and Regehr and Brooks, participants only practiced classifying each training stimulus five times before the transfer phase began. In the current experiment, eight of the training items were each classified five times before the transfer phase began (low practiced items) but the remaining eight training items were each classified 40 times before the transfer phase began. This extra practice was used to ensure that each item would be highly familiar to participants and aid in the transition to a fast classification procedure.

To confirm that the extra practice was leading to the development of a faster, less effortful classification procedure, a power function was fit to the training response times. As argued by Logan (1988, 1992) and Newell and Rosenbloom (1981), acquiring a skill in a task often leads to improvements in response times that follow a power function of the number of practice trials. This power function is considered to be indicative of a transition to a faster more automatic strategy for performing a task. Using a least squares criterion to determine the best fitting function, a power function of the form  $RT = A + Bx^{-c}$  (where  $A$  is the asymptote,  $B$  is the difference between the response time on the first trial and the asymptote,  $x$  is the trial number and  $c$  is the learning rate) was fit to the training times for each practice condition.

The predictions for this classification task are relatively straightforward. If extra practice facilitates a transition to retrieval of prior instances as a basis for classification, then it may also facilitate similarity-based retrieval for novel items, because once there is a sufficiently fluent stock of instances in memory, similarity-based retrieval may occur automatically. If so, then the difference in the proportion of rule-inconsistent responses between BT and GT conditions should be larger for the high practiced stimuli compared to stimuli that were not as highly practiced .

### *Method*

#### *Participants*

Fifteen students taking an introductory psychology course participated for course credit.

#### *Materials*

As shown in Figure 2, the stimuli consisted of drawings of imaginary animals that varied on five binary dimensions: body shape (rounded or angular), neck length (short or long), body markings (no spots, spots), leg length (short or long) and number of legs (two, six). Appendix A contains the informational structure of the stimuli that were used in this experiment. Eight items were created, and as shown in Appendix A, the combination of features that represented each of these items was chosen so that any one feature only occurred four times across the set of eight items and was not consistently paired with any other feature. This set of items was designated as the training items. Transfer items were then created by changing the value of the feature that represented body markings. This manipulation resulted in a set of eight transfer stimuli that were very similar to the training items. A second set of eight pairs of training and transfer items was created by repeating the first set of eight training and transfer pairs but reversing which member of the pairs was a training item and which was a transfer item. That is, the items designated as training stimuli in the first set of eight pairs became the transfer

stimuli in the second set of eight pairs and the items designated as transfer stimuli in the first set of eight pairs became the training stimuli for the second set of eight pairs. Although the training stimuli from the first set of eight pairs had the same structure as the transfer stimuli from the second set of eight pairs, and vice versa, the perceptual manifestation of the features was changed so that different animals were created. For example, compare the Low Practiced training item at the top left of Figure 2 with the High Practiced transfer item at the top right of Figure 2. Both animals have the same structure, 00000 but the perceptual characteristics of the features are different so that the animals look different from one another. Changing the perceptual manifestation of the features in this fashion made it possible to alternate the two sets of stimuli between the practice conditions so that each set served equally often as the high practiced items and the low practiced items.

The stimuli were categorized as builders or diggers according to a two out of three feature additive rule that used body shape, body markings, and number of legs as the relevant features. An animal was defined as a builder if it had builder values on at least two of the three relevant dimensions with all other animals being diggers. Two different rules were used to categorize the animals and were chosen to ensure that no one feature was predictive for the categories. The features associated with the builder category for each of these two rules were: any two of an angular body, six legs, or spots for the first rule and any two of an angular body, two legs, or spots for the second rule.

The combination of two classification rules and alternating the sets of 16 pairs of stimuli between training and transfer conditions ensured that each stimulus served once as a GT item, once as a BT item, once as a training item similar to a GT item, and once as a training item similar to a BT item. Four different counterbalancing orders were required to meet these constraints.

### ***Procedure***

The training stimuli were presented one at a time on the screen of a Macintosh Quadra 800 computer. Participants' responses were made by pressing one of the two keys on the keyboard, one designating a builder response and one a digger response. Response time was measured as the time elapsed between the presentation of the stimulus on the screen and the key press.

Each participant was told they would see a set of drawings of imaginary animals and their task was to place each animal into its correct category. To help them, the classification rule was presented on the screen at the beginning of the training phase and participants were asked to remember it. After signalling that they were ready, participants were presented with the training stimuli, consisting of a block of nine training stimuli with each block repeated 40 times. The order of items within a block was randomized with each repetition and always consisted of the eight high practice training items and one of the eight low practice training items. The low practice training items were chosen in a pseudorandom fashion with the constraint that an individual item was only shown five times during the training phase. For each presentation, the animal stayed on the screen until classified by the participant at which point the screen cleared and the animal was replaced, one second later, with the next trial. Feedback, consisting of the computer beeping and presenting the correct category label above the animal, was given if participants misclassified the item according to the rule. The animal and the category label stayed on the screen for one second before being replaced, one second later, with the next trial.

Participants were then presented with the list of transfer stimuli consisting of eight blocks of 10 stimuli. The order of items within a block was randomized and always consisted of seven of the eight high practice training items, one low practice training item, one high practice transfer item and one low practice transfer item. The low practice



training item, the GT item and the BT item were randomly chosen with the constraint that each item was only shown once within the transfer list. The high practice training items were also randomly chosen with the constraint that each item was only shown seven times within the transfer list. After the entire transfer list was presented once, the list was repeated in a new random order. Participants were not notified when the transfer phase began nor was feedback given for any item that was misclassified. Each item stayed on the screen until it was classified at which point it was replaced, one second later, with the next trial.

### *Results*

Although 16 participants were required to complete four rotations through the counterbalancing orders only 15 participants completed this study. Also, the second repetition of transfer stimuli for two participants was lost due to technical difficulties.

#### *Training*

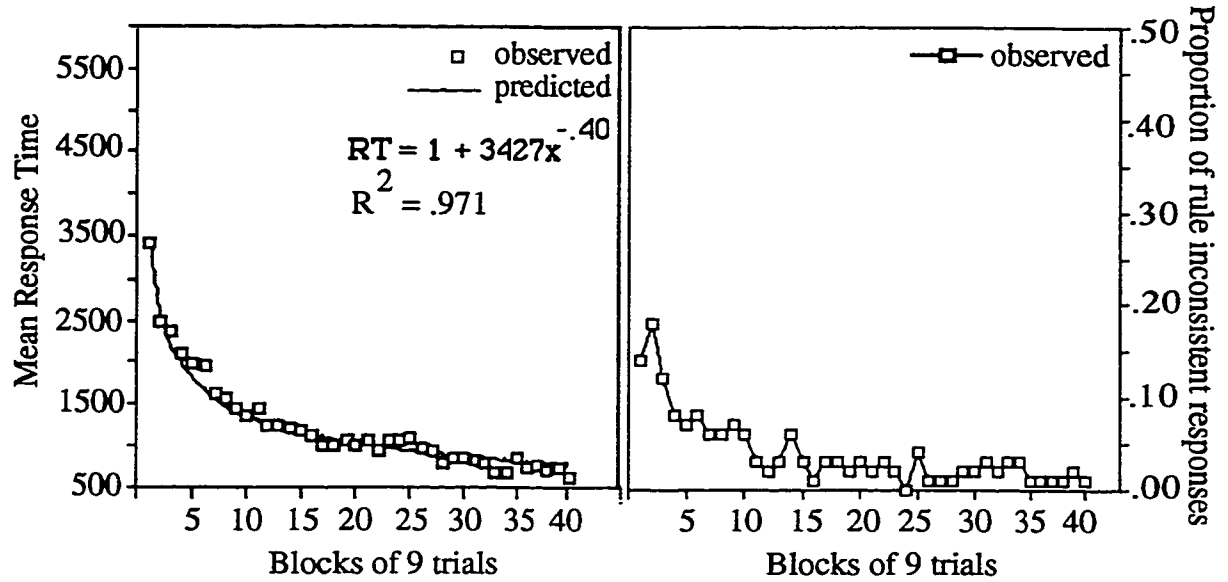
Figure 3 contains a plot of the mean response times, the corresponding fitted power functions, and the proportion of rule-inconsistent responses for each block of nine training trials. It appears, looking across the training blocks, that participants got faster and made fewer errors with practice.<sup>1</sup> Much of this benefit, however, occurred for high

---

1. A correct answer is defined as anything that is consistent with the classification rule, and an error is anything that is inconsistent with the classification rule. Classification errors can occur for a number of reasons and the significance of these errors depends on the stimulus type. For example, an error during training can occur because of guessing, because the participant did not use the rule correctly, because the participant misjudged a relevant feature, or because a participant misremembered a prior response. During transfer, errors on Old items and errors on GT items can occur for the same reasons that errors do at training. Because BT items and GT items were designed to be comparable with each other, errors on BT items also occur for the all same reasons as errors on GT items except BT items have an additional factor that can produce errors. Recall that rules and retrieval lead to opposite responses on BT items, and therefore if a participant makes a similarity-based response on a BT item it will not be consistent with the rule and will be considered an error. To separate errors due to similarity from errors due to other factors, the BT items are always compared to GT items and any difference between the two is attributed to the effect of similarity.

Because a similarity-based response is not really an error in the same sense as the other types of errors, the term "rule-inconsistent response" was used to describe any response not in accordance with a rule. This term still reflects a classification error on training items and on GT items, but it was felt that rule-

### High Practiced



### Low Practiced




Figure 3. Mean response times with corresponding fitted power functions, and mean error rates for high practiced and low practiced training stimuli plotted as a function of blocks of training trials (Experiment 1).

inconsistent response better captures the intended meaning of what is meant when an error occurs while classifying a BT item.

practiced items because they appear to be classified faster and are less error prone than the low practiced items. To verify these impressions, the mean response times for each block of nine training trials and the mean proportion of rule-inconsistent responses were analyzed. Presumably, if participants were learning a skill then both response times and errors should drop with practice. Two 2 x 40 ANOVAs were conducted with practice (high, low) and training block (training blocks 1 to 40) treated as repeated measures variables. Participants responded faster across trials, as shown by the significant main effect of training blocks,  $F(39,546)=9.98$ ,  $MSe = 755353$ ,  $p < .001$ . This benefit of practice was more pronounced for the high practiced stimuli than it was for the low practiced stimuli, as shown by the main effect of practice,  $F(1,14)=23.49$ ,  $MSe = 8529967$ ,  $p < .001$  and by the significant interaction between practice and the training blocks,  $F(39,546) = 2.554$ ,  $MSe = 660290$ ,  $p < .001$ . Participants also made fewer errors with practice, as shown by the significant main effect of training blocks,  $F(39,546) = 1.89$ ,  $MSe = .04$ ,  $p < .01$ , and the finding that participants made fewer errors in the high practice conditions than they did in the low practice condition  $F(1,14) = 9.32$ ,  $MSe = .08$ ,  $p < .01$ .

To determine if performance met the criteria of becoming faster and less effortful, the training response times for each practice condition were plotted as a function of each block of nine trials. Power functions were then fit to these plots using a least squares criterion to determine the best fitting power function. As shown in Figure 3, the best fitting power function for the low practice condition was of the form,  $RT = 73 + 2445x^{-.09}$ , but this function did not provide a very good fit between the observed and predicted scores,  $R^2 = .128$ . The best fitting power function for the high practice condition was of the form,  $RT = 1 + 3427x^{-.40}$  and did provide a good fit between the observed and predicted scores,  $R^2 = .971$ . These fits suggest that performance was fast and less effortful but only when participants had practiced

classifying individual stimuli a number of times.

### ***Transfer***

Analyses were performed on both the proportion of responses that were inconsistent with the rule and on participants' median response time for correct (rule consistent) responses made during the transfer phase. Two 2 x 3 ANOVAs were performed with two repeated measures variables: practice (high, low), and transfer type (Old, GT, BT). A summary of the results appears in Table 3.

***Rule inconsistent responses.*** The extra practice did not significantly increase the proportion of rule-inconsistent responses,  $F(1,14) < 1$ , nor did it interact with transfer type,  $F(2,28) = 1.53$ ,  $MSe = .01$ ,  $p > .10$ . There was a difference, however, in the proportion of rule-inconsistent responses across transfer types,  $F(2, 28) = 29.06$ ,  $MSe = .01$ ,  $p < .001$ . A subsequent set of paired t-tests were conducted after collapsing over the practice conditions and revealed that, whereas GT items ( $M = .03$ ) were less prone to a rule-inconsistent response than BT items ( $M = .23$ ),  $t(14) = 5.58$ ,  $p < .001$ , there was no difference between GT items and Old items ( $M = .03$ ),  $t(14) < 1$ .

***Response times.*** Median response times for correct (rule consistent) responses revealed a significant main effect of practice,  $F(1,14) = 4.74$ ,  $MSe = 167713$ ,  $p < .05$ . As shown in Table 3, participants responded faster to the high practiced items than they did to the low practiced items. There was also a significant main effect of transfer type,  $F(2,28) = 25.87$ ,  $MSe = 299351$ ,  $p < .001$ , and the interaction between practice and transfer type was also significant,  $F(2,28) = 5.86$ ,  $MSe = 307479$ ,  $p < .01$ . To explore the interaction further, a series of planned, paired t-tests were conducted. Because of the number of comparisons involved, however, a Modified Bonferroni Test for Planned Comparisons was used to prevent increasing the familywise error rate. The level of significance for these comparisons was set at  $p = .03$ . For low practiced stimuli, Old

Table 3.  
Mean responses to high practiced and low practiced transfer stimuli (Experiment 1)

Measure	<u>low practiced</u>		<u>high practiced</u>	
	Mean	SD	Mean	SD
Proportion of rule inconsistent responses				
Old	0.06	(0.06)	0.01	(0.01)
GT	0.03	(0.07)	0.03	(0.09)
BT	0.20	(0.16)	0.26	(0.22)
Median response time (rule consistent)				
Old	1426	(649)	707	(190)
GT	1492	(750)	1402	(832)
BT	1950	(750)	2196	(1200)

Note. low practiced = each training item seen 5 times before transfer  
high practiced = each training item seen 40 times before transfer  
GT = good transfer items.  
BT = bad transfer items.

items did not differ in response time from their corresponding GT items,  $t(14) < 1$ . There was a difference, however, between GT and BT items  $t(14) = 4.91$ ,  $p < .01$ , with BT items being slower. For the high practiced stimuli, response times for Old items were faster than the GT items,  $t(14) = 3.73$ ,  $p < .01$ , and responses to GT items were also faster than responses to BT items,  $t(14) = 2.60$ ,  $p < .03$ .

### *Discussion*

The purpose of this experiment was to determine whether the transition to a fast, retrieval-based classification procedure occurs relatively automatically given sufficient practice, despite participants also being given a rule. If so then this use of retrieval should extend to both familiar training stimuli and also novel but similar-to-old stimuli. When mean response times for the training items were plotted as a function of blocks of training trials, it was found that a power function provided a good fit of the high practiced items. This result, combined with the drop in errors that occurred with practice, confirms

that participants had acquired some degree of skill at classifying the items. This skill was specific to the high practiced items, however, as a power function did not provide a good fit of the low practiced training times.

During the transfer task, participants made more rule-inconsistent responses when classifying BT items than GT items. This result is consistent with other studies that have used this GT/BT procedure (Allen & Brooks, 1991; Regehr & Brooks, 1993; Neal et al., 1995) and suggests that participants were relying on retrieval to some degree. However, the difference between GT and BT items was not as large as would be expected if the transition to a fast classification procedure reflected the exclusive use of retrieval. This smaller than expected difference suggests that participants were continuing to rely on the rule for classifying the novel transfer stimuli.

An additional finding was that there was little difference in the proportion of rule-inconsistent responses for BT items between the two practice conditions. In fact, the difference between GT and BT items was within the same range as the other studies that have used this procedure despite the extra practice (Allen & Brooks, 1991; Regehr & Brooks, 1993). This finding is surprising because if extra practice leads to a faster, less effortful classification procedure, and if this procedure reflects an increased reliance on retrieval, then the extra practice should have produced an increase in rule inconsistent responses for BT items. The nonsignificant difference in rule inconsistent responses that occurred between the practice conditions may suggest that the extra practice had no effect on facilitating similarity-based retrieval. This conclusion is not definitive, however, because there may not have been enough power to detect such a difference if it existed. Cohen's  $d$ , which is a measure of effect size, was calculated for the difference in rule-inconsistent responses between low practice BT items and high practice BT items. This value was fairly small,  $d = .24$ , and suggests that the experiment did not have enough power to detect an effect. This lack of power is not crucial, however, because it was

expected that the extra practice would facilitate the reliance on retrieval and produce a large difference between GT and BT items rather than the small difference that was found.<sup>2</sup>

Although the extra practice may not have had much of an effect on how the novel transfer stimuli were classified, it had a large effect when classifying training items. There was a significant drop in response time when classifying the high practiced training items compared to training items that were low practiced. Additionally, the difference in response time when classifying GT items compared to Old items is larger in the high practiced condition than the low practiced condition, despite the number of rule-inconsistent responses being comparable for these items. Apparently, given sufficient practice, GT items are treated differently than Old items. This suggests that while extra practice may not facilitate the use of retrieval to classify novel but similar-to-old instances, it does benefit the recognition of prior instances, producing performance that is considerably faster than what was obtained without the extra practice.

Finding no increase in the difference between GT and BT items suggests that extra practice, by itself, may not be a strong enough inducement for retrieval-based processes to generalize to novel items. The next experiment was designed to strengthen the similarity manipulation even further in hopes of creating conditions conducive for similarity-based responding.

2. It is difficult to say exactly how large the difference in rule inconsistent responses between GT and BT items should be with extra practice. There is some data, however, that at least serves as a rough guide. For example, Allen and Brooks (1991) and Regehr and Brooks (1993) found that when participants were not given a rule at training and therefore had to rely on similarity-to-old to classify the novel transfer stimuli, roughly 80% of the BT items were given a rule- inconsistent response. In a pilot study using stimuli similar to those described in Experiment 2 of this thesis, 90% of the BT items were given rule-inconsistent responses when participants were not given a rule at training. If a fast, less effortful classification procedure reflects an increased reliance on retrieval and a decreased reliance on the rule, then it is likely that the extra practice should lead to an increase in rule inconsistent responses that are within the same range as those found in these studies.

