PERCEPTUAL MANIFESTATIONS OF AN ANALYTIC STRUCTURE: THE PRIORITY OF HOLISTIC INDIVIDUATION

By

GLENN REGEHR, B.A.

A thesis

Submitted to the School of Graduate Studies

in Partial Fulfillment of the Requirements

for the Degree

Doctor of Philosophy

McMaster University

PERCEPTUAL MANIFESTATIONS OF AN ANALYTIC STRUCTURE

DOCTOR OF PHILOSOPHY (1992) McMASTER UNIVERSITY

(Psychology) Hamilton, Ontario

TITLE: Perceptual Manifestations of an Analytic Structure: The Priority of Holistic Individuation

AUTHOR: Glenn Regehr, B. A. (University of Waterloo)

SUPERVISOR: Professor Lee R. Brooks

NUMBER OF PAGES: (ix), 105

ABSTRACT

For many natural categories, there is idiosyncratic information contained in the perceptual stimuli that is not likely to be represented in an analytic description of those stimuli, and thus is likely to be ignored during attempts to establish an abstracted, probabilistic description of the category. This idiosyncratic perceptual information is useful, however, in the identification of distinct individuals within a category and, therefore, will be important in more nonanalytic, exemplar based classification procedures. The present work addresses two such sources of perceptual information: 1) Feature Individuation, the specific manner in which each separate feature is manifested in an individual item, and 2) Holistic Individuation, the manner in which the item's separate features cohere into an individuated gestalt whole. Data from Experiment Sets 1 and 2 indicate that the presence of either source of perceptual information will have an impact on classification behaviour. Furthermore, the extent of this impact is, at least in part, a function of whether subjects are provided with the classification rule prior to training or are expected to induce the classification rule through repeated feedback during training. Data from Experiment Set 3 indicate that if both sources of perceptual information are present, holistic individuation of the stimuli has priority. The presence of individuated features is not sufficient to produce exemplar based transfer if an item's holistic individuality is altered, but will affect responding if the holistic nature of the stimuli is eliminated during the transfer phase of the experiment. Finally, data from Experiment Set 4 indicate that learning about the two types of nonanalytic information

does not appear to be incompatible with learning about the analytic information contained in the stimuli. Learning to use these various sources of information for the purposes of classification, however, are somewhat antagonistic.

ACKNOWLEDGEMENTS

My time at McMaster has been both an excellent learning experience and downright fun, and virtually all members of the McMaster psychology community have had a part in making this true. This is especially true for my supervisor, Dr. Lee Brooks, who has always worked to ensure that I got the most out of my graduate education. He has treated me as a colleague and as a friend, and I look forward to maintaining my relationship with him in both these regards for many years to come. I also thank the members of my advisory committee, Dr. Larry Jacoby, Dr. Rolfe Morrison and Dr. Ian Begg, for their comments and insights. I have benefitted from my interactions with each of them.

I thank my friends and fellow occupants of the first floor lab wing, all of whom have contributed positively to my experience at McMaster both socially and academically. A special thanks goes to Phil Higham, who always forced me to think about why I believed the things I believed rather than just letting me accept them as truths. In the time we shared an office, we didn't agree on much, but we had a hell of a good time disagreeing.

And finally, I thank my best friend, Maja, for everything.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Introduction

For purposes of comparison and categorization, sets of stimuli can be described in terms of abstract features and dimensions. Evidence from the perception and categorization literature indicates that subjects are quite adept at such description. If the task set out by the experimenter requires it, subjects are both willing and able to decompose each item within a stimulus set into a collection of component parts and to evaluate each separate part on some dimension that varies systematically across the set. Even for sets in which the stimuli are said to be "unitary" (as opposed to "analyzable"; Shepard, 1964), subjects will identify and make use of dimensions of systematic variation (see for example Foard & Kemler Nelson, 1984). As Lockhead (1966, p. 104) states, "almost certainly Ss can analyze any set of multidimensional stimuli."

Ideally, the end result of this description process will be an organized, structural representation of the stimulus set that is well suited to analytic activities such as hypothesis testing. In order to be economical for these purposes, the dimensions of variation that are identified in the description process should have certain properties. First, the dimension of variation itself should be identified easily using familiar verbal description (e.g. size, shape). Second, the potential values on the dimension should be described discretely (often dichotomously), again using familiar verbal descriptions (e.g. large/small, round/square). Finally, the values on a given dimension should be selected such that they are more or less equally distributed across the stimulus set (e.g. 50-50 for dichotomous dimensions). Because of these special properties, the structural representation that results from a subject's analysis of the stimulus set might be thought of as resembling the arrays of 1's and 0's that are used to describe most stimulus sets in the literature (see, for example, Table 1).

Even under circumstances in which the subject is not thought to be actively analyzing the stimuli, however, it is often assumed that each stimulus within a set is being represented in terms of an analytic structure. L. B. Smith (1989; Smith & Evans, 1989), for example, suggests that although a represented object will be experienced consciously as a unitary entity, the object is nonetheless being represented as a set of constituent dimensions. Similarly, Estes (1986) makes the assumption that exemplars with which an individual has had experience are stored in the form of a vector of feature values. In other words, a subject's memory representation of a series of items would, again, take a form much like that seen in Table 1. Even when it appears that no systematic dimensions of variation exist, Estes argues, methods of multidimensional scaling can generate descriptions of the stimulus set in terms of dimensions in a psychological space, dimensions that can in turn be represented in an array.

In principle, a dimensional description could be theoretically neutral, simply a descriptive device necessary to allow the expression of whatever theory is being proposed. In practice, however, these dimensions tend to be treated by experimenters as real entities, as separate analytic dimensions on which subjects can selectively perform cognitive operations such as attention and/or differential weighting (for a similar argument, see J. D. Smith, 1989, in his discussion of weak vs.

strong holistic assumptions). Within the concept formation literature, the idea that the stimulus is represented as a set of separate, functionally independent dimensions has largely been treated as uncontroversial. What has been debated, instead, is the nature of the cognitive operations performed on the representation.

Traditional models of concept formation have emphasized the ability to summarize, across experience with individual members of a category, a general set of these separate features that are predictive of the category as a whole. This general set of predictive features, whether in the form of a rule (Bruner, Goodnow & Austin, 1956), a prototype (Franks & Bransford, 1971; Posner & Keele, 1968), or a feature frequency list (Hayes-Roth & Hayes-Roth, 1977), is then used as the basis for subsequent classification decisions. According to these models, then, category learning (and subsequent categorization) is analytic in at least two senses (Brooks, 1983). First, as mentioned earlier, it is analytic in its treatment of the stimulus, in that the stimulus is broken into a set of separate, independent components and is represented in terms of its analytic structure. Second, it is analytic in the process of concept formation in that an attempt is being made to analyze, across a variety of exemplars, the features that are relevant to category membership.

As an alternative to this position, there has been a growing emphasis on people's use of exemplar based, episodic classification procedures. The episodic approach suggests that much of our knowledge about a category is represented in the form of specific exemplars of the category. That is, items with which a person has had experience maintain their autonomy in memory, and category membership is simply an additional property of the individual item. Thus, categorical decisions are based not on the presence of features that are deemed analytically relevant to category membership, but on an item's similarity to specific previously encountered members of the category.

A number of models of episodic classification have been proposed (see, for example, Estes, 1986; Hintzman, 1984; Medin & Schaffer, 1978; Nosofsky, 1984, 1986; L. B. Smith, 1989). Although there are differences in the exact instantiation of each, all the models assume that the similarity between two items is some (non-linear) reciprocal function of the differences in their dimensional structures. Thus, according to these models, the classification of a new item is based on the number of separate features it shares with each previously encountered item. This approach to modelling episodic classification can be considered nonanalytic in its representation of the category, because it is specific item information, rather than analytic category level information that is supporting the process of categorization. It might be considered analytic, however, in its treatment of the stimulus, because the similarity judgement on which categorization is based is a function of the number of independent analytic features shared by the stimuli (see J. D. Smith, 1989).

Evidence suggests, however, that reference to the analytic dimensional structure of a stimulus set alone fails to capture an important source of variance found in subjects' classification behaviour. Allen and Brooks (1991), for example, found that after a short training session, subjects attempting to classify a set of cartoon animals (composed of 5 binary dimensions) showed evidence of instance based, episodic transfer whereas subjects presented with analytically equivalent verbal descriptions of the animals' dimensional structures did not. I propose that this discrepancy exists because the analytic structure fails to capture the individuality inherent in a perceptual stimulus. Unlike most experimental stimuli, many natural stimulus sets contain perceptual information that is idiosyncratic to individual stimuli within the set. This idiosyncratic information cannot easily be described (either by experimenters or by subjects) in terms of a small number of evenly distributed dimensions and, therefore, it is not likely to play an important part in subjects' analytic efforts during classification training Yet, such idiosyncratic information is likely to be important for learning and remembering individual stimuli. Because, as is argued by models of exemplar-based classification, the memory for an individual will often include information about category membership, this idiosyncratic information can be important for the purposes of classification. Furthermore, it can act not only as a basis of classification for old items, but also as a basis for the generalization of categorization decisions to stimuli that are judged similar to them.

As mentioned earlier, however, analytic effort is pervasive. Thus, although many natural stimuli are highly amenable to being dealt with in a nonanalytic, exemplar based manner during categorization tasks, this does not preclude analysis of these stimuli. As a result, two sources of information are available to subjects during the classification of many natural stimuli: *systematic* variations across the stimuli, which are likely to support analytic processes, and *idiosyncratic* variations in the stimuli, which are likely to support nonanalytic processes.

The major purpose of the research in this paper is to explore the stimulus variation that controls the relative contributions of analytic and nonanalytic processes during categorization. Within each experiment, the relevant stimulus sets were constructed from a single analytic structure similar to that shown in Table 1. As can be seen in this table, there are five binary-valued dimensions with equal frequencies of occurrence for each value on each dimension. This structure, however, may be manifested in a wide variety of perceptual forms, as can be seen most clearly in the comparison of the stimuli in Panels A and D of Figure 1. The dimensional structure of both these sets is represented in Subset 1 of Table 1. Although the dimensional structure of the two sets are identical there is a clear difference in the individuality of the items within a set. The animals in Panel A appear to be little more than the perceptual analog of the analytic 1's and 0's. The most striking characteristic of the set is the fact that, although any two animals chosen at random are clearly different from each other, there is a general uniformity that seems to exist across the set as a whole. By contrast, the animals in Panel D are highly varied perceptually and appear to be separate individuals. Although the analytic structures for this set are identical to those for Panel A, the manner in which each analytic structure manifests itself is perceptually distinctive.

The research presented here will address two sources of perceptual individuation that are not captured in the analytic structure of an item: 1) *Feature individuation*, the unique manner in which each separate feature is manifested in an individual item and 2) *Holistic individuation*, the manner in which the separate features composing an item cohere into an individuated whole. Each of these sources is represented as a dimension in the 2 x 2 format of Figure 1. Collectively, I will refer to these idiosyncratic variations as *nonanalytic variation* in contrast to the systematic analytic variation that is being kept more or less constant across the four stimulus sets. It is predicted that systematic variations in the dimensional structure underlying each stimulus set will capture subject's analytic activities, whereas idiosyncratic variations in the perceptual manifestations of that structure (as found in Panel D) will control the contribution of similarity to old items as a basis of categorization.

Table 1. The analytic structure of the stimuli used in Experiment 1. The various perceptual forms for the items in Subset 1 may be seen in panels A and D of Figure 1 (examples of the items from Subset 2 will be presented in Figure 2). Items from Subsets 1 and 2 alternately served as training and transfer items.

Item Number	Body Shape	Neck Length	Spots	Number of Legs	Leg Length
1.1	0	0	0	0	0
1.2	0	0	0	1	1
1.3	0	1	1	0	1
1.4	0	1	1	1	0
1.5	1	0	1	0	1
1.6	1	0	1	1	0
1.7	1	1	0	0	0
1.8	1	1	0	1	1
2.1	0	0	1	0	0
2.2	0	0	1	1	1
2.3	0	1	0	0	1
2.4	0	1	0	1	0
2.5	1	0	0	0	1
2.6	1	0	0	1	0
2.7	1	1	1	0	0
2.8	1	1	1	1	1
	0 = Round 1 = Angular	0 = Short 1 = Long	0 = No Spots 1 = Spots	0 = Two 1 = Six	0 = Short 1 = Long

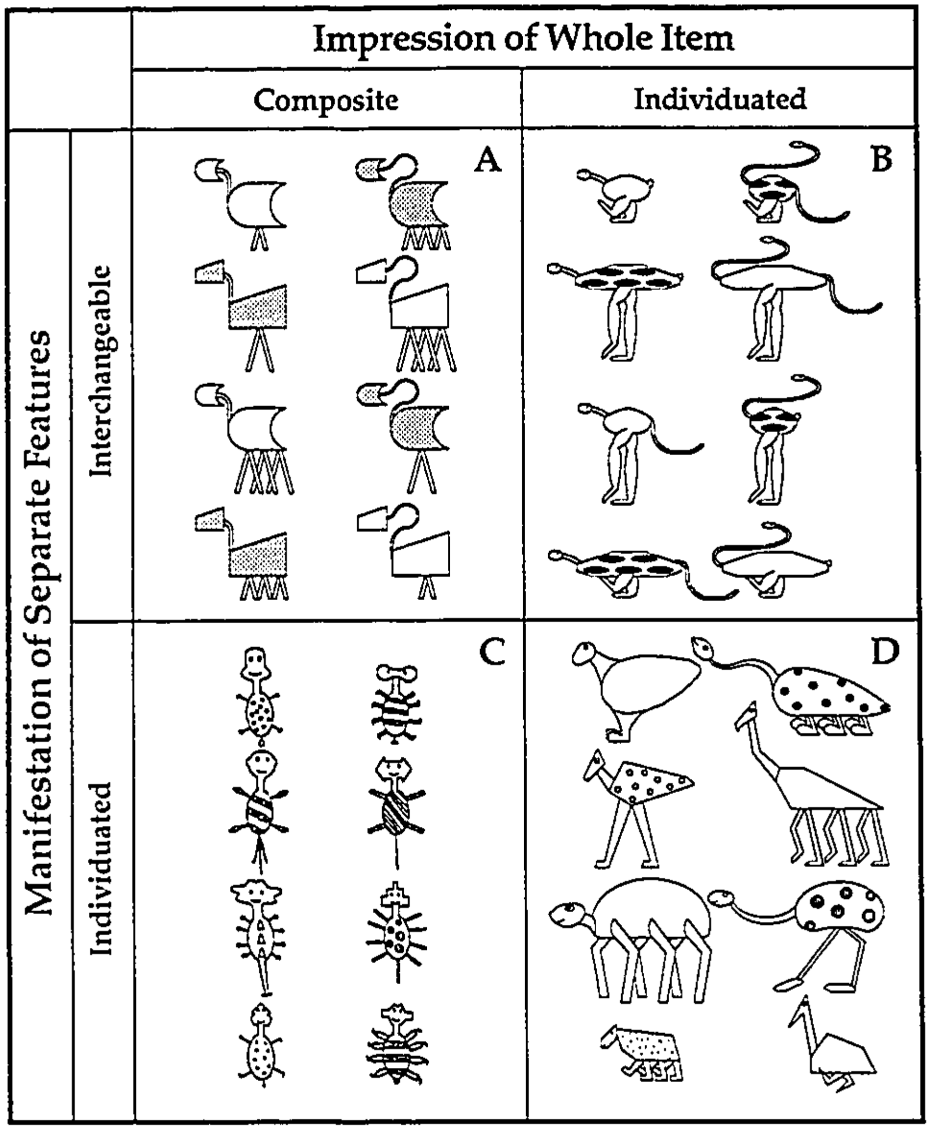

Figure 1. Stimuli for Experiment Sets 1 and 2. A similar logical structure of binary-valued dimensions underlies each set of animals. However, two aspects of the perceptual manifestation of the underlying structure are varied factorially. One factor, the perceptual manifestation of separate features, is whether or not a feature occurs in identical form in different items. The other factor, the impression of the item as a whole, is the extent to which an item's features cohere into an individuated whole.

Feature Individuation: The Unique Manifestation of Separate Features

In the vast majority of laboratory generated categories, the analytic structure is sufficient to identify each specific stimulus in the set. That is, if two items have the identical structure on dimensions that are convenient for analysis, then they are the same perceptual stimulus. For most categories outside the lab, however, a given value on an analytic dimension may have many perceptual manifestations. This phenomenon is most obvious when comparing exemplars from different subordinate categories. For example, eagles, hummingbirds and penguins all have wings, and in the analytic structure used for characterizing birds as opposed to other animals, each would receive a value of "1" on that dimension. Yet one hardly gets the impression that the wings are interchangeable perceptually.

The distinction between perceptually interchangeable and perceptually unique features is reflected in the two rows of Figure 1. For the two sets in the top half of the figure, features within a set are fully interchangeable. Within either set, for example, every pair of short legs is identical. Thus, for both sets, the separate features of a given animal provide no information beyond that present in the item's analytic structure. For the two sets in the bottom half of the figure, however, the perceptual manifestation of each separate feature is unique to an animal. One need only look at a single feature to identify an individual.

The unique variations across items within a category are often differences that, if a person were asked, could be described verbally. In the normal course of discussion about categorization, however, they are not likely to be mentioned. When conveying information about classification to a novice, for example, one is likely to describe only the analytic structure, trusting the novice to learn, through personal experience, the wide range of perceptual variation that the analytic description was meant to capture. Obviously, any attempt to describe all these variations would make the rules hopelessly unwieldy.

Yet, these unique variations in the perceptual manifestation of the analytic features will have an important impact on the degree of similarity between items within a category. Because of these perceptual variations, two items with the identical analytic structure might appear quite different, and two items with some differences in their analytic structures might appear more similar. Thus, the degree to which two items are similar in the perceptual manifestation of a feature (or set of features) could have an impact on exemplar based classification behaviour. This impact, however, would be invisible at the level of the analytic structure.

