THE CONTRIBUTION OF SPECIAL EXPERTISE TO THE IDENTIFICATION OF PICTURES

By

John Gerard Baker, B. Sc.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree
Doctor of Philosophy

McMaster University

SPECIAL EXPERTISE IN PICTURE IDENTIFICATION

DOCTOR OF PHILOSOPHY (Psychology)

McMaster University
Hamilton, Ontario

TITLE: The Contribution of Special Expertise to the Identification of Pictures

AUTHOR: John Gerard Baker, B. Sc. (University of Lethbridge)

SUPERVISOR: Professor L. R. Brooks

NUMBER OF PAGES: viii, 93

ABSTRACT

Determining how people identify objects and pictures in real world settings is one of the major challenges in psychology. Under normal conditions, objects and pictures are subject to an extraordinary amount of variability in perceptual characteristics. Within object categories, there are a wide variety of objects, and for any given object. there is a wide variety of possible viewing conditions. The ease with which people overcome the variability in the surface characteristics of objects and pictures is the phenomenon of interest here.

Current theories of object identification stress robust identification procedures. By the robustness view, identification proceeds by identifying reliable stimulus characteristics, and assigning category on the basis of these characteristics. Idiosynchratic, stimulus specific features are not considered a normal path to object identification. In abstractive theories of concept learning, this view of object identification is often implicit, but sometimes it is an explicit design feature of the theory. In contrast, episodic views of concept learning stress the use of prior items in categorization. Implicit in such a view is that adventitious stimulus characteristics that have been used previously to good advantage will be exploited in object identification. This is akin to Newell and Simon's view of expertise in which the expert adapts prior solutions to new problems rather than apply the same algorithmic solution again. This adaptation of prior solutions to new problems will be termed special expertise.

It is widely recognized that robust identification procedures do

not exist for all objects and viewing conditions. Therefore, normal object and picture identification is a mixture or hybrid of robust procedures and special expertise. However, it is necessary to place some empirical constraints on the contribution of special expertise to determine whether special expertise is a special case reserved for unusual perceptual problems or whether it is a normal part of object perception.

The following series of experiments assesses the contribution of special expertise to picture identification by replicating and extending prior work. Additional work is reported that clarifies factors potentially important to special expertise. Implications for theories of concept learning will be discussed.

ACKNOWLEDGEMENTS

A thesis is the result of the contributions of many people. My supervisory committee has been of invaluable benefit to me in my work. I owe a great debt of gratitude to my supervisor, Dr. Lee Brooks. He has been a source of a great many insights, a constant source of encouragement and support, and has been tireless in his efforts to assure I got the most out of my graduate education. I would like to thank Dr. Larry Jacoby. Our discussions have been very productive and rewarding. His expertise and insight has broken conceptual roadblocks on many occasions, and has contributed immeasurably to the work reported here. I would also like to thank Dr. Betty Ann Levy for her interest, enthusiasm, and encouragement.

I would also like to thank Dr. John Vokey for introducing me to cognitive psychology and inspiring me to pursue graduate studies. He is and will always remain my role model as an academic. I am greatful to my student colleagues. Two warrant special mention: Scott Allen and Jane Collins. They have been good friends and colleagues. If in my future career I meet others of their caliber, I will indeed be fortunate.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

One of the central challenges facing cognitive psychology is the question of how people identify natural objects and classify them into conventional naming categories. People can identify the objects about them very quickly, accurately, and without apparent effort. The central issue in the work reported here is what kind of knowledge we must have about object categories to produce this kind of easy identification performance.

An important consideration for any theory of object identification is that the conditions of presentation of real world objects are subject to considerable variability. First, there is tremendous variety within object categories. Chairs, dogs, and trees can each have a variety of very different appearances. Second, there is a great deal of variability in the manner of presentation of each individual object. Lassie in profile presents a very different set of perceptual properties than does Lassie bounding through tall grass towards the viewer. The changes in orientation are irrelevant to the categorical status of the object, but greatly affect the information available for the purposes of identification. In some way, the system of representation must deal with these sorts of changes in perceptual characteristics.

In addition, people often do not use all the information potentially available in any given presentation. Consider the series of pictures in Figure 1. This series of pictures mimics a common kind of identification problem. As the object emerges from behind an occlusion, more visual information becomes available. At some point, typically before full presentation of the object, identification becomes possible. This kind of identification problem is used in the following experiments.