## CHAPTER 4

### *Experiment 2: The effect of many similar training neighbors*

If the transition to a fast classification procedure is based on retrieval, then providing participants with multiple similar training instances, rather than just repeating the same training instance, may facilitate generalization around the training instances. As an analogy, our ability to classify dogs may be better if we learned to classify ten different collies at training compared to just seeing Lassie over and over again. Evidence consistent with this view comes from a phenomenon in category learning called the category size effect, in which increases in the number of similar members in a category leads to better classification of both prototypes and other novel members of the category. This effect has been attributed to the influence of similarity, either to a prototype or to prior instances. For example, Homa et al. (1981) concluded that the category size effect occurs because large categories make it easier to abstract a prototype to which novel items can then be compared. Hintzman (1986) and Nosofsky (1988) reasoned that the category size effect occurs because with large categories there is an increased likelihood that a novel exemplar will be very similar to at least one of the known exemplars.

In order to increase the size of the categories, it was necessary to introduce a degree of variation in the training stimuli. This was done by altering the perceptual characteristics of each training instance so that, although the underlying structure was the same, perceptually, no two training instances were identical. Introducing variation within the training instances may also benefit the use of similarity-based retrieval because several studies have found that variation improves transfer performance. Schmidt and Bjork (1992), in a review of practice effects, concluded that training with a variety of situations rather than just one situation facilitates the transfer of that skill to novel



situations. For example, practice throwing a basketball into a net from a variety of distances leads to better transfer of that skill to an unpracticed distance than does repeatedly throwing the basketball from the same distance. This benefit of variation at training has been found in a number of classification studies as well (Posner & Keele, 1968; Homa, 1984; Hintzman, 1986; Nosofsky, 1988). For example, Posner and Keele (1968) trained one group of participants with medium distortions of four category prototypes and another group with low distortions of the same four category prototypes. They found that the group that learned with the medium distortions showed better transfer to novel stimuli that were highly distorted versions of the prototype. Homa and Vosburgh (1976) found a similar result of variation but also found that variation interacted with category size. They had two groups of participants learn to classify dot patterns into three categories with each category having either three, six or nine training exemplars. One group of participants was trained using a variety of distortions of the training exemplars, whereas the other group was trained with only low level distortions of the training exemplars. Homa and Vosburgh found that participants who trained with the low level distortions of the training exemplars only produced a transfer benefit when classifying novel patterns from the small category. Those participants who trained with a variety of distortions of the training exemplars produced a transfer benefit when classifying novel items from the large categories.

Another reason also exists for why variation in the training exemplars should benefit transfer. The introduction of new perceptual manifestations of features in the training stimuli might accustom participants to variation in what is essentially the same animal. Thus, increased variation might allow participants to loosen their criteria for what they consider similar and increase the likelihood that novel transfer stimuli will benefit from their similarity to prior instances. As an analogy, classifying ten different Golden Retrievers, ten different Collies, and ten different Miniature Poodles as dogs may make it

easier to then classify a Cocker Spaniel as a dog.

As shown in Figure 4, the underlying structure of the animals used in this experiment consisted of six binary features: tail length (short or long), neck length (short or long), leg length (short or long), body shape (rounded or angular), body markings (stripes or spots), and number of legs (two or six). Two different manipulations were used to increase the category size and to introduce more perceptual variation within the category members. First, two pairs of similar animals were created that differed from one another on only 2 of the 6 underlying binary features. This ensured that the pairs were more similar to one another than they were to any other category member. For example, consider the pair of training animals located at the top of the upper panel of Figure 4. The underlying structure of the builder located on the left is represented as 001000 and the underlying structure of the builder located on the right is 010000. Thus the members of the pair have four of the six features in common and are similar to one another. Additional training exemplars were created by changing the postures of each pair of animals and then by changing the perceptual manifestation, but not the informational value, of one of the relevant dimensions. For example, consider the "builder" located in the top left corner of Figure 4. Every animal located below this particular animal has the same underlying structure as the original (001000), although perceptually each animal is different. The animals located under the other member of the pair (010000), also have the same underlying structure, although perceptually they are different from one another.

Creating training exemplars in this fashion ensured that each category was larger than those used in Experiment 1 and also contained more variation than existed in Experiment 1. It is expected that the addition of multiple similar training instances combined with extensive practice, individuated stimuli, and a large proportion of training items in the transfer list, will be sufficient to induce participants to rely on similarity-based retrieval when classifying the novel transfer stimuli. If so, then the difference in the proportion

**Rule:** Builders have any combination of two of the three following features, short tail, long neck, long legs, otherwise they are diggers.

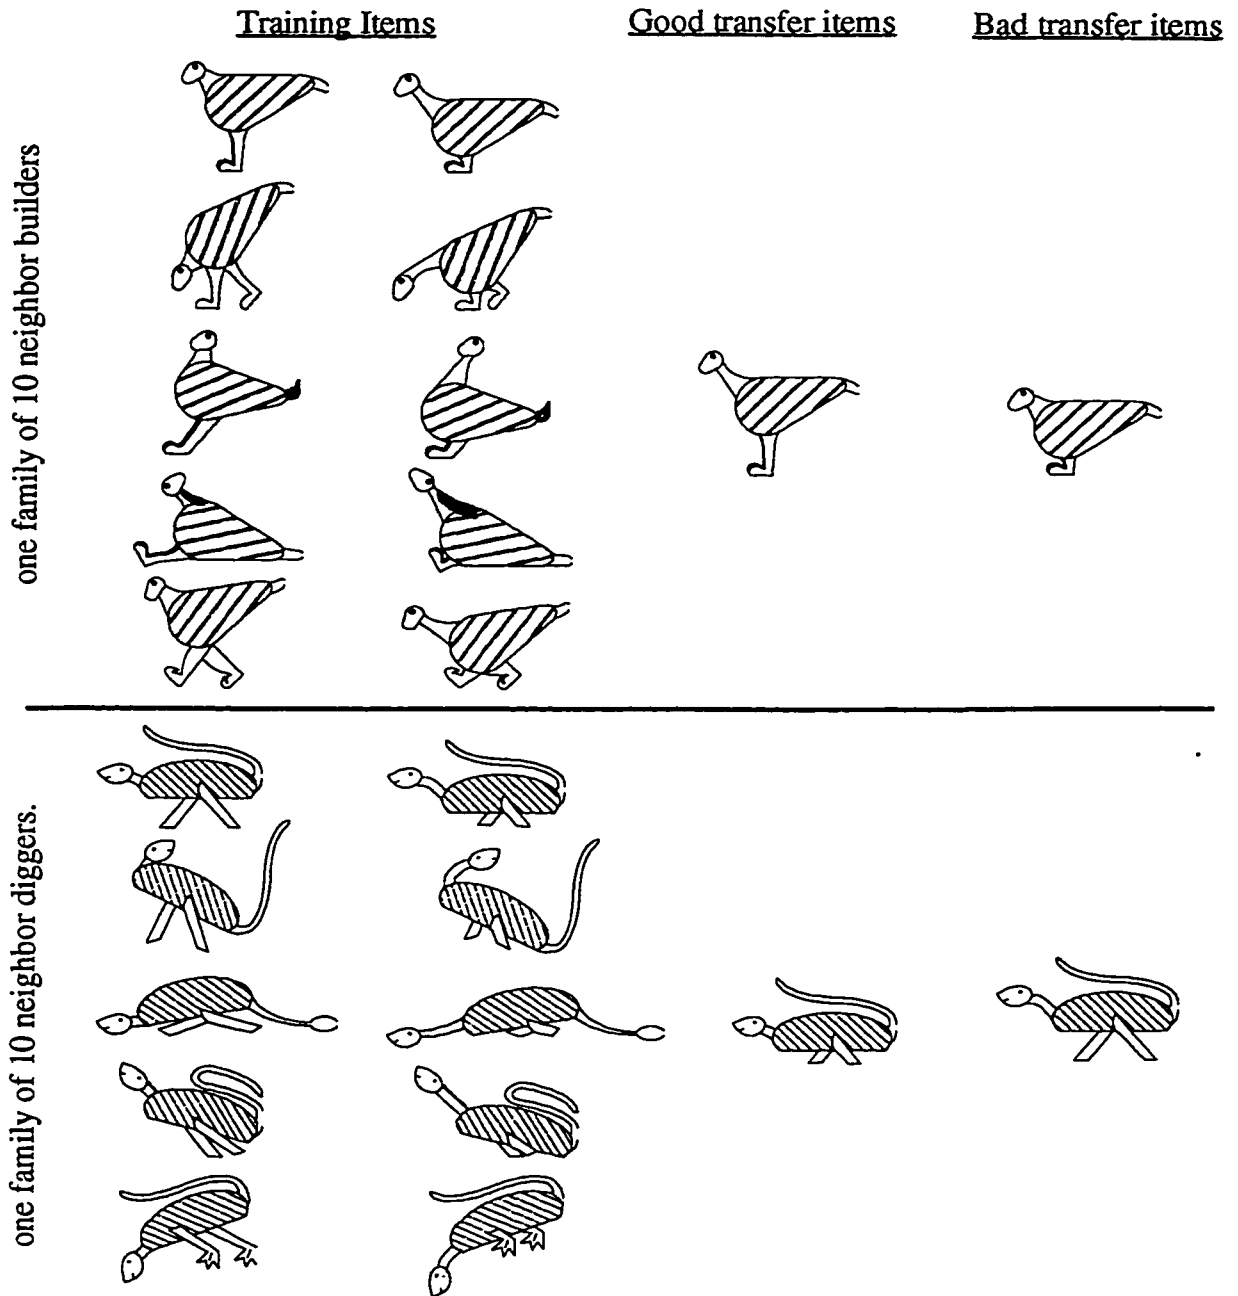


Figure 4. Examples of two families of stimuli used in the multiple neighbor experiment (Experiment 2).

of rule-inconsistent responses between BT and GT conditions should be quite large.

### *Method*

#### *Participants*

Twelve students taking an introductory cognitive psychology course participated for course credit.

#### *Materials*

The stimuli were line drawings of imaginary animals that varied on six binary dimensions: tail (short or long), neck length (short or long), leg length (short or long), body shape (rounded or angular), body markings (stripes or spots), number of legs (two or six). Appendix B contains the informational structure of these stimuli. Eight families of stimuli were created and consisted of ten perceptually similar stimuli plus a GT and BT item. These families first started as a quartet of stimuli and consisted of two similar training stimuli that differed from each other on two of the three relevant features, a GT item that differed from the training pairs by one relevant feature, and a BT item that also differed from the training pairs by one relevant feature. The combination of features that represented each item was chosen so that no one feature occurred on all members of a category. To create the entire family, four variations of each training stimulus was created by changing the overall posture of the animal but not the underlying structure. As shown in Figure 4, this produced eight new variations of training stimuli in each family with each variation containing the same underlying structure as the original pair. In addition, on six of these new variations, the perceptual manifestation of one relevant dimension was changed (i.e. a short tail but a different looking short tail).

The stimuli were categorized as builders or diggers according to a two out of three feature additive rule that used tail length, neck length and number of legs as the relevant features. An animal was defined as a builder if it had builder values on at least two of the three relevant dimensions, all others being diggers. Three different rules were used to

categorize the animals, requiring three different counterbalancing orders. The features associated with the builder category for each of these three rules were: any two of a short tail, long neck, or long legs for the first rule; any two of long tail, long neck, or short legs for the second rule; and any two of short tail, short neck, or short legs for the third rule.

### ***Procedure***

The training stimuli were presented, one at a time, on the screen of a Macintosh LC 475 computer. Participants' responses were made by pressing one of the two keys on the keyboard, one designating a builder and one a digger response. Response times were measured as the time elapsed between the presentation of the stimulus on the screen and the key press. Participants were told that they would see a set of drawings of imaginary animals and their task was to place each animal into its correct category. They were not given any instructions emphasizing speed or accuracy. To help participants with their classification, the two out of three feature rule was presented concurrently in a separate window at the bottom of the computer screen and remained there for the entire training phase.

Participants were presented with the training stimuli, consisting of a list 80 items that contained 10 similar training items in each of the eight families. This list was repeated four times, with the order of stimuli randomized within each repetition and for each participant. For each presentation, the stimulus stayed on the screen until classified by the participant. Feedback, which consisted of the computer beeping and presenting the correct category label underneath the item, was given if participants misclassified the item according to the rule. If feedback was not necessary, the animal disappeared as soon as the participant responded and was replaced with the next trial one second later. If feedback was necessary, the picture of the animal and its category label stayed on the screen for one second at which point it disappeared and was replaced with the next trial one second later.

At the beginning of the transfer phase, a message appeared on the computer screen saying that the participant was doing fine and was to continue without feedback. Participants were allowed a short break of no more than 30 seconds and were then asked to continue. No mention was made of the new animals, nor were the rules presented on the screen. In addition, feedback was no longer given for incorrect classifications and participants were not given any explicit instructions to respond fast or accurately. Participants were presented with a list of transfer stimuli that consisted of the 80 training items plus the 8 good transfer and 8 bad transfer items. The list was presented in a random order, and then the entire list was presented again in a different random order for each participant. Each item stayed on the screen until the participant classified it, at which point it disappeared and was replaced, one second later, by the next trial.

## *Results*

### *Training*

Figure 5 contains the mean response times, the corresponding fitted power functions, and the mean proportion of rule-inconsistent responses for each block of eight training trials. Participants appeared to respond faster and make fewer rule-inconsistent responses as the number of trials increased. To verify these impressions, the mean response times for each block of eight training trials and the mean proportion of rule-inconsistent responses for each block of eight training trials were analyzed. Two 2 x 40 ANOVAs were conducted with training block (training blocks 1 to 40) treated as a repeated measures variable. These analyses confirmed that participants responded faster with practice, as shown by the significant main effect of training blocks,  $F(39,429)=30.86$ ,  $MSe = 408972$ ,  $p < .001$ , and that they made fewer errors with practice, as shown by the significant main effect of training blocks,  $F(39,429) = 4.69$ ,  $MSe = .007$ ,  $p < .001$ .

To determine if performance had switched to a fast classification procedure, the training response times were plotted as a function of the number of blocks of eight trials.




Figure 5. Mean response times with the corresponding fitted power function and mean proportion of error plotted as a function of blocks of training trials when there are multiple similar training neighbors (Experiment 2).

A power function was then fit to these average times using a least squares criterion to determine the best fitting function. The best fitting power function was of the form,  $RT = 0 + 5958x^{-.39}$  and provided a good fit between the observed and predicted scores,  $R^2 = .969$ .

### **Transfer**

Analyses were performed on both the proportion of rule-inconsistent responses and on participants' median response time for correct (rule consistent) classifications made during the transfer phase of the experiment. These analyses involved two one way ANOVAs with one repeated measures variable, transfer type (Old, GT, BT). A summary of the results appears in Table 4.

**Rule-inconsistent Responses.** An analysis of the proportion of rule-inconsistent responses revealed that there was an effect of transfer type,  $F(2,22) = 21.10$ ,  $MSe = .011$ ,

Table 4.  
Mean responses to transfer stimuli that are similar to multiple training instances  
(Experiment 2).

Measure	Mean	SD
Proportion of rule inconsistent responses		
Old	0.02	(0.02)
GT	0.01	(0.02)
BT	0.26	(0.18)
Median response time (rule consistent)		
Old	1248	(519)
GT	1148	(544)
BT	1988	(1181)
<u>Note.</u> GT = good transfer items BT = bad transfer items		

$p < .001$ . Subsequent planned paired t-tests showed that participants were more likely to misclassify BT items than GT items,  $t(11) = 4.71$ ,  $p < .01$ , but there was no difference in performance between Old items and GT items,  $t(11) = 1.40$ ,  $p > .10$ .

**Response Times.** Participants appear to be influenced by the type of item as shown by the significant main effect of transfer type,  $F(2,22) = 14.39$ ,  $MSe = 175757$ ,  $p < .001$ . Planned paired t-tests revealed that participants took longer to classify BT items than GT items,  $t(11) = 4.01$ ,  $p < .01$ . GT items were also responded to faster than Old items,  $t(11) = 2.47$ ,  $p < .05$ .

### Discussion

The purpose of this study was to determine if an increase in the size of the category and in the variation within that category would help facilitate similarity-based retrieval when classifying novel stimuli. When mean response times for the training items were plotted as a function of blocks of training trials, a power function provided a good fit of training times. Combined with the decrease in errors that occurred across the training



trials, it would appear that participants had acquired some degree of skill at classifying the items, despite the increase in the size of the categories and the variation within the categories.

During the transfer task, participants made more rule-inconsistent responses to BT items than GT items suggesting that they relied on similarity to some degree. The difference between GT and BT items, however, was not as large as expected and was within the same range as the difference in Experiment 1. This smaller than expected difference, combined with the presence of slower responses to BT items relative to GT items, suggests that participants had maintained their use of the classification rule despite the use of experimental manipulations designed to facilitate retrieval. An additional result was also curious. The response times required to classify Old items in this experiment were actually slower than the response times required to classify GT items. This is puzzling especially when one considers that participants had already classified the Old items several times and that the difference between Old and GT items was in the opposite direction in Experiment 1. A potential explanation for this result could be that studying so many similar training instances may have impaired participants' ability to recognize individual training stimuli, which then caused them to take longer to respond than they might have otherwise. In any case, the small difference in rule-inconsistent responses between GT and BT items is puzzling and not what was expected if participants were relying on retrieval.

The intent of this experiment, and of Experiment 1, was to create conditions that would be conducive to using retrieval as a basis of classification and then see to what degree retrieval-based processes were used to classify both familiar stimuli and also novel but similar-to-old stimuli once performance had become fast and effortless. As such, the conditions that were used are consistent those found in everyday category learning. Individual instances were highly distinctive and the novel stimuli were highly

similar to the training instances. In addition, the categories consisted of a large number of old and similar-to-old items, the cost of an error is low, and there is no pressure to explicitly justify answers. The factors that were manipulated in the experiments (i.e. practice and multiple similar neighbors) have been found, by other researchers, to facilitate the use of retrieval.