A similar point was made by Medin, Dewey and Murphy (1983) who suggested that exemplars are composed of category level information plus idiosyncratic information. One source of idiosyncratic information that they identified is the unique realization of category level information in specific exemplars. Medin, Wattenmaker and Hampson (1987) focused on this feature level idiosyncratic information when creating a set of perceptually distinctive artificial stimuli for their studies. The stimuli in Panel C of Figure 1, in fact, are adaptations of the stimuli used by Medin and his colleagues (Medin et al., 1987, Experiment 2b & 3).

Holistic Individuation: The Individuated Nature of the Item as a Whole

A second source of perceptual individuation, also likely to be lost in any analytical description of the stimulus, is related to the Gestalt perceptual tradition. Early Gestalt psychologists proposed that, for many stimuli, the perceived whole is different than the sum of its component parts. That is, the nature of the object as a whole cannot be captured by an analysis of its components (Kohler, 1929). In the extreme expression of this holistic form of perceptual individuation, the separate features are subsumed in the whole and have no immediate psychological status, referred to as strongly holistic processing by Kemler Nelson (1989).

The degree to which a holistic source of individuation is present in a stimulus set is manipulated in the columns of Figure 1. The animals in Panel B, for example, like those in Panel A, have interchangeable features. The features provide no unique information in themselves. Yet, for the animals of Panel B, each set of features seems to cohere into an individuated item. From these animals, one almost gets a feeling of "personality" that is simply lacking in the composite animals of Panel A. By contrast, the animals in Panel C, like those of Panel D, each possess unique perceptual manifestations of the analytic features. At the level of the separate features, each animal is unique. In the animals of Panel C, however, there is a certain consistency of construction across the set. The bodies are extremely similar, and the neck, tail and legs protrude from the same place on the body in each animal. These consistencies result in an overall feeling of sameness across the set, a feeling that the items are perceptually composite in their construction. To the extent that any "personality" exists in these animals, it likely arises from the unique perceptual manifestation of the head rather than as a result of any coherence of the features into an individuated whole.

The notion that the holistic coherence of a stimulus is important in classification has been examined directly in a number of contexts.

Using Saltz' (1972) notion of "boundedness", for example, Modigliani (1971, 1974; Modigliani & Rizza, 1971) examined the effect of stimulus integration in a standard classification paradigm. In one set of studies (Modigliani, 1971), for example, the stimuli were comprised of four binary dimensions (shape of flowerpot, shape of bloom, shape of leaves, and direction of stem). However, some subjects were trained and tested on stimuli that were organized into coherent pictures of flowers in pots, whereas other subjects were trained and tested on stimuli in which the same features were scattered randomly about the card. Modigliani first trained subjects to use one of the binary dimensions to classify stimuli into one of two categories. For the test phase, he gave the subjects the possibility of using an inbetween, "uncertain" category in addition to the two used in training. He then began altering the "logically" irrelevant dimensions: adding new values to the pre-existing dimensions, removing one or more of the pre-existing dimensions, adding new previously unseen dimensions, etc. Because the relevant dimension was unchanged in the new stimuli, it was possible to continue placing these stimuli into the appropriate category according to the previously induced rule. If, however, the categorically irrelevant alterations had an effect on subjects' categorization behavior, then stimuli containing these alterations would be placed into the new (grey) category, despite the presence of the defining characteristic for the previously learned categories. Modigliani found that alterations in the logically irrelevant dimensions had a far more disruptive effect on classification for subjects viewing the coherent stimuli than for subjects viewing the randomly scattered stimuli. Thus, when the stimuli were coherent, subjects became uncertain about the classification as analytically irrelevant changes were made.

Using a similar manipulation, Brooks (1976) produced an artificial alphabet that was designed such that words could be constructed either in "discrete" form or in "glyphic" form. A discrete construction was a horizontal listing of the characters that made up the word (much like the lexical constructions used in English spelling). A glyphic construction was the same set of "letters" in the same order, but presented vertically rather than horizontally. Both at the level of the analytic structure and at the level of the separate perceptual features, therefore, the two constructions contained an equal amount of information about the word. However, the letters were designed such that the horizontal constructions appeared visually as a list of discrete symbols whereas the vertical constructions integrated into visually distinct wholes. When subjects were asked to respond rapidly with a unique response to each item (i.e. to pronounce each word) the glyphic constructions produced faster responding than the discrete constructions both initially and after considerable practice. These results suggest that the integrated visual organization of the glyphic stimulus set was providing information at the level of whole items that was not being captured by the set of discrete features alone.

Clearly, then, the holistic quality of the items within a stimulus set can have an impact on subjects' classification behaviour. The point being asserted by Brooks and his colleagues (Allen & Brooks, 1991; Brooks, 1978, 1987, 1990; Jacoby & Brooks, 1984; Whittlesea, 1987; Vokey & Brooks, 1991) is that this holistic, nonanalytic quality can result in sufficient individuation to create a basis for the classification of the item itself and new items that are judged similar to it. Because this source of individuation will be a property of the item as a whole rather than a function of the independent features that compose the item, however, its presence or absence will, at best, be inconsistently represented by the analytic structure.

The Coordination of Analytic and Nonanalytic Resources

Smith and Medin (1981) have argued that if instances of a concept share a large number of properties, then that concept will likely be represented in probabilistic form, whereas if the instances share few properties, the concept is likely to be represented in exemplar form. This implication that the representation of a concept is *either* probabilistic *or* exemplar based, suggests that the acquisition and use of analytic and nonanalytic sources of information are reciprocally related. A similar implication is made in discussions in which holistic processing is characterized as being developmentally or cognitively more primitive than is analytic processing. It has been suggested, for example, that younger subjects or subjects acting under heavy cognitive load are more likely to process stimuli holisitically and correspondingly less likely to process the stimuli analytically (for a critique of these ideas with respect to categorization, see J. D. Smith, 1989, and Smith & Shapiro, 1989). An expectation of reciprocal relations between the two is further implied by experimental procedures in which a subject's choice of transfer stimuli is interpreted as being consistent with *either* analytic *or* holistic processing of the stimuli (Kemler Nelson, 1984; Smith & Kemler Nelson, 1984).,

The idea that these two modes of processing are in a compensatory relationship likely arises, at least in part, from the assumption that the processing of a stimulus, whether analytic and holistic, is performed on a dimensional structure. If, as implied by the notion of strong holistic processing, the separate features have no immediate psychological status, then concurrent analytic processing of these separate features cannot take place. By contrast, I propose that the information that most easily supports nonanalytic processing of a stimulus may not be well represented in the analytic structure of the stimulus. Instead, for many categories outside the laboratory, this information is found in the specific perceptual manifestation of the analytic structure. Thus, expert classifiers may use overlapping properties in the items' analytic structures to produce a useful probabilistic representation of the concept, while using the individuating perceptual manifestations of those analytic structures for nonanalytic, exemplar based categorization procedures.

This proposal does not negate the possibility that the acquisition and use of analytic and nonanalytic processes tends to be in a compensatory relationship. It is possible, for example, that applying or discovering an analytic rule might tend to "break up" a stimulus so that the holistic properties of the stimulus are less apparent. Under these circumstances, then, a person might be less likely to acquire and make use of this holistic information during categorization. On the other hand, it is possible that, with practice, responding to the holistic or distinctive feature properties of the stimuli is faster and less effortful than is rule application. Thus, if a beginner is given a rule for categorizing unfamiliar stimuli, the featural or holistic distinctiveness of the stimuli would not be exploited until a great deal of practice had occurred, at which time the analytic processing would have less control. Clearly, then, it is possible to imagine situations in which the acquisition and use of analytic and nonanalytic information would be in a compensatory relationship. As formulated here, however, this is an empirical issue rather than a logical necessity.

In fact, research that demonstrates at least some conditions in which analytic and nonanalytic categorization occur concurrently has been reported by Allen and Brooks (1991). Subjects in their experiments were given an easy, perfectly predictive classification rule, followed by training in applying that rule to a set of practice items. On a subsequent transfer test, the accuracy and speed of classifying new items was strongly affected by similarity to previously seen items. New items that were similar to old items in the same category (good transfer items) were responded to approximately as rapidly and accurately as the old items themselves. New items that were similar to old items in the opposite category (bad transfer items) produced slower response times and a much higher rate of error. If classification was solely on the basis of a speeded application of the rule, then one would not expect more errors to the bad transfer items than to the good transfer items; in both cases perfect classification could be achieved by attending to the relevant analytic features. On the other hand, the rule had not simply been abandoned. The subjects made only 45% errors on the bad transfer items rather than closer to 100% as would be expected if they were relying only on similarity to prior exemplars. Thus, in these experiments, knowledge of the rule helped to maintain accuracy, but it did not prevent facilitation or interference from item-specific and context-specific knowledge.

The research presented in the present thesis used the same general method as that employed by Allen and Brooks to examine how the information available in the analytic structure and the information available in the perceptual manifestation of that structure contribute to performance on a classification task. This was done by systematically manipulating the presence or absence of each source of perceptual individuation within the stimulus set presented to subjects, as depicted in Figure 1.

In addition, I varied whether or not the subjects were given the classification rule before training. There is good reason to believe that whether or not the subjects have a rule could affect the way in which they organize the stimuli. In the studies presented by Allen and Brooks (1991), one group of subjects were not given the rule but were given training trials with feedback on the correct category for each of the training items. These "No Rule" subjects made over 85% errors on the bad transfer items, indicating they were classifying the transfer stimuli largely by analogy to specific training items. But in addition to responding to the holistic or distinctive feature properties of the items, these subjects also reported trying to analyze the stimuli to discover a classification rule. It is possible that comparable No Rule subjects would respond differently to the various types of stimuli used in this thesis than would subjects who knew the classification rule from the beginning. Regardless of how these subjects actually do respond, however, the Rule and the No Rule groups represent two important conditions for categorization studies. The No Rule (rule discovery) procedure is by far the most common method used in studies of classification. But, as Allen and Brooks (1991) have argued, providing individuals with a set of classification rules reflects many training conditions found in natural teaching situations. Whether or not subjects are provided with the rule for classification, therefore, was another manipulated variable in the following studies.

CHAPTER 2

The Initial Demonstration

The first study was an attempt to replicate the Allen and Brooks (1991) finding of episodic transfer effects using a new set of materials, and to provide a clear contrast case in which there are no episodic transfer effects. Allen and Brooks contrasted their perceptual stimuli (drawings) with a set of verbal feature lists. When the same analytic information was given during both training and test in the form of a list of verbal descriptions, there was no sign of episodic transfer effects. Apparently, at least when the subjects were given a sufficient classification rule, there were no effects of similarity to old items unless those items were perceptually and mnemonically distinctive. In Experiment 1 my interest lay in producing a set of *perceptual* stimuli, namely the drawings in Panel A, in which the episodic transfer effect would be clearly absent. It was with this contrast in mind that the stimuli in Panels D and A from Figure 1 were created.

Materials

The stimuli were line drawings of imaginary animals. The analytic structure of the stimuli included five orthogonal binary dimensions: body shape (round or angular), neck length (long or short), spots (present or absent), number of legs (two or six), and leg length (long or short). A set was composed of sixteen stimuli constructed in pairs (one member for each of two subsets of items). Within each pair, the two stimuli differed only on the analytic dimension of spots. Across the pairs, each stimulus differed from every other on at least two analytic dimensions. The analytic structure of the sixteen animals composing a set may be seen in Table 1.

The stimuli in panels A and D are the perceptual constructions resulting from the analytic structures of subset 1 (the top half) of Table 1. The set of animals shown in panel D were designed with the intention that both sources of perceptual distinctiveness be present. In each animal the perceptual manifestation of the features was intended to be unique. Furthermore, each animal was meant to cohere an individuated whole. In contrast, the animals in panel A were designed such that neither source of perceptual distinctiveness was present. Here the features are completely interchangeable and, although the features clearly combine into animals, each stimulus gives the appearance of a composite structure rather than giving any strong sense of individuality.

The preliminary step is to determine if the addition of perceptual distinctiveness to the animals of panel D fundamentally altered their analytic affordances. It is possible that the most natural analytic contrasts among these animals are not the analytic dimensions that were intended but rather some new set of dimensions. Alternatively, the addition of perceptually idiosyncratic information might have changed the extent to which people rely on any consistent dimensions. With these stimuli, the subjects might use some "family resemblance" structure, involving overall similarity, that does not consistently use any dimensions. In Experiment 1A, subjects were told that the animals with which they were presented could be divided into two categories and were asked to give any potential bases that they could see on which this classification could be made. In Experiment 1B, following Medin, Wattenmaker and Hampson (1987), subjects were asked to group the stimuli into whatever categories seemed natural to them. Medin et al., using stimuli similar to those in panel C, found that subjects tended to idiosyncratically center on a single analytic dimension to form their categories, a strategy that I will call "analytic simplification." The question for Experiments 1A and 1B is whether the addition of the perceptually distinctive information in panel D changed the dimensions available to the subjects or their tendency to use any one dimension. In other words, are the analytic dimensions from which the stimuli were originally created also the ones used by the subjects? Subsequently, a group that was not given the classification rule (Experiment 1C) and a group that was (Experiment 1D) were tested to see if the stimuli in Panels A and D differ in the extent to which the nonanalytic properties of the items contribute to classification performance.

Experiment 1A: The reported analytic affordances of the stimuli

Method

Subjects: The subjects were 16 undergraduate students from McMaster University participating for course credit.

Procedure: Each of sixteen stimuli (the eight shown in panel A of Figure 1 and the eight shown in panel D of Figure 1) was placed on a separate 3 x 5 inch file card. Each subject received either the eight animals from set A or the eight animals from set B. The eight cards were laid in front of a subject in a random arrangement and the subject was told:

> The eight animals in front of you can be divided into two groups. If you were told that you had to figure out how the animals divided into the two categories, what would you do?

What would you test as potential reasons for dividing the animals? Please name anything you can think of.

During responding, subjects were repeatedly prompted to name any additional methods for dividing the animals that they could think of. Subjects' responses were recorded.

Results

The responses from subjects viewing the perceptually uniform set of animals were extremely consistent. Six of the eight subjects identified the five dimensions present in the analytic structure of the stimuli and no others. The remaining two subjects identified the five dimensions, but in addition identified the shape of body and shape of head separately (these two dimensions are perfectly correlated). One of these two subjects subsequently recognized the confound and indicated that it would not be worthwhile checking the head and body separately.

Of the eight subjects viewing the perceptually distinctive set of animals, seven identified all five dimensions of variation from the analytic structure; the eighth subject failed to identify the length of neck as a potential dimension. Three subjects also identified shape of head (again, subjects' descriptions of this dimension were fully confounded with angularity of body). Two additional sources of systematic variation were identified (each by one subject): the presence or absence of a mouth, and whether or not the eye was filled.

Discussion

Generally, the two sets of stimuli do not seem to differ radically in their analytic affordances. For both sets, all five of the original dimensions seem apparent to subjects. Furthermore, the tendency for subjects to find additional sources of systematic variance in the perceptually distinctive stimuli is small. In this experiment, however, subjects were not actually asked to classify the stimuli. Instead, subjects were simply encouraged to treat the items as analytically as possible. It would be worthwhile, therefore, to explore whether subjects continue to use these dimensions when instructions to analyze the stimuli are replaced with instructions to classify the stimuli.

Experiment 1B: The analytic affordances used during free classification

The design of Experiment 1B was consistent with the general design presented by Medin et al. (1987). That is, subjects were presented with a set of stimuli and were asked to divide them into two categories in whatever way seemed natural to them. Medin et al. found that subjects tended to focus on a single analytic dimension when forming their categories. If subjects viewing the current stimuli also engage in this kind of analytic simplification, the free sorting technique should reveal the dimensions that subjects actually use for the purposes of classification.

Method

Subjects: The subjects were 32 undergraduate students from McMaster University participating for course credit.

Procedure: Each of sixteen stimuli (the eight shown in panel A of Figure 1 and the eight shown in panel D of Figure 1) was placed on a separate 3 x 5 inch file card. A subject received either set A or set D in a randomly scrambled order. The subject was asked to lay out the examples, look them over carefully, and place the animals into two categories in any way that seemed appropriate. No restriction was placed on the number of animals that were in each category. Following the categorizations, subjects were asked to describe the criteria by which they classified the stimuli.

Results

In both sets of animals, 15 of 16 subjects reported using a single analytic dimension to classify the animals into the two categories. The breakdown of features reported for both the perceptually uniform animals of panel A and the perceptually distinctive animals of panel D may be seen in Table 2.

Table 2. The frequency with which each feature was named as the sole criterion for classification when no feedback was provided (Experiment 1B). Perceptually Uniform and Perceptually Distinctive conditions refer to panels A and D of Figure 1.

Perceptual Form	Body Shape	Neck Length	Spots	Number of Legs	Leg Length	Multiple Features
Uniform	7	2	6	0	0	1
Distinctive	5	1	3	4	2	1

The remaining subject viewing the perceptually uniform animals reported using a conjunctive combination of body shape and spots such that one category was defined as requiring both a round body and no spots. The remaining subject viewing the perceptually distinctive set reported using a combination of the number of legs and the length of neck, although the exact instantiation of this combination was unclear. With the exception of one stimulus, this subject placed all stimuli into groups consistent with number of legs. Similarly, one of the subjects viewing the perceptually uniform animals reported using neck length as the criterion, but placed one stimulus with a long neck in the inappropriate category. With the exception of these two subjects, all items were categorized in a manner consistent with the subject's verbal report.

Discussion

Consistent with Medin et al. (1987), the present results indicate that subjects who are asked to classify items without the benefit of feedback will tend to use strongly analytic strategies. That is, the tendency is to use a single analytic dimension as the criterion for classification. Furthermore, this analytic strategy is maintained despite the presence of perceptually idiosyncratic properties. What is more relevant to the present discussion, however, is that subjects viewing the perceptually distinctive items used only dimensions that were present in the original analytic structure. It appears that no new analytic dimensions emerged in these stimuli during actual classification.