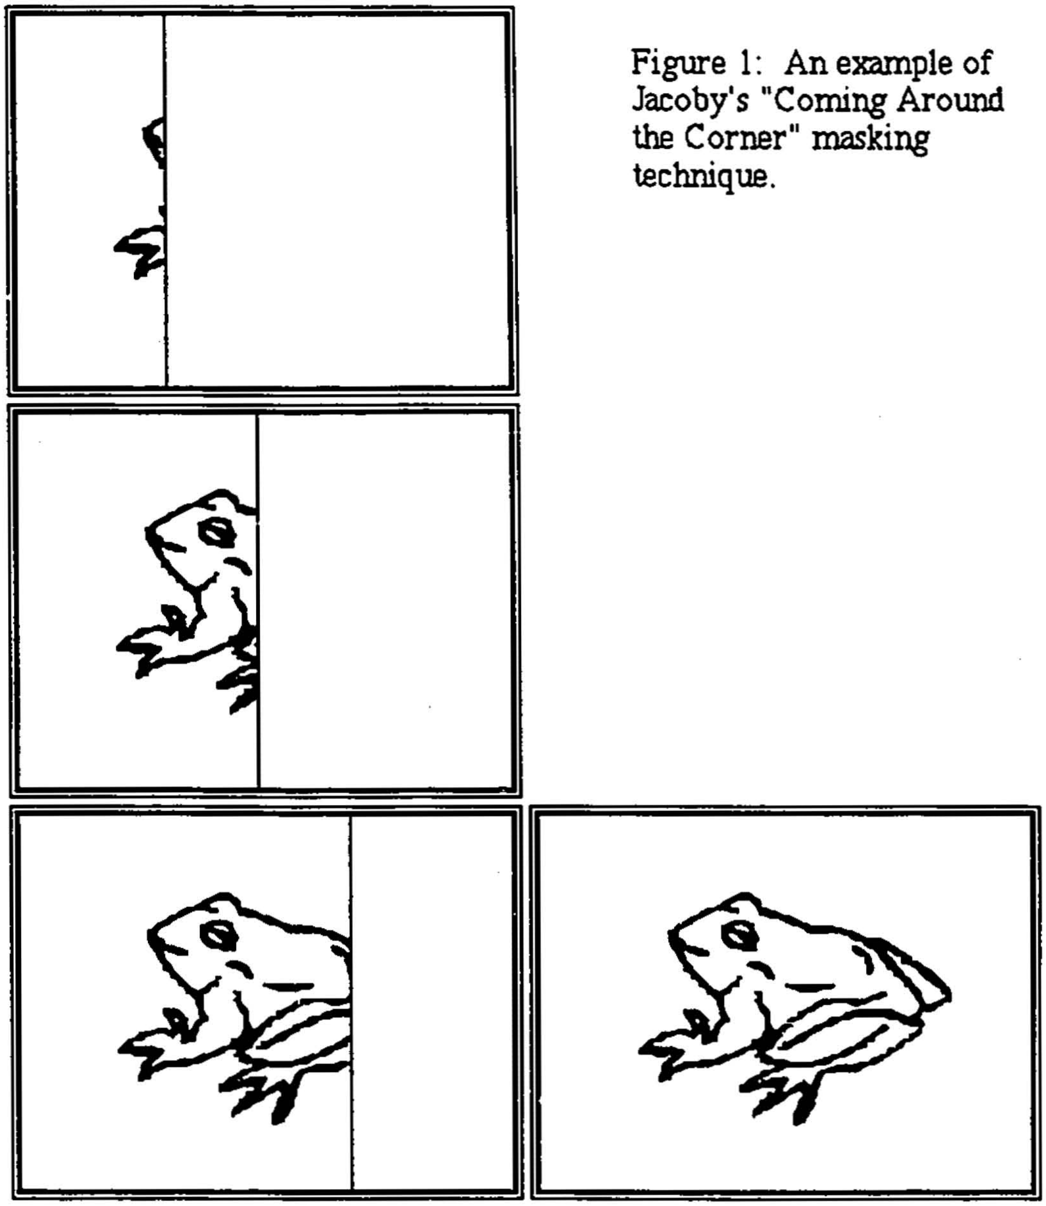

Figure 1: An example of Jacoby's "Coming Around the Corner" masking technique.

The assumption behind these experiments is that the identification of drawings under conditions of occlusion and variable appearance makes demands of the cognitive system similar to those routinely exploited in operating in the world, and do not represent unusual skills or knowledge.

The identification of objects without complete visual information is more common than might be obvious. Most objects are partially or, more commonly, completely opaque and therefore information on the far side of an object is occluded from view and cannot be used to identify the object. In fact, it is difficult to think of objects for which occlusion is not a problem. Despite the near universality of occlusion, variations in the viewpoint of the observer, and variations among different members of the same category, object identification is fast and accurate. Clearly, the cognitive system has some way of identifying objects despite the general problem of perceptual variability.

There are two general approaches to dealing with the issue of perceptual variability. The most common approach is to make identification primarily dependent upon robust characteristics: characteristics that are largely invariant across individuals within a category and across different views of an individual. That is, identification would not rely upon stimulus specific or highly variable perceptual features. This approach of relying on "necessary and sufficient properties" has been termed the Classical view by Medin and Smith (1984), arguing that historically it has been the dominant approach in philosophical and psychological discussions of category identification.

One possible explanation for the intuition that stimulus specific properties are not a normal part of object identification is that this intuition results from our use of an alphabetic writing system. In an alphabetic system, words are constructed from the essentially independently recombinable components, letters. However, the alphabetic writing system may be a poor model of how we represent object categories. In describing an object, one refers to an object as having legs, a body, a nose, eyes, etc. as if each of these object components was generic in the same sense the letter 'e' is generic. However, it is impossible to render these sorts of object components generically. Dogs have dog noses and cows have cow noses. One could not assemble a cow by choosing four legs from all possible legs, a nose from all possible noses, a tail from all possible tails, and so on. There is a kind of perceptual variability not reflected in component names. While relying only on invariances for identification logically would deal with the general issue of variability, there are no apparent invariances in the immediately available stimulus.