The findings from Experiment 1 and Experiment 2 suggest that participants are able to adopt a fast classification procedure during training and there is some evidence that this fast procedure reflects the use of a retrieval-based process. However, there was also a puzzling tendency for participants to continue to use the rule when classifying novel stimuli despite the presence of the conditions designed to facilitate the use of similarity-based retrieval. The idea that faster, less effortful classification performance reflects the fading away of the classification rule and an increase reliance on retrieval appears to be questionable.

## CHAPTER 5

### *Experiment 3: The effect of a deadline*

In the first two experiments, the presence of rule-inconsistent responses on BT items demonstrates some reliance on retrieval-based categorization procedures when classifying novel items. However, the idea that fast classification performance reflects the fading away of the classification rule has received little support. Neither extra practice nor increasing the number of similar neighbors at training seems to produce an increased reliance on similarity-based retrieval when classifying novel but similar-to-old items. Instead, participants seem willing to maintain their use of the classification rule when classifying the transfer stimuli. It is possible that this continued reliance on the classification rule occurs because the experimental conditions are still not conducive to using similarity-based retrieval. It was decided to introduce a factor that has already been shown to lead to increases in the use of similarity, even when classifying stimuli that encourage a classification procedure based on single feature rules. Using a response deadline, Smith and Kemler-Nelson (1984) found that when participants were instructed to respond quickly they increased their reliance on similarity relative to a condition in which they were allowed to respond at their own pace.

The introduction of a response deadline should also prove useful as a manipulation check to see if the experimental conditions are conducive for using similarity. According to a classification account derived from Logan's Instance theory, rule application is slow and effortful while retrieval is fast and effortless. If this assumption applies under the current experimental conditions, then encouraging participants to respond quickly should make it hard to apply the rule, thus allowing similarity to play a larger role and produce an increase in the proportion of BT errors.

## *Method*

### *Participants*

Seventeen students taking an introductory cognition psychology course participated for course credit.

### *Materials*

The stimuli were line drawings of imaginary animals similar to those used in the prior experiment and are shown in Figure 6. Eight triplets of stimuli were created and each triplet consisted of a training item, a GT item and BT item. The GT and BT items were both equally similar to their corresponding training stimulus, differing by one relevant feature. The classification rules were identical to those used in Experiment 2, and the informational structure of the stimuli is shown in the Appendix C.

### *Procedure*


The procedure for this experiment was identical to Experiment 2 except for the following changes. The training list consisted of a block of eight training stimuli and each block was repeated 18 times. After this initial round of training, participants were told that their task was to continue to classify the stimuli as builders or diggers but now they were to beat a deadline imposed by the computer. The deadline was set for each participant and was determined by calculating their median response times for the 17th and 18th blocks of training trials. Participants were not told how fast they were to respond but were allowed to practice beating the deadline for two additional presentations of the set of eight training items. If they responded slower than the deadline, a tone sounded immediately after their response. Participants also continued to receive feedback on classification errors which consisted of the computer beeping and placing the category label under the animal.

The procedure for the transfer phase was also identical to Experiment 2 with the exception that participants were presented with a transfer list that consisted of the set of

---

**Rule:** Builders have any combination of two of the three following features, short tail, long neck, long legs, otherwise they are diggers.

Training items




Builder



Builder

Good transfer items



Builder



Builder


Bad transfer items



Digger



Digger



Digger



Digger




Builder



Digger



Digger



Builder

---

Figure 6. Examples of some of the stimuli used in the deadline experiment (Experiment 3) and the corresponding rule-consistent response for each stimulus.

eight training items, plus eight GT items and eight BT items. The list of 24 items was then repeated and the items within each repetition were presented in a different random order for each participant. Participants were told that they would be presented with the animals they had just classified but in addition there would be some new animals that resembled the ones they had just classified. They were also told to continue to respond as fast as possible. Feedback was no longer given for any item that was misclassified according to the rule but the tone continued to sound if a participant's response was not faster than their deadline. Each item stayed on the screen until a response was made.

### *Results*

Two participants were replaced as their response deadlines were approximately two seconds and were judged to be too slow both for a speeded classification task and in relation to the other participants.

### *Training*

Figure 7 contains the mean response times and proportion of rule-inconsistent responses for each block of eight training trials. As expected, participants responded faster and made fewer errors as the number of trials increased. To confirm these expectations, the mean response times and the proportion of rule-inconsistent responses for the trials before the response deadline was introduced (the first 18 blocks of eight training trials) were analyzed. Two 2 x 18 ANOVAs were conducted with training block (training blocks 1 to 18) treated as a repeated measures variable. These analyses confirmed that participants responded faster with practice, as shown by the significant main effect of training blocks,  $F(17,238)=36.11$ ,  $MSe = 385931$ ,  $p <.001$ , and that they made fewer errors with practice, as shown by the significant main effect of training blocks,  $F(17,238) = 8.56$ ,  $MSe = .007$ ,  $p <.001$ . Responses for the final two blocks of training, when participants were under time pressure, were not analyzed. It should be noted, however, that the mean response deadline across participants was 703 ms and, as




Figure 7. Mean response times with the corresponding fitted power function and mean proportion of errors plotted as a function of blocks of training trials when there is a response deadline (Experiment 3). Note that the response deadline was imposed on the 19th and 20th blocks of trials therefore the power function was only fitted to the first 17 blocks of training trials.

shown in Figure 7, the average response times for these final two blocks of training were faster than the response times for previous blocks of training ( $M = 741$  ms and  $M = 571$  ms respectively). It should also be noted that with the decrease in speed there was a corresponding increase in rule-inconsistent responses. As participants began to respond faster, they also made more rule-inconsistent responses. The mean proportion of rule-inconsistent responses for the 19th block of training was .06, and for the 20th block of training the mean was .17. Apparently participants found it difficult to maintain accuracy when they were responding fast.

To determine if participants had switched to a faster, less effortful strategy during the

initial 18 blocks of training trials, the mean training response times were plotted as a function of each block of eight trials. A power function was fit to these average times using a least squares criterion to determine the best fitting function. As shown in Figure 7, the best fitting power function was of the form,  $RT = 0 + 4561x^{-.52}$  and provided a good fit between the observed scores and those predicted by the power function,  $R^2 = .975$ .

### ***Transfer***

Analyses were performed on both the proportion of rule-inconsistent responses and on the participants' median response time for correct (rule consistent) responses made during the transfer phase. Two one way ANOVA's were performed with one repeated measures variable: transfer type (Old, GT, and BT). A summary of the results appears in Table 5.

***Rule-inconsistent responses***. The proportion of rule-inconsistent responses differed significantly across transfer types,  $F(2, 28) = 32.44$ ,  $MSe = .02$ ,  $p < .001$ . A subsequent set of planned paired t-tests was conducted and revealed that participants made fewer rule-inconsistent responses when classifying GT items than BT items,  $t(14) = 6.55$ ,  $p < .001$ . Participants also made fewer rule-inconsistent responses when classifying GT than Old items,  $t(14) = 2.68$ ,  $p < .05$ .

***Response times***. The mean response deadline, across subjects, was 703 ms. Median response times for correct (rule consistent) responses revealed a marginal effect of transfer type,  $F(2, 28) = 2.84$ ,  $MSe = 3360$ ,  $p < .08$ . A subsequent paired t-test revealed that although participants were trying to respond as fast as possible, BT items were still responded to more slowly than GT items,  $t(14) = 2.30$ ,  $p < .05$ .

### ***Discussion***

This experiment was intended to check if the assumptions of an account of classification, based on Logan's Instance theory, held in the current design. If rule



Table 5.  
Mean responses to transfer stimuli when a response deadline is imposed at transfer (Experiment 3).

Measure	Mean	SD
Proportion of rule inconsistent responses		
Old	0.21	(0.14)
GT	0.13	(0.14)
BT	0.49	(0.15)
Median response time rule consistent responses		
Old	560	(104)
GT	567	(108)
BT	607	(142)
Mean Response Deadline 703 ms		
<u>Note.</u> the response deadline was set for each participant and was based on their training times.		
GT = good transfer items		
BT = bad transfer items		

application is more effortful than similarity-based retrieval, then stressing rule application by making participants respond as fast as possible should make it hard to use the rule and facilitate the use of similarity. There were several findings of note. Introducing a response deadline produced an increase in the number of rule-inconsistent responses for BT items. Unfortunately, the increase to 49% was still not as high as would be expected if the rule was not being used. Furthermore, the proportion of rule-inconsistent responses increased relative to the previous experiments, but this increase in errors was not disproportionately higher for BT items because the number of rule-inconsistent responses also went up for the Old items and the GT items. In fact, participants were surprisingly error prone when classifying the Old items and produced a large number of rule-inconsistent responses after the deadline was introduced at training and during the transfer task. The proportion of rule-inconsistent responses for Old items during the

transfer task was actually larger than it was for GT items. Responding quickly appears to not only makes participants more error prone but it may selectively disrupt how the Old items are classified.

In an attempt to understand how participants were classifying the stimuli, extensive protocols were taken from participants. One of the questions they were asked was to estimate how many of the animals, classified in the last part of the experiment, had been ones they had studied. The mean estimate, across participants, was that 54% of the transfer items were old. Recall that only a third of the transfer items were actually old, therefore participants were overestimating the number of training stimuli that were present in the transfer phase. In addition, participants were also asked to describe how they were classifying both the old and new animals in the transfer phase. Curiously, nine participants said they could recognize the training items and did not need to use the classification rule to classify these items. Twelve of the 15 participants said they were trying to use the rule as fast as possible when classifying animals they judged to be new. The protocols suggest that old versus new may have a special status for participants with most participants trying to recognize old stimuli and use the rule on new stimuli. This latter conclusion, in which most participants use the rule on items they judged to be new, is surprising and suggests that despite the emphasis on speed, participants continue to use the classification rule.

When the findings from the first three experiments and the protocols are considered together, an interesting method of coordinating rules and retrieval emerges, one that is more complex than was originally assumed. Perhaps, under the current experimental conditions, participants are performing an old/new discrimination before classifying the stimuli, relying on a faster memory retrieval process when classifying old items, and applying the slower classification rule when classifying novel items. This strategy would account for the fast classifications of old animals given extensive practice and the

continued reliance on the rule when classifying some of the novel animals. In addition, participants claimed to be using the rule when classifying novel items but some of the BT items were given rule-inconsistent responses suggesting that participants are relying on retrieval to some degree. Perhaps when a similarity-based classification is made for novel stimuli, some proportion of these responses actually reflect participants failing to discriminate the novel transfer items as new. In essence, while the difference between GT and BT items may still reflect a retrieval-based classification process, this process may actually consist of false recognitions rather than similarity-based retrieval.

Performing an old/new discrimination before classifying the stimuli also accounts for some of the more puzzling findings that were obtained in the first three experiments. For example, in the low practice condition of Experiment 1, the small difference in response times between Old items and GT items may have occurred because participants did not have enough training to easily tell old from new. In Experiment 2 and in Experiment 3, participants' performance on Old items was worse than their performance on GT items. This finding may have occurred because the high degree of variation within the training stimuli of Experiment 2 and the instructions to respond quickly in Experiment 3 made the old/new discrimination harder for old items than it was for novel items.

The next two chapters were designed to provide converging evidence for the hypothesis that participants are doing an old/new discrimination before classifying the stimuli. The experiments in Chapter 5 manipulated the stimuli so that the difficulty of discriminating old from new varied, and this range in difficulty influenced how participants classified the stimuli. By explicitly asking participants to decide whether a stimulus was old or new, the experiments in Chapter 6 were designed to provide a more sensitive measure of the degree that novel transfer stimuli are being falsely recognized. Evidence consistent with the hypothesis that participants rely on memory to classify items they believe are old and rules to classify items they believe are new was found in

both Chapters.

## CHAPTER 6

The experiments in this chapter were designed to provide supporting evidence for the hypotheses that participants perform a recognition check before classifying the stimuli and that at least some of the evidence for similarity-based retrieval can be reinterpreted as false recognition of novel stimuli. The logic behind the experiments in this chapter is that if novel transfer stimuli are falsely recognized as old, then making the old/new discrimination harder should have the additional effect of increasing the proportion of rule-inconsistent responses on BT items. If this logic is true, then an interesting issue is how participants decide whether something is old or new. Obviously if a stimulus is completely different from the other training stimuli, then it should be an easy task to decide that the item is new. The stimuli used in this thesis, however, were designed to be very similar to the training items and therefore may not signal their novelty to the same degree as a completely new item. Perhaps the presence of some kind of novel feature, on an otherwise familiar animal, signals the novelty of the item. If so, then designing the stimuli so that the novel feature is less obvious should make it harder to perform an old/new discrimination. If novel transfer stimuli are falsely recognized as old, then a more difficult old/new discrimination should increase the proportion of rule-inconsistent responses on BT items. In Experiment 4, two variables, multiple prior neighbors and a decrease in novelty cues, were deliberately confounded with each other in order to make the old/new discrimination more difficult and thereby demonstrate a relative increase in bad transfer errors. In Experiment 5, multiple prior neighbors and a decrease in novelty cues were unconfounded from each other in order to determine their relative contributions to participants' performance.

***Experiment 4: The role of novelty detection***

In order to make the old/new discrimination harder it was decided to reduce the novelty inherent within the transfer stimuli. The procedure for creating stimuli was similar to that used in Experiment 2. As shown in Figure 8, the underlying structure of the animals that were used in this experiment consisted of six binary features: tail length (long or short), body markings (stripes or spots), number of legs (2 or 6), body shape (rounded or angular), neck length (short or long) and leg length (short or long). Two conditions were created, one in which novelty was reduced and the other in which the novel cue was present. First, pairs of similar animals were created that differed from one another on two of the six underlying binary features but were always in the same category according to the rule. For example, consider the first pair of training items in the bottom panel of Figure 8. The underlying structure of the builder on the left is 010000 and the underlying structure of the builder on the right is 001000. Therefore, each member of the pair has 4 of the 6 features in common and is also in the same category. The GT and BT items were created by recombining the relevant features of the training pairs so that both transfer items differed from the training pairs by one feature but it was a feature that had already been seen in the context of a similar training item. For example, the GT item associated with the pair of builders in Figure 8 has the underlying structure 011000. This item differs from both training stimuli by one relevant feature, but notice that the feature that differs is one that occurred on the other member of the training pair. The GT item has a different number of legs relative to the first member of the training pair but it has the same number as the second member of the pair. In addition, the GT item has different body markings relative to the second member of the training pair, but these markings had occurred on the first member. Thus, this "2 neighbor/novelty reduced" GT item does not have a feature that was not already seen in the context of a similar animal; as such the novelty of the animal is reduced. The same method was used

**Rule :** Builders have any combination of two of the following three features, a short tail, spots, and six legs, otherwise they are diggers.




Figure 8. Examples of some of the stimuli used in the novelty reduced experiment (Experiment 4) and the corresponding rule-consistent response for each stimulus. The 1 Neighbor/Novelty present transfer stimuli differ from their corresponding training item by one feature that is also novel in the context of the animal. The 2 Neighbor/Novelty reduced transfer stimuli also differs from their corresponding similar training pairs by one feature but that feature is not novel in the context of the animal because it occurred on one of the similar training animals.

to create the 2 neighbor/novelty reduced BT item. The BT item associated with the pair of builders in Figure 8 has the structure 000000. This differs from both the training stimuli by one feature but it is not a feature that had not already been seen on one of the training animals. For example, the BT item has different body markings than the first member of the pair but has the same body markings as the second member of the pair. In addition, the BT item has a different number of legs relative to the second member of the pair but the number of legs is identical to that used on the first member of the training pair. Thus, this 2 neighbor/novelty-reduced BT item also does not have a feature that was not seen in the context of a similar animal.

Creating stimuli with a novel cue simply involved not presenting one of the members of the training pair. The GT and BT items still differed from their respective training item by one relevant feature but this feature had not been seen before in the context of the animal. For example, consider the GT builder located in the top panel of Figure 8. This item differs from its similar training item by one feature but note that the feature that is different, the tail length, had not been seen before on a similar animal.

The BT item is also one feature away from the training item and differs in the number of legs that it has. Note, however, that these 2 legs were not seen on a similar item and thus, the "1 neighbor/novelty present" GT and BT stimuli both contain a novelty cue.

The procedure used in this experiment was also similar to Experiment 2 and was designed to promote the acquisition of both rule-based and retrieval-based classification. To promote rule-based classification, participants trained with a perfectly predictive classification rule and were given feedback if they made any responses that were inconsistent with this rule. To promote retrieval-based classification, each training item consisted of individuated features, there was a high degree of similarity between training and transfer items, and to help reinstate the study context there was a large number of training items repeated during the transfer phase.



If participants are trying to discriminate old items from new items before classifying the stimuli then they should find it difficult to decide if the 2 neighbor/novelty reduced stimuli are old or new. If so, and if participants do falsely recognize some novel items as old, then the difference between BT and GT conditions should be larger when novelty is reduced relative to when it is present.

### *Methods*

#### *Participants*

Twenty four students taking an introductory psychology course participated for course credit.

#### *Materials*

As shown in Figure 8, the stimuli were line drawings of imaginary animals similar to those used in the previous experiments and varied on six binary dimensions: tail (short or long), body markings (stripes or spots), number of legs (two or six), body shape (rounded or angular), neck length (short or long), and leg length (short or long). Appendix B contains the informational structure of these stimuli. Eight quartets of perceptually similar stimuli were created and consisted of 2 training stimuli that differed from each other on 2 of the 3 relevant features, a GT item that differed from the training pairs by one feature and a BT item that also differed from the training pairs by one relevant feature. The GT and BT items were created by recombining the relevant features of the training pairs so that although the GT item and the BT item both contained a novel feature, neither item had a feature that had not already been seen on a similar training item. The eight quartets were then divided into two groups of four quartets, consisting of an equal number of builder and digger animals. For one of these groups, one of the pairs of training stimuli was removed leaving four triplets of stimuli consisting of one training item and a GT and BT item that differed from the similar training stimulus by one feature. Thus, these 1 neighbor/novelty-present transfer stimuli all contained a feature

that had not been seen during training in the context of a very similar item.