Experiment 1C: The use of nonanalytic affordances in "No Rule" training

Although subjects tend to use only the analytic affordances of the stimuli during a classification task in which they create the categories, in many natural (and experimental) situations the subject's job is not to create the categories, but to learn them. In the majority of experimental designs this learning arises by way of feedback about subjects' decisions. Experiment 1C was performed to determine the degree to which this category learning procedure induces subjects to make use of the nonanalytic information to support generalization to new items. As in the experiments by Allen and Brooks (1991), the transfer phase included items that were similar to old items in the same category (good transfer items) and new items that were similar to old items in the opposite category (bad transfer items). Because these items were balanced for all of the analytic dimensions, differential accuracy on these items is an index of whether the subjects were using similarity to old items, be it on the basis of distinctive individual features or holistic properties, to support categorization.

Method

Subjects: The subjects were 32 undergraduate students from McMaster University participating for course credit.

Design: The sixteen animals composing a set were divided into two categories (builders and diggers) on the basis of a three-feature additive rule. That is, three of the five dimensions were designated as categorically relevant dimensions and two of the dimensions were designated as categorically irrelevant dimensions. For the three categorically relevant dimensions, one of the values was designated a builder feature, and the other value was designated a digger feature. If an animal possessed a majority of builder features, it was a member of the builder category. If it possessed a majority of digger features, it was a member of the digger category. All eight possible combinations of the binary values for the relevant features were present within each subset of eight animals, so a given feature classification was consistent with the classification of the animal in 75% of the cases. The values of the two irrelevant dimensions appeared equally often in animals of each category and therefore were definitionally nondiagnostic.

Four different rules were used to categorize the animals. The features associated with builders for each of these four rules were:

1) Long legs, Angular body, Spots present

2) Short legs, Long neck, Spots present

3) Six legs, Angular body, Spots present

4) Two legs, Long neck, Spots present

These rules were chosen in order to counterbalance four of the five dimensions across relevance to classification. However, the presence or absence of spots (the only dimension that separates the two members of a stimulus pair) is always a relevant dimension. This creates four critical conditions:

1) Good Old (GO): A training item whose twin transfer item will be in the same category.
2) Bad Old (BO): A training item whose twin transfer item will be in the opposite category.
3) Good Transfer (GT): A transfer item seen for the first time in the test phase that (according to the three feature additive rule) is in the same category as its twin training item.
4) Bad Transfer (BT): A transfer item seen for the first time in the test phase that (according to the three feature additive rule) is in the opposite category to its twin training item.

Examples of the four conditions may be seen for each set of animals in Figure 2. The four rules and the counterbalancing of training and transfer subsets ensured that each of the sixteen stimuli appeared equally often in each of the four critical conditions. For the purposes of analyses, the GO and BO items were treated as a single set of eight Old items.

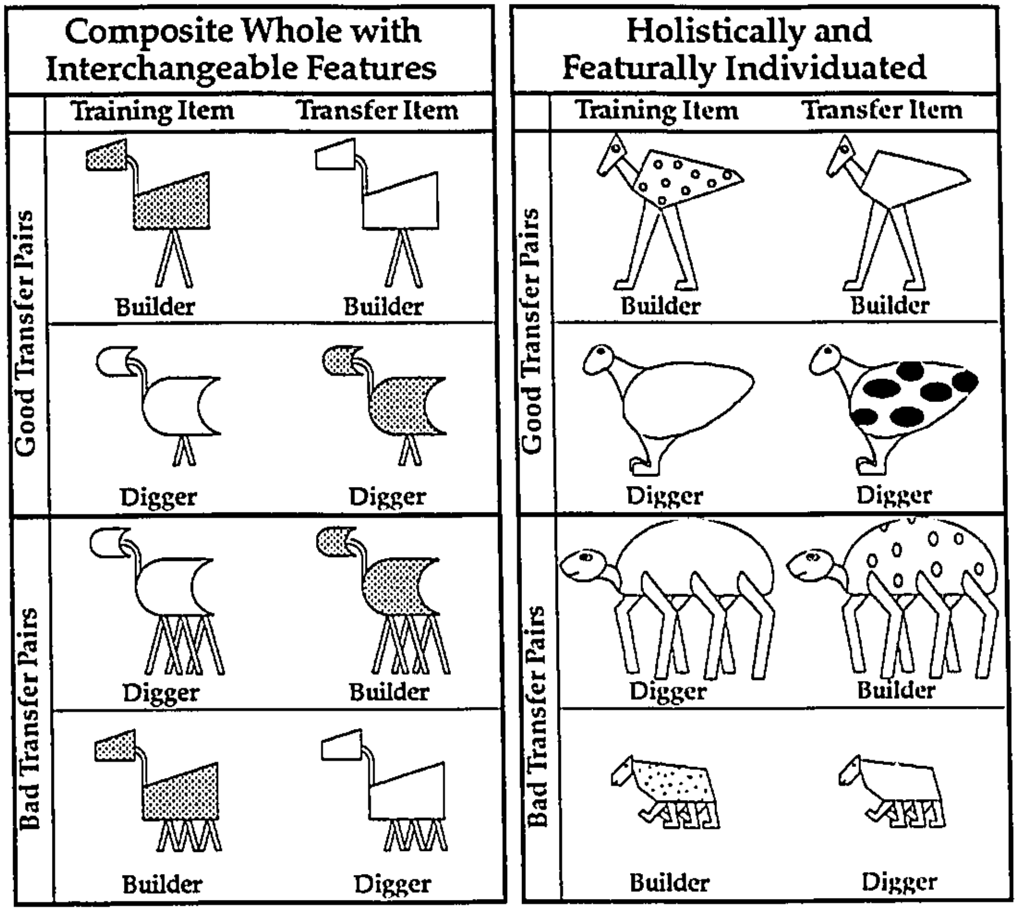

Figure 2. Examples of the transfer conditions used in all of the experiments (shown here for each of the stimulus sets used in Experiment 1). In this example, the category labels are assigned according to the categorization rule: "At least two of long legs, angular body and spots indicates a builder". For Good Transfer pairs the transfer item is in the same category as its training twin according to the analytic rule. For Bad Transfer pairs, the transfer item is in the opposite category to its training training twin according to the analytic rule. A difference in errors and classification time between Good and Bad Transfer items is interpreted as showing an effect of similarity to old items.

As argued by Allen and Brooks (1991), differences in performance on GT and BT items can shed light on the nature of the information that subjects acquire and make use of during classification training. The GT items operate largely as a baseline condition. Perceptually, each of these new items is very similar to an Old item that, according to the analytic rule, is in the same category. If the GT items are classified as rapidly and as accurately as the Old items, then we have some indication that subjects can generalize their classification knowledge from the Old items, on which they have had training, to new items, with which they have had no prior experience. The rapid and accurate classification of GT items, however, indicates little about the type of classification knowledge being generalized. For GT items, as for Old items, the analytic and nonanalytic sources of information are leading to the same category decision. Thus, the use of either analytic or nonanalytic sources of information will result in accurate classification.

By contrast, BT items are designed such that the analytic and nonanalytic sources of information are in opposition. Perceptually, each of the BT items, like each of the GT items, is very similar to a single Old item. Yet, because of the change in the analytic dimension of spots, each BT item is in the category opposite to the Old item it resembles. Thus, nonanalytic categorization procedures based on similarity will lead to one decision, whereas the analytic application of the rule will lead to the opposite decision. This conflict between the analytic and nonanalytic categorization processes is likely to be manifested both in the accuracy with which BT items are classified (according to the analytic rule) and in the speed with which accurate decisions are made. Thus, if subjects are using analytic categorization procedures exclusively, then the BT items should be classified as rapidly and as accurately as the GT items. If subjects are using nonanalytic categorization procedures exclusively, then BT items will be placed consistently in the category opposite to that dictated by the rule (the error rates will be very high compared to GT items), but the process of categorization should be as rapid for the BT items as for the GT items. If both nonanalytic sources of information and analytic sources of information are having an impact on classification behavior, however, then the classification of BT items should be slower and only somewhat more error prone than the classification of GT items.

Procedure: The materials were presented on the screen of an Apple MacIntosh II. Subjects' categorical decisions were made by pressing one of two keys on the MacIntosh keyboard and the elapsed time from the presentation of the stimulus until the response of the subject was measured. Each subject was tested individually. Subjects were either in the perceptually uniform group or the perceptually distinctive group, which differed only in the type of items that subjects were asked to classify. Each subject was told that there were two categories of items and that the subject's job was to try to place each item into the appropriate category. It was suggested to subjects that they would likely just be guessing to begin with, but that the computer would be correcting their decisions so that, with time, they would become better at placing items in the correct categories. These instructions were meant to be as neutral as possible with respect to analytic or nonanalytic methods of classification. The subject was then presented with a set of eight training items (the GO and BO items) repeated five times. The order of items was randomized within each repetition. For each of the 40 trials, subjects were to classify the item as quickly as possible without sacrificing accuracy. Feedback was given after each classification.

The transfer phase contained the eight old items and the eight transfer items. Each item was seen once. The order of the items was random except that a transfer item was always separated from its old twin by at least two items. Subjects were informed that the eight old items would now be supplemented by eight new items and were asked to continue classifying as they had been all along. No feedback was given during this phase.

Results

Analyses were performed both on the proportion of errors and on the subjects' median response times[1] for correct classifications made during the transfer phase of the experiment. An error in classification was identified as the failure to place the animal in the category consistent with the three-feature additive rule. These analyses included two 2 x 3 ANOVAs with one between-subjects factor, perceptual form (Uniform vs. Distinctive), and one within-subjects factor, transfer type (Old vs. GT vs.

---

[1]As noted by Allen and Brooks (1991), there is a confound in the comparison of Good Transfer and Bad Transfer items, in that half of the Good Transfer items possess all three features consistent with their category membership, whereas this is true for none of the Bad Transfer items. However, their findings were the same whether only items with two consistent features or all the items were included in the analysis. Analysis of data from the experiments presented throughout this paper repeatedly failed to find significant differences between Good Transfer items possessing three features consistent with category membership and Good Transfer items possessing only two features consistent with category membership. Thus, the analyses presented throughout this paper will include all items in the interest of obtaining more stable estimates.

BT). A summary of the data is presented in Table 3. A criterion p value of 0.05 was used for all analyses.

Table 3. Mean responses to the perceptually uniform stimuli (panel A of Figure 1) and the perceptually distinctive stimuli (panel D of Figure 1) when no explicit rule was provided to aid subjects' classifications (Experiment 1C).

PROPORTION ERRORS

PERCEPTUAL FORM	OLD	GT	BT	GT - BT Effect
Perceptually Uniform	.320 (.144)	.469 (.196)	.328 (.220)	N. S.
Perceptually Distinctive	.109 (.128)	.187 (.194)	.766 (.232)	*

MEDIAN RESPONSE TIMES (ms)

PERCEPTUAL FORM	OLD	GT	BT	
Perceptually Uniform	1173 (721)	1140 (770)	1091 (628)	N. S.
Perceptually Distinctive	847 (520)	1123 (676)	1090 (480)	N. S.

Errors: Analysis of the error data revealed no main effect of perceptual form, $F(1,30)=1.06$, MSe=0.05, but there was a main effect of transfer type, $F(2,60)=35.12$, MSe=0.04 and an interaction of perceptual form with transfer type, $F(2,60)=15.67$, MSe=0.04. To explore these results further, two separate one-way ANOVAs were performed, one for each of the two perceptual forms. Analysis of errors in the perceptually uniform condition revealed no significant effect of transfer type, $F(2,30)=2.52$, MSe=0.04. A subsequent single sample t-test, which collapsed across the three transfer types, revealed that proportion of errors for subjects viewing the perceptually uniform stimuli was significantly below chance, $t(15)=3.33$. For the perceptually distinctive condition, there was an effect of transfer type, $F(2,30)=53.22$, MSe=0.04. A subsequent set of paired t-tests revealed no reliable difference between Old and GT items, $t(15)=1.77$, but did indicate a significantly greater proportion of error rates on BT items than on GT items, $t(15)=6.60$.

Response Times: An analysis of the correct response times revealed no main effect of perceptual form, $F(1,30)<1$, and no main effect of transfer type, $F(2,60)=1.39$, MSe=85387. There was, however, a significant interaction of perceptual form with transfer type, $F(2,60)=3.72$, MSe=85387. Subsequent one way ANOVAs performed for each of the perceptual forms, however, revealed no significant effects of transfer type either for the perceptually uniform condition, $F(2,30)=1.48$, MSe=48405, or the perceptually distinctive condition, $F(2,30)=2.98$, MSe=122369.

Discussion

Frequently, subjects faced with stimuli that contain no contrastive perceptual information beyond that found in analytic description showed evidence of analytic classification behavior. That is, the accuracy rate on the 16 items in the transfer phase was reliably better than would be predicted by chance. Furthermore, these subjects seemed to apply this analytic, category level information consistently across all items (old and new alike). These conclusions are lent support by anecdotal evidence collected during post-experimental interviews. During these interviews, subjects were asked the basis for their categorical decisions during the transfer phase. All subjects presented with the perceptually uniform items reported attempts, during training, to establish category membership on the basis of a single analytic feature. It should be noted that although no single feature is perfectly predictive, the three-feature additive rule used in this study ensures that the use of a single relevant feature in isolation will result in 75% accuracy. Although all subjects indicated a lack of satisfaction with any single feature, a number of subjects viewing the perceptually uniform stimuli reported using a single-feature rule in the transfer phase regardless of this dissatisfaction. Consistent with these reports, a review of the data indicated that for nine of the sixteen subjects in this condition, 100% of their classification responses could be accounted for by reference to a single analytic feature. For seven of these subjects, the feature chosen was one of the three features relevant to category membership. When these nine subjects were removed from the analysis, the the seven remaining subjects showed a mean error rate of 0.482, which did not differ reliably from chance. Data from this condition, then, support the findings of the early literature on hypothesis testing, which indicate that when the stimuli are perceptually uniform, subjects' attempts at classification tend to be conscious and analytic (e.g., Bruner, Goodnow & Austin, 1956).

By contrast, when stimuli possessed perceptual distinctiveness (both in the unique perceptual manifestations of the separate features and in the form of holistic individuation), this distinctiveness seems to have had a strong impact on classification behaviour. First, memory for the category of old items was quite good. Furthermore, compared to the perceptually uniform condition, there was a huge rise in the error rate for BT items and a corresponding drop in the error rate for GT items. These results are consistent with the use of a nonanalytic classification strategy based on the specific similarity of the new item to a previously seen item. These conclusions are again supported by the post-experimental interview. When subjects viewing the perceptually distinctive animals were asked how they classified animals in the transfer round, the majority indicated that they classified the item by placing it in the same category as the old item it resembled. When asked how old items were classified, subjects generally indicated that they simply memorized the items during training. Consistent with these post-experimental statements, none of the subjects viewing the perceptually distinctive stimuli displayed a response pattern that could be explained with reference to a single feature.

This is not to say, however, that the perceptually distinctive stimulus condition eliminated all efforts at analysis. When explicitly asked, all subjects indicated that their classification attempts in the early training rounds were characterized by attempts to establish category membership on the basis of a single analytic feature. As a result, these subjects could, when asked, identify many of the analytic features that were varying and could, in many cases, identify at least one feature that was correlated with category membership.

Thus, the findings of this study are, again, consistent with those of Medin et al. (1987). First, conscious analytic decisions about category membership seem to be based not on a "family resemblance" principle, but rather on very simple, single dimensional rules. In fact, these data suggest that this tendency toward analytical simplification is sometimes powerful enough to persist even when repeated feedback indicates that a single feature is insufficient for accurate classification. Furthermore, anecdotal evidence from the post-experimental interviews suggests that this tendency for analytic behaviour to focus on single features in isolation is present both in perceptually uniform and perceptually distinctive stimuli. The present data also suggest, however, that when idiosyncratic perceptual information is useful as an aid to classification, this information is likely to take precedence over subjects' tentative analytic information.

Experiment 1D: The use of nonanalytic affordances in "Rule" training

Allen and Brooks (1991) have shown that the effect of similarity to prior training items also occurs when subjects' analytic categorical information is in no way tentative. Prior to training, they provided subjects with the three feature additive rule that would result in perfect classification. Following a fairly short training session (that is, a short period of practice in applying the analytic rule) subjects showed evidence of nonanalytic, single item transfer. In Experiment 1D, I will attempt to replicate this finding with the perceptually distinctive animals of panel D. Furthermore, I will attempt to contrast this nonanalytic effect in the perceptually distinctive animals with a clear lack of the effect in the perceptually uniform stimuli of panel A.

Method

Subjects: The subjects were 32 undergraduate students from McMaster University participating for course credit.

Procedure: The procedure of Experiment 1D was identical to that of Experiment 1C except that, prior to the presentation of animals in the training round, subjects were presented with the three-feature additive rule that would result in perfect classification.

Results

Analyses were performed both on the proportion of errors and on the subjects' median response times for correct (rule consistent) responses made during the transfer phase of the experiment. These analyses involved two 2 x 3 ANOVAs with one between-subjects factor, perceptual form (Distinctive vs. Uniform) and one within-subjects factor, transfer type (Old vs. GT vs. BT). A summary of the data is presented in Table 4.

Errors: An analysis of error data revealed that subjects viewing the perceptually distinctive stimuli were, in general, more likely to make errors in classification than subjects viewing the perceptually uniform stimuli, $F(1,30)=18.78$, MSe=0.02. Furthermore, there was a significant main effect of transfer type, $F(2,60)=5.00$, MSe=0.03. These main effects, however, must be interpreted in light of a significant interaction between perceptual form and transfer type, $F(2,60)=4.30$, MSe=0.03. A series of one-tailed paired t-tests confirmed the earlier predictions. For the perceptually uniform animals, Old items showed no advantage over GT items, $t(15)=0.62$, and GT items showed no advantage over BT items, $t(15)=0.00$. For the perceptually distinctive animals, again, there was no advantage for Old items over GT items, $t(15)=0.19$. The classification of the distinctive BT items, however, was significantly more error prone than the classification of the distinctive GT items, $t(15)=2.11$.