A related version of the robust features approach is to have identification depend upon invariant "deep" features related to surface features by specifiable transformations. This approach is analogous to the use of transformational grammars in linguistics. As an example of this approach, Biederman and his colleagues (Biederman, 1987; Biederman and Ju, 1988) propose a theory of object identification based on generic object components called geons. Early perceptual processing recovers information from the perceptual array necessary to define the geons of an object, and discards surface variability. Geons are volumetric components created by generalized cones. A generalized cone is the volumetric form created by sweeping a two dimensional form along an axis through the third dimension. For example, a simple cylinder is created by sweeping a circle along a straight axis for some distance at right angles to the plane of the circle.

Generalized cones can have a variety of shapes. The axis through which the two dimensional form is swept may be straight or curved. For example, an arc is created by sweeping a circle 180° along a circular axis. In addition, the two dimensional form can change shape as it moves along that axis. Any simple three dimensional form can be described as a generalized cone or geon, and any complex three dimensional form can be described as a collection of these geons. Since objects are three dimensional, any object can be described in terms of geons. For example, an arc connected to the top of a cylinder could represent a bucket, while an arc connected to the side of a cylinder could represent a cup. See Figure 2 for an illustration of geon descriptions of objects. These collections of geons can then be thought of as a lexicon of objects. By this view, each object category is represented by a small number of geons, ranging from one to several.

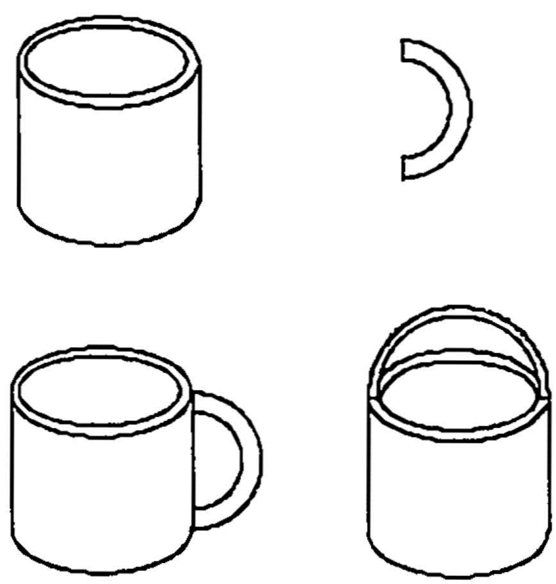

Figure 2: An illustration of how geons could be used to represent objects. A simple cylinder and an arc can be combined to form a cup or a pail. Each of these objects could be defined by the same two geons and their points of attachment.

The attraction of geons is their robustness. Across changes in orientation, all that changes in the geon description is the aspect ratio of the geons[1], and not the geons themselves. Therefore, geon descriptions are robust across changes in orientation. An important consequence of this scheme is that the information used in the identification of objects does not include information subject to change with changes in the object's orientation. Further, Biederman and Ju (1988) argue that properties likely to vary between category members, such as color and texture, are not used in normal object identification or "primal perceptual access" for objects. Thus, Biederman's suggestion is that objects can be described as a combination of independent geons in much the same way words are formed from independent letters. By this proposal, the robustness of geons and their independence of variable qualities such as color and texture, makes the robust feature approach a feasible way of dealing with perceptual variability.

A different principle for making the robust knowledge approach to the problem of perceptual variability work is represented by the prototype view of category identification. By this view, categories are represented by a host of characteristic features. Each feature is characteristic of the category in that it occurs more often in that category than in alternative categories. Because there are a large number of them and no one characteristic is necessary, different views of the object will generally provide a sufficient number of features to allow identification. The conjunction of these characteristics is the prototype, and identification normally proceeds by comparing a new object to this knowledge structure.

[1]The aspect ratio of a geon is the ratio of the length of the generalized cone over the average width of the two dimensional form swept through that length. For the example of a cylinder, the aspect ratio is the length over the diameter.

Prototypes are robust across conditions of perceptual variability by containing enough redundancy to cover most viewing conditions.

The alternative to the robust knowledge approach to solving the perceptual variability problem is the special expertise approach. This approach is implied by the exemplar view of concept learning in which concept knowledge is represented by prior instances (Brooks, 1978; Medin and Schaffer, 1978). By this view, identification proceeds by exploiting adventitious stimulus characteristics. Imagine a difficult identification problem, such as a shape looming out from a snowstorm. At some point, identification of the looming object will become possible. This experience will now be a resource that can be used in a subsequent identification problem. When confronted by an object looming out of another snowstorm, the object may be identified by consulting the prior item.

This continuity between episodic memory and identification has been developed by Jacoby and Brooks (1984). They argue that the facilitative effects of prior exposure on an identification task can be profitably discussed in terms of episodic memory. For example, Jacoby (1983a) demonstrated that a single prior exposure to a common word can facilitate performance on a perceptual identification task at least five days later. Effects of this durability cannot be attributed reasonably to transient priming[2], but rather reflect a long term learning, or episodic memory. To further demonstrate that the facilitation resulted from episodic memory, Jacoby (1983b) demonstrated that the details of processing substantially affected transfer performance: a finding not easily reconciled

[2]Priming is commonly used in two very different senses: as a description of an empirical finding, and as an explanation for that finding. When referring to priming, I refer to its explanatory, rather than descriptive sense.

with a simple priming approach. This dependence of transfer on the details of a prior processing episode has been extended to the domain of picture identification by Jacoby, Baker, and Brooks (1989). This extension is critical to the current thesis and will be described subsequently in greater detail.