The stimuli were categorized as builders or diggers according to a two out of three feature additive rule that used tail length, body markings and number of legs as the relevant features. An animal was defined as a builder if it had builder values on at least two of the three relevant dimensions, all others being diggers. Three different rules were used to categorize the animals. The features associated with the builder category for each of these three rules were: any two of short tail, spots, or six legs for the first rule, any two of long tail, spots, or two legs for the second rule and any two of short tail, stripes, two legs for the third rule.

Stimuli alternated between the 1 neighbor/novelty present and 2 neighbor/novelty reduced conditions and, combined with the three classification rules, required six different counterbalancing orders.

### ***Procedure***

The procedure was identical to that used in Experiment 2 with the exception that the training list consisted of 12 training stimuli: four 1 Neighbor items, and four pairs of 2 neighbor items. This list of 12 stimuli was repeated 20 times, with the order of items within a list presented in a different random order for each participant.


The procedure for the transfer phase was also identical to Experiment 2 with the exception that participants were presented with a list of transfer items consisting of the twelve training items, each repeated five times, plus the eight good transfer items and eight bad transfer items. The 76 trials were presented in a different random order for each participant, and then the entire transfer list was repeated in a different random order. Each item stayed on the screen until it was classified.

### ***Results***

#### ***Training***

Figure 9 contains the mean response times and mean proportion of rule-inconsistent

### 1 neighbor/novelty present



### 2 neighbor/novelty reduced

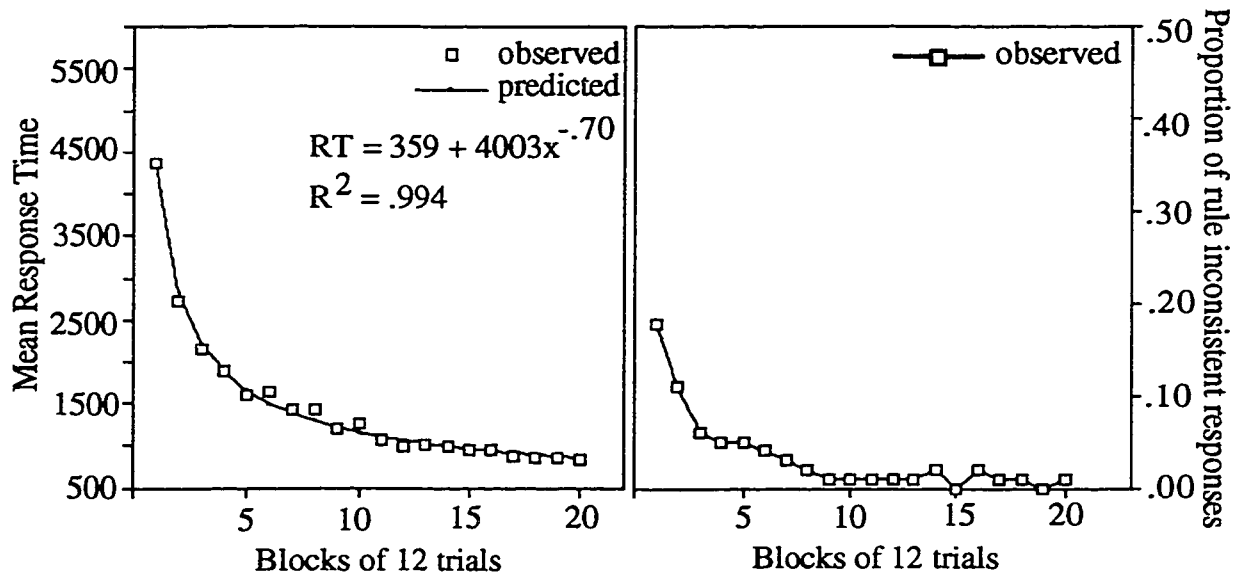


Figure 9. Mean response times with the corresponding fitted power functions and mean error rates for the novelty present and novelty reduced conditions plotted as a function of blocks of training trials (Experiment 4).

responses for each block of twelve training trials. It appears that participants responded faster and made fewer rule-inconsistent responses as the number of trials increased. This finding suggests that participants were learning the categories. In addition, there does not appear to be much difference in response time between the two novelty conditions. To confirm these impressions, the mean response times for each block of twelve training trials and the proportion of rule-inconsistent responses were analyzed. Two  $2 \times 20$  ANOVAs were conducted with novelty (present, reduced) and training block (training blocks 1 to 20) treated as a repeated measures variable. These analyses confirmed that participants responded faster with practice, as shown by the significant main effect of training blocks,  $F(19,437)=53.30$ ,  $MSe = 766378$ ,  $p < .001$ . There was no difference between the novelty conditions,  $F(1,23)=1.13$ ,  $MSe = 364306$ ,  $p > .10$  but there was an interaction between trials and the novelty condition,  $F(19,437)=1.78$ ,  $MSe = 211623$ ,  $p < .05$ . The interaction occurred because there was a bigger difference between 1 Neighbor/novelty present items and 2 Neighbor/Novelty reduced items in the early training blocks than there was during the latter training blocks. An analysis of the proportion of rule-inconsistent responses confirmed that as participants got faster at the task they also made fewer errors. There was a significant main effect of training blocks,  $F(19,437) = 12.42$ ,  $MSe = .008$ ,  $p < .001$ , but no other difference between the novelty conditions,  $F < 1$ , and no interaction between novelty and trials.

To confirm that the faster response times reflected a transition to a faster classification procedure, mean training response times for each novelty condition were plotted as a function of each block of 12 trials and power functions were fit to these plots using a least squares criterion to determine the best fitting function. As shown in Figure 9, the mean responses times for the 1 neighbor/novelty-present condition were fit quite well by a power function,  $R^2 = .997$ , as were the response times for the 2 neighbor/novelty reduced condition,  $R^2 = .994$  suggesting that participants were

classifying the stimuli quickly and effortlessly.

### ***Transfer***

Analyses were performed on both the proportion of rule-inconsistent responses and on participants' median response time for correct (rule consistent) responses made during the transfer phase. Two 3 x 2 analyses of variance were performed with two repeated measures variables: transfer type (Old, GT, BT) and novelty (1 neighbor/novelty present, 2 neighbor/novelty reduced). A summary of the results appears in Table 6.

***Rule-inconsistent responses.*** Participants made more rule inconsistent responses when classifying the 2 neighbor/novelty reduced stimuli ( $M=.15$ ) than they did when classifying the 1 neighbor/novelty present stimuli ( $M=.09$ ),  $F(1,23) = 12.09$ ,  $MSe = .01$ ,  $p < .01$ . There was also a significant main effect of transfer type,  $F(2,46) = 43.34$ ,  $MSe = .04$ ,  $p < .001$ , and the interaction between transfer type and novelty was also significant,  $F(2,46) = 13.32$ ,  $MSe = .01$ ,  $p < .001$ . To explore this interaction further, two sets of planned t-tests were conducted. However, to prevent an increase in the familywise error rate a Modified Bonferroni Test for Planned Comparisons was used to set the significance level. The level of significance for these comparisons was set at  $p = .02$ . When the novelty cue was present, participants made more rule-inconsistent responses to BT items than GT items,  $t(23) = 4.39$ ,  $p < .001$ , but there was no difference between Old items and GT items,  $t(23) = 1.48$ ,  $p > .10$ . A similar result was observed when novelty was reduced. Participants made more rule-inconsistent responses to BT than GT items,  $t(23) = 7.09$ ,  $p < .001$  and there was no difference between Old items and GT items,  $t(23) < 1$ . The main source of the interaction is due to a difference in performance on BT items. A subsequent paired t-test comparing the proportion of rule-inconsistent responses on BT items shows that participants were more likely to make a rule-inconsistent response to BT items when novelty was reduced than when it was present,  $t(23) = 3.67$ ,  $p < .02$ .

Table 6  
Mean responses to transfer stimuli when a novelty cue is present versus when a novelty cue is reduced (Experiment 4).

Measure	1 Neighbor Novelty present		2 Neighbor Novelty reduced	
	Mean	SD	Mean	SD
Proportion of rule inconsistent responses				
Old	0.02	(0.04)	0.01	(0.02)
GT	0.01	(0.04)	0.01	(0.04)
BT	0.23	(0.26)	0.42	(0.28)
Median response time (rule consistent)				
Old	1024	(295)	1040	(342)
GT	1399	(441)	1048	(403)
BT	1873	(521)	1879	(555)

Note. GT = good transfer items  
BT = bad transfer items

**Response times.** The data from two participants were dropped from the analysis of correct (rule consistent) response times, one because he/she had no correct responses on any BT items, the other because he/she had no correct responses on 2 Neighbor/novelty-reduced BT items. For the remaining 22 participants, there was a difference in the median correct (rule consistent) response times between the two novelty conditions. Participants took longer to respond to 1 neighbor/novelty present stimuli ( $M=1432$ ) than they did to 2 neighbor/novelty reduced stimuli ( $M=1322$ ),  $F(1,21) = 9.53$ ,  $MSe = 41,580$ ,  $p < .01$ . The median response times also differed significantly across transfer type,  $F(2,42) = 63.90$ ,  $MSe = 134,826$ ,  $p < .001$ , and the interaction between transfer type and novelty was significant,  $F(2,42) = 13.82$ ,  $MSe = 34,763$ ,  $p < .001$ . To explore the interaction, two sets of planned paired t-tests were conducted. A Modified Bonferroni Test for Planned Comparisons was used to set the level of significance for these comparisons at  $p = .03$ . Participants took longer to

respond to BT items than GT items when the novelty cue was present,  $t(21) = 7.12$ ,  $p < .001$ , and when it was reduced,  $t(21) = 7.65$ ,  $p < .001$ . In addition, when the novelty cue was present, GT items were slower than Old items,  $t(21) = 5.48$ ,  $p < .001$ , but when novelty was reduced Old items and GT items did not differ,  $t(21) < 1$ .

This pattern of results, in which Old items were categorized faster than GT items when there was a novel cue present, and in which there was little difference between Old and GT items when novelty was reduced, is consistent with the hypothesis that participants performed an old/new discrimination. When there was a novel feature on the transfer item, participants slowed down to classify it, but when there was not a novel feature on the transfer item, participants responded relatively fast. This latter results could suggest that participants had not noticed the new items and were treating them as if they were old. Further converging evidence supporting the idea that some novel items are falsely recognized as old can be found by comparing the response times for rule-inconsistent responses when classifying BT items to correct response times when classifying Old items. If rule-inconsistent responses on BT items are due to participants treating the items as Old then there should be little difference between correct responses on Old items and rule-inconsistent responses on BT items. To perform this analysis, only the data from participants who made rule-inconsistent responses on both the 2 neighbor/novelty reduced and 1 neighbor/novelty present BT items were analyzed. To meet this constraint, data from seven participants were dropped, two for not making any rule-inconsistent responses items at all, four for not making any rule-inconsistent responses when the novelty cue was present, and one for not making any rule-inconsistent responses when novelty was reduced. Paired t-tests revealed that when the novelty cue was present, the remaining 17 participants took longer to classify BT items ( $M=1439$  ms) than Old items ( $M=943$  ms),  $t(16) = 2.33$ ,  $p < .05$ , but when novelty was reduced there was little difference in response time for rule-inconsistent responses on

BT items ( $M=896$  ms) and correct responses on Old items ( $M=919$  ms),  $t(16) < 1$ .

### *Discussion*

This experiment was designed to provide evidence consistent with the hypothesis that participants perform an old/new discrimination before classifying the transfer stimuli and that some of the evidence for similarity-based classification is actually due to participants failing to discriminate an item as new and treating it as if it was old. In this experiment, the old/new discrimination was made more difficult by reducing the novelty of some of the transfer stimuli. If participants are doing an old/new discrimination before classifying the stimuli, then this modification should make their task more difficult. In addition, if novel items are occasionally falsely recognized as old, then a more difficult old/new discrimination may also lead to a corresponding increase in the number of rule-inconsistent responses when classifying BT items. As predicted, the rule-inconsistent responses when classifying BT items in the 2 neighbor/novelty reduced group were nearly double those in the 1 neighbor/novelty present group.

The response time data was also consistent with the idea that some novel items were falsely recognized, especially when the old/new discrimination was difficult. Classification performance was roughly equivalent between the 1 Neighbor Old items and the 2 Neighbor Old items on both the training and transfer tasks. Differences between the novelty conditions occurred, however, when participants classified the novel transfer items. When the novelty cues were reduced, the 2 Neighbor/novelty reduced GT response times were within the same range as the 2 Neighbor Old items. On the other hand, when the novelty cues were present, GT times were substantially slower than 1 Neighbor Old items. Additionally, correct responses to Old items were equivalent to rule-inconsistent responses on BT items when novelty was reduced but correct responses on old items were faster than rule-inconsistent responses on BT item when a novel cue was present. These findings would be expected if some of the 2 neighbor/novelty reduced



transfer items were being falsely recognized as Old items and a more time consuming rule application procedure was used to classify items believed to be new, especially those stimuli with a novel feature.

***Experiment 5: Unconfounding novel cues from number of neighbors***

A potential problem with the interpretation of the data from the previous experiment is that the presence of a novel feature was confounded with the number of similar neighbors. Whenever a transfer item was similar to one training instance, it also had a novel feature. Whenever a transfer item was similar to two training instances, the novelty of that item was reduced. Confounding the number of neighbors and the degree of novelty in this fashion makes it difficult to determine if the increase in rule-inconsistent responses for BT items in Experiment 4 occurred because of the reduction in novelty cues or because of the number of similar neighbors. In hindsight, the same confound was also present in Experiment 2. In that experiment, two pairs of similar stimuli were used as the training stimuli and the transfer items did not contain any features that had not already been seen in the context of a similar training item. Experiment 2 had the additional confound that there were multiple variations of these two training stimuli, which may also have influenced how participants classified the stimuli. In summary, because of the design of Experiment 4 and of Experiment 2, it is difficult to determine what role similarity to multiple neighbors and degree of novelty play in influencing participants' performance.

The goal of Experiment 5 was to unconfound the number of similar training instances and the presence of a novel feature. The stimuli and the design for this experiment were identical to those used in Experiment 4 except that all transfer stimuli, both one and two neighbor, now contained a novel feature. As shown in Figure 10 and in Appendix B, the underlying structure of the stimuli was identical to that used in Experiment 4. One set of transfer items was similar to a single training example and differed from that training

**Rule :** Builders have any combination of two of the following three features, a short tail, spots, and six legs, otherwise they are diggers.




Figure 10. Examples of some of the stimuli used in the experiment in which the number of similar neighbors and the degree of novelty is unconfounded from one another (Experiment 5), and the corresponding rule-consistent response for each stimulus. The 1 Neighbor/Novelty present transfer stimuli differ from their corresponding training item by one feature, which is also novel in the context of the animal. The 2 Neighbor/Novelty present transfer stimuli also differ from their corresponding similar training pairs by one feature. For example, although a training item and a transfer item may both have spots, the pattern of spots is different relative to each other. Thus the 2 Neighbor/Novelty reduced transfer stimuli are similar to two training items but also contain a novel feature.

item by one relevant feature that was also novel in the context of the animal. The other set of transfer stimuli were similar to two training examples and were created by recombining the features of the training stimuli. The difference between these 2 Neighbor transfer stimuli and those used in Experiment 4 is that the perceptual manifestation of the feature that was different was changed so that it was novel in the context of the experiment. For example, consider the first pair of training items in the bottom panel of Figure 10. The structure of the builder on the left is 010000 and the structure of the builder on the right is 001000. The GT item, associated with this pair of builders, has the structure 011000 and therefore has 5 of the 6 features in common with each training item. Additionally the GT item was created by recombining the features of the training stimuli so that , at the informational level, there is no one feature that had not occurred on a similar training animal. Notice, however, that there is a different pattern of spots on the GT item relative to the training items. Although both the GT item and a training item have spots, the style of spots are different from one another. Thus, this "2 neighbor/novelty present" GT item is very similar to 2 prior training stimuli but also has a novel feature. The same method was used to create the 2 neighbor/novelty present BT item. The BT item associated with the pair of builders in Figure 10 has the structure 000000 and differs from both the training stimuli by one feature. It was also created by recombining the features of the training stimuli to make a novel item that is very similar to the training stimuli. Notice, however, that the BT item has a different set of 2 legs relative to the other training item with 2 legs. Thus, this "2 neighbor/novelty present" BT item is also similar to 2 prior training stimuli but contains a novel feature.

Creating stimuli in this fashion serves to unconfound the effects of similarity to a number of neighbors from the degree of novelty present within the transfer stimuli. If similarity produced the increased number of rule-inconsistent responses found when classifying BT items in Experiment 4, then there should be a comparable level found

here. If, on the other hand, participants use the rule on items they perceive to be novel, then the presence of a novel feature on all of the transfer stimuli should alert participants to the existence of that novelty and there should be a corresponding decrease in the number of rule-inconsistent responses when classifying 2 Neighbor BT stimuli.

### *Method*

#### *Participants*

Twenty five students taking an introductory psychology course participated for course credit.

#### *Materials*

The stimuli used for this experiment were identical to those used in Experiment 4 with the exception that the 2 neighbor transfer stimuli now contained a new relevant feature. As shown in Figure 10, this was done by changing the perceptual manifestation of one of the relevant features so that all the 2 Neighbor transfer stimuli contained a feature that was novel in the context of the animal.

#### *Procedure*

The procedure for this experiment was identical to Experiment 4.

### *Results*

One participant's data was replaced because he/she made errors on more than 25% of the training trials.