Table 4. Mean responses to the perceptually uniform stimuli (panel A of Figure 1) and the perceptually distinctive stimuli (panel D of Figure 1) when the rule resulting in perfect classification was provided prior to classification training (Experiment 1D).

PROPORTION ERRORS

PERCEPTUAL FORM	OLD	GT	BT	GT - BT Effect
Perceptually Uniform	.031 (.056)	.047 (.101)	.047 (.136)	N. S.
Perceptually Distinctive	.094 (.109)	.109 (.153)	.328 (.212)	*

MEDIAN RESPONSE TIMES (ms)

PERCEPTUAL FORM	OLD	GT	BT	
Perceptually Uniform	863 (295)	894 (420)	924 (292)	N. S.
Perceptually Distinctive	894 (324)	1018 (414)	1295 (396)	*

Response Times: Analysis of subjects' median response times for correct responses revealed a very similar pattern to that found in the error data. Although there was no significant main effect for perceptual form, $\underline{F}(1,30)=1.77$, MSe=415645, there was a significant main effect for transfer type, $\underline{F}(2,60)=6.31$, MSe=70311. Again, however, the main effect must be interpreted in light of an interaction between perceptual form and transfer type, $F(2,60)=3.51$, MSe=70311. A series of one tailed paired t-tests revealed a pattern of results consistent with the error data. For the perceptually uniform animals that were correctly categorized, there was no significant difference in response times between Old items and GT items, $t(15)=0.41$ and no significant difference in response times for GT items and BT items, $t(15)=0.52$. For perceptually distinctive animals that were classified correctly, the difference in response times between Old items and GT items was not reliable, $t(15)=1.66$. However, the correct classification of BT items was significantly slower than the correct classification of GT items, $t(15)=1.78$.

Discussion

Consistent with Allen and Brooks (1991), Experiment 1D provides evidence that, despite subjects' awareness of a perfectly predictive analytic rule for categorization, idiosyncratic information in the perceptual manifestation of the analytic structure can lead to more nonanalytic, exemplar based categorization. If an item was perceptually similar to a previously learned item that, according to the analytic rule, was in the opposite category, then efforts to classify the new item were slower and more error prone. As in the Allen and Brooks data, it appears that the rule for classification was still present and operating. The error rate on BT items was substantially lower in Experiment 1D than in Experiment 1C, in which no rule was provided for subjects. However, the use of the rule seems to have been modified to accommodate the idiosyncratic information provided in the perceptually distinctive stimuli.

These effects cannot be a function merely of items' similarity at the level of the analytic structure. Evidence for similar effects in the perceptually uniform animals of panel A is clearly lacking. In these animals, there were no reliable differences between GT and BT items either in classification accuracy or in classification speed. In fact, there were not even any reliable differences between the transfer items and the items on which subjects had been given previous classification training.

This set of studies, then, demonstrates the importance of perceptually idiosyncratic information in classification behaviour. Because the studies were designed specifically to highlight this importance, however, the perceptually distinctive stimuli were created to be perceptually individuated in both the unique manner in which the separate analytic features were manifested and the individuated manner in which the animals cohered into wholes. The perceptually uniform stimuli lacked both these sources of individuation. As a result of this confound, these studies have not addressed the relative contribution of each source of perceptual distinctiveness to the phenomenon of instance based categorization observed.

One possibility is that the effect of each is mediated by the task demanded of the subject. Using pseudowords as stimuli, Whittlesea (1987) has shown that if subjects are encouraged to treat training items holistically as pronounceable words, then later classification of new items follows a nonanalytic, exemplar based pattern of responding. When training encourages subjects to treat the items as a series of independent cues to category membership, however, the classification of new items is far more analytic. Similarly, Smith and Shapiro (1989) used the rapid pronunciation of pseudowords presented on the left or right half of the page as an incidental learning condition. Consistent with Whittlesea, they found that these conditions resulted in less analytic categorization strategies.

By analogy, in the studies presented above, it is possible that subjects provided with an explicit rule were being encouraged to break up the stimulus in a way that subjects receiving no explicit rule were not. If this is the case, then it is possible that the nonanalytic pattern of responding on perceptually distinctive items' in the "No Rule" experiment was being supported by the items' holistic distinctiveness whereas the nonanalytic intrusions in the "Rule" experiment were more a function of the unique manifestation of the analytically relevant features.

CHAPTER 3

Orthogonalizing the two sources of perceptual individuation

To examine the effect of each source of perceptual individuation on classification behavior and to explore how each interacts with the two training conditions, the effect of each source of perceptual distinctiveness will be examined in isolation. We will return, therefore, to the 2 x 2 matrix of Figure 1 and fill in the off-diagonal panels. These remaining two panels complete the attempt to orthogonalize the two types of nonanalytic information. The two experiments to be presented here will look at the interaction of each of these two forms of perceptual information with the two training conditions that have been discussed.

Experiment 2A: The effect of holistic individuation

Materials

Stimuli were a single set of 16 animals constructed using the same procedure as that described in Experiment 1, except that the dimension of leg number (two vs. six) was replaced by the dimension of tail length (short vs. long). Thus, all animals had two legs that systematically varied across animals in length only. The logical combinations of features (0's and 1's) however, remained identical to that displayed in Table 1. Perceptually, these animals were designed to have interchangeable features, but to nonetheless cohere into individuated wholes. One of the two subsets composing the 16 stimuli that resulted from these constraints can be seen in panel B of Figure 1.

Method

Subjects: Subjects were 32 undergraduate students at McMaster University participating for course credit.

Design: The design was identical to that of Experiments 1C and 1D except that the features associated with builders in each of the four rules were altered to:

1) Long legs, Angular body, Spots present
2) Short legs, Angular body, Spots present
3) Long legs, Round body, Spots absent
4) Short legs, Round body, Spots absent

Procedure: Both the presentation of stimuli and design of the training and transfer rounds were identical to that described in Experiments 1C and 1D. Subjects were in one of two training conditions: the "Rule" training condition and the "No Rule" training condition. The groups differed only in the instructions they received prior to the training phase. As in Experiment 1C, subjects in the No Rule group were told that there were two categories of items and that they were to try to place items in the appropriate categories. It was suggested that they would likely be guessing at the beginning, but that the computer would be correcting their decisions so that, with practice, they would become better at classifying items. Consistent with Experiment 1D, subjects in the Rule group were presented with the three-feature additive rule that would result in perfect classification of builders and diggers, and were told that they would be given practice in applying the rule to a set of training stimuli.

Results

Analyses were performed on the proportion of errors and on the median response times for correct categorizations made by subjects in the transfer phase of the experiment. These analyses included two 2 x 3 ANOVAs with one between-subjects factor, training condition (Rule vs. No Rule), and one within-subjects factor, transfer type (Old vs. GT vs. BT). A summary of the data is presented in Table 5.

Table 5. Mean responses to the holistically individuated stimuli with perceptually interchangeable features (panel B of Figure 1) under conditions of either "Rule" or "No Rule" training (Experiment 2A).

PROPORTION ERRORS

TRAINING CONDITION	OLD	GT	BT	GT - BT Effect
No Rule	.156 (.132)	.146 (.225)	.500 (.261)	*
Rule	.052 (.084)	.063 (.113)	.083 (.123)	N. S.

MEDIAN RESPONSE TIMES (ms)

TRAINING CONDITION	OLD	GT	BT	
No Rule	802 (273)	853 (296)	850 (287)	N. S.
Rule	909 (262)	904 (274)	1136 (369)	*

Errors: Analysis of the error data revealed that subjects in the No Rule condition were, in general, more error prone than subjects in the Rule condition, $\underline{F}(1,30)=23.75$, MSe=0.03. Furthermore, there was a significant main effect of transfer type, $\underline{F}(2,60)=10.26$, MSe=0.03. These main effects, however, must be interpreted in light of a significant interaction between training condition and transfer type, $\underline{F}(2,60)=7.63$, MSe=0.03. A subsequent set of paired t-tests indicated that, for the Rule training condition, there was no significant difference in the error rates of Old vs. GT items, $\underline{t}(15)=0.35$, and no significant difference in the error rates of GT vs. BT items, $\underline{t}(15)=0.52$. For the No Rule training condition, there was also no significant difference in the error rates of Old vs. GT items $\underline{t}(15)=0.16$. The classification of BT items, however, was significantly more error prone than classification of GT items in the No Rule training condition, $\underline{t}(15)=4.77$.

Response Times: Analysis of subjects' response times revealed no main effect of training condition, $F(1,30)=1.56$, MSe=248782. There was a significant effect of transfer type, $F(2,60)=6.62$, MSe=19588. However, this main effect must be interpreted in the light of a significant interaction between training and transfer type, $F(2,60)=4.66$, MSe=19588. A subsequent set of paired t-tests indicated that, for the No Rule training condition, there was no reliable difference in the speed of responding to Old vs. GT items, $t(15)=1.49$, and no reliable difference between GT and BT items, $t(15)=0.12$. For the Rule training condition, there was no significant difference in the speed of responding Old vs. GT items, $t(15)=0.18$. The classification of BT items was significantly slower than the classification of GT items, $t(15)=3.22$.

Discussion

Subjects viewing these holistically individuated items showed evidence of episodic effects in classification. In the No Rule condition, BT items were more prone to errors than the GT items, suggesting some use of the similarity to previously learned exemplars as a source of classification information. It should be noted that of the 16 subjects in the No Rule condition, three subjects showed patterns of responding consistent with classification on the basis of a single-feature rule, indicating a tendency in some subjects to treat these stimuli as a set of independent, recombinable parts. When these subjects were removed from the analysis, however, the remaining subjects showed on average an even greater discrepancy between the error rates on GT and BT items (the error rate for GT items dropped to 0.096 and the error rate for BT items increased to 0.577). Generally, then, the effects here do not seem to be as powerful as the effects seen in the stimuli that possessed both sources of idiosyncratic perceptual information, but some evidence of an effect is present.

In the Rule condition, the effect is apparent in subjects' response times. Although all items were classified with approximately equal accuracy (Old, GT and BT items alike), accurate classification of the BT items was slow in comparison to the other items. This slowing of accurate BT responses suggests that, again, the episodic effect was present, but in a weaker form than previously seen. It appears that the conflict between the classification rule and the holistic similarity to old items was making subjects hesitate when they were classifying the BT items. Thus, under both training conditions, similarities in the holistic quality of items seems to have had an effect on subjects' classification behaviour, despite the lack of any unique perceptual information in the manifestations of the items' independent features. Experiment 3 will address the issue of whether it is most appropriate to think of such effects as being due to holistic properties or due to relational properties between the features that are substantially less encompassing than suggested by the term holistic.

A tentative hypothesis previously entertained was that, with these stimuli, knowledge of the rule would eliminate the response to the holistic quality that was evident in the No Rule group. Because each feature was identical in every animal in which it occurred, it might be easy for the subjects to search for the exact perceptual appearance of the features named in the rule, thereby breaking up the overall appearance of the animal. Because the subjects in the No Rule group did not know which features to search for, they might have been more sensitive to the overall organization of the item. Although the pattern of errors indicates less nonanalytic generalization in the Rule than in the No Rule group, the inconsistent results with the response times prevent taking this as support for the hypothesis.

Experiment 2B: The effect of perceptually unique features

Materials

The stimuli for this experiment were adapted from the stimuli of Medin et al. (1987). Consistent with the stimuli created by Medin and his colleagues, the animals in this study vary systematically on only four analytic binary dimensions: head shape (round vs. angular), body markings (stripes vs. spots), leg number (four vs. eight), and tail length (short vs. long). As before, two subsets of eight animals were used (one as the training set and one as the transfer set). For one subset (shown in panel C of Figure 1) five animals were duplicates of the animals used by Medin et al. (1987). The remaining three animals were altered in the dimension of tail length in order to produce the appropriate analytic structure for the set. Each of these animals has unique manifestations of the isolated features. Yet, as mentioned earlier, there is a consistency of construction across the set. I suspected that this uniformity in construction would undermine the tendency for the features in a given animal to produce an overall feeling of individuality. Instead, the set in general gave the impression of being composite in nature. The matched counterpart for each of these animals in the second subset was created by altering only the analytic dimension of body markings.

Method

Subjects: Subjects were 32 undergraduate students from McMaster University participating for course credit.

Procedure: The design and procedure were identical to that of experiment 2A except that the features associated with the builder category in each of the four rules were:

1) Eight legs, Angular head, Spots
2) Four legs, Angular head, Spots
3) Eight legs, Round head, Stripes
4) Four legs, Round head, Stripes

The dimension of tail length was the single definitionally irrelevant dimension.

Results

Analyses were performed on the proportion of errors and on the median response times for categorizations made by subjects in the transfer phase of the experiment. These analyses included two 2 x 3 ANOVAs with one between-subjects factor, training condition (Rule vs. No Rule), and one within-subjects factor, transfer type (Old vs. GT vs. BT). A summary of the data is presented in Table 6.

Table 6. Mean responses to the composite stimuli with perceptually unique features (panel C of Figure 1) under conditions of either "Rule" or "No Rule" training (Experiment 2B).

PROPORTION ERRORS

TRAINING CONDITION	OLD	GT	BT	GT - BT Effect
No Rule	.148 (.131)	.141 (.182)	.688 (.255)	*
Rule	.063 (.079)	.078 (.151)	.078 (.120)	N. S.

MEDIAN RESPONSE TIMES (ms)

TRAINING CONDITION	OLD	GT	BT	
No Rule	926 (307)	1013 (362)	1001 (366)	N. S.
Rule	903 (261)	911 (236)	928 (201)	N. S.

Errors: Analysis of the error data showed that subjects in the No Rule condition were, in general, more error prone than subjects in the

Rule condition, $F(1,30)=55.33$, $MSe=0.03$. Furthermore, there was a significant main effect of transfer type, $F(2,60)=21.49$, $MSe=0.04$. These main effects, however, must be interpreted in light of a significant interaction between training condition and transfer type, $F(2,60)=20.32$, $MSe=0.04$. A subsequent set of paired t-tests indicated that, for the Rule training condition, there was no significant difference in the error rates of Old vs. GT items, $t(15)=0.02$, and no significant difference in the error rates of GT vs. BT items, $t(15)=0.00$. For the No Rule training condition, there was also no significant difference in the error rates of Old vs. GT items $t(15)=0.01$. The classification of BT items, however, was significantly more error prone than classification of GT items in the No Rule training condition, $t(15)=4.93$.

Response Times: Analysis of subjects' response times revealed no reliable differences either in the main effects or the interaction.

Discussion

Subjects in the No Rule condition were clearly making use of item-specific information in the classification of these animals. Accuracy on Old items and on GT items was quite good, but BT items were consistently being placed in the category opposite to that predicted by the three-feature additive rule. Furthermore, none of the sixteen subjects in the No Rule training condition showed a classification pattern consistent with the use of a single analytic feature in the transfer round. By contrast, when subjects were supplied with the rule that would provide perfect classification, the pattern of responding suggests that subjects were able to ignore this item-specific information and respond exclusively on the basis of the analytic structure of the stimuli. Neither the error rates nor the response times suggest that subjects were experiencing any difficulty in classifying animals in which the analytic rule was in conflict with the perceptual similarity to previously trained items.

Taken together, the data from Experiments 2A and 2B suggest a pattern of results opposite to the pattern that was suggested as a possibility earlier. First, a number of pieces of evidence suggest that, under No Rule training conditions, the unique manifestation of individual features rather than the holistic quality of the individual items seems to have had a stronger influence on responding. A small subset of subjects viewing the holistic animals, for example, made exclusive use of the analytic affordances of the stimuli in their classification attempts, but no subjects viewing the featurally individuated animals used this analytic simplification approach. Furthermore, although a post hoc analysis of the data did not reveal the difference to be statistically significant, subjects viewing the featurally distinctive animals showed, on average, a slightly stronger BT effect than subjects viewing the holistically individuated animals (even after the subjects displaying analytic simplification strategies had been removed from the holistically individuated condition).

On the other hand, under Rule training conditions, it was the effects of the attempt at holistic individuation rather than the manipulation of unique separate features that appeared more persistent. Even though subjects were told to pay attention to specific features, to treat the stimuli as a set of parts, the presence of idiosyncratic parts alone was not sufficient to produce similarity based responding for this group. Instead, it was the stimuli designed to possess holistic individuation rather than the stimuli thought to possess only unique perceptual manifestations of the separate features that seemed to show nonanalytic intrusions on purely analytic performance. Although the effect was weak in the animals designed to possess only the holistic source of perceptual individuation (in that it was present only in the response times), there was no effect at all for the animals possessing only featural individuation.

Conclusions based on the comparison of these two sets of stimuli, however, are tentative. The intent for the stimuli in Panel C was to produce a set of items in which there was clear featural individuation but no effective holistic differentiation. However, it cannot be assumed that the contrast actually achieved was this pure. These stimuli, for example, might have had even less overall individuation if the bodies were five times larger relative to the appendages. Because the shape of the bodies is the same for all of the animals, such a change might have perceptually emphasized the commonality of the animals and perceptually isolated the legs from the tails and in turn from the heads. Under these conditions the contrast between the legs on the different animals would not have been diminished at the level of separate features, but there would be less of a tendency for the particular combination of legs to cohere with the tail and the head to produce an overall impression for the animal. Possibly with such stimuli, there would have been no effect of similarity to old, suggesting that overall similarity rather than distinctiveness of individual features is crucial for exemplar based transfer.

Similarly, performance with the animals of Panel B might not be a pure result of their holistic properties. One could imagine, for example, that the ratio of leg length to body length was itself a feature. If so, then this could make a distinctive feature that was not based on enough information about the item to refer to it as holistic. There was certainly no indication in subjects' reports that they were using such featural interactions in their explicit attempts to develop a rule for categorization. However, the manipulation is not sufficient to implicate the holistic structure as the only source of less explicit, nonanalytic intrusions on strictly analytic classification behaviour. At best, it can be concluded that the perceptual distinctiveness of the features used by subjects for the purposes of analysis is insufficient to account for the signs of exemplar-based generalization that we obtained in Experiments 1C and 1D.