A way of thinking about the special expertise approach is as an adaptation of Newell and Simon's (1972) view of expertise to the perceptual domain. By the Newell and Simon view, the expert does not have to solve the same problems over the way they were before. Rather, a prior solution is adapted to the current problem. In terms of object identification, an expert parser is one that does not have to solve the same difficult identification problems over and over the same way. Instead, prior solutions can be adapted directly.

A very basic test of a view is that of plausibility. Would it make sense to design a system that relied upon special expertise? If a parser (either grammatical or perceptual) was confronted with the same problem several times, would the problem be treated exactly the same way every time? If it was, then the parser would have failed to capitalize on the tendency of the world to repeat larger patterns. This illustrates one attractive property of the episodic view. Generalization from a restricted cohort of prior events represents a tailoring of the knowledge base to the redundancy of the world.

Special expertise can be used to account for automaticity: the phenomenon in which the performance of a task on a particular set of stimuli becomes faster and (typically) more accurate with practice. Logan (1988, 1990) has advanced a theory in which automaticity results from the retrieval of prior events rather than speeding up or improving the efficiency of an algorithm. The operation of a parser can be thought of as an algorithmic solution to a problem. However, by Logan's view, the parser or algorithm is used less with practice, and the retrieval of prior events comes to dominate as automaticity develops. This is a clear example of the special expertise approach.

Another implication of the special expertise view is that if prior experiences with an object category are diverse, then there is likely to be at least one similar prior experience for a given test stimulus. Thus, the diversity of prior experience should be an important component of generalization. That is, it is more likely that a person will be able to adapt a prior solution to a given test stimulus if the person has a variety of prior experiences. A result reported by Bartram (1974) speaks to this hypothesis. Bartram used a picture naming task in which subjects were required to quickly name several pictures of each of several common object categories. If subjects had named several different examples of an object category, they were much quicker to name a new example of that category than if they had previously named the same number of pictures of a single category member. Such a finding cannot be accommodated by a priming explanation, as the number of prior exposures to the category are equal across conditions. Further, the effects of prior exposure on speeded picture naming have been demonstrated to be very long lasting. Mitchell and Brown (1988) found that a single prior naming of a picture can speed subsequent naming of that picture for at least six weeks. This result suggests that naming pictures is subject to special expertise.

The stress on either robustness or special expertise represents extreme positions. By stressing robustness, models of object identification rely upon categorically relevant invariances. In stressing special expertise, the attempt is to rely upon adventitious stimulus properties. A more likely position is that some aspects of both are relevant to object identification. However, hybrid models can be underspecified. It is necessary, then, to determine some empirical constraints upon the contribution of special expertise to object or picture identification. The goal of the work reported here is to document the extent to which special expertise contributes to the identification of drawings.

The specific question asked in the following research is what kind of information do people use to identify drawings under conditions of distortion. Conditions of visual distortion, such as those in the 'coming around the corner' masking technique, represent a kind of variability in perceptual information that can be characteristic of object identification in the world. The way in which this question is approached is to examine the consequences of specific prior experiences on subsequent identification performance. By examining the consequences of specific prior experiences in detail, it will be possible to examine what is learned, and what is used subsequently to identify pictures.

There are three potential outcomes to this general line of research. First, it may be that picture identification is independent of any single prior experience. The identification of a given stimulus might depend upon a fixed procedure, and that procedure changes sufficiently slowly that a single prior experience will not substantially change the amount of visual information necessary. Since cognitive psychology is replete with examples of the effects of specific knowledge and expectation in performance of cognitive tasks, it is most unlikely that the current experimental preparation will act otherwise.

The question then becomes one of the nature of the effects of a specific prior experience. By the prototype view, a prior experience might serve to prime the prototype, thereby reducing the amount of stimulus information necessary for identification, provided the stimulus is presented before the prototype loses its activation. The second type of outcome from prior presentation, then, would be temporary priming of the category. On the other hand, the episodic view suggests that the details of presentation will be important to the amount of information necessary for identification. If a subsequent event presents the same perceptual problem as that previously experienced, then one should show maximum benefit from the prior experience. However, by changing the stimulus or by changing the conditions of distortion, the benefit should be reduced. The third type of outcome would be durable facilitation that is graded according to the match in details between the prior presentation and the subsequent test item.

There are several key results from prior research using the picture identification task. The prior presentation of an item facilitates subsequent identification of that item through distortion. Further, having previously seen the item through distortion facilitates performance more than having seen the intact item (Gollin, 1960; Jacoby, Baker, and Brooks, 1989; Snodgrass, 1987, 1990). An extension of this finding was reported by Snodgrass (1987). She found that if subjects were previously shown fragmented pictures, test performance was better if the same fragments were shown in test than if different fragments were shown. These findings suggest that the basis of facilitation is not general to the category or even the category member, but rather are specific to how that item was encountered. This processing dependency is a prominent feature of the episodic view of concept learning, and central to the special expertise view.