#### *Training*

Figure 11 contains the mean response times and proportion of rule-inconsistent responses for each block of twelve training trials. As expected, participants responded faster and made fewer errors as the number of trials increased. In addition, there does not appear to be much difference in performance between the two novelty conditions. To confirm these impressions, the mean response times, and the proportion of rule-inconsistent responses for each block of twelve training trials were analyzed. Two

---


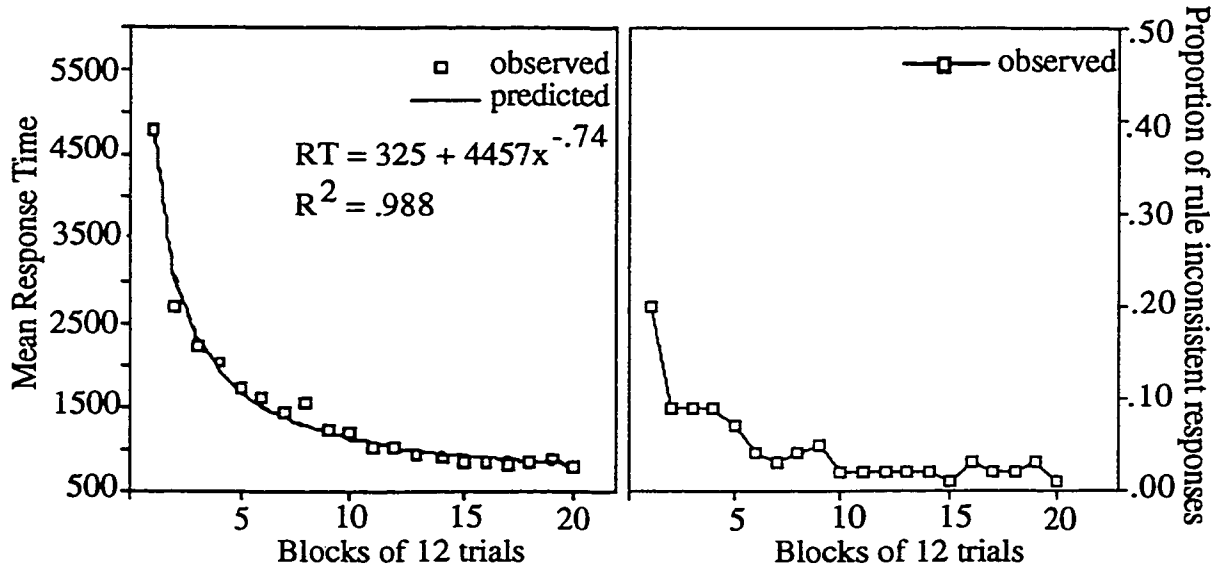
**1 neighbor/novelty present**

**2 neighbor/novelty present**


Figure 11. Mean response times and fitted power functions plotted as a function of blocks of training trials for training items in which the number of similar neighbors is unconfounded from the novelty alert (Experiment 5).

---

2 x 20 ANOVAs were conducted with number of neighbors (one, two) and training block (training blocks 1 to 20) treated as a repeated measures variable. These analyses confirmed that participants responded faster with practice, as shown by the significant main effect of training blocks,  $F(19,437)=53.30$ ,  $MSe = 766360$ ,  $p < .001$ . There was no difference between the novelty conditions,  $F(1,23)=1.13$ ,  $MSe = 364176$ ,  $p > .10$ , but there was an interaction between trials and the novelty condition  $F(19,437)=1.78$ ,  $MSe = 211644$ ,  $p < .05$ . The interaction occurred because there was a bigger difference between 1 Neighbor/Novelty present and 2 Neighbor/Novelty present in the early rounds of training than there was toward the latter rounds. An analysis of the proportion of rule-inconsistent responses revealed a significant main effect of training blocks,  $F(19,437) = 10.40$ ,  $MSe = .009$ ,  $p < .001$  and occurred because participants made fewer errors with increased practice.

To determine if the performance met the criteria of becoming fast and less effortful, the mean training response times for each novelty condition were plotted as a function of each block of twelve training trials. Power functions were then fit to the plots using a least squares criterion to determine the best fitting function. As shown in Figure 11, the best fitting power function for the 1 Neighbor/Novelty present condition was of the form  $RT = 335 + 4660x^{-.79}$  and provided a good fit of the training times,  $R^2 = .992$ . The best fitting power function for the 2 Neighbor/Novelty present stimuli was of the form  $RT = 325 + 4457x^{-.74}$  and also provided a very good fit of the training times,  $R^2 = .988$ .

### ***Transfer***

Analyses were performed on both the proportion of rule-inconsistent responses and the participants' median response time for correct (rule consistent) classifications made during the transfer phase of the experiment. These analyses involved two 3 x 2 ANOVAs with two repeated measures variables: transfer type (Old, GT, BT) and neighbor

(1 neighbor/novelty present, 2 neighbor/novelty present). A summary of the results appears in Table 7.

**Proportion of rule-inconsistent responses.** The proportion of rule-inconsistent responses differed significantly across transfer types,  $F(2,46) = 31.05$ ,  $MSe = .01$ ,  $p < .001$ . However, there was no effect of number of neighbors  $F(1,23) < 1$ , nor was the interaction between transfer type and number of neighbors significant,  $F(2,46) < 1$ . Collapsing across the novelty condition, two paired t-tests showed that participants made more errors on BT items ( $M=0.16$ ) than GT items ( $M=0.02$ ),  $t(23) = 5.45$ ,  $p < .001$ , and there was no difference in errors between old items ( $M=0.01$ ) and good transfer items,  $t(23) = 1.68$ ,  $p > .10$ .

**Response time.** Although participants took longer to classify 2 neighbor/novelty present transfer items ( $M=1442$  ms) than 1 neighbor/novelty present transfer items ( $M=1381$  ms), the difference only approached significance,  $F(1,23) = 3.41$ ,  $MSe = 40,357$ ,  $p < .08$ . Response times differed significantly across transfer types,  $F(2,46) = 82.01$ ,  $MSe = 126,881$ ,  $p < .001$ , and the interaction between transfer type and number of neighbors was not significant,  $F(2,46) < 1$ . Paired t-tests conducted after collapsing across number of neighbors showed that participants took longer to classify BT items than GT items ( $M=1914$  ms and  $M=1327$  ms respectively),  $t(23) = 7.09$ ,  $p < .001$ , and they took longer to classify GT items than Old items ( $M=1327$  ms and  $M=994$  ms respectively),  $t(23) = 5.51$ ,  $p < .001$ .

This pattern of results, in which Old items were classified faster than GT items, is consistent with the hypothesis that the presence of a novel feature on GT items alerts participants to the novelty of the item and produces a slowing in response time. Further converging evidence favoring the idea that some of the novel items are falsely recognized as old can be found by comparing the response times for rule-inconsistent responses when classifying BT items to correct response times when classifying Old items. If

Table 7  
Mean responses to transfer stimuli when a novelty cue and number of neighbors is unconfounded from each other Experiment 5).

Measure	1 Neighbor Novelty present		2 Neighbor Novelty present	
	Mean	SD	Mean	SD
Proportion of rule inconsistent responses				
Old	0.00	(0.01)	0.01	(0.01)
GT	0.01	(0.04)	0.02	(0.06)
BT	0.16	(0.13)	0.16	(0.18)
Median response time (rule consistent)				
Old	971	(224)	1014	(242)
GT	1270	(450)	1384	(512)
BT	1901	(421)	1926	(452)

Note. GT = good transfer items  
BT = bad transfer items

rule-inconsistent responses on BT items are due to participants treating the items as Old then there should be little difference between these two measures. The data from participants were removed from this analysis if they did not make any rule-inconsistent responses on BT items. Twelve participants fit this criteria and were removed: two for not making any rule-inconsistent responses when classifying 1 neighbor/novelty present items, six for not making any rule-inconsistent responses when classifying the 2 neighbor/novelty present, and four for not making any rule-inconsistent responses on any BT items. Two paired t-tests were conducted on the data from the remaining 12 participants. Rule-inconsistent responses for BT items were as fast as correct responses for Old items. More specifically, the time required to make a rule-inconsistent response for 1 neighbor/novelty present BT items ( $M = 889$  ms) was as fast as the time required to



correctly classify 1 neighbor/novelty present Old items ( $M=892$  ms),  $t(11) < 1$ . Additionally, although rule-inconsistent responses on 2 neighbor/novelty present BT items ( $M=1325$  ms) appeared to take longer than correct responses to old items ( $M=890$  ms), the difference only approached significance  $t(11) = 2.01$ ,  $p < .08$ .

### *Discussion*

The purpose of this experiment was to investigate the separate contribution of multiple similar training neighbors and the presence of novelty cues in promoting a retrieval-based classification procedure. Experiment 4 showed that when the two were confounded there was an increased reliance on retrieval when classifying the BT stimuli. It is of theoretical interest to determine if this increased reliance occurred because the transfer items were similar to multiple neighbors, or if it occurred because participants falsely recognized more stimuli when the novelty cues were reduced.

In Experiment 5, novelty was held constant by placing a new feature on all transfer stimuli while the number of similar neighbors varied. Participants continued to make rule-inconsistent responses on BT items suggesting that there was some reliance on retrieval. However, there was no difference between one neighbor and two neighbor conditions, which suggests that the presence of a novel feature on a stimulus was an important factor in determining whether a rule or retrieval was used. In addition, if the means from Table 6 and 7 are compared, it appears that the difference in rule-inconsistent responses between GT and BT items actually gets smaller in Experiment 5. These two findings, no difference in the proportion of rule inconsistent responses between novelty conditions and a smaller difference in the number of rule-inconsistent responses between GT and BT items, are particularly interesting in light of the fact that the only difference between Experiment 4 and Experiment 5 was that the 2 Neighbor transfer stimuli in Experiment 4 did not have a novel feature, but did have one in Experiment 5. A subsequent cross experiment analysis of the rule-inconsistent responses for BT items was

conducted in order to confirm these observations. A 2 x 2 ANOVA was conducted on the proportion of rule-inconsistent responses for BT items. Neighbors (one, two) was a repeated measures variable and experiment (Experiment 4, Experiment 5) was a between subject variable. There was a significant interaction between neighbors and experiment  $F(1,46) = 9.41$ ,  $MSe = .02$ ,  $p < .01$  with the difference between 1 Neighbor and 2 Neighbor BT items greater in Experiment 4 than it was in Experiment 5.

In summary, Experiment 4 showed that the proportion of rule-inconsistent responses on BT items increases when the degree of novelty is reduced and Experiment 5 showed that this proportion decreases when a novel cue is present. This pattern of results is consistent with the hypothesis that participants are performing an old/new discrimination before classifying the stimuli. The more difficult this old/new discrimination is, the more likely participants are to falsely recognize an item as old. The easier the old/new discrimination, the better participants are at recognizing an item as new. The pattern of data observed here is also consistent with the hypothesis that at least some of the evidence for similarity-based retrieval is actually due to participants falsely recognizing new items as old.

## CHAPTER 7

### *Experiment 6: Making Old/New Explicit*

The previous two experiments provide converging evidence for the idea that participants perform an old/new discrimination before classifying the stimuli and that rule-inconsistent responses on BT items may reflect false recognitions rather than similarity-based retrieval. The evidence for false recognition is still circumstantial, however, and consists mainly of an increase in fast similarity-based responses when novelty is reduced relative to when a novel cue is present. The goal of this experiment was to provide a more explicit measure of how often novel transfer stimuli are falsely recognized.

The training for this experiment was identical to that used in the prior experiments but the transfer task was modified. Rather than classifying the transfer stimuli as builders or diggers, participants were asked to classify the transfer stimuli as old builders, old diggers, new builders, or new diggers. In addition, to determine if a false recognition effect is limited only to novel items that are similar-to-old, an extra condition was added to the transfer phase. This extra condition consisted of completely novel transfer stimuli that did not resemble the training items. Figure 12 contains an example of two quartets of stimuli, each of which consists of an Old item, a 1 Neighbor/Novelty present GT item, a 1 Neighbor/Novelty present BT item and a completely novel transfer item. Performance on the completely novel items should be determined almost exclusively by rule application and should therefore provide a measure of performance that is neither supported nor hindered by specific retrieval. For the other transfer items, it was expected that if participants falsely recognize some of the novel transfer stimuli as old then there should be a high proportion of rule-inconsistent responses for BT items that participants

---

**Rule:** Builders have any combination of two of the following three features; short tail, long neck, long legs, otherwise they are diggers.




Figure 12. Examples of some of the stimuli used in the experiment in which the old/new discrimination was made explicit (Experiment 6) and the corresponding rule consistent response for each stimulus. Note that these are all examples of 1 Neighbor/Novelty cue present stimuli. In addition, while the completely novel item has the identical underlying structure as the training item, the manifestation of each feature was changed producing a brand new animal.

---

think are old and a low proportion of rule-inconsistent responses for BT items that participants think are new.

### *Method*

#### *Participants*

Thirty students taking an introductory cognition course participated for course credit.

#### *Materials*

As shown in Figure 12, the stimuli for this experiment were similar to those used in

the previous experiments and consisted of imaginary animals that varied on six binary dimensions: tail (short or long), neck length (short or long), leg length (short or long), body shape (rounded or angular), body markings (stripes or spots), and number of legs (2 or 6). Eight quartets of 2 neighbor/novelty reduced stimuli were created in the same fashion as the 2 neighbor/novelty reduced stimuli in Experiment 4. Appendix B contains the informational structure for these stimuli. In addition, eight triplets of 1 Neighbor/Novelty present stimuli were created in the same fashion as the 1 Neighbor/Novelty present stimuli in Experiment 4. Appendix C contains the informational structure for these stimuli. Finally, a third type of transfer item was added to the transfer phase. As shown in Figure 12, these novel stimuli had the same underlying structure as their respective training stimuli but the perceptual manifestation of each feature was changed. For example, consider the set of 1 Neighbor/Novelty present stimuli presented at the top of Figure 12. Both the training item and the completely novel transfer item have the same underlying structure, 111010, but the manifestation of each feature was different on the novel item producing a stimulus that is completely novel in the context of the experiment. Creating novel stimuli in this fashion, produced eight pairs of completely novel transfer items in the 2 neighbor/novelty reduced condition and eight completely novel transfer items in the 1 Neighbor/novelty present condition.

The stimuli were categorized as builders or diggers according to a two out of three feature additive rule that used tail length, neck length, and leg length as the relevant features. An animal was defined as a builder if it had builder values on at least two of the three relevant dimensions, all others being diggers. Three different rules were used to categorize the animals. The features associated with the builder category for each of these three rules were: any two of short tail, long neck, long legs; any two of long tail, long neck, short legs; and any two of short tail, short neck, short legs. Each rule required a

separate counterbalancing order.

### ***Procedure***

The procedure for the training phase was identical to Experiment 2 with the following exception. One group of participants studied 20 blocks of eight, 1 neighbor/novelty present items and were tested with a transfer list that consisted of the eight training animals, eight GT items, eight BT items and eight Novel transfer items. A second group of participants studied 20 blocks of sixteen, 2 neighbor/novelty reduced training items and were tested with a transfer list that consisted of 16 training items, eight GT items, eight BT items and 16 Novel transfer items.

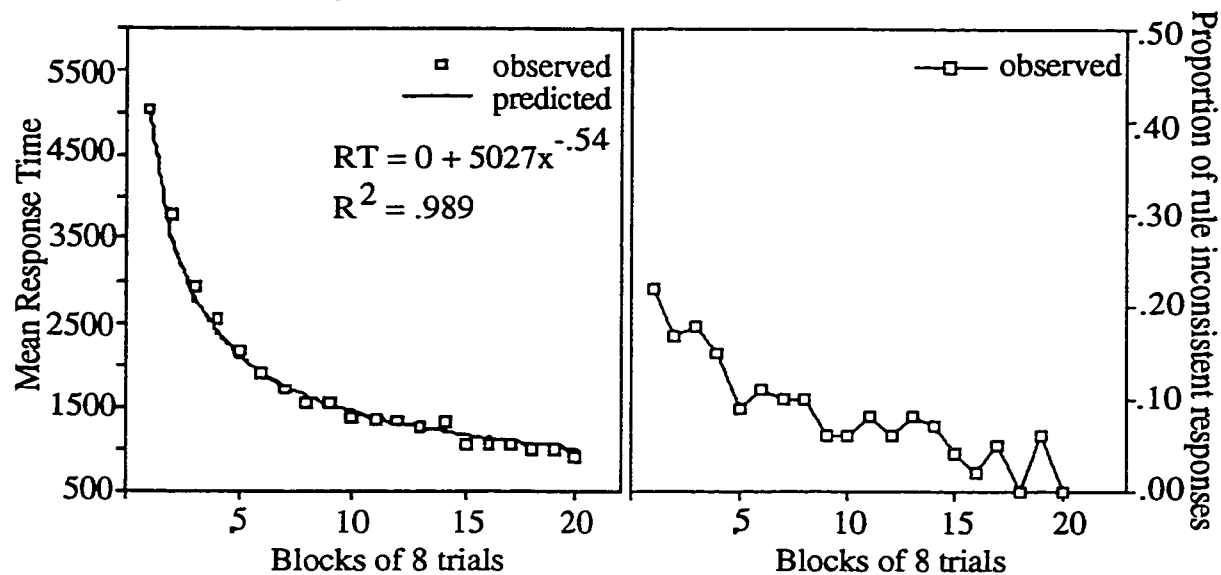
### ***Results***

#### ***Training***

Figure 13 contains the mean response times and proportion of rule-inconsistent responses plotted as a function of each block of training trials. To verify that participants responded faster and made fewer errors with practice, the mean response times and the proportion of rule-inconsistent responses for each block of training trials was analyzed. Two 2 x 20 ANOVAs were conducted with training blocks (1 to 20) treated as a repeated measures variable and novelty (present, reduced) treated as a between subject factor. These analyses confirmed that participants responded faster with practice as shown by the significant main effect of trials,  $F(19,532) = 58.98$ ,  $MSe = 507977$ ,  $p < .001$ . No other effects approached significance. Participants also made fewer rule-inconsistent responses with practice as shown by the significant main effect of trials,  $F(19,532) = 14.02$ ,  $MSe = .007$ ,  $p < .001$ . No other effects approached significance.