CHAPTER 4

The Priority of Holistic Individuation

The main purpose of the current experiment is to show that something more like the overall shape of the animal is critical to obtaining the bad transfer effects that are being used as an index of exemplar-based transfer. To accomplish this, I used the stimuli shown in Figure 3. For these stimuli, I used only the dimensions that are grouped at the center of each animal as relevant dimensions: namely, number of legs, body shape, and spots. It is possible to alter the overall impression given by the animal by changing only the categorically irrelevant dimensions, as is shown by the difference in the necks, heads and tails between Training Set A and Training Set B. In this comparison, the perceptually individuated qualities of the relevant features remain the same, which means that on these dimensions the two training sets have exactly the same relation to the Transfer Set that will be used after both sets of training items. It is assumed that giving a person a rule will concentrate the person's attention on the named dimensions, which would mean that maintaining the perceptual identity on these dimensions should be especially important for the subjects in the Rule group. In principle, it is possible for the Rule group to concentrate on the center of the animal, as they are directed to do by the rule. If so, and if the observed effect on the Bad Transfer items is due to a summation of information from the distinctive features, then there should be at least some part of this effect evident in the transfer from Training Set B to the Transfer Set. On the other hand, if the holistic nature of the items is particularly important for producing the bad transfer effect, then there will be substantially less transfer from Training Set B than from Training Set A. This question will be explored in Experiment 3B.

The second purpose of Experiment 3 is to probe the relationship between featural and holistic individuation. If there is no sign of a bad transfer effect between Training Set B and the Transfer Set, then we have evidence that altering the holistic individuality of the stimuli critically interferes with the ability to obtain this sign of exemplar based transfer. Yet it says little about manner in which this interference occurs. For example, it is possible that, in the presence of holistic individuation, subjects simply fail to learn the categorical relevance of the distinctive information found in the separate features. Alternatively, it might be that subjects will learn the categorical assignments of the distinctive features despite the presence of holistic individuation, but that whatever effect is exerted by the subjects' knowledge of individuated features is dependent on *not* experiencing conflict with the holistic identity of the animals. In Experiment 3C, I will examine whether the presence of holistic individuation interferes with acquiring knowledge about the categorical assignments of distinctive separate features during learning, or interferes with the use of this knowledge during later classification decisions.

Experiment 3A is a preliminary experiment to demonstrate that I did not alter the basic bad transfer effect when I created the new set of stimuli. This is an important baseline against which to evaluate the apparent null effect of maintaining the perceptual distinctiveness of the relevant features in Experiment 3B.

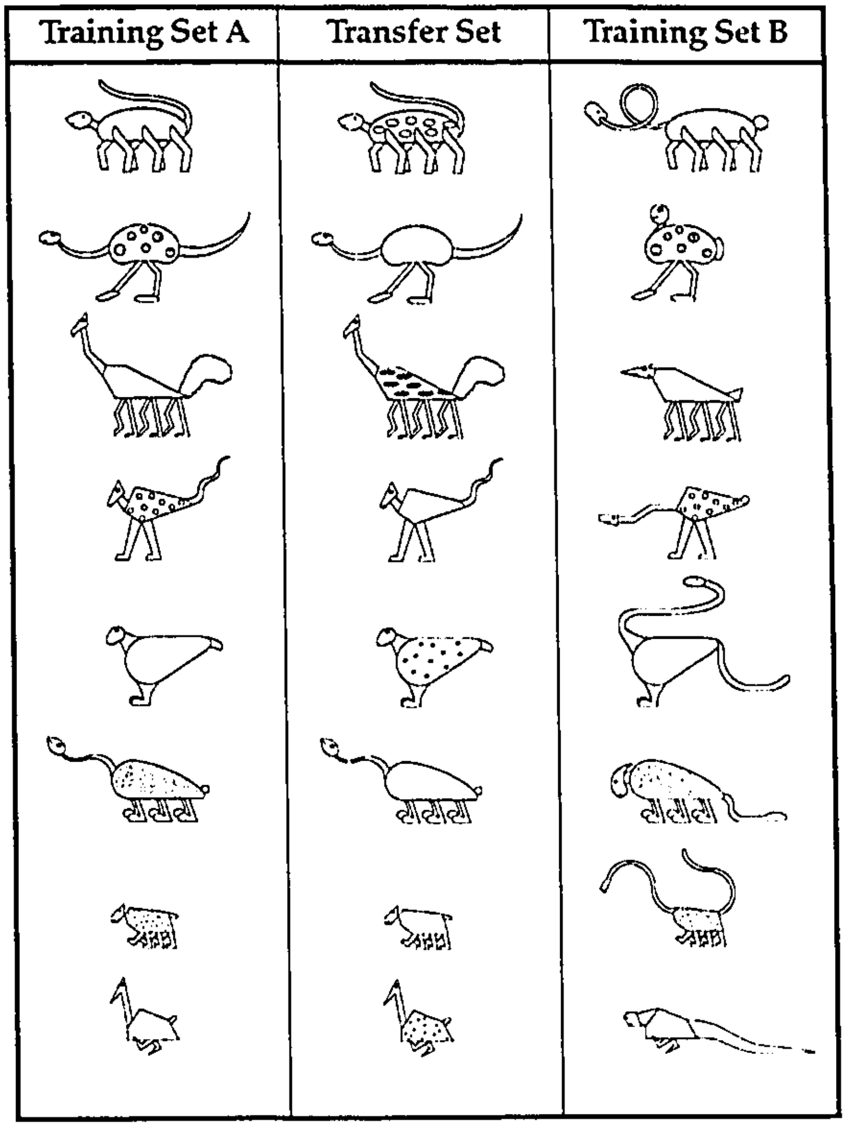

Figure 3. Stimuli used in Experiment 3A (Training Set A and the Transfer Set) and Experiment 3B (Training Set B and the Transfer Set). In both experiments, the perceptually distinctive aspects of two of the relevant features (number of legs and shape of bodies) stayed the same between the training set and the transfer set. However, because of alteration of the irrelevant features (tail and neck), the overall appearance of each animal changed between Training Set B and the Transfer Set but not between Training Set A and the Transfer Set.

Experiment 3A: A replication of Experiments 1C and 1D

Method

Materials: The stimuli are a single set of 16 animals adapted from the perceptually distinctive set used in Experiment 1. These stimuli are identical to the earlier set except that the length of the animals' legs was altered (made more consistent across the animals) and a tail was added to each animal. Thus for this set, the dimension of leg length (long vs. short) was replaced by the dimension of tail length (long vs. short). The analytic structure of 1's and 0's remained identical to that seen in Table 1. The resulting set of 16 stimuli are Training Set A and the Transfer Set of Figure 3.

Subjects: Subjects were 32 undergraduate students from McMaster University participating for course credit.

Procedure: The design and procedure for Experiment 3A were identical to that of Experiment 2A. Thus, subjects were in one of two groups (Rule or No Rule) that differed only in the instructions received prior to the training phase of the experiment. In Experiment 3A, however, the features associated with builders in each of the four rules were altered to:

1) Six legs, Angular body, Spots present
2) Two legs, Angular body, Spots present
3) Six legs, Round body, Spots absent
4) Two legs, Round body, Spots absent

Thus, the dimensions of neck length and tail length were definitionally irrelevant in all conditions.

Results

Analyses were performed on the proportion of errors and on the median response times for correct categorizations made by subjects in the transfer phase of the experiment. These analyses included two 2 x 3 ANOVAs with one between-subjects factor, training condition (Rule vs. No Rule), and one within-subjects factor, transfer type (Old vs. GT vs. BT). A summary of the data is presented in Table 7.

Table 7. Mean responses to the stimuli composing Training Set A and the Transfer Set of Figure 3 under conditions of either "Rule" or "No Rule" training (Experiment 3A).

PROPORTION ERRORS

TRAINING CONDITION	OLD	GT	BT	GT - BT Effect
No Rule	.094 (.107)	.097 (.153)	.812 (.194)	*
Rule	.102 (.131)	.125 (.183)	.344 (.256)	*

MEDIAN RESPONSE TIMES (ms)

TRAINING CONDITION	OLD	GT	BT	
No Rule	849 (375)	897 (293)	878 (292)	N. S.
Rule	811 (311)	875 (300)	1300 (369)	*

Errors: Analysis of the error data revealed that subjects in the No Rule condition were, in general, more error prone than subjects in the Rule condition, $\underline{F}(1,30)=21.55$, MSe=0.04. Furthermore, there was a significant main effect of transfer type, $\underline{F}(2,60)=88.00$, MSe=0.03. These main effects, however, must be interpreted in light of a significant interaction between training condition and transfer type, $\underline{F}(2,60)=23.23$, MSe=0.03. A subsequent set of paired t-tests indicated that, for the Rule training condition, there was no significant difference in the error rates of Old vs. GT items, $\underline{t}(15)=0.68$. The classification of BT items, however, was significantly more error prone than classification of GT items following Rule training, $\underline{t}(15)=3.22$. Similarly, for the No Rule training condition, there was no significant difference in the error rates of Old vs. GT items $\underline{t}(15)=0.06$. Again, however, the classification of BT items was significantly more error prone than classification of GT items, $\underline{t}(15)=11.04$. A subsequent between groups t-test indicated that the difference in error rates between GT and BT items was significantly smaller in the Rule condition than in the No Rule condition, $\underline{t}(30)=5.29$.

Response Times: Analysis of subjects' response times revealed no main effect of training condition, $F(1,30)=1.33$, MSe=262441. There was a significant effect of transfer type, $F(2,60)=21.27$ MSe=27940. However, this main effect must be interpreted in the light of a significant interaction between training and transfer type, $F(2,60)=19.49$ MSe=27940. Consistent with previous No Rule conditions, a set of paired t-tests indicated no reliable differences in the speed of responding either between Old and GT items, $t(15)=0.99$, or between GT and BT items, $t(15)=0.38$. For the Rule training condition, there was no significant difference in the speed of responding to Old vs. GT items, $t(15)=1.60$. The accurate classification of BT items, however, was significantly slower than the accurate classification of GT items, $t(15)=5.60$.

Discussion

Clearly, altering the stimuli from panel D of Figure 1 to include a tail and altering the rules from Experiment 1 to include only legs, body and spots did not have a fundamental impact on subjects' responses to the stimuli. The pattern of results for these animals is extremely similar to the pattern for the perceptually distinctive animals of Experiment 1 when trained under similar conditions. Both the error rates and the response time data replicate the findings for the perceptually distinctive stimuli of Experiments 1C (as seen in Table 3 on page 31) and 1D (as seen in Table 4 on page 37).

Experiment 3B: The effect of altering holistic individuation

This set of stimuli provides the opportunity to test whether the episodic effects seen in the Bad Transfer items are a function of the unique manifestation of the isolated features or a function of the more holistic individuation of the animals. To distinguish these possibilities, Training Set B of Figure 3 was given to a Rule and to a No Rule group. As indicated in the introduction to Experiment 3, on the basis of definitionally relevant features alone, which are possible to concentrate on as a group because they are clustered in the center of the animals, these new animals bear the same relation to the Transfer Set as do the animals in Training Set A both in analytic structure and in the perceptual manifestation of these features. If the observed effect on the Bad Transfer items is due to a summation of the distinctive features, then when subjects are given a rule, the Bad Transfer effect should be critically dependent on maintaining perceptual similarity of the relevant features from training to transfer items. Under these circumstances, the alteration of definitionally irrelevant features should not decrease errors made on Bad Transfer items. However, alteration of the definitionally irrelevant features will alter the holistic nature of the items. Thus, if holistic individuation affects classification, then the alteration of the irrelevant features will affect the frequency of instance based errors in the BT items.

Method

Subjects: Subjects were 32 undergraduate students from McMaster University participating for course credit.

Procedure: The design and procedure for Experiment 4B were identical to that of Experiment 3A.

Results

Analyses were performed on the proportion of errors and on the median response times for categorizations made by subjects in the transfer phase of the experiment. These analyses included two 2 x 3 ANOVAs with one between-subjects factor, training condition (Rule vs. No Rule), and one within-subjects factor, transfer type (Old vs. GT vs. BT). A summary of the data is presented in Table 8.

Errors: Analysis of the error data revealed that subjects in the No Rule condition were, in general, more error prone than subjects in the Rule condition, $\underline{F}(1,30)=38.01$, MSe=0.02. Furthermore, there was a significant main effect of transfer type, $\underline{F}(2,60)=5.06$, MSe=0.03. These main effects, however, must be interpreted in light of a significant interaction between training condition and transfer type, $\underline{F}(2,60)=3.80$, MSe=0.03. A subsequent set of paired t-tests indicated that, for the Rule training condition, there was no significant difference in the error rates of Old vs. GT items, t(15)=0.28, and no significant difference in the error rates of GT vs. BT items, t(15)=0.37. For the No Rule condition, Old items were significantly less error prone than the GT items, t(15)=3.56, and no reliable difference in error rates was found between GT items and BT items, t(15)=0.17.

Table 8. Mean responses to the stimuli composing Training Set B and the Transfer Set of Figure 3 under conditions of either "Rule" or "No Rule" training (Experiment 3B).

PROPORTION ERRORS

TRAINING CONDITION	OLD	GT	BT	GT - BT Effect
No Rule	.102 (.068)	.313 (.233)	.297 (.245)	N. S.
Rule	.039 (.060)	.047 (.101)	.063 (.112)	N. S.

MEDIAN RESPONSE TIMES (ms)

TRAINING CONDITION	OLD	GT	BT	
No Rule	1095 (442)	1320 (492)	1258 (409)	N. S.
Rule	1194 (404)	1237 (326)	1271 (340)	N. S.

Response Times: Analysis of the response times revealed a main effect for transfer type, $\underline{F}(2,60)=3.49$, MSe=50079, but no main effect for training condition, $\underline{F}(1,30)<1$, and no significant interaction, $\underline{F}(2,60)=1.33$, MSe=50079. Collapsing across training condition, a subsequent set of paired t-tests revealed no reliable difference between GT and BT items, $\underline{t}(31)=.26$, but responses to Old items were significantly faster than responses to the Transfer items, $\underline{t}(31)=2.63$.

Discussion

Evidence from Experiment 3B suggests that for both the Rule and the No Rule training conditions, the items' holistic individuation, rather than the unique manifestation of independent features, was essential for the bad transfer effect. When the holistic individuation was changed through the alteration of definitionally irrelevant features, all evidence of exemplar based transfer disappeared.

Subjects classifying the animals following Rule training were extremely accurate in the classification of all stimuli (Old, GT and BT stimuli alike). Furthermore, response times for the Rule condition were clearly slower in the transfer phase of Experiment 3B than in the transfer phase of Experiment 3A, suggesting that the perceptually altered transfer items produced a general context of unfamiliarity, which resulted in more deliberate rule application in the final round.

In the No Rule condition, responding to Old items was fairly accurate (at least as accurate as responding to other perceptually distinctive stimuli viewed under No Rule training conditions). Thus, for the familiar old items, the memory for specific cases appears to have been influential in classification. In the transfer items, however, GT and BT items were classified with equal accuracy. This pattern is not consistent with the pattern expected under conditions of episodic transfer. The level of accuracy for the transfer items, however, was better than the 50% accuracy predicted by chance. It was, in fact, quite similar to the level of accuracy achieved by subjects viewing the perceptually uniform stimuli under the No Rule conditions of Experiment 1C. It appears, therefore, that the classification of perceptually familiar old items was aided by memory for specific exemplars, but the classification of the perceptually novel transfer items was supported by some (albeit incomplete) analytic category level information, similar to the analytic simplification observed in Experiment 1C.

Thus, when the individual as a whole became unfamiliar, subjects responded as if no similarity existed between the new individual and the earlier training items. Subjects failed to take advantage of the perceptual and analytic similarities that continued to exist despite the loss of the holistic similarity. It appears, therefore, that when holistic similarity is present, it is this form of similarity that supports exemplar based responses.

Experiment 3C: The effect of disrupting holistic individuation

Based on the data from Experiment 3B, it appears that holistic individuation interferes with the tendency for unique features to control responding. These data, however, do not shed much light on whether this interference is at the level of acquisition or response control. It is possible, for example, that in the presence of holistic individuation, subjects fail to notice the unique nature of the individual features out of which the whole is emerging. If this were true, then in the presence of holistic individuation, the specific features would have no separate association with the object. The only contribution of a feature to the individual would be in the manner in which it melds with other features to create the whole. Undoubtedly, the unique manifestation of the specific features will contribute to the holistic distinctiveness of the animal. Anecdotal evidence from everyday life, however, suggests that the uniqueness of separate features is not entirely lost in the gestalt of the whole. Children, for example, are often told that the have their mother's eyes or their father's nose. Statements of this sort suggest that separate features are recognized as belonging to the individual. Thus, it is likely that subjects in Experiment 3B *could* have identified the perceptually similar features as belonging to a previous individual item. They did not, however, because the altered holistic nature of the stimulus suggested that it was a new, very different individual. To examine whether the subjects learned anything about the associations of the unique features with the categories, another condition was tested. For this condition, stimuli in the transfer phase were composed of features scattered randomly about the screen (see Figure 4). Consistent with Modigliani (1971, 1974), it can argued that such a manipulation will destroy the holistic nature of the stimuli while leaving the unique nature of the separate features intact.