One of the most diagnostic tests of whether a finding results from priming or episodic memory is the durability of the effect. If the effect is short lived, then priming would seem a good explanation since priming is supposed to be short lived, while if the effect is long lived then episodic memory would seem the better explanation since episodic memory is not supposed to be short lived. Parkin and Streete (1988) found that prior presentation facilitated subsequent identification of fragmented versions for at least two weeks. Further, they did not detect any decrement in facilitation across those two weeks. This further suggests that facilitation in the task is special expertise rather than priming of a stable identification procedure.

Essential to the special expertise view is that of generalization: the process by which prior solutions are adapted to new problems. Prior presentation of an item facilitates identification of other members of the same category (Jacoby, Baker, and Brooks, 1989; Vokey, Baker, Hayman, and Jacoby, 1986). A result reported by Jacoby, Baker, and Brooks (1989) is especially relevant to this point. Not only did they find that prior presentation of an item facilitated identification of other members of that name category, but the amount of facilitation depended upon how that prior item had been processed. If the prior item had been viewed through distortion, there was relatively less facilitation on same name items than if the item had been viewed intact. This is important because it suggests that performance on same name items is under control of the same processes as those controlling performance on previously seen items.

The experiments reported here extend the demonstrations of the contribution of item memory to identification performance. To avoid the potential problem that the findings discussed above result from specific experimental preparations, the demonstrations will be extended to a variety of experimental designs and arrangements. This is done in order to buttress the assertion that the use of item memory in picture identification is sufficiently robust and ubiquitous that it should be considered a normal component of picture and object identification.

In the following experiments, three rather different kinds of masking procedures will be employed. The masking procedure illustrated in Figure 1, called the 'coming around the corner' mask, mimics a commonplace perceptual problem. This masking procedure is used to lend ecological validity to the experiments reported here. To enhance the case that special expertise is a normal part of object identification, the conditions of presentation ought to be a simulation of naturally occurring distortion.

The second kind of masking used here is referred to as 'checkerboard masking.' This form of masking is similar to that used first by Gollin and later by Snodgrass. This kind of masking has some analytic advantages in that it is easy to manipulate, and is easier to relate to current theories of object identification. Tversky and Hemenway (1984) have suggested that parts of objects might be a privileged type of information. Representations of object categories might be in terms of object components, and identification may proceed by the identification of constituent parts. Parts are typically small in comparison to the object. It has been demonstrated that the effectiveness of a mask depends upon the match between the spatial frequency of the signal and that of the mask

(Legge, Cohen, and Stromeyer, 1978). While no attempt has been made here to match spatial frequencies in detail, the checkerboard masking procedure was designed such that the mask and object parts should be of approximately the same order. The third kind of masking is called random dot masking, first applied by Jacoby to the picture identification task. In this form of masking, the stimulus emerges slowly from a display of random dots. Because the mask is of a high spatial frequency, it preferentially masks fine detail. This is similar in consequence to defocussing or viewing an object through fog or snow. Demonstrating the same pattern of results across mask types would further illustrate the robustness and ubiquity of the use of item memory in object identification.

The issue of generalization from prior episodes to new is approached in two ways. First, specificity of transfer is examined. A consequence of the episodic view is that the details of a prior episode are important to the amount of visual transfer. The test display can be thought of as a retrieval cue. In accordance with the encoding specificity principle (Tulving and Thomson, 1973), the success of retrieval, and hence the amount of transfer, should depend upon the match in cues between the test display and the prior episode. The dependence of the magnitude of transfer on the match of visual detail, or specificity of transfer, is documented.

Replicating studies by Gollin (1960), Snodgrass (1990, 1987), and Jacoby, Baker, and Brooks (1989), two kinds of encoding events are used here: viewing intact pictures and viewing fragmented or masked pictures. This manipulation has been shown to affect the magnitude of transfer in picture identification. Having seen a masked version of the picture allows earlier identification than would having seen the intact picture. This general finding has been extended to the case of same versus different fragments between training and test. The demonstration of dependence between visual detail and transfer would be strong support for special expertise in the picture identification task.

Chapter 2: Specificity in the picture identification task

The purpose of the following two experiments is to document the contribution of special expertise in the picture identification task. As argued in Chapter 1, specificity of transfer is consistent with the encoding specificity principle, thereby supporting the argument that task performance relies upon item memory. To this end, the dependence of facilitation on the details of prior encoding will be examined.

There are three design features in the following experiments that extend prior research. First, the generality of previous results is demonstrated by using two masks with different perceptual characteristics than those used previously. These masks will be described subsequently. The second feature is an integrated training/test series instead of discrete training and test phases. In this modest way, the experimental conditions emulate the kinds of demands placed on people in most real world identification problems. Under normal circumstances, people encounter new items embedded in a context of familiar ones rather than a long series of new items followed by another series of items just seen. By using the integrated training/test series, it is likely that subjects will not adopt a special strategy to deal with the task demands of separate phases. The third feature is that the effects of same versus different picture fragments is more extensively investigated than in the results reported by Snodgrass (1987). In addition, these experiments will replicate the contrast between viewing intact versus fragmented pictures previously reported by Gollin (1960), Snodgrass (1987), and Jacoby, Baker, and Brooks (1989).