To verify that this improvement with practice also reflected a transition to a faster, less effortful classification procedure, the training response times for each novelty condition were plotted as a function of each block of training trials and a power function was fit to this data. As shown in Figure 13, the best fitting power function for the

### 1 neighbor/novelty present



### 2 neighbor/novelty reduced

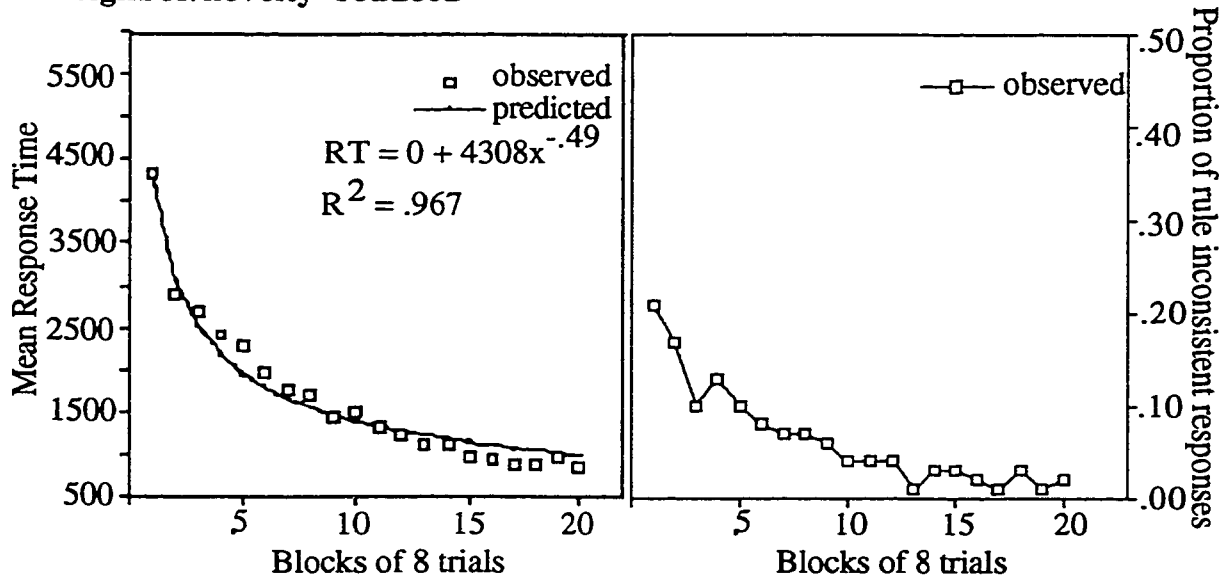


Figure 13. Mean response times and fitted power functions plotted as a function of blocks of training trials for training items in which the old/new discrimination, at transfer, was made explicit (Experiment 6).

1 Neighbor/Novelty present condition was of the form  $RT = 0 + 5027x^{-.54}$  and provided a good fit of the training times,  $R^2 = .989$ . The best fitting power function for the 2 Neighbor/Novelty reduced condition was of the form  $RT = 0 + 4038x^{-.49}$  and also provided a good fit of the training times  $R^2 = .967$ .

### *Transfer*

Table 8 contains the proportion of items that participants called old. As shown in the table, participants do quite well at correctly discriminating the Old items as "old" and the completely novel items as "new" but appear to have problems with the GT and BT stimuli. When there is a novel feature on the transfer animals, roughly a quarter of the GT and BT items were falsely recognized as being old animals. When the degree of novelty was reduced, the proportion of old responses increased with participants falsely recognizing 58% of the BT items and 70% of the GT items.

This pattern of means, in which novel but similar-to-old transfer stimuli are called old, especially when the novelty cues were reduced, supports the hypothesis that participants are performing an old/new discrimination and when that discrimination is made difficult by reducing the cues that signal novelty there is an increase in false recognitions. Of more interest, however, is the likelihood of making a rule-inconsistent response given that an item was judged as "old" or "new". A rule-inconsistent response for BT items judged as "old" would reflect the false recognition of that item while a rule-inconsistent response for a BT item judged as "new" would reflect the use of similarity-based retrieval. Unfortunately these comparisons require the analysis of two dependent measures simultaneously, leaving the analysis open to potential biases from subject selection effects (Hintzman, 1980). To compensate for these effects, it was decided to analyze only data from participants whose performance was roughly equivalent on the old/new discrimination. For both novelty conditions, a false alarm score



Table 8  
 Proportion of old responses to transfer stimuli when the old/new discrimination was made explicit (Experiment 6).

Proportion of Old responses	1 Neighbor <u>Novelty present</u>		2 Neighbor <u>Novelty present</u>	
	Mean	SD	Mean	SD
Old	0.81	(0.18)	0.85	(0.18)
GT	0.27	(0.20)	0.70	(0.26)
BT	0.28	(0.21)	0.58	(0.30)
Novel	0.02	(0.04)	0.05	(0.13)

Note. GT = good transfer items  
 BT = bad transfer items

was calculated by collapsing over the proportion of old responses on GT and BT items. Participants who had similar patterns of hits versus false alarms were then grouped together. In the 1 Neighbor/Novelty present condition, the data from two participants were dropped from the analysis; one participant had a strong bias to respond "new" and the other participant did not respond "old" on any BT items. The performance of the remaining 13 participants was roughly equivalent with a hit rate over 50% for each participant and a consistent difference between hits and false alarms. In the 2 Neighbor/Novelty reduced condition, the data from one participant was dropped from the analysis because he/she had more false alarms than hits. Of the remaining 14 participants, there appeared to be two groups. One group of four participants, tended to call all the Old, GT and BT items "old", while the other group of 10, appeared to be able to discriminate old from new items to some degree. In subsequent analyses these groups are referred to as the "high false alarm group" and the "low false alarm group" respectively.

An analysis of the probabilities of making a rule-inconsistent response, conditional on

whether an item was judged as "old" or "new", was conducted. For each novelty condition, paired t-tests were used to analyze the probability of making a rule-inconsistent response for each type of transfer item. A summary of both the overall mean proportion of rule-inconsistent responses and also the probability of making a rule-inconsistent response conditionalized on whether the item was judged as "old" or "new" is presented in Table 9.

*1 neighbor/novelty present* . There are a couple of patterns of note displayed in Table 9. First, participants have a tendency to produce fewer rule-inconsistent responses when classifying 1 Neighbor/Novelty present BT items ( $M=.30$ ) than the 2 Neighbor/Novelty reduced BT items ( $M=.35$  for the low false alarm group and  $M=.53$  for the high false alarm group). This pattern of means is similar to that found in Experiment 4, which also used 1 Neighbor/Novelty present and 2 Neighbor/Novelty reduced items. Considering that there were a number of procedural differences between the two experiments, this replication is important. A second pattern of note is that a large proportion of BT items ( $M=.63$ ) were given rule-inconsistent responses when they were judged to be "old". If the BT item is judged as "new", then the proportion of BT items given rule-inconsistent responses drops ( $M= .18$ ). This difference in the likelihood of a rule-inconsistent response depending on whether the item was judged as "old" or "new" has several implications. First, it signals the importance of the old/new distinction, because participants are more likely to make a rule-inconsistent response if they also think a BT item is Old. If the BT item is judged as new, then participants are more likely to apply the rule than make a similarity-based response. Second, it suggests that both false recognition and similarity-based retrieval occur. The presence of rule-inconsistent responses on items judged to be "old" is what would be expected with false recognition. The presence of rule-inconsistent responses for items judged to be "new" is what would be expected with the use of a similarity-based classification procedure. Third, even when

**Table 9.**  
Overall proportion of rule-inconsistent responses and the probability of a rule-inconsistent response conditional on the old/new discrimination when that discrimination was made explicit (Experiment 6).

Novelty	Proportion of rule-inconsistent responses	Proportion of a rule-inconsistent responses conditional on the old/new discrimination			
		<u>Old Responses</u>		<u>New Responses</u>	
		Mean	SD	Mean	SD
1 neighbor/novelty present					
	Old	0.07	0.07 (0.10)	0.04	(0.14)
	GT	0.03	0.01 (0.03)	0.02	(0.04)
	BT	0.30	0.63 (0.27)	0.18	(0.22)
	Novel	0.11	0.00 (0.00)	0.07	(0.08)
2 neighbor/novelty reduced (low false alarm group)					
	Old	0.00	0.01 (0.02)	0.00	(0.00)
	GT	0.01	0.01 (0.00)	0.01	(0.04)
	BT	0.35	0.65 (0.37)	0.19	(0.24)
	Novel	0.10	0.07 (0.21)	0.09	(0.09)
2 neighbor/novelty reduced (high false alarm group)					
	Old	0.06	0.05 (0.03)	0.00	(0.00)
	GT	0.02	0.02 (0.03)	0.00	(0.00)
	BT	0.58	0.60 (0.23)	0.13	(0.25)
	Novel	0.01	0.00 (0.00)	0.13	(0.09)

**Note.** the overall proportion of a rule-inconsistent responses reflects the proportion of responses irrespective of whether participants judged the item to be "old" or "new" low false alarm group were participants with few false recognitions high false alarm group were participants with a high number of false recognitions

BT items are judged as "old" there remains a large proportion of BT items for which the rule is applied. This result suggests that participants are able to use a retrieval-based classification procedure that produces false recognitions and similarity-based responses, and also a rule-based classification procedure that helps maintain accuracy.

To verify the above conclusions, planned paired t-tests were conducted to contrast

performance for the three types of novel transfer items. Given that the transfer item was judged as "old", any difference between the BT items and the other two novel transfer items indicates the prevalence of false recognitions. Given that an item was judged as "new", any difference between the BT items and the other transfer stimuli indicates the occurrence of similarity-based retrieval.

When an item was judged as "old", participants were more likely to make a rule-inconsistent response when classifying BT items than GT items,  $t(12) = 57.85$ ,  $p < .001$ , suggesting some use of a retrieval-based strategy. More importantly, these rule-inconsistent responses were limited to BT items because participants were also more likely to make rule-inconsistent responses when classifying BT items relative to the Novel baseline items,  $t(12) = 8.49$ ,  $p < .001$ . This pattern, in which the likelihood of making an "old", rule-inconsistent response when classifying a novel but similar-to-old stimulus is greater than the likelihood of making a rule-inconsistent response on a completely novel stimulus, is consistent with the hypothesis that participants were falsely recognizing some of the similar-to-old transfer stimuli.

A different pattern of results occurs for those items judged as "new". Planned paired t-tests contrasted the probability of making a rule-inconsistent responses for items judged as "new" for each of the three types of transfer items. Participants were more likely to make rule-inconsistent responses when classifying BT items than GT items,  $t(12) = 2.74$ ,  $p < .05$ , a finding consistent with similarity-based retrieval. However, it is not clear that all of the difference between GT and BT items can be attributed to similarity-based retrieval because the difference between BT items and completely novel items only approached significance  $t(12) = 1.95$ ,  $p < .10$ .

**2 neighbor/novelty reduced.** The means for the low false alarm group and the high false alarm group are presented in Table 9. There are two patterns of note displayed in the Table 9. First, as discussed above, participants were more likely to give a

rule-inconsistent response when classifying the 2 Neighbor/Novelty reduced BT stimuli relative to the 1 Neighbor/Novelty present BT items. This pattern replicates those found in Experiment 4 which used comparable stimuli. A second pattern of note is that in both the low false alarm group and the high false alarm group, the proportion of BT items judged to be "old" and also given a rule-inconsistent response was quite high ( $M = .65$  and  $M = .60$  respectively). When the BT items were judged as "new" the proportion of rule-inconsistent responses decreased for both the low false alarm group ( $M = .19$ ) and the high false alarm group ( $M = .13$ ). This difference in the likelihood of a rule-inconsistent response depending on outcome of the old/new discrimination has several implications. First, it signals the importance of the old/new distinction. If participants think a BT item is old, then they are more likely to make a rule-inconsistent response than a rule consistent response. If they think an item is new then participants are more likely to apply the rule than make a similarity-based response. These results also suggest the presence of both false recognition and similarity-based effects. The presence of rule-inconsistent responses on items judged to be "old" is what would be expected if false recognition occurred. The presence of rule-inconsistent responses for items judged to be "new" is what would be expected with the use of a similarity-based classification procedure. In addition, the presence of rule-consistent and rule-inconsistent responses on all BT items continues to suggest that participants are using a mix of retrieval-based knowledge and a rule-based knowledge when classifying the stimuli.

An analysis of the conditional probabilities was conducted to support these conclusions. A set of planned paired t-tests was conducted in which performance for the three types of novel transfer items were contrasted with each other. Given that the stimuli were judged as "old", any differences between the BT items and the other two novel transfer items indicate the prevalence of false recognitions. If judged as "new", then any difference between the BT items and the other transfer stimuli reflects the use of

similarity-based retrieval. Participants in the low false alarm group were more likely to make rule-inconsistent responses to BT items than GT items,  $t(9) = 5.50$ ,  $p < .001$ . This similarity effect was limited to items that resembled training stimuli because participants also made more rule-inconsistent responses to BT items than Novel items,  $t(9) = 3.94$ ,  $p < .01$ . This pattern, in which the likelihood of making an "old", rule-inconsistent response for a novel item that resembles a training item is greater than the likelihood of making that response on a novel item that does not resemble a training item, is consistent with the hypothesis that participants were falsely recognizing some of the similar-to-old transfer stimuli as new.

A different pattern of results occurred for rule-inconsistent responses for items judged as "new". For the low false alarm group, participants made more rule-inconsistent response to BT items than GT items  $t(9) = 2.55$ ,  $p < .05$ , suggesting the possible use of similarity when classifying items participants believe are novel. It is not clear, however, how much of this difference between GT and BT stimuli can be attributed to similarity-based retrieval because the difference between BT items and novel items was not significant  $t(9) = 1.04$ ,  $p > .10$ .

No analysis was conducted on the high false alarm group because there were only four participants. The means presented in Table 9 suggest that the pattern of results would be similar to the low false alarm group. Participants appeared to be more likely to make a rule-inconsistent responses to a BT item judged to be "old" than a GT item judged to be "old" or a Novel item judged to be "old". In addition, while there is a difference between BT and GT items that were judged to be "new", the likelihood of making a rule-inconsistent response does not appear to differ between the BT and Novel transfer items.

### *Discussion*

The goal of this experiment was to make the role of the old/new discrimination in classification more explicit, thus providing a measure of the degree to which novel

transfer stimuli are falsely recognized. There were two findings of interest. First, a number of novel, but similar-to-old, transfer stimuli were falsely recognized as old, especially when the cues that signalled novelty were reduced. A second finding was that the probability of making a rule-inconsistent response on a BT item was greater when that item was also judged to be "old" than when it was judged to be "new". These findings support the idea that participants perform an old/new discrimination before classifying the stimuli. If they believe an item is new they tend to apply the classification rule. If they believe an item is old they tend to make a retrieval-based response, but this response may be more a reflection of false recognition than similarity-based retrieval. There was also some evidence that not all rule-inconsistent responses on BT items are due to false recognition. A small proportion of BT items were given rule-inconsistent responses but also judged to be "new"; a result that would be expected if participants were also relying on similarity-based retrieval. Unfortunately, this conclusion is tempered by the additional finding that just as many completely Novel stimuli were also given rule-inconsistent responses and therefore it may not be justified to attribute these particular rule-inconsistent responses to similarity-based retrieval.

In summary, when the use of an old/new discrimination procedure was made explicit, strong evidence was found for the hypothesis that participants rely on retrieval if an item is believed to be old and a classification rule if an item is believed to be new.

## CHAPTER 8

### *General Discussion*

When classifying a novel object people are able to use both rule-based and retrieval-based processes to determine the category for the object. Given that either process can be used to classify the same object, the overall issue addressed in this thesis is how we coordinate our use of these two types of knowledge. This thesis investigated a method of coordination, in which there is a gradual transition from relying on rule-based processes to relying on retrieval-based processes as one becomes practiced at classifying members of a category. This transition makes a great deal of intuitive sense because many conditions in everyday category learning are conducive to relying on retrieval-based processes. For example, often similarity and rule-based procedures give the same answer, retrieval is faster and less effortful than the explicit application of a rule, most everyday categories consist of a large number of old and similar-to-old items, and categories are structured so that items that are similar to one another usually belong to the same category. All of these conditions occurred in the thesis experiments and therefore should provide favorable conditions for a transition to retrieval to occur. Given that it was assumed that this transition would occur relatively automatically, the experiments were designed to test an extreme view of this transition, in which it was assumed that a transition to retrieval-based processes would occur even when participants were given a simple, predictive, and easy to use classification rule.

To confirm that participants were relying on a faster, less effortful classification procedure at the end of training, response times for the training tasks were plotted as a function of the number of trials, and a power function was fit to the resulting plots. As described by Logan (1988, 1992) and Newell and Rosenbloom (1981), learning times that are well fit by a power function reflect the transition to a procedure that is fast and



effortless. The experiments in this thesis showed that power functions provided excellent fits of the learning times, confirming that participants had acquired some degree of skill at classifying the stimuli.

The underlying basis of this faster, less effortful classification skill was assessed using a transfer task in which participants classified a set of stimuli that were novel but also similar to the stimuli with which they were trained. Participants made more responses that were inconsistent with the classification rule and took longer to classify novel items that were similar to, but in the opposite category as, specific training items (BT items) than they did when classifying novel items that were similar to, and in the same category as, specific training items (GT items). These findings suggest that participants relied on retrieval of similar instances to some degree. However, because the proportion of rule-inconsistent responses was never close to the 100% that would be expected if participants had abandoned the rule and were relying solely on similarity, application of the rule evidently was still useful.