Materials

During training, subjects viewed the stimuli from Training Set B of Figure 3 five times each. In the transfer phase, these eight stimuli were supplemented with the eight stimuli from the Transfer Set of Figure 3. In this experimental condition, however, the features composing each of the sixteen stimuli in the transfer phase were no longer arranged to create whole animals. Instead, for each stimulus, the five features were scattered

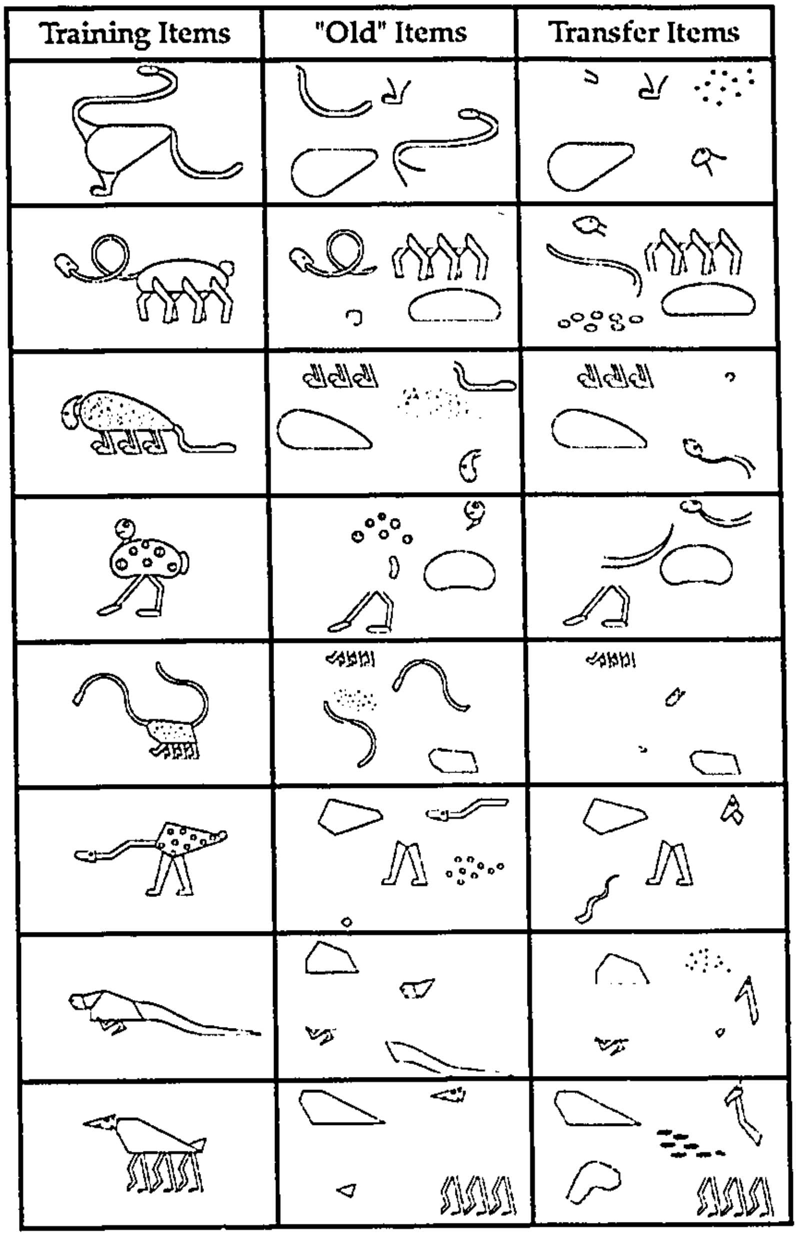

Figure 4. Stimuli used in Experiment 3C. On the basis of separate features alone, both the Training items and the "Old" items (seen in the transfer phase) are identical to the items of Training Set B in Figure 3, and the Transfer Items are identical to the items of the Transfer Set in Figure 3. For these stimuli, however, the items seen at transfer have been broken apart in order to disrupt the holistic individuation.

randomly around the screen (as seen in Figure 4). Thus, the separate pieces of analytic and perceptual information were present in the stimulus, but no holistic component existed.

Method

Subjects: Subjects were 8 undergraduate students from McMaster University participating for course credit.

Procedure: The procedure was identical to the No Rule condition of Experiment 3B.

Results

As in the previous No Rule conditions reported in this thesis, subjects' response times were not particularly informative and will not be discussed here. The error rates for the Old, GT and BT items may be seen in Table 9. A set of paired t-tests revealed no significant difference between the error rates of Old and GT items, $\underline{t}(7)=0.26$, but a significant increase in errors for the BT items over GT items, $\underline{t}(7)=2.96$.

Table 9. Mean responses to the stimuli composing Training Set B and the Transfer Set of Figure 3 when the holistic nature of the stimuli has been disrupted as depicted in Figure 4 (Experiment 3C).

	PROPORTION ERRORS			
	OLD	GT	BT	GT - BT Effect
TRAINING CONDITION				
No Rule	.203 (.132)	.219 (.248)	.656 (.229)	*

Discussion

The disruption of the holistic nature of the animals in this No Rule training condition does seem to have hurt performance on Old items at least to some degree. The 20% error rate is higher than has generally been seen on perceptually distinctive stimuli. Interestingly, this manipulation, which seems to have hurt the classification of Old items, provides evidence of an episodic transfer effect when none was apparent in Experiment 3B. It appears that when subjects are faced with a set of randomly scattered parts, the similarity of a few separate features is sufficient to influence responding. Thus the presence of holistic individuation during training does not interfere dramatically with the acquisition of knowledge about the unique perceptual manifestation of separate features. The continued presence of holistic individuation during transfer, however, does appear to interfere with the use of this knowledge about unique features.

CHAPTER 5

Factors affecting the acquisition of analytic and nonanalytic knowledge

Consistent with Allen and Brooks (1991), evidence from the studies presented to this point suggest that subjects can possess analytic and nonanalytic knowledge concurrently. The evidence is most apparent for stimuli that contain both forms of perceptually idiosyncratic information (the animals of Panel D and of Experiment 3). Even under circumstances in which a perfectly predictive analytic rule was provided prior to experience with these stimuli, subjects learned about the nonanalytic perceptual information available in the stimuli and made use of this nonanalytic knowledge for the purposes of classification. The acquisition of the nonanalytic information, however, did not appear to be associated with a loss of access to the analytic, category level information that was provided prior to training. If, as in Experiment 3B, subjects were subsequently asked to classify items under conditions of clear novelty, they could use the previously provided analytic information quite effectively. Thus, possession of analytic knowledge did not eliminate the acquisition (and subsequent use) of nonanalytic information, and the acquisition of this nonanalytic information did not necessitate the loss of any analytic knowledge.

Similarly, evidence from No Rule conditions suggests that whether or not nonanalytic perceptual information was available (and acquired) had little effect on the degree to which analytic information was acquired. When a perfectly predictive rule for classification was not provided, subjects acquired and made use of the available idiosyncratic perceptual information for the purposes of classification. This acquisition of nonanalytic idiosyncratic information, however, did not seem to affect subjects' ability to identify and assess the systematic analytic variations present in the stimuli. If the idiosyncratic information was subsequently removed (or rendered less obvious as in Experiment 3B), then subjects' response patterns indicated that they had acquired just as much analytic knowledge about the stimulus set as subjects who had no idiosyncratic perceptual information available during training. Again, the acquisition of nonanalytic knowledge did not seem to decrease the acquisition of analytic knowledge.

It is interesting to note, however, that regardless of the presence or absence of idiosyncratic perceptual information, not a single subject in the No Rule training condition actually discovered the three-feature additive rule that would have resulted in perfect classification. Instead, the systematic application of an analytically generated rule was almost exclusively in the form of a single-feature rule. Thus, any improvement in classification that resulted from the analytic assessment of the stimulus set resulted from the identification and use of a single feature that was partially correlated with category membership. The finding that subjects viewing perceptually distinctive stimuli and subjects viewing perceptually uniform stimuli were equally effective in discovering analytically relevant information, then, indicates little more than the fact that each group could identify the separate dimensions of systematic variation in the stimulus set. It does not say much about the relative ability of the two groups to discover a more complicated, multi-feature analytic rule.

The strong tendency toward analytic simplification seen in these studies may be due, at least in part, to the fact that subjects were placed into circumstances in which analytic simplification had some use. That is, the use of a single relevant feature in isolation resulted in 75% accuracy. If this moderate level of success encouraged the use of analytic simplification, then it might well have interfered with the discovery of the more complicated multi-feature rule that would have resulted in perfect classification. The following experiments, therefore, used a two-feature "exclusive or" (XOR) rule to establish category membership. The advantage of the XOR rule is that no single dimension in isolation is correlated with category membership. Any attempt to use the strategy of analytic simplification results in an accuracy rate that is no better than chance performance. Thus, this rule provides the opportunity to assess whether subjects continue to use analytic simplification in the face of feedback that indicates it is not an effective strategy. Furthermore, to the extent that subjects abandon this strategy, it will be possible to assess subjects' ability to acquire the XOR classification rule, and to assess the extent to which this ability is affected by the perceptual nature of the stimulus set with which subjects are presented.

The second issue addressed in this set of studies involves the effect of processing on the information acquired by subjects during classification training. As mentioned earlier, the No Rule instructions provided for subjects in the previous experiments were carefully designed to avoid biasing the subjects either toward analytic or nonanalytic classification strategies. Although post-experimental interviews suggested that subjects, whether viewing perceptually uniform or perceptually distinctive animals, interpreted the instructions (at least initially) as an expectation that they would attempt to discover the rules of classification, it was never specifically indicated that such analysis was vital. Thus, when no analytic solution was readily apparent, subjects viewing the perceptually idiosyncratic stimuli might have abandoned the analytic strategy in favour of a more nonanalytic instance based classification procedure. If this was the approach that subjects adopted, then the acquisition of analytic and nonanalytic knowledge would not be concurrent, but serial processes. Subjects would have first acquired whatever analytic information they could, then subsequently learned the appropriate category for each individual.

The following experiments, therefore, contrasted the standard, "neutral" No Rule training instructions of the previous studies with a set of No Rule instructions that stress the use of analytic strategies. In this analytic instruction condition, subjects viewing perceptually distinctive stimuli will be told that rule discovery is vital because, following training, they will be asked to classify a set of new (unfamiliar) animals that follow the same categorical rules. In this condition, there is a clear indication that learning the category membership of individuals will not help in the subsequent transfer round and, as a result, there is a much stronger injunction to continue treating the animals as a set of separate, analytically abstracted parts. Thus, this instruction provides the opportunity to assess the acquisition of nonanalytic categorical information about individuals despite a highly analytic approach to the category learning task. Furthermore, because the explicit task for subjects in this condition is the discovery of the analytic rule, these instructions are likely to produce a state of mind similar to that found in subjects who view the perceptually uniform stimuli following "neutral"

No Rule instructions. Thus, the uniform condition provides a useful baseline for examining the effect of nonanalytic perceptual variation on subjects' analytic ability to discover a more complicated, multi-feature classification rule.

Materials

The stimuli were adaptations of the perceptually uniform and perceptually distinctive animals used in Experiment 1. Each of these animals was given a tail that could be either long or short. Thus, for these studies, the analytic structure of the stimuli included six orthogonal binary dimensions. As before, a set was composed of sixteen stimuli constructed in pairs (one member for each of two subsets of items). However, for this set, the two stimuli within four of the eight pairs differed on the analytic dimension of spots, and the two stimuli within the remaining four pairs differed on the analytic dimension of tail length. Across pairs, each stimulus differed from every other on at least two analytic dimensions (the analytic structure of the sixteen stimuli composing a set may be seen in Table 10). As in Experiment 1, this single analytic structure was manifested in two perceptual sets: a distinctive set containing both sources of perceptually idiosyncratic information, and a uniform set, containing neither source (examples of the two sets are depicted in Figure 5).

Across subjects, four different rules were used to categorize the animals. The combinations of features associated with builders for each of the four rules were:

1) Angular body and Spots *OR* Round body and No spots
2) Angular body and Long tail *OR* Round body and Short tail

3) Long legs and Spots *OR* Short legs and No spots

4) Long legs and Long tail *OR* Short legs and Short tail

Table 10 The analytic structure of the stimuli used in Experiments 4A and 4B. The items from Subsets 1 and 2 alternately served as training and transfer items.

Item Number	Body Shape	Spots	Tail Length	Leg Length	Neck Length	Number of Legs
1.1	0	0	0	0	0	0
1.2	0	1	1	0	0	1
1.3	1	0	0	1	0	1
1.4	1	1	1	1	0	0
1.5	0	0	1	1	1	1
1.6	0	1	0	1	1	0
1.7	1	0	1	0	1	0
1.8	1	1	0	0	1	1
1.1	0	1	0	0	0	0
1.2	0	0	1	0	0	1
1.3	1	1	0	1	0	1
1.4	1	0	1	1	0	0
1.5	0	0	0	1	1	1
1.6	0	1	1	1	1	0
1.7	1	0	0	0	1	0
1.8	1	1	1	0	1	1

Using these rules, half of the items in the transfer set differed from their training twin on a category relevant feature and half of the items differed from their training twin on a category irrelevant feature. Because, with an XOR rule, a change in a relevant feature always results in a change of category membership, a transfer item differing from its training twin on an analytically irrelevant feature acted as a Good Transfer item and an item differing from its training twin on an analytically relevant feature acted as a Bad Transfer item.

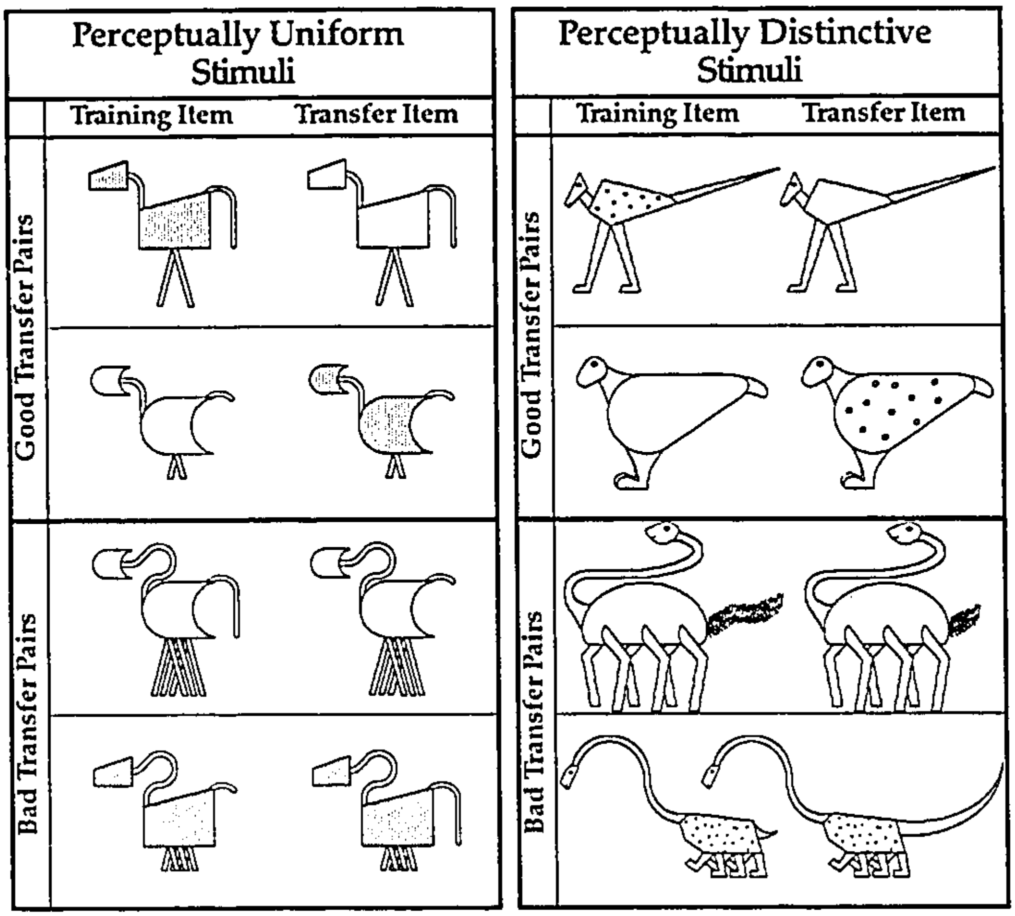

Figure 5. Examples of the transfer conditions used in Experiments 4A and 4B (shown here for each of the stimulus sets used in the experiments). In this example the categorization rule is: "The presence of EITHER long legs and a long tail OR short legs and a short tail indicates a builder". For this rule, a change in spots will not affect category membership so the pairs in the top half of the figure are Good Transfer pairs. However, a change in tail length will result in a change in category membership so the pairs in the bottom half of the figure are Bad Transfer pairs.

Experiment 4A: Interfering with the acquisition of nonanalytic knowledge

Experiment 4A examined the degree to which a strong analytic bias interferes with learning about the category membership of individual items during training. The No Rule conditions of the previous experiments suggest that, when viewing perceptually distinctive stimuli, five training rounds provide sufficient exposure to learn the category membership of the individuals (this finding is replicated using the adapted stimuli in the present experiment). However, pilot data suggested that five training rounds would be insufficient for subjects to discover the XOR rule that is determining categorization (a finding that is confirmed in the perceptually uniform condition of the present experiment). It was assumed, therefore that subjects who received strongly analytic instructions prior to viewing perceptually distinctive animals would fail to discover the XOR rule within the training period. What was assessed in the transfer phase of Experiment 4A was the extent to which these subjects learned about the category membership of the individuals.

Method

Subjects: Subjects were 36 undergraduate students from McMaster University participating for course credit.