Experiment 1a

This experiment introduces the 'coming around the corner' masking technique, first suggested by Jacoby. As argued in Chapter 1, this mask mimics a common form of visual distortion.

The major premise of the special expertise position is that old solutions are adapted to new problems. An implication of this is that the memorability of a prior item may have consequences upon transfer performance. If two items differ in memorability, then transfer from those items may differ. To investigate this possibility, some subjects were required to make concurrent memory judgments about the item they had just identified. If these subjects demonstrate poorer item memory for a given item, it may well be that they will also show poorer transfer. A further advantage of this procedure is that a special orientation to the past may be produced by the requirement of concurrent memory judgments, and consequently produce changes in.transfer performance.

Method

Subjects: The subjects were 32 McMaster University students enrolled in an introductory Psychology course who participated for partial course credit. Subjects were run individually.

Materials: The pictures used were 36 black on white line drawings of familiar objects, each from a different name category. Objects were drawn by the experimenter, taken from prior experiments by Nelson, Reed and Walling (1976) and Jacoby or taken from such sources as children's coloring books. The drawings were digitized using a Computer Stations Dithertizer II image digitizer, a video camera, and an Apple II+ computer. The pictures were displayed on an Apple //e computer with a 14 inch (diagonal) monochrome monitor. The pictures were digitized such that the major axis of the object (horizontal, vertical or diagonal) extended approximately 80% of the screen in that direction.

Design: Each of the 36 line drawings was presented twice in an interleaved training/test series. First and second presentations occurred intermixed in a 72 trial series. The average lag between first and second presentations was seven, and ranged between six and nine. On first presentation, half the pictures were presented intact on the screen for five seconds. Subjects were required to name the picture out loud, and were instructed to study the picture for as long as it remained on the screen. The other half of the first presentations started with a blank display. Picture information (pixels) was added to the screen starting at one side of the object. The appearance of the display was as if the picture was being occluded by a blank page, and that blank page was slowly being drawn away to the side. The description given the subjects was that the item was "coming around a corner."

All of the second presentations involved clarification of the picture. For those items previously seen intact (not requiring clarification), half were clarified from the right and half from the left. To eliminate the direction of clarification (with respect to the computer screen) as a valid cue to the picture's identity, half of the pictures were mirror reversed between first and second presentations. Each picture has a normal left/right orientation as stored in the computer. When referring to directions of clarification, the direction is with respect to the object in its normal orientation, and not the computer screen or subject. With the direction of clarification defined with respect to the picture this means that the direction the illusory blank page moved on the screen was reversed as well.

There were two groups in the experiment. One group was required to answer a question from memory about the item. On second presentations, and after the item had been displayed on the screen for its allotted time, the message "This was the second presentation of this item. Did this item 'come around the corner' on its FIRST presentation? (Y/N)." Subjects were to answer the question by pressing the "Y" or "N" keys of the computer. The other group was not required to make this judgment, and no message was presented after presentation of the item.

Sixteen counterbalance orders were constructed that rotated items through the cells of the design across subjects.

Procedure: At the outset of the experiment, subjects were told that they were to name a series of pictures as rapidly and accurately as possible, and that some of these objects would be presented as if they were "coming around a corner." At the beginning of a trial, the prompt "Press a key" was displayed, allowing the subject to pace the intertrial interval. If the picture was to be presented in the clear, the full picture was displayed, and the subject named it. If the picture was to be clarified, a blank screen was presented. The computer automatically added picture information to the display working horizontally across the screen. The picture was revealed at such a rate that the mask would be completely removed in approximately 30 seconds. As the pictures are of varying width, the amount of time required to fully clarify a picture ranged about 12 to 25 seconds.

Subjects were instructed to press the space bar of the computer to stop the clarification process as soon as the they thought they knew the identity of the object. They then verbally stated their hypothesis. If the subject was incorrect, the subject was instructed to press the space bar of the computer to resume the clarification process. If the subject was correct, the subject was instructed to press the ESCAPE key of the computer. Once the ESCAPE key was pressed, the picture was presented fully clarified for five seconds. Those pictures presented in the clear on their first presentation were displayed for the same five seconds. Subjects were instructed to study the picture for as long as it remained on the screen.

The correctness of an hypothesis was assessed using the naming criterion that if a name was offered that could be applied to that object, but no other item in the list, it was considered correct. For example, "bunny" would be considered a correct name for a picture of a rabbit but "animal" would not, given the presence of other animals in the list. If subjects offered a name that applied to more than one item in the list, such as "animal," they were instructed to be more specific in their naming.

In the group with concurrent memory judgments, a judgment was required after the second presentation of an item. After the five second presentation of the fully clarified picture, the computer printed the message that the trial had been a second presentation, and asked if that object had "come around the corner" on its first presentation. The subjects responded by pressing either the "Y" or "N" key of the computer.

The datum resulting from a trial was the lateral position of the mask at the time of identification. The horizontal offset of the picture from the edge of the display was subtracted, and the difference was divided by the width of the drawing. The resulting value is a proportion measure. On those trials where the subject allowed the mask to be removed past the far edge of the object, the dependent measure was truncated at 1.0. These proportions were converted to percentages and submitted to planned t tests. As the design is necessarily incomplete[3], a more conventional analysis of variance was not possible. A criterion p value of 0.05 was used for all statistical tests.