This pattern of transfer results was maintained despite the introduction of factors that were thought to facilitate the use of retrieval. Experiment 1 showed that extensive practice benefited the classification of training items. These items were classified considerably faster than training items that were not extensively practiced and were also classified considerably faster than novel transfer stimuli. However, this extensive practice did not increase the likelihood of making a rule-inconsistent response when classifying BT stimuli. Experiment 2 showed that increasing the number of similar training stimuli from one to ten, and introducing variation within these training stimuli, did not produce an increased reliance on retrieval. If anything, it reduced the benefit of extra practice for classifying training stimuli that was found in Experiment 1. In Experiment 3, a response deadline was imposed at transfer and resulted in an increase in the number of rule-inconsistent responses when classifying BT items. It also resulted in

an increase in rule-inconsistent responses in general. The findings from the first three experiments were interpreted as suggesting that memory retrieval does play an important role in the transition to a fast classification procedure, at least for old items. A transition to a faster, less effortful strategy, however, does not necessarily mean that similarity is generalized to novel items, nor does it necessarily mean that the rule is abandoned. Instead, participants appear to adopt a strategy such that if something is sufficiently similar to a familiar prior instance it is treated as if it is that instance, whereas if something is sufficiently different, it is treated as being new and the classification rule is applied.

The remaining three experiments were designed to provide converging evidence for this “retrieve if believed old, rule if believed new” interpretation. Experiments 4 and 5 showed the importance of novelty detection for coordinating rules and retrieval. When the cues that signalled novelty were reduced, there was an increase in the number of rule-inconsistent responses for BT items, whereas when cues that signalled novelty were present, the number of rule-inconsistent responses on BT items remained relatively low. In Experiment 6, the old/new discrimination was made explicit, and evidence was found that participants falsely recognized a number of novel transfer stimuli rather than generalizing their use of similarity. Although the findings from these latter three studies do not rule out the idea that there is some residual amount of similarity-based retrieval that is used to classify novel items, they do support the claim that a majority of the retrieval-based performance in this situation simply involves treating new items as if they are old.

The intention of this thesis was to borrow ideas from the skill acquisition literature to account for the transition from slow, effortful performance to faster, less effortful performance when learning a category. There are two broad intuitions from this literature that can be applied to category learning. The first intuition is based on Anderson's ACT-R

model (Anderson 1993; also Anderson, 1982; Anderson & Singley, 1989) and suggests that initial classification is based on the deliberate use of analogies or general classification rules. Given sufficient practice, participants eventually form a production rule, at which point their classification performance becomes fast and less effortful. This fast rule-based strategy will transfer to novel stimuli if the underlying rule-based procedures are consistent with how participants classified the training stimuli. If the procedure changes then participants revert back to the slower rule. A classification framework based on this intuition is not consistent with the findings in this thesis. Structurally, the novel GT and BT stimuli were equally similar to the training items, differing in only one feature. Therefore, the underlying rule-based procedures required to classify the transfer stimuli ought to have been consistent with the procedures required to classify the training stimuli, and transfer to the novel but-similar-to old stimuli should have been fast and rule-based. This did not occur, however, because fast responses on BT items tended to be consistent with similarity.

Note that this shortcoming of the ACT-R model does not imply that the model is incorrect, but only that the experimental conditions used in this thesis may not have been conducive to the use of a production rule. The ACT-R model was initially derived to account for performance in tasks that are relatively complex, such as computer programming, solving logic problems, or solving alpha-arithmetic problems. These types of tasks may require the execution of sets of steps, and performance on such tasks may be well described by the ACT-R model. Stimuli that lend themselves to direct perceptual processing, like those used in the present studies, may be relatively easy to process and to remember. Consequently, classification may be biased toward reliance on memory rather than the execution of a set of steps. It is also possible that the ACT-R model is not consistent with the findings in this thesis because the amount of repetition in these experiments led to an increased reliance on retrieval-based processing. For example,

Anderson et al. (1997) found that repeating specific examples during learning led to a reliance on instance retrieval whereas presenting multiple variations of specific training problems, without repeating specific problems, led to the development of a production rule. Anderson et al. concluded that skilled performance in a cognitive task represents the mixed use of memory retrieval and production rules.

The second intuition from the skill acquisition literature is based on Logan's Instance theory (Logan, 1988, 1992). According to this view, initial classification is also based on the explicit use of a slow effortful classification rule. However, given sufficient practice, retrieval of instances from memory begins to form the basis of responding, eventually resulting in memory-based classification that is fast and effortless. As Logan instantiated the theory, the range of similarity in which memory for a prior item will transfer to a novel instance is extremely narrow. In most cases, even a new instance that is similar to a training instance does not benefit from that similarity and requires the application of the rule for classification. This assumption, that the use of retrieval is limited only to prior items, is problematic because similarity-based retrieval is a critical component to many models of classification and memory. It was reasoned that due to the importance of similarity-based retrieval in memory, and because Logan's Instance Theory relies on retrieval from memory, it should be possible to extend the theory to account for the classification of both old instances and novel but similar-to-old instances.

The findings in this thesis were inconsistent with this extension of an Instance theory account of classification. Novel but similar-to-old items only benefitted from their similarity if they were falsely recognized as being old; if participants realized the items were new, they slowed down and applied the rule. This occurred despite the experimental conditions being designed to create conditions that should have been conducive to similarity generalizing to novel stimuli. The findings, however, were not inconsistent with an Instance Theory account of classification if a false recognition

component is added to the model rather than a generalization process. That is, if the assumption is added to Logan's Instance theory that novel but similar-to-old stimuli can be falsely recognized, then a classification framework based on Logan's original instantiation of the Instance theory is sufficient to account for the findings in this thesis.

***An assessment of the relations between similarity and false recognition***

According to formal exemplar-based models such as Medin and Schaffer's (1978) Context model and Nosofsky's (1988) Generalized Context Model, similarity-based retrieval is an important process for both recognition and classification of novel items. How this similarity-based retrieval process is used, however, differs between the two tasks. In classification, novel stimuli that are similar to previously learned instances often benefit from that similarity and are classified faster and more accurately than novel stimuli that are not similar to previously learned instances. In addition, the more similar a novel stimulus is to a previous instance of a category, the more likely it will benefit from that similarity. This benefit in performance for novel but similar-to-old items is attributed to a similarity-based retrieval process in which similarity is determined by comparing a novel item to other instances of a category that are stored in memory. This use of similarity is also called generalization. In a recognition task, novel stimuli that are similar to studied instances tend to be falsely recognized as old more often than novel stimuli that are not similar to studied instances. In addition, the more similar a novel stimulus is to a studied instance, the more likely it will be falsely recognized. This increase in false recognitions is attributed to a similarity-based retrieval process in which similarity is determined by comparing the novel item to all items in memory.

All of the experiments reported in this thesis provided evidence of similarity-based retrieval responding. A number of BT items were placed in the same category as the training instances they were similar to despite participants having learned a rule that would have led them to put the BT items in the opposite category. This use of similarity

should have been consistent with generalization because the thesis experiments were classification tasks. However, the similarity-based retrieval process that was used appeared to be more consistent with false recognition. This finding is surprising because of the assumption made by exemplar-based models that specific uses of similarity are linked to specific tasks; generalization to classification tasks and false recognition to recognition tasks. Finding that categorization is dependent on recognition could suggest that the distinction between the two uses of similarity-based retrieval is not as crucial as is assumed by exemplar-based models.

If the link between a specific use of similarity and a specific task is not crucial, then it is possible that some of the evidence for generalization that has been found in past studies may have occurred because the novel stimuli were being falsely recognized. If so, then this may have an effect on how the findings should be interpreted. For example, Palmeri (1997) described a set of findings that were interpreted as being due to the use of a generalization process to classify novel stimuli. He trained participants to count the dots on sets of spatial patterns and then presented them with novel patterns that were distortions of the training patterns. Palmeri found that participants responded faster to novel shapes that were similar to the training shapes than they did to novel shapes that were not similar to the training shapes. Additionally, participants responded faster to moderately similar shapes than shapes that had a low degree of similarity to the training shapes. Palmeri attributed these similarity effects to a similarity-based retrieval process and described how a model called the Exemplar-Based Random Walk model (Palmeri, 1997; Nosofsky & Palmeri, 1997) would account for these results.

The Exemplar-Based Random Walk model combines aspects of Nosofsky's Generalized Context Model with aspects of Logan's Instance theory. Like the Generalized Context Model, categories consist of stored exemplars, and like the Instance Theory, when a stimulus is to be classified, these exemplars race against each other to be

retrieved from memory. The Exemplar-Based Random Walk model differs from these prior theories, however, in that the exemplars race to be retrieved with speeds proportional to their similarity to the presented item. The more similar a probe is to a stored exemplar, the more likely the stored exemplar will be retrieved. Furthermore, each retrieval of an exemplar only provides incremental evidence to drive a random walk, and it is only after sufficient evidence has accumulated that a response is made.

Given that the Exemplar-Based Random Walk model described by Palmeri (1997) is similar in spirit to the Instance-based classification framework described in this thesis, it may be subject to some of the same difficulties. That is, there is no reason to believe that the same false recognition effects, present in the thesis experiments, were not also present in Palmeri's generalization data. After all, the more similar a spatial pattern is to a learned pattern, the more likely it could be either falsely recognized or retrieve a similar prior instance. Distinguishing between the two uses of similarity may provide an additional constraint to the Exemplar-Based Random Walk model.

If generalization is no longer linked exclusively to classification and if false recognition is no longer linked exclusively to recognition, then it becomes crucial to determine the conditions that might lead a participant to adopt one use of similarity over the other. A recent experiment by Neal (unpublished data) highlights this importance. Using a procedure similar to that reported by Neal et al. (1995), Neal had participants listen to descriptions of individuals who were applying for a bank loan. The participants were given a two out of three feature classification rule, and after an initial training phase, they were presented with GT and BT descriptions to classify. Neal found a difference between the two transfer conditions, with participants making more rule-inconsistent responses on the BT descriptions than the GT descriptions. Of more interest is that Neal also had participants identify the descriptions as old or new immediately after they classified them. He found virtually error free performance with

participants falsely recognizing only about 1% of the novel transfer stimuli and correctly recognizing about 99% of the training items. This finding, in which there are rule-inconsistent responses for BT descriptions without evidence of false recognitions, suggests that Neal's similarity effect was more akin to generalization than false recognition.

A current topic of interest, therefore, is to explore the differences in experimental design that led to false recognitions in the thesis experiments and generalization in Neal's experiment. One possibility is that memorial differences between pictorial and verbal stimuli alter the use of similarity. The stimuli in this thesis were all pictorial items, whereas Neal's stimuli were all verbal descriptions. This difference in stimuli may be important because Paivio (1971) found that pictorial stimuli are easier to remember than verbal stimuli. Easy to remember stimuli, combined with the large number of repetitions that were used in this thesis, may have encouraged participants to treat the task as a memory test. When novel stimuli were then presented in the transfer phase, participants tried to discriminate old from new items, much as they would in a recognition test. The verbal stimuli that Neal used may be more difficult to remember thus preventing a recognition strategy from developing.

It is also possible that rule application is considerably easier for pictorial stimuli than verbal stimuli. With pictorial stimuli, it is easy to quickly scan a stimulus to check if the requisite features are present. Listening to a verbal description requires a participant to keep the description in mind until the right features are read aloud. This ease in rule application for pictorial stimuli may allow participants to encode additional item-specific knowledge that later encourages a recognition strategy.

A third possibility for the differences between Neal's work and that reported in this thesis comes from recent work on object and social categorization. Using stimuli with the same informational structure, Wattenmaker (1995) found that if the features were



manifested in terms of an object, like a hammer, classification and sorting was based on the presence of specific defining features. If the features were manifested in terms of descriptions of people, classification and sorting was based on the linear separability of the stimuli. Wattenmaker concluded that social categories require participants to use background knowledge to interpret the features while classification of object categories require participants to look for specific features. The stimuli used in Neal's studies are analogous to Wattenmaker's social categories and may require background knowledge to perform the task. The stimuli used in this thesis may be analogous to Wattenmaker's object categories and may encourage participants to learn about the individual features. If so, then it is possible prior knowledge may be an important factor for the use of similarity-based retrieval.

#### ***A reassessment of the relations between rule-based and retrieval-based classification***

In the experiments that were reported, participants maintained their use of a classification rule, even under conditions that were designed to encourage the use of retrieval. This robust finding raises questions about the role of rule-based and retrieval-based knowledge in classification. Ever since Rosch (Rosch, 1975; Rosch & Mervis, 1975; Mervis & Rosch, 1981) showed that categories in the real world are ill-defined, there has been a backlash against rule-based models, and the belief that simple rules don't exist in the real world has become popular. Many of the subsequent classification models that were developed had no role for rule-based knowledge, and some, like Medin and Schaffer's (1978) Context Model, were able to show that by changing the weight given to specific features when calculating the similarity between items, findings analogous to rule-based classification could occur without the explicit use or storage of a rule.

In recent years, however, there has been some dissatisfaction with the idea that simple rules play only a minor role in categorization. For example, Medin and colleagues

(Medin, Wattenmaker & Hampson, 1987; Ahn & Medin, 1992; also Regehr & Brooks, 1995) showed that when participants were asked to sort arrays of stimuli into two categories that have a family resemblance structure, participants had a strong tendency to sort the category on the basis of a single defining feature. This is not what would have been expected if similarity-based "family resemblance" strategies had been natural for the subjects. Nosofsky and colleagues (Nosofsky, Palmeri, & McKinley, 1994; Palmeri & Nosofsky, 1995; also Ward & Scott, 1987) have also questioned the apparent lack of support for rule-based classification and have developed the "rules plus exception" (Rulex) model of classification. According to this model, classification involves active hypothesis testing in an attempt to find a sufficient classification rule, using memory only to store exceptions to that rule when one is found. Much of the data that apparently supported similarity-based models, such as the Generalized Context Model, was found to be fit better by the Rulex model. The present findings (see also Allen & Brooks, 1991; Regehr & Brooks, 1993) are consistent with this trend toward an increased role for rules in classification.

How should the relations between rule-based and retrieval-based classification be resolved? The current materials were originally designed to present an extreme example of the relation between rule-based and similarity-based categorization: if a rule that is simple and easy to apply yields, with experience, to the use of prior episodes, then we should entertain the possibility that reliance on prior episodes is a very common means of classification. Given what is known about episodic memory it certainly seems feasible that people would remember the many variants of a single category, and save the use of a rule or other explicit knowledge for dealing with special problems. One advantage of this changed basis of classification would be greater efficiency in retrieval, a possibility that was confirmed in the present work by the consistently faster reaction times to Old than new (GT and BT) items. However, the continued reliance on rules for items believed to

be new and false recognition of items believed to be old suggests that the rules to episodes transition for similarity-based retrieval is, at the least, subject to more complicated conditions than was anticipated. Again, the research by Neal suggests that such conditions can be found; their generality is subject to further investigation.

The rules to episodes transition also addresses an additional issue: it might provide one essential reason why people are able to maintain the belief that categories in the world are simpler than they really are (the psychological essentialism referred to by Medin & Ortony, 1989; see also Malt, 1990). That is, people might start with a simple, explicit rule that is useful for initially categorizing instances and, with experience, gradually rely on prior episodes rather than the original rule. By no longer relying on the original rule, they might not notice when contradictions between that rule and the category occur, which they are successfully categorizing by other means. When participants are later challenged to think explicitly about the identification criteria for that category, they can retrieve the original rule and, in the context of having had no trouble with categorizing instances of that category, misattribute their success to that rule. An example of this type of misattribution is demonstrated in Brooks, Squire-Graydon and Wood (1998). When participants were asked if there is a simple single feature rule that all bottles have in common, the majority of participants confidently answer that there was. This confidence in the sufficiency of a rule is particularly interesting considering that when asked if their rule included baby bottles and pill bottles but excluded jars, glasses, and cartons, virtually all of the participants admitted they were mistaken.

If the use of prior episodes does not contribute to the belief that categories in the world are simpler than they really are, then it is hard to imagine how rule-based knowledge could lead to this belief with natural categories, but not produce the same belief with categorization tasks in the laboratory. As shown by Nosofsky et al. (1994; also Palmeri & Nosofsky, 1995; Ward & Scott, 1987), participants are aware of the

complexity of the categories and know when a rule does not work. In the experiments reported in this thesis, participants also demonstrated that they were well aware of the complexities of the stimuli in that they were surprisingly good at knowing when a single feature had been changed on an otherwise familiar stimulus.

One approach to investigating the difference between common laboratory tasks and everyday categorization learning was provided by Brooks et al. (1998). Their argument was that most laboratory tasks that use family resemblance categories do not result in the everyday "simpler than it is" belief because the stimuli and the tasks encourage direct attention to identification procedures during training. In order to simulate the learning that occurs under many natural circumstances, they developed a procedure called diverted analysis, in which the participants' analytic abilities are diverted from the way in which the stimuli are identified to the use to which those stimuli are to be put. Participants were given a rule for movement on a chess board for each of two (family resemblance) categories of items. The emphasis throughout the training period was on the consequences of movement ("who won?") rather than the identification procedure. Categorization was supported initially by providing labels for the items and only gradually withdrawing this support after familiarity with the items increased. After this procedure, participants stated with high confidence that "all members of one category had some feature that no members of the other category had." This same proposition was given drastically lower endorsement after conventional category learning procedures, like those used in this thesis, were used or when the participants were instructed to memorize the category of the items. Brooks et al. suggest that the latter manipulations led to an increased awareness of the complexity of the stimuli while the diverted analysis procedure led participants to assume the stimuli consisted of simple structures.

Returning to the thesis, the initial goal was to examine whether a memory-based framework of skill acquisition can account for the pattern of results that occur when

practiced classifiers are asked to classify both old and novel but similar-to-old stimuli. Under the current experimental conditions, designed to facilitate retrieval of similar instances from memory as a basis of classification, participants adopted a “retrieve if believed old, rule if believed new” strategy. If an item was new they relied on the classification rule, whereas if an item was old, or was new but participants failed to discriminate it as new, then participants relied on fast retrieval-based processes for classification. It is possible, therefore, that under some circumstances, transfer of a learned skill to novel situations is simply due to false recognition and not similarity-based retrieval. This possibility needs to be borne in mind when evaluating evidence for instance-based categorization. Finally, the more general proposition, in which explicit rule-based methods of classification routinely yield to instance or episodic bases of classification, is clearly subject to more complex conditions than was originally assumed.