Procedure: Subjects were in one of three conditions. For two of the conditions (the Standard/Uniform and the Standard/Distinctive conditions) the procedure was identical to that of Experiment 1C. That is, each subject was told that there were two categories of items and that the subject's job was to learn to place items into the appropriate categories. It was suggested to subjects that they would likely be guessing to begin with, but that the computer would be correcting their decisions so that, with time, they would become better at placing items into the appropriate categories. These instructions were meant to be as neutral as possible with respect to analytic or nonanalytic methods of classification. Following five training rounds, during which feedback was provided after each classification, subjects were informed that the eight old items would now be supplemented by eight new items and were asked to classify the sixteen items as accurately as they could. No feedback was provided during this transfer phase of the experiment. As in Experiment 1C, the Standard/Uniform and Standard/Distinctive conditions differed only in the perceptual manifestation of the items subjects were asked to classify (perceptually uniform or perceptually distinctive). For the third condition (the Analytic/Distinctive condition) the presentation of stimuli was identical to the other two conditions: five training rounds of eight items with feedback followed by one transfer round of sixteen items with no feedback. Subjects in this group viewed the perceptually distinctive stimuli seen in the Standard/Distinctive condition. For this group, however, the instructions were altered. Prior to the five training rounds, subjects in this condition were informed that each animal belonged to one of two categories and that their job was to figure out how to place items into one category or the other. It was suggested to subjects that, at the beginning, they would likely be guessing about category membership, but that the computer would be correcting their guesses so that hopefully by the end of the training round they would be able to figure out the rules that made a digger a digger and made a builder a builder. Following training, subjects were asked to describe their efforts to learn the rules of categorization, then were given the following instructions:

> In this last round, I am not interested in the rules that place items in one category or the other. For each of the eight animals that you have been looking at, the computer has told you five times what category the animal belongs in. For this round I just want you to remember for me what category the computer kept telling you that the animal belonged in. You will see each animal (or an animal that looks very similar to it) twice. Again, for each, just do your best to place the old (or similar) animal in the category that the computer was telling you to earlier.

The subject was then given the sixteen animals of the transfer round. No feedback was provided during this round.

Results

As in previous No Rule conditions, subjects' response times are not particularly informative and will not be discussed here. Analyses were performed on the proportion of errors made during the transfer phase of the experiment. For ease of comparison, in all three conditions an "error" in classification was identified as the failure to place the animal in the category consistent with the XOR rule (although this is not, in fact, an error for the Analytic/Distinctive group who received memory instructions in the transfer round). Thus, as in the previous experiments, the difference between GT errors and BT "errors" is a measure of subjects' use of nonanalytic, episodic classification procedures (as subjects in the Analytic/Distinctive condition were asked to do during the transfer phase). A summary of the data is presented in Table 11.

Table 11 Mean error rates for subjects viewing the Uniform or Distinctive stimuli of Figure 5 following either Standard or Analytically biased "No Rule" training instructions (Experiment 4A).

	PROPORTION ERRORS			
	OLD	GT	BT	GT - BT Effect
TRAINING CONDITION				
Standard/Uniform	.490 (.100)	.500 (.337)	.500 (.184)	N. S.
Standard/Distinctive	.072 (.124)	.167 (.195)	.792 (.234)	*
Analytic/Distinctive	.385 (.155)	.395 (.225)	.625 (.291)	*

A 3 x 3 ANOVA with one between-subjects factor, training condition (Standard/Uniform vs. Standard/Distinctive vs. Analytic/Distinctive) and one within-subjects factor, transfer type (Old vs. GT vs. BT) revealed a significant main effect of training condition, $\underline{F}(2,33)=5.56$, MSe=0.04, and a significant main effect of transfer type, $\underline{F}(2,66)=22.49$, MSe=0.05. These main effects, however, must be interpreted in light of a significant interaction between training condition and transfer type, $\underline{F}(4,66)=9.44$, MSe=0.05. A set of paired t-tests performed for each of the training conditions indicated that for the Standard/Uniform condition, there was no significant difference in the error rates of Old vs. GT items, $\underline{t}(11)=0.07$, and no significant difference in error rates between GT vs. BT items, $\underline{t}(11)=0.00$. For the Standard/Distinctive condition, there was also no significant difference in the error rates of Old vs. GT items, $t(11)=1.33$. For this condition, however, the classification of BT items was significantly more error prone than the classification of GT items, $t(11)=6.27$. Subjects in the Analytic/Distinctive condition showed a pattern of results similar to that of the Standard/Distinctive condition in that there was no significant difference in the error rates of Old vs. GT items, $t(11)=0.14$, but a significantly greater proportion of "errors" on BT items than on GT items, $t(11)=2.93$. A subsequent between groups t-test, however, indicated that the difference in "error" rates between GT and BT items was significantly smaller in the Analytic/Distinctive condition than in the Standard/Distinctive condition, $t(22)=3.12$.

Discussion

Data from the Standard/Uniform condition indicate quite clearly that the XOR rule was not easy to discover. Following five rounds with the eight training stimuli, not one of the subjects in this condition discovered the combination of features that would result in perfect classification. By contrast, data from the Standard/Distinctive condition replicated the finding from previous studies that five exposures to each of eight perceptually distinctive individuals can be sufficient to allow learning of the individuals' category membership (memory for the category of Old items was quite good). Furthermore, the low error rates for GT items and the correspondingly high error rates for BT items suggest that this nonanalytic source of knowledge was generalized to the classification of perceptually similar transfer items. These data provide the backdrop for a discussion of performance in the Analytic/Distinctive condition.

Consistent with the earlier prediction, post-training interviews with subjects in the Analytic/Distinctive condition indicated that these subjects were also unable to discover the analytic XOR rule that would have resulted in perfect classification. At the same time, however, data from the transfer round suggest that these subjects learned less about the category membership of individual items than did subjects who were not biased toward analytic strategies. These subjects were, in fact, only 60% accurate both in their memory for the category of old items and in their ability to classify the perceptually similar transfer items. Thus, the acquisition of nonanalytic knowledge as a source of classification information is clearly not independent of the manner in which the stimulus is processed. When subjects are encouraged to treat perceptually distinctive stimuli as sets of analytically abstracted parts, acquisition of the nonanalytic perceptual information as an aid to classification is diminished (although not eliminated).

Experiment 4B: Interfering with the acquisition of analytic knowledge

Experiment 4A indicates that a strong analytic bias can reduce subjects' tendency to learn about the category membership of the individuals, but it says little about whether the presence of nonanalytic information will interfere with the acquisition of a more complicated analytic rule. As in previous experiments most subjects in Experiment 4A (whether viewing perceptually uniform or perceptually distinctive items) could identify the dimensions of systematic variation in the stimulus set, but none discovered the multi-feature rule that would have resulted in perfect classification of the stimuli. Experiment 4B, therefore, repeated the three training conditions of Experiment 4A, except that subjects were given ten training rounds with the eight animals. As will be shown, this increase in training exposure allowed approximately one third of the subjects in the Standard/Uniform condition to discover the XOR rule that results in perfect classification. It is this baseline against which subjects in the Analytic/Distinctive condition could be compared in order to assess the effect of idiosyncratic perceptual variation on the acquisition of a multi-feature analytic rule. Because there is reason to believe (from the previous studies) that subjects will tend to override tentative analytic hypotheses in the face of contradictory nonanalytic information, the assessment of analytic learning by subjects in the Analytic/Distinctive condition was performed using the perceptually uniform set of stimuli in the transfer round of this experiment. However, the increased exposure to the stimuli also provides greater opportunity for subjects in the Analytic/Distinctive condition to learn the category membership of the individuals. Thus, following the transfer phase in which these subjects were tested using the perceptually uniform stimulus set, an assessment of their memory for the category of previously viewed items was performed.

Method

Subjects: Subjects were 48 undergraduate students from McMaster University participating for course credit.

Procedure: Subjects were in one of three conditions. The procedure for two of the conditions (the Standard/Uniform and the Standard/Distinctive conditions) was identical to that of Experiment 4A except that the training phase was increased from five rounds to ten rounds. Subjects in the third condition (the Analytic/Distinctive condition) were given the analytically biased training instructions as in Experiment 4A, but were also informed that following training with the perceptually distinctive animals, they would be asked to classify a new set of items that followed the same rules, but looked quite different. Thus, the emphasis on analytic attempts at rule discovery was further strengthened. Following ten training rounds with the perceptually distinctive animals, subjects in the Analytic/Distinctive condition were asked to classify the sixteen perceptually uniform animals without feedback. Subjects were then presented with each of the perceptually distinctive Old items paired its previously unseen training twin and were asked, for each pair, to identify the Old item that they had classified in the training phase. Finally, subjects were given the memory instructions of Experiment 4A and were asked to classify the eight perceptually distinctive items that they had viewed during training. No feedback was provided during either of these two memory tests.

Results

Subjects in the Standard/Distinctive condition showed a pattern of results consistent with previous No Rule conditions. No subjects in this condition either discovered the XOR rule of classification or made use of analytic simplification strategies in the transfer round. The mean error rates for subjects in this condition were: 0.11 (s.d.=0.14) for Old items, 0.20 (s.d.=0.21) for GT items, and 0.78 (s.d=0.22) for BT items.

The method of classification used by subjects in the remaining two conditions (viewing perceptually uniform stimuli in the transfer phase) fell into one of three general strategies. Some subjects discovered the XOR rule during the training round and subsequently made use of this rule during the transfer round. These subjects could verbalize the rule, and were 100% accurate in their classifications during the transfer round. Other subjects made use of analytic simplification. Again, each of these subjects could easily state the single feature rule used to classify the stimuli in the transfer round, and 100% of their classifications could be accounted for using the single feature identified. The remaining subjects used less systematic strategies. When these subjects attempted to describe their classification criteria, the classification "rules" were, in general, poorly articulated and inconsistently applied in the transfer round. The accuracy rates for these subjects did not differ reliably from chance responding. The frequency with which subjects in the Standard/Uniform and the Analytic/Distinctive conditions made use of the XOR rule, analytic simplification, or some other (less systematic) method of classification is presented in Table 12. An L^2 statistic (Kennedy, 1983) comparing the proportion of subjects in each condition making use of the XOR rule during transfer revealed that subjects in the Standard/Uniform condition were more likely to discover and make use of the correct analytic rule in the transfer round than were subjects in the Analytic/Distinctive condition (L^2=4.96). For the remaining subjects who failed to discover the XOR rule, however, there was no significant difference between the two groups in the proportion of subjects who made use of analytic simplification strategies (L^2=0.03).

When subjects in the Analytic/Distinctive condition were asked to select each Old, previously seen item from its perceptual twin distractor, however, subjects were quite accurate. The mean error rate of 0.11 (s.d.=0.14) was significantly below the level of 0.50 predicted by chance, t(15)=11.14 Furthermore, when these subjects were asked to classify the eight previously seen distinctive animals on the basis of memory, the mean proportion of errors (0.13, s.d.=0.17) was, again, significantly below chance, t(15)=8.71.

Table 12 Number of subjects viewing the Uniform in training following Standard "No Rule" instructions or viewing Distinctive stimuli in training following Analytically biased "No Rule" instructions who used the appropriate XOR rule, an inappropriate single feature rule, or some other strategy when viewing perceptually uniform stimuli in the transfer round (Experiment 4B).

	CLASSIFICATION STRATEGY		
TRAINING CONDITION	XOR Rule	Analytic Simplification	Other
Standard/Uniform	6	5	5
Analytic/Distinctive	1	8	7

Discussion

The evidence from subjects viewing the perceptually uniform stimuli indicates that a substantial proportion of subjects were able to discover the XOR classification rule. This seems to happen quite rarely, however, for subjects viewing stimuli that also possess unsystematic, idiosyncratic perceptual variation. One potential reason for this difficulty in discovering the critical combination of analytic features is the possibility that subjects viewing the perceptually distinctive stimuli have the additional task of identifying the dimensions of systematic variation. It is possible, for example, that chance orderings of items will produce dimensions of local systematicity. For three or four consecutive items, a subject might entertain the possibility that the presence or absence of feet is critical to classification, then subsequently realize that only one of the animals is lacking feet. A large number of such local systematicities might interfere with the discovery of the general rule for classification.

Several pieces of evidence, however, suggest that this is not the main reason for subjects' failure to identify the XOR rule. First, in post-experimental interviews, only one subject indicated using a dimension of variation during training that was not present in the uniform animals viewed in the transfer round (whether the animal was looking up or down). Thus, consistent with data reported earlier, it appears that subjects were quite adept at identifying the dimensions of systematic variation that were built into the perceptually distinctive stimuli. More convincing evidence, however, comes from spontaneous reports given by subjects at the end of the training round when they were told that they would now have to use the rule they had discovered to classify a different set of animals. At that point, a large number of subjects reported (often with some embarrassment) that they did not discover any rule, but simply learned the categories of the individual items. Consistent with these reports, eleven of the fifteen subjects who failed to discover the XOR rule in the Analytic/Distinctive condition were nonetheless 100% accurate in classifying the perceptually distinctive stimuli by the tenth round of training. It appears, therefore, that the nonsystematic perceptual variation in the distinctive stimuli did interfere with the discovery of the XOR rule. However, it does not appear to have had its effect by disguising the systematic dimensions of variation that subjects required for rule discovery. Instead, the idiosyncratic variation seems to have decreased rule discovery by providing a simpler method of classification, thereby allowing subjects to abandon analytic efforts prior to successful discovery of the rule.

Evidence from the Analytic/Distinctive condition also confirms the finding of Experiment 4A that analysis of the stimuli interfered with the acquisition of knowledge about the category membership of individuals. Consistent with Experiment 4A, the four subjects in this condition who were not 100% accurate in classifying the perceptually distinctive stimuli by the end of training (i.e. subjects who continued to try to find the analytic rule but failed) were only 62% accurate in their classification of Old perceptually distinctive items when they were later asked to classify these items using memory. In fact, these four subjects were the four least accurate subjects on this memory measure. Thus, again, there is evidence that a continued analytic bias interferes with learning the category membership of specific individuals. It is interesting to note, however, that although these subjects were quite poor at identifying the category membership of Old items, they were no less accurate than the other subjects in identifying which items were Old and which were the modified training twins (the mean error rate for the the four subjects on this memory measure was 0.10). Thus, subjects who were treating the stimuli analytically were nonetheless learning about the nonanalytic, idiosyncratic information necessary to identify individuals, but that this learning of individuals did not include learning about category membership.

Experiments 4A and 4B, then, indicate that the acquisition of analytic and nonanalytic category information are, at least to some degree, reciprocal. Analytic treatment of the stimulus set interferes with the nonanalytic tendency to learn about the category membership of individuals, and the presence of nonanalytic information (which promotes learning about individuals' category membership) interferes with the learning of more complicated analytic category information. The interference in both directions, however, appears to have its effect at the level of category learning. Analytic treatment of the stimuli decreases learning about the category membership of individuals, but it does not interfere with the learning of the individuals themselves. On an old/new recognition task, even subjects who had failed to learn each item's category by the end of training showed that they had learned the individual items (despite the fairly difficult discrimination required). Similarly, subjects viewing the perceptually distinctive stimuli were less likely to discover the general rule for categorization, but these subjects were able to identify the analytic structure of the individuals within the set. Thus, acquisition of analytic and nonanalytic information about individuals need not be reciprocal processes. The reciprocity may instead occur in learning to use these sources of information for the purposes of classification.

CHAPTER 6

General Discussion

There are three central assertions in this thesis. (1) For most natural stimuli, the dimensional structure that describes the subjects' analytic activity is not sufficient to describe the perceptual distinctiveness that supports concurrent exemplar-based categorization. (2) Holistic individuation (the distinctiveness of whole items) is not in general produced by increasing the distinctiveness of individual features and even has some priority over individual feature distinctiveness in supporting exemplar-based categorization. (3) Use of the two types of nonanalytic distinctiveness, holistic individuation and featural individuation, is not necessarily in a reciprocal relation with use of analytic knowledge. These three themes will be discussed in turn.

Perceptual Manifestations of Analytic Structure

The principal manipulation in these experiments was to vary the perceptual manifestation of a more or less identical analytic structure. There were several lines of evidence demonstrating that the dimensional variation called "analytic" was used by the subjects faced with an analytic task; the explicit dimensional descriptions elicited from the subjects in Experiment 1A, the free two-category sortings produced in Experiment 1B, the single dimension solutions (analytic simplifications) produced by subjects viewing the nondistinctive stimuli of Panel A in Figure 1, and the free reports of analytic activity given by subjects after the transfer phase in the relevant experiments all support the original designation of the analytic structure. However, the different perceptual manifestations of this analytic structure differed widely in their tendency to support transfer based on similarity to particular training instances. Clearly, different aspects of the stimuli are important for supporting these two different types of processing.

Materials and tasks such as the current ones, in which both analytic and nonanalytic categorization processes are occurring, provide a problem for some concept learning models. Many models calculate similarity between two stimuli as some function of the overlap in their dimensional structure (e.g. some function of the number of 1's and 0's shared in an abstract description such as that in Table 1). This same dimensional structure is also used to describe whatever analytic activity is occurring. But on the basis of the current results, there is no reason to believe that, except for some special stimulus sets, the same aspects of stimulus variation are important for both types of processing. This is particularly evident for the holistic variation, which cannot even be nested into the same description. That is, holistic variation cannot be treated as simply a more concrete or distinctive level of the analytically useful structure. For stimuli in which such effects occur, two different descriptions of the stimuli seem to be required.

I suspect that there is nothing rare about stimuli and tasks that require two different descriptions of the stimuli. Most natural stimuli show perceptual variety and holistic integrity at least comparable to that evident in the perceptually distinctive stimuli (e.g. Panel D of Figure 1). What is unusual about the perceptually distinctive stimuli found here is that they also show such obvious analytic variation. However, in many natural situations with less obvious analytic dimensions, people either have a prior established theory or have received deliberate instruction that contains a useful analytic description. Both the strengths and the limits of such analytic descriptions are shown especially in areas that require a combination of textbooks and years of practice, such as dermatology, radiology, and histology (see Brooks, Norman & Allen, 1991; Allen, Norman & Brooks, 1992; Norman, Brooks, Coblentz & Babcook, 1992). Experimental models that contain the affordances for both analytic and nonanalytic processes would provide a useful tool for investigating an interesting class of natural problems.

The Priority of Holistic Individuation

The manipulations in this paper distinguished between featural individuation (whether or not a feature occurred in identical form in different items) and holistic individuation (the extent to which the separate features of an item cohere into an individuated whole). Logically, individuated features could be sufficient to produce analogy to individual items; a new item possessing several features that previously had been uniquely associated with a particular training item might be categorized by analogy with that item. However, the current evidence demonstrates that for the conditions in this paper this is not sufficient.