Results and Discussion

The first question to be addressed is regarding the general effects of prior experience. Does having seen a picture before improve subjects' ability to identify the items on second presentation? In order to assess the impact of prior presentation, performance on those first occurrences requiring clarification (NEWS) were compared to second presentations (OLDS). New items required 40.6% clarification, while Olds required only 23.2%, $t(31) = 17.181$. Clearly, the prior presentation of an item greatly improves performance in the clarification task. This effect of prior presentation replicates the findings of other researchers, i.e. Jacoby, Baker, and Brooks (1989), Parkin and Streete (1989), Snodgrass (1987, 1990), Vokey, Baker, Hayman, and Jacoby (1986).

While consistent with the episodic view argued here, a simple effect of prior presentation on clarification performance does not differentially support either the prototype or episodic view. By either view, the prior presentation of an item should affect performance. By the episodic view, the prior presentation of the item results in an encoded episode. This prior identification is adapted to aid the identification of the item on its second

---

[3]The dependent measure is the percentage of the picture's horizontal extent present at the time of identification. On first presentations, half of the items were viewed intact, and never clarified. Since no data can be collected for these items, no complete orthogonal factor structure can be used to describe the experiment.

presentation. By the prototype view, prior presentation of a category member could act to prime the category. Because the category is primed, subsequent identification of a category member would be improved.

Table 1: Experiment 1a collapsed across the grouping factor of concurrent memory judgments. Entries are in percent clarified at correct identification.

New:	40.637
Old:	23.149
Subdividing Olds into...	
Previously Clarified:	22.002
Previously Viewed Clear:	24.297
Subdividing Previously Clarified into...	
Same Information:	20.394
Different Information:	23.609

There is reason to believe that priming effects should be short lived, though there is no principled answer to the question of just how short lived they should be. By the priming explanation, priming results from the temporary activation of a knowledge structure. If this activation is not temporary, the prior presentation constitutes a change in the knowledge structure itself, thereby defeating the goal of relying on stable representations. What is subject to some debate is how long lived the effects may be. In the current experiment, at least six items occur between first and second presentations. This lag is clearly longer than that observed in classic priming demonstrations such as lexical decision, and will be extended to 24 hours in Experiment 3.

The major contrast between the episodic and priming explanations is the generality of the effects of prior presentation. By the priming account, facilitation depends upon the activation of the knowledge structure underlying the identification of an item. This knowledge structure is general to the category, and represents features characteristic of the category. Therefore, facilitation should be ahistorical. That is, all that will effect the amount of facilitation is the amount of activation given the knowledge structure, and the amount of time that knowledge structure has had to decay. The exact details of how that knowledge structure came to be activated to the level it was in the first place ought not have an effect on the magnitude of that facilitation.

By the episodic view, however, transfer depends upon the retrieval of the prior event. In terms of encoding specificity (Tulving and Thomson, 1973), if those cues available at the time of encoding are also available at the time of retrieval, then retrieval is likely to be successful. Thinking of the visual display as a collection of cues, then, retrieval should be successful if there is a large overlap in cues across the two presentations. In contrast to the priming account, the episodic account suggests that transfer should be sensitive to the exact history of each item.

The critical contrast between the episodic and prototype accounts in this experiment, then, is the dependency of facilitation upon the details of the prior encounter. One test of this dependency is the effect of different kinds of prior exposure. Subjects required 24.5% clarification to identify those items seen clear on their first presentation, but only 22.0% to identify those items previously clarified, $\underline{t}(31) = 3.473$. Subjects required reliably less clarification to identify a second presentation if the item had been clarified on its first presentation. This replicates the results of Jacoby, Baker, and Brooks (1989) using a different procedure.

There is an alternative account for the advantage of previously clarified items over previously viewed clear. The previously viewed clear items were seen for five seconds on their first presentation. Previously clarified items were viewed for the same five seconds plus the time required to identify the item. In addition, the effort invested during the first presentation is substantially greater on previously clarified items. It could be that the amount a knowledge structure is primed depends upon the amount of time or effort spent working with the prime stimulus. With either of these provisions, the prototype view can be made to account for the observed data.

A simple "time spent" or "effort of encoding" explanation is insufficient to explain the advantage for previously clarified items. For half the previously clarified items, the same stimulus details were presented on both presentations, and for the remainder, different details were displayed on the two presentations. Within the population of items clarified on first and second presentations, which equates items for time spent and effort expended, the effects of the match of visual detail between presentations can be examined. If the visual details mismatched between presentations, subjects required 23.6% clarification. If the details matched, subjects required only 20.4% clarification, $\underline{t}(31) = 2.861$. Thus, subjects required reliably less clarification to identify the item if the visual details available on the second presentation matched those available at the time of encoding.

This result is consistent with the encoding specificity principle, and lend support to the episodic view. It is at odds with a prototype account, as by the prototype account, the amount of facilitation depends upon two factors: the initial level of activation and the delay between prime and target. The delay is equal for the same information and different information trials. Therefore, the only remaining source of differential priming effects is the initial level of activation. However, the activation of the knowledge structure resulting from the first presentation cannot depend upon the details of the second presentation.