## CHAPTER 9

### *References*

- Ahn, W. K., & 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.
- Allen, S. W., & Jacoby, L. L. (1990). Reinstating study context produces unconscious influences of memory. *Memory & Cognition, 18*, 270-278.
- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review, 89*, 369-406.
- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ:Erlbaum.
- Anderson, J. R., & Fincham, J. M. (1994). Acquisition of procedural skills from examples. *Journal of Experimental Psychology: Human Learning and Memory, 20*, 1322-1340.
- Anderson, J. R., & Fincham, J. M. & Douglass, S. (1997). The role of examples and rules in the acquisition of a cognitive skill. *Journal of Experimental Psychology: Human Learning and Memory, 23*, 932-945.
- Andrews, S. (1989). Frequency and neighborhood size effects on lexical access: Activation or search? *Journal of Experimental Psychology: Human Learning and Memory, 15*, 802-814.
- Bellezza, F. S. (1984). Reliability of retrieval from semantic memory: noun meanings. *Bulletin of the Psychonomic Society, 22*, 377-380.
- Blessing, S. B., & Anderson, J. R. (1996). How people learn to skip steps. *Journal of Experimental Psychology: Human Learning and Memory, 22*, 576-598.
- Brooks, L. (1976). Visual pattern in fluent word identification. In A.S. Reber & D.L.

- Scarborough (Eds.), *Toward a psychology of reading*. Hillsdale, NJ: Erlbaum.
- Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. E. Rosch & B. Lloyd (Eds.) *Cognition and categorization*. (pp. 169-211) Hillsdale, NJ: Erlbaum.
- Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, *120*, 278-287.
- Brooks, L. R., Squire-Graydon, R. & Wood, T. J. (submitted). The role of inattention in everyday concept learning: Identification in the service of use.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study in thinking*. New York: Wiley.
- Carlson, R. A., & Lundy, D. H. (1992). Consistency and restructuring in learning cognitive procedural sequences. *Journal of Experimental Psychology: Human Learning and Memory*, *18*, 127-141.
- Charness, N. & Campbell, J. I. D. (1988). Acquiring skill at mental calculation in adulthood: A task decomposition. *Journal of Experimental Psychology: General*, *117*, 115-129.
- Cho, J. R., & Mathews, R. C. (1996). Interactions between mental models used in categorization and experiential knowledge of specific cases. *The Quarterly Journal of Experimental Psychology*, *49A*, 572-595.
- Elio, R., & Anderson, J. R. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 397-417.
- Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD : Erlbaum.
- Gentner, D., Ratterman, M. J., & Forbus, K. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, *25*,

- 524-575.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1-67.
- Haygood, R. C., & Bourne, L. E. (1965). Attribute and rule learning aspects of conceptual behavior. *Psychological Review*, 72, 175 - 195.
- Hinton, G. E., McClelland, J. L., Rummelhart, D. E. (1986). Distributed representations. In D. E. Rummelhart & J. L. McClelland (Eds. ) *Parallel distributed processing: Explorations in the microstructure of cognition: Volume I. Foundations* (pp. 77-109). Cambridge, MA: MIT Press.
- Hintzman, D. L. (1980). Simpson's Paradox and the analysis of memory retrieval. *Psychological Review*, 87, 398-410.
- Hintzman, D. L. (1986). Schema Abstraction in a multiple trace memory model. *Psychological Review*, 93, 411-428.
- Homa, D., Sterling, S. & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 418-439.
- Homa, D., & Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 322-330.
- Hull, C. L. (1920). Quantitative aspects of the evolution of concepts. *Psychological Monographs*, 28, (1, whole No. 123).
- Jacoby, L. L. & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. H. Bower (Ed. ) *The psychology of learning and motivation* (pp. 1-47). New York: Academic Press.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.



- Kemler-Nelson, D. G. (1984). The effect of intention on what concepts are acquired. *Journal of Verbal Learning and Verbal Behavior*, 23, 734 - 759.
- Kemler-Nelson, D. G. (1989). The nature and occurrence of holistic processing. In B. E. Shepp & S. Ballesteros (Eds.) *Object perception: Structure & process..* (pp. 357-386). Hillsdale, N.J.: Erlbaum.
- Krushcke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22-44.
- Lassaline, M. E., & Logan, G. D. (1993). Memory-based automaticity in the discrimination of visual numerosity. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19, 561-581.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.
- Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms. *Cognitive Psychology*, 22, 1-35.
- Logan, G. D. (1992). Shapes of reaction-time distributions and shapes of learning curves: A test of the instance theory of automaticity. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 883-914.
- Logan, G. D., & Stadler, M. A. (1991). Mechanisms of performance improvement in consistent mapping memory search: Automaticity or strategy shift? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 17, 478-496.
- Malt, B. C. (1990). Features and beliefs in the mental representation of categories. *Journal of Experimental Psychology : Learning, Memory, and Cognition*, 12, 479-488.
- Martin, R. C., & Caramazza, A. (1980). Classification in well-defined and ill-defined categories: Evidence for common processing strategies. *Journal of Experimental Psychology: General*, 109, 320-353.

- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology : General*, 109, 320-353.
- McCloskey, M. E. & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory & Cognition*, 6, 462-472.
- Medin, D. L. (1986). Commentary on "memory and retrieval processes in category learning." *Journal of Experimental Psychology: General*, 115, 373-381.
- Medin, D. L., Dewey, G. I., Murphy, T.D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 607-625.
- Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds. ), *Similarity and analogical reasoning*. New York: Cambridge University Press.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.
- Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 7, 355-368.
- Medin, D.L. & Smith, E.E. (1981). Strategies and classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 241 - 253.
- Medin, D. L., Wattenmaker, W. D. & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, 19, 242-279.
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, 32, 89-115.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative

- information. *Psychological Review*, 89, 609-626.
- Neal, A., Hesketh, B., & Andrews, S. (1995). Instance-based categorization: Automatic versus intentional forms of retrieval. *Memory & Cognition*, 23, 227-242.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed. ). *Cognitive skills and their acquisition* (pp. 1-55). Hillsdale, N. J. Erlbaum.
- Nosofsky, R. M. (1986). Attention, Similarity, and the Identification-Categorization Relationship. *Journal of Experimental Psychology:General*, 115, 39-57.
- Nosofsky, R. M. (1988). Exemplar-based accounts of the relationships between classification, recognition, and typicality. *Journal of Experimental Psychology:Learning, Memory and Cognition*, 14, 700-708.
- 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., Palmeri, T. J., & Mckinley. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.
- Nosofsky, R. M., Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104, 266-300.
- Paivio, A. (1971). Imagery and verbal processes. New York : Holt
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology:Learning, Memory and Cognition*, 23, 324-354.
- Palmeri, T. J, & Nosofsky, R. M. (1995). Recognition memory for the exceptions to the category rule. *Journal of Experimental Psychology:Learning, Memory and Cognition*, 21, 548-568.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353-363.

- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 77, 317-327.
- Reber, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning & Memory*, 2, 88-94.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118, 219-235.
- Reber, A. S., & Allen, R. (1978). Analogy and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition*, 6, 189-221.
- Reber, A. S., Kassin, S. M., Lewis, S., Cantor, G. (1980). On the relationship between implicit and explicit modes of learning a complex rule structure. *Journal of Experimental Psychology: Human Learning & Memory*, 6, 492-502.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382-407.
- Regehr, G. & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure. *Journal of Experimental Psychology : General*, 122, 92-114.
- Regehr, G. & Brooks, L. R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology : Learning, Memory, Cognition*, 21, 347-363.
- Riesbeck, C. K., & Schank, R. C. (1989). *Inside Case-based Reasoning*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Rips, L. J., Shoben, E.J. & Smith, E. E. (1973). Semantic distances and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1-20.
- Rosch, E. E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192-233.
- Rosch, E. E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-603.

- Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Human Learning and Memory*, *13*, 629-639.
- Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Human Learning and Memory*, *15*, 456-468.
- Rueckl, J. G. (1990). Similarity effects in word and pseudoword repetition priming. *Journal of Experimental Psychology: Human Learning and Memory*, *16*, 374-391.
- Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science*, *3*, 207-216
- Singley, M. K. & Anderson, J. R. (1989). *The transfer of cognitive skill..* Cambridge, MA. : Harvard University Press.
- Smith, J. D. & Kemler-Nelson, D. G. (1984). Overall similarity in adults's classification: The child in all of us. *Journal of Experimental Psychology: General*, *113*, 137-159.
- Vokey, J. R. & Brooks, L. R. (1992). The salience of item knowledge in learning artificial grammars. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *18*, 324-344.
- Ward, T. B., & Scott, J., (1987). Analytic and holistic modes of learning family-resemblance concepts, *Memory & Cognition*, *15*, 42-54.
- Ward, T. B., Foley, C. M., & Cole, J. (1986). Classifying multidimensional stimuli: Stimulus, task and observer factors. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 211 - 225.
- Wattenmaker, W. D. (1991). Learning modes, feature correlations, and memory-based categorization. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *17*, 908-923.

- Wattenmaker, W. D. (1995). Knowledge structures and linear separability: Integrating information in object and social categorization. *Cognitive Psychology, 28*, 274-328.
- Whittlesea, B. W. A. (1987). Preservation of specific experiences in the representation of general knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13*, 465-472.
- Whittlesea, B. W. A., Brooks, L. R., & Westcott, C. (1994). After the learning is over: Factors controlling the selective application of general and particular knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 259-274.
- Whittlesea, B. W. A., & Dworkin, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General, 116*, 224-248.

### Appendix A

The analytic structure of the stimuli used to compare high practiced stimuli to low practiced stimuli. (Experiment 1)

---

items	<u>Training items</u>					<u>Transfer items</u>				
	A	B	C	D	E	A	B	C	D	E
1.1	0	0	0	0	0	0	0	1	0	0
1.2	0	0	0	1	1	0	0	1	1	1
1.3	0	1	1	0	1	0	1	0	0	1
1.4	0	1	1	1	0	0	1	0	1	0
1.5	1	0	1	0	1	1	0	0	0	1
1.6	1	0	1	1	0	1	0	0	1	0
1.7	1	1	0	0	0	1	1	1	0	0
1.8	1	1	0	1	1	1	1	1	1	1
2.1	0	0	1	0	0	0	0	0	0	0
2.2	0	0	1	1	1	0	0	0	1	1
2.3	0	1	0	0	1	0	1	1	0	1
2.4	0	1	0	1	0	0	1	1	1	0
2.5	1	0	0	0	1	1	0	1	0	1
2.6	1	0	0	1	0	1	0	1	1	0
2.7	1	1	1	0	0	1	1	0	0	0
2.8	1	1	1	1	1	1	1	0	1	1

---

Note. Items 1.1 - 1.8 and 2.1 - 2.8 alternately served as high practiced or low practiced items.

A = Body shape (0 = round 1 = angular)

B = Neck length (0 = short 1 = long)

C = Body markings (0 = no spots 1 = spots)

D = Leg length (0 = short 1 = long)

E = Number of legs (0 = two 1 = six)

### Appendix B

The analytic structure of the stimuli used when multiple similar training instances were required. (Experiments 2, 4, 5 & 6)

Rule	<u>Training item 1</u>						<u>Training item 2</u>						<u>GT items</u>						<u>BT items</u>							
	A	B	C	D	E	F	A	B	C	D	E	F	A	B	C	D	E	F	A	B	C	D	E	F		
1.1	0	0	1	0	0	0	0	1	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0			
1.2	0	1	0	0	1	0	1	1	1	0	1	0	0	1	1	0	1	0	1	1	0	0	1	0		
1.3	0	1	0	1	0	1	1	1	1	1	0	1	0	0	1	1	1	0	1	1	1	0	1	0		
1.4	0	0	1	1	1	1	0	1	0	1	1	1	1	0	1	1	1	1	1	0	0	0	1	1	1	
1.5	1	0	1	0	0	0	1	1	0	0	0	0	0	1	0	0	0	0	0	1	1	1	0	0	0	
1.6	0	0	0	0	1	0	1	0	1	0	1	0	1	1	0	0	0	1	0	0	1	0	1	0	0	
1.7	0	0	0	1	0	1	1	0	1	1	0	1	0	1	0	0	1	0	1	0	1	0	1	0	1	
1.8	1	0	1	1	1	1	1	1	0	1	1	1	1	1	0	0	1	1	1	1	1	1	1	1	1	
2.1	0	1	0	0	0	0	1	1	1	0	0	0	0	1	1	0	0	0	0	0	0	1	0	0	0	
2.2	0	1	0	0	1	0	1	0	0	0	1	0	1	1	0	0	1	0	1	0	0	0	0	1	0	
2.3	0	1	0	1	0	1	1	0	0	1	0	1	0	1	1	0	1	0	1	0	0	0	1	0	1	
2.4	0	1	0	1	1	1	1	1	1	1	1	1	1	1	0	1	1	1	1	0	1	1	1	1	1	
2.5	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	1	0	0	0	0	1	0	0	0	0	
2.6	0	1	1	0	1	0	1	0	1	0	1	0	1	0	0	1	0	1	0	1	1	1	0	1	0	
2.7	0	1	1	1	0	1	1	0	1	1	0	1	0	0	1	1	0	1	0	1	1	1	1	0	1	
2.8	0	0	0	1	1	1	1	0	1	1	1	1	1	0	0	1	1	1	1	1	1	0	0	1	1	1
3.1	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0
3.2	0	0	1	0	1	0	0	1	0	0	1	0	1	0	0	0	0	1	0	0	0	1	1	0	1	0
3.3	0	0	1	1	0	1	0	1	0	1	0	1	0	0	0	1	0	1	0	1	0	1	1	1	0	1
3.4	0	1	0	1	1	1	1	0	0	1	1	1	1	0	0	0	1	1	1	1	1	0	1	1	1	1
3.5	0	1	1	0	0	0	1	0	1	0	0	0	0	1	1	1	0	0	0	0	0	1	0	0	0	0
3.6	1	0	1	0	1	0	1	1	1	0	0	1	0	1	1	1	0	1	0	1	1	0	0	0	1	0
3.7	1	1	0	1	0	1	1	0	1	1	0	1	0	1	1	1	1	0	1	0	1	0	0	1	0	1
3.8	1	0	1	1	1	1	1	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	1	1	1	1

Note. rule = type of classification rule

#### Experiment 4 & 5

A = Tail length (0 = short 1 = long)  
 B = Body markings (0 = stripes 1 = spots)  
 C = Number of legs (0 = two 1 = six)  
 D = Body shape (0 = round 1 = angular)  
 E = Neck length (0 = short 1 = long)  
 F = Leg length (0 = short 1 = long)

#### Experiment 2 & 6

A = Tail length (0 = short 1 = long)  
 B = Neck length (0 = short 1 = long)  
 C = Leg length (0 = short 1 = long)  
 D = Body shape (0 = round 1 = angular)  
 E = Body markings (0 = stripes 1 = spots)  
 F = Number of legs (0 = two 1 = six)



### Appendix C

The analytic structure of the stimuli used when a deadline was imposed at transfer (Experiment 3) or when 1 Neighbor/Novelty present transfer stimuli were used (Experiments 4,5, and 6).

Rule	<u>Training items</u>						<u>GT items</u>						<u>BT items</u>						
	A	B	C	D	E	F	A	B	C	D	E	F	A	B	C	D	E	F	
1.1	0	1	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	
1.2	1	1	1	0	1	0	0	1	1	0	1	0	1	1	0	0	1	0	
1.3	1	1	1	1	0	1	0	1	1	1	0	1	1	1	0	1	0	1	
1.4	0	0	1	1	1	1	0	1	1	1	1	1	1	0	0	1	1	1	
1.5	1	0	1	0	0	0	1	0	0	0	0	0	0	1	1	1	0	0	0
1.6	0	0	0	0	1	0	1	0	0	0	1	0	1	0	0	1	0	1	0
1.7	0	0	0	1	0	1	1	0	0	1	0	1	1	0	0	1	1	0	1
1.8	1	1	0	1	1	1	1	0	0	1	1	1	1	1	1	1	1	1	1
2.1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	1	1	0	0	0
2.2	1	0	0	0	1	0	1	1	0	0	1	0	1	0	0	0	0	1	0
2.3	1	0	0	1	0	1	1	1	0	1	0	1	1	0	0	0	1	0	1
2.4	0	1	0	1	1	1	1	1	0	1	1	1	1	0	1	1	1	1	1
2.5	0	0	0	0	0	0	0	0	1	0	0	0	0	1	0	0	0	0	0
2.6	0	1	1	0	1	0	0	0	1	0	1	0	1	1	1	1	0	1	0
2.7	0	1	1	1	0	1	0	0	1	1	0	1	1	1	1	1	0	1	1
2.8	1	0	1	1	1	1	0	0	1	1	1	1	1	1	0	0	1	1	1
3.1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0
3.2	0	0	1	0	1	0	0	0	0	0	1	0	1	0	1	1	0	1	0
3.3	0	0	1	1	0	1	0	0	0	1	0	1	1	0	1	1	1	0	1
3.4	0	1	0	1	1	1	0	0	0	1	1	1	1	1	1	0	1	1	1
3.5	0	1	1	0	0	0	1	1	1	0	0	0	0	0	0	1	0	0	0
3.6	1	1	0	0	1	0	1	1	1	0	1	0	1	1	0	0	0	1	0
3.7	1	1	0	1	0	1	1	1	1	1	0	1	1	0	0	1	0	1	1
3.8	1	0	1	1	0	1	1	1	1	1	1	1	1	0	0	1	1	1	1

Note. rule = classification rule

#### Experiment 4 & 5

A=Tail length (0 = short, 1 = long)  
 B=Body markings (0 = stripes, 1 = spots)  
 C=Number of legs (0 = two, 1 = six)  
 D=Body shape (0 = round, 1 = angular)  
 E=Neck length (0 = short, 1 = long)  
 F=Leg length (0 = short, 1 = long)

#### Experiment 3 & 6

A=Tail length (0 = short, 1 = long)  
 B=Neck length (0 = short, 1 = long)  
 C=Leg length (0 = short, 1 = long)  
 D=Body shape (0 = round, 1 = angular)  
 E=Body markings (0 = stripes, 1 = spots)  
 F=Number of legs (0 = two, 1 = six)