Experiments 3B and 3C were designed to test the relationship between holistic and featural individuation. In the transfer phase, the two analytically *irrelevant* features were altered substantially. For the transfer stimuli of Experiment 3B, this meant that the holistic quality of each stimulus was substantially altered as well. The effect of this alteration was the elimination of episodic transfer in subjects' classifications. In Experiment 3C, the same features were altered in the same way. Because the transfer stimuli were broken apart, however, the holistic nature of the stimuli was not altered, it was entirely destroyed. Under these circumstances, the episodic transfer effect returned. It appears, therefore, that when both holistic and unique feature information are available in the training stimuli, both sources of information are acquired under No Rule training conditions.

Yet, if both sources are present at the time of transfer, the holistic individuation of the stimulus seems to have a stronger influence on classification. If, as in Experiment 3B, the holistic nature of the stimulus is altered, the resulting stimulus is an unfamiliar individual. In this context of unfamiliarity, subjects fall back on an analytic pattern of responding. If, however, the holistic nature of the stimulus is entirely disrupted, the stimulus is no longer an individual at all, but a set of parts. Under these circumstances, the focus seems to shift to the separate features, and the perceptual similarity of some of these features to the old training items is sufficient to produce more episodic, similarity based responding. Thus, the presence of holistic individuation in the training stimuli does not seem to interfere with the *acquisition* of knowledge about the items' perceptually unique features. The continued presence of holistic individuation during transfer, however, does appear to interfere with the *use* of this knowledge about unique features.

Perhaps this priority of the whole is not surprising in the context of No Rule training. Subjects were simply being presented with a series of individuals and asked to learn the category membership of these individuals. Although there is evidence that subjects attempted to treat these individuals analytically, there was no specific injunction from the experimenter that they do so. This was not the case, however, under the

Rule training conditions. Here subjects were given explicit instructions regarding their treatment of the stimuli. They were told to pay attention to specific features, that only these features were relevant to classification. The specific features identified by the rule remained perceptually unaltered in the transfer stimuli. Despite this strong encouragement to treat the stimuli as a set of parts, however, it was again the holistic nature of the stimuli rather than the unique perceptual manifestation of the separate features that seemed to intrude on analytic performance.

This priority of the whole is probably influenced by the way the subjects think about the stimuli. The items were presented in training as whole animals, not as sets of disassembled parts. Consequently the subjects were likely to think of them as whole animals and fail to react to transfer animals that seemed to be essentially different for whatever reason. It is expected that if subjects had been asked to categorize the legs rather than the whole animals, then holistic priority would have accrued to the legs rather than the whole animal despite the continued figural connectedness of the whole animal. That is, the holistic properties are probably at least partly relative to the way the items are framed (for a review of framing effects in concept learning, see Barsalou, 1991).

Finally, it is worth noting that although holistic integration in this paper was produced by manipulating the form of the stimuli, similar effects can be produced by varying the way the subjects process exactly the same stimuli. Whittlesea and his colleagues (Whittlesea, 1987; Whittlesea & Cantwell, 1988; Whittlesea & Brooks, 1989; Whittlesea, Brooks & Westcott, 1992) have shown large variations in the extent to which a stimulus supports exemplar-based generalization by varying whether the subjects treated a set of pseudowords as individuals or as data for judgments about the category as a whole.

The Nonreciprocal Relation between Analytic and Nonanalytic Resources

The evidence in this thesis indicates that neither the acquisition nor the use of analytic and nonanalytic information are necessarily antagonistic to one another. In the Rule conditions, subjects started the training phase with an analytic rule that was sufficient for perfect categorization. Nonetheless, this did not prevent them from acquiring and using sufficient nonanalytic knowledge to produce a bad transfer effect. Allen and Brooks (1991) have demonstrated that this effect is robust across experimental variations such as prompting the subjects for accuracy and warning them about the presence of "trick" bad transfer items. The subjects in the No Rule conditions used in the present thesis also learned about and made use of both analytic and nonanalytic knowledge. Post-experimental interviews with the No Rule subjects viewing perceptually distinctive stimuli indicated strong analytic tendencies. As with subjects viewing the perceptually uniform animals, these subjects attempted to isolate the dimensions relevant to category membership. As seen in Experiment 3B, subjects used this analytic knowledge to classify perceptually novel items with a moderate level of accuracy (a level essentially equal to that achieved by subjects viewing the perceptually uniform stimuli under No Rule training conditions). Thus, it appears that the No Rule subjects' acquisition of knowledge about the individuality of specific cases did not preclude their acquisition of knowledge about the analytic information contained in those cases. Rather, when distinctive perceptual information was available in the stimulus, subjects used it to supplement their incomplete analytic knowledge, thereby improving their accuracy in the classification of old and perceptually similar new items. When this information was not apparent, the analytic information was used exclusively and accuracy in classification was limited to the adequacy of this analytic knowledge.

The assertion that both analytic and nonanalytic knowledge are acquired and used during classification is consistent with the data reported by Vokey and Brooks (1991). They presented subjects with a set of letter strings to memorize, all of which were generated using a complicated, artificial grammar. Following this "training" phase, subjects were informed of the existence of the grammar and were then presented novel letter strings. The subjects' task was to indicate whether each of the new items obeyed the grammar. These new strings varied both on their grammaticality (whether or not they were consistent with the grammar) and on their similarity to a previously studied item (one letter away vs. several letters away). Vokey and Brooks found that when grammaticality and similarity-to-old were orthogonalized experimentally, they became independent, noninteracting sources of variance in subjects' classification behaviour. Again, the analytic and nonanalytic sources of information were not in a reciprocal relationship. Rather, subjects were learning about, and making use of both sources of information during the experiment. A similar additive relationship between the factors of context-free similarity (a measure of the structure of the set) and specific memorability was found for recognition memory by Vokey and Read (1992).

Logan (1988) has proposed a model of automaticity that makes explicit use of the notion that both episodic and analytic sources of information are available for the performance of tasks. He used tasks such as basic arithmetic in which he sees two methods for producing an answer: the purposeful application of an algorithm, and the instance based memory for one's previous solution to the problem. Performance on a problem is determined by a race between the two processes. As practice continues, the number of instances increases, thereby increasing the probability that the fastest access of an instance will be fast enough to beat the algorithmic solution. According to Logan's model, automatic performance is the end result of a large number of episodes. Under these circumstances, an episodic solution will almost always beat the algorithmic solution. Thus, automatic performance is equated with instance based processes whereas purposeful, reflective performance is equated with more the more analytic application of the algorithm. Allen and Brooks (1991) have questioned the use of a straightforward race model to describe the pattern of results found in experiments similar those presented here. Logan's general approach, however, is quite consistent with the one being presented here. Both the analytic and nonanalytic sources of information are acquired during training. In Logan's model, competition between the two is not at the level of acquisition, but at the level of response control.

The hypothesis that analytic and nonanalytic sources of information are acquired independently with experience is also being asserted in the memory literature. Jacoby (1991; see also a more general review in Jacoby, Ste-Marie, & Toth, 1992) has developed a technique to isolate the effects of conscious recollection and unreflective familiarity on subjects' classification of items as "old" or "new" in a recognition memory paradigm. Using the technique, Jacoby has shown that more intensive processing of an item (experiencing the item as an anagram that must be solved as opposed to simply reading the item) increases not only the probability that the item will be recollected, but also the probability that the item will feel familiar. That is, a manipulation that increases one source of memory does not necessarily produce a decrease in the other. The processes described by Jacoby in the framework of the memory literature may map very well onto the analytic/nonanalytic distinction being used in the categorization literature. If this is true, then his process dissociation technique may become a powerful tool in our attempts to isolate and examine the two processes in categorization, and may provide converging evidence for their independence.

This emphasis on methods that allow independent variation of analytic and nonanalytic information is also the major difference between the current work and that of Kemler Nelson (1989) and J. D. Smith (1989). The distinctions made in this thesis are closely related to distinctions made by those authors. Kemler Nelson distinguishes three different senses in which processing can be nonselective; that is, three different ways in which judgments of similarity can depend relatively unselectively on many aspects of the stimulus. These differ in the extent to which judgments of overall similarity result from independent processing of constituent stimulus attributes. One extreme is *analytic* processing in which the stimulus attributes are independently processed and then the similarity judgment is constructed from these already identified attributes. In this form of processing, the calculation of similarity is clearly secondary to the processing of the stimulus attributes. On the other extreme is *strong holistic* processing in which processing the individual properties is not a part of the apprehension of overall similarity. The model of such strong holistic processing is the integral or configural processing discussed by Garner (1974) in which the underlying dimensions do not have psychological primacy in judging distance between two stimuli. In between is *weak holistic* processing, in which the overall similarity computation is clearly primary but in which specification of individual properties occurs as a part of this computation.

Using Kemler Nelson's distinctions, it possible that what I refer to here as holistic individuation corresponds to weak holistic processing. Because the interest lay in facilitating concurrent analytic processing, I did not use stimuli that are likely to support strong holistic processing (in which the component dimensions lose all priority). Thus, within my work, reference to holistic processing does not necessarily imply that the subject is being totally nonselective. Because of the high availability of the analytic dimensions, a subject's assessment of what characterizes an item as a whole may not equally weight all component properties nor even utilize them exhaustively. Similarly, the whole point of the perceptual manipulations was to demonstrate the insufficiency of a strictly analytic computation of similarity. Because of the analytic affordances that were built into the current stimuli, it is not possible to use the specific free and directed sorting methods that Kemler Nelson and J. D. Smith used to diagnose these various forms of processing. However, it is clear that, descriptively, their distinctions are very similar to those being presented here. Indeed, the current materials could provide some interesting extensions to the ideas described by J. D. Smith (1989) and Smith and Shapiro (1989) on the relation between holistic and analytic processing.

Finally, consistent with Lockhead (1970), it is important to point out that the degree to which an item will be individuated by the perceptual manifestation of its analytic structure must be a function of the set in which the item is presented. Imagine, for example, that a single item from the perceptually uniform set of panel A were to replace its analytic equivalent in panel D. The item would now be perceptually different from all other items in the set both in the manifestation of its separate features, and in the impression that it projects as a whole. Under these circumstances, it is plausible to assume that the item would show episodic effects. It important to note, however, that one would expect these effects to emerge only if one also swapped the appropriate transfer item. To the extent that this is true, the perceptual similarity between two items within a given stimulus set is not well predicted by the analytic structure alone.

Conclusions

This set of studies has served to highlight the importance of the perceptual array in the classification of items within a stimulus set. Whether faced with perceptually uniform or perceptually distinctive stimuli, subjects can induce and make use of abstract categorical information. What is of interest, however, is the fact that subjects challenged with perceptually distinctive stimuli will often make use of this item-specific information in addition to the analytic information. In fact, this tendency to make use of perceptually idiosyncratic information is sufficiently powerful that it will often persist even under circumstances in which subjects possess analytic information that is perfectly predictive of category membership.

It should be stressed that this perceptual idiosyncrasy is not merely at the level of the item's separate features. Clearly, the unique perceptual manifestation of an item's separate features affect classification behaviour. Often, however, its effect on classification behaviour is in its contribution to the individuated nature of the item as a whole. It is the presence of this holistic individuation that seems to affect subjects' application of an explicit classification rule most strongly.

The finding that item classification is a function both of abstract, category level information and of item specific information is becoming common in the literature (see, for example, Allen & Brooks, 1991; Brooks, 1987, 1990; Jacoby & Brooks, 1984; Malt, 1989; Medin et al., 1983; Medin & Ross, 1989; Ross, 1987; Whittlesea & Brooks, 1988). Such findings have led to the description of classification behaviour as a hybrid of both rule based and item based knowledge. Inevitably, such descriptions lead to questions regarding the relative availability and the coordination of the two bases of classification. The point here is that the answer to this question is not a general one. Instead, conclusions of this sort must include a statement about the perceptual composition of the stimulus set being classified.

REFERENCES

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., Norman, G.R., & Brooks, L.R. (in press). Experimental studies of learning dermatological diagnosis: The impact of examples. *Teaching & Learning in Medicine.*

Barsalou, L.W. (1991). Frames, concepts and conceptual fields. In, A. Lehrer and E Kittay (Eds.), *Frames, Fields, and Contrasts.*

Brooks, L. R. (1976). Visual pattern in fluent word identification. In A. Reber & D. Scarborough (Eds.), *Toward a psychology of learning* (pp 143-181). Hillsdale, N. J.: Erlbaum.

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

Brooks, L. R. (1983). *The insufficiency of analysis.* Unpublished manuscript. McMaster University, Hamilton, Ontario, Canada.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp 141-174). Cambridge: Cambridge University Press.

Brooks, L. R. (1990). Concept formation and particularizing learning. In P. Hanson (Ed.), *Information, language, and cognition: Vancouver studies in cognitive science. Vol. 1*, Vancouver: UBC Press.

Brooks, L.R., Norman, G.R., & Allen, S.W. (1991). The Role of Specific Similarity in a Medical Diagnostic Task. *Journal of Experimental Psychology: General, 121*, 278-287.

Bruner, J., Goodnow, J., & Austin, G. (1956). *A study of thinking*. New York: John Wiley.

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

Foard, C. F., & Kemler Nelson, D. G. (1984). Holistic and analytic modes of processing: The multiple determinants of perceptual analysis. *Journal of Experimental Psychology: General, 113*, 94-111.

Franks, J.J. & Bransford, J.D. (1971). Abstraction of visual patterns. *Journal of Experimental Psychology, 90*, 65-74.

Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.

Hayes-Roth, B. & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning and Verbal Behavior, 16*, 321-338.

Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments, and Computers, 16*, 96-101.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513-541.

Jacoby, L. L. & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory, Vol. 18*, 1-47.

Jacoby, L. L., Ste-Marie, D. & Toth, J. P. (1992). Redefining automaticity: Unconscious influences, awareness and control. In A. D. Baddeley & L. Weiskrantz (Eds.), *Attention, Selection, Awareness and Control. A Tribute to Donald Broadbent.*

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 and process.* Hillsdale, N. J.: Erlbaum.

Kennedy, J. J. (1983). *Analyzing qualitative data.* New York: Praeger.

Kohler, W. (1929). *Gestalt psychology.* New York: Liveright.

Lockhead, G. R. (1966). Effects of dimensional redundancy on visual discrimination. *Journal of Experimental Psychology, 72*, 95-104.

Lockhead, G.R. (1970). Identification and the form of multidimensional discrimination space. *Journal of Experimental Psychology, 85*, 1-10.

Logan, G.D. (1988). Toward an instance theory of automatization. *Psychological Review, 95*, 492-527.

Malt, B. C. (1989). An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning Memory and Cognition, 15*, 539-555.

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. & Ross, B. H. (1989) The specific character of abstract thought: Categorization, problem solving and induction. In R. Sternberg (Ed.), *Advances in the Psychology of Human Intelligence, Vol. 5*, New York: Academic Press.

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

Medin, D. L., Wattenmaker, W. D. & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology, 19*, 242-279.

Modigliani, V. (1971). On the conservation of simple concepts: Generality of the affirmation rule. *Journal of Experimental Psychology, 87*, 234-240.

Modigliani, V. (1974). Abstraction of unidimensional concepts from larger conceptual systems. *Journal of Experimental Psychology, 102*, 341-343.

Modigliani, V. & Rizza, J. P. (1971). Conservation of simple concepts as a function of deletion of irrelevant attributes. *Journal of Experimental Psychology, 90*, 280-286.

Norman, G.R., Brooks, L.R., Coblentz, C.K., & Babcook, C.J. (in press). The interdependence of feature and category identification in diagnostic radiology. *Memory & Cognition*

Nosofsky, R. M. (1984). Choice, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: Learning, Memory and Cognition, 10*, 104-114.

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

Posner, M.I. & Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77,* 353-363.

Ross, B. H. (1987) This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13,* 629-639.

Saltz, E. (1972). *Learning and conceptual behavior.* Homewood, Ill.: Dorsey Press.

Shepard, R.N. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology, 1,* 54-87.

Smith, E.E. & Medin, D.L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Smith, J. D. (1989). Analytic and holistic processes in categorization. In B.E. Shepp & S. Ballesteros (Eds.), *Object perception: Structure and process.* Hillsdale, N. J.: Erlbaum.

Smith, J. D. & Kemler Nelson, D. G. (1984). Overall similarity in adults' classification: The child in all of us. *Journal of Experimental Psychology: General, 113,* 137-159.

Smith, J. D. & Shapiro, J. H. (1989). The occurrence of holistic categorization. *Journal of Memory and Language, 28,* 386-399.

Smith, L. B. (1989). A model of perceptual classification in children and adults. *Psychological Review, 96,* 125-144.

Smith, L. B. & Evans, P. (1989). Similarity, identity, and dimensions: Perceptual classifications in children and adults. In B. E. Shepp & S. Ballesteros (Eds.), *Object perception: Structure and process.* Hillsdale, N. J.: Erlbaum.

Vokey, J. R. & Brooks, L. R. (1991). The salience of item knowledge in learning artificial grammars. *Journal of Experimental Psychology: Learning, Memory and Cognition*, Manuscript submitted for publication.

Vokey J. R. & Read, J. D. (in press). Familiarity, memorability and the effect of typicality on the recognition of faces. *Memory & Cognition.*

Whittlesea, B. W .A. (1987). Preservation of specific experiences in the representation of general knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13*, 3-17.

Whittlesea, B. W. A. & Brooks, L. R. (1988). Critical influence of particular experience in the perception of letters, words and phrases. *Memory & Cognition, 15*, 465-472.

Whittlesea, B.W.A., Brooks, L.R. & Westcott, C. (1992). Utilization of analytic and nonanalytic processing: Factors affecting reliance on general and particular knowledge. Submitted manuscript.

Whittlesea, B.W.A., & Cantwell, A.L. (1987). Enduring influence of the purpose of experiences: Encoding-retrieval interactions in word and pseudoword perception. *Memory and Cognition, 15*, 465-472.