Part of the rationale for this experiment is to examine the effects of concurrent recognition memory judgments. Part of the purpose was to test whether episodic effects in picture identification are due to an explicit orientation towards memory, thereby invoking a reflective, memory based strategy in task performance. Two way analyses of variance paralleling the t tests reported above were performed to examine the influence of concurrent memory judgments. In every test involving the grouping factor of Presence/Absence of concurrent memory judgments, both as a main effect and in an interaction, the $F(1,30)$ was less than 1.3 ($F(1,30)$ critical value is 4.17), and the p values were greater than 0.25. Requiring concurrent memory judgments during the clarification task had no detectable effect on subjects' performance. The role of orientation to the past in performance of the picture identification task will be examined in further detail in later experiments.

Another planned analysis was to compare the amount of transfer on those items for which subjects correctly answered the memory question to those which subjects did not. If subjects are incorrect, the subjects might not have as good a memory for that item, and the amount of transfer might be less. Unfortunately, very few errors were made. Most of the subjects made no errors on at least one of the old picture types (previously viewed clear, previously clarified revealing the same information, and previously clarified revealing different information). Therefore, the analysis was impossible.

Experiment 1b

The results of Experiment 1a support the contention that memory for prior items is used in the clarification task. The masking technique used in Experiment 1a matches some ecologically ubiquitous conditions of visual distortion. The purpose of the current experiment is to replicate the basic findings of Experiment 1a using a kind of distortion potentially more relevant to current theories of concept representation and object identification. The masking technique used here, checkerboard masking, was aimed at the masking of object components, or in Tversky and Hemenway's (1984) terms, parts.

The effectiveness of a mask has been demonstrated to depend upon the match of spatial frequencies between mask and target. Because parts of an object are typically small compared to the object itself, constructing a mask of somewhat higher spatial frequency than the very low frequency 'coming around the corner' mask might reveal effects not observed in Experiment 1a. Biederman (Biederman, 1987; Biederman and Ju, 1988) has offered a theory of object perception in which objects can (usually) be described by a few geons, ranging from one in the case of a beachball to several in the case of a quadrupedal animal. With multiple geons per object, each of those geons must be of a higher spatial frequency (smaller) that the object itself. While the spatial frequencies of the pictures themselves were not measured, the spatial frequency of the mask was increased by a factor of eight. This should keep the mask and the object parts or geons being masked in the same general magnitude. The ultimate goal of this line of research will be to specifically mask particular parts of objects, but this is beyond the capacity of the current computer system. As a result, the checkerboard mask was used as a first approximation.

Method

Subjects: Subjects were 18 McMaster University students enrolled in an introductory Psychology class. Subjects were run individually. None had participated in the previous experiment.

Materials: The pictures used were the same 36 black on white line drawings of familiar objects used in Experiment 1a. The pictures were displayed on an Apple //e computer with a 14 inch (diagonal) monochrome monitor.

Design: The design of the experiment is very similar to that in Experiment 1a. Subjects were presented each of the 36 pictures twice in an interleaved training/test series. The number of items between the two presentations of an item ranged from six to nine, with an average lag of about seven. Upon first presentation, half the items were viewed in the clear and half were clarified. Upon second presentation, all items were clarified. For those items clarified on first presentation, half were clarified revealing the same picture information, and half were clarified revealing different picture information. A set of counterbalances were constructed that rotated items through each condition in the experiment across subjects.

Procedure: At the outset of a trial, the prompt "Press a key..." was presented on the screen. Upon the subject's doing so, the trial began. If the item was to be presented in the clear, the item was presented for five seconds. Subjects were required to name the picture aloud, and were instructed to study the picture for as long as it remained on the screen.

The masking used in this experiment was the checkerboard clarification technique. The screen of the computer was divided into an eight by eight grid. At the outset of a clarification trial, a nearly blank screen was displayed. When subjects pressed a key, pixels from the picture were added to the display in alternate squares. Lines of picture information were added starting at the tops of the squares and working down. On first presentation trials requiring clarification, one of the two subsets of squares was clarified. Either the set of squares including the top right (Information subset A) or the set including the top left (Information subset B) could be clarified on the first presentation of an item. At 50% clarification, then, all the picture information coincident with half of the squares would be displayed, and the picture information coincident with the alternate squares would be blank.

If an item was being presented for the second time, one standard test mask was used. The use of one standard test mask eliminates the mask per sé as a cue to the picture's identity independent of the picture information available. The test mask was the set of squares containing the top left square, but was shifted a half square to the right. On second presentation, the picture was shifted laterally one half square behind the mask. If the picture is shifted right, Information subset A was presented, but if shifted left, subset B was presented.

The subject controlled the addition of picture information to the display, thereby allowing subjects to pace the rate of clarification. Every time the subject pressed a key, approximately 2% of the picture information was added to the screen. Subjects were instructed to press the key until they had an idea as to what the picture might be. Subjects stated their hypotheses and received verbal feedback. The correctness of an hypothesis was assessed using the same naming criterion as that used in Experiment 1a. After having correctly named the picture, the subject was instructed to press the ESCAPE key of the computer. Doing so presented the fully clarified picture on the screen for five seconds.

The datum resulting from a trial was the number of key presses necessary for identification. The number of key presses is linearly related to the percentage of picture information present on the screen, or clarification. On some trials, the subject would fail to identify the item before the picture had been clarified to 50%. If the clarification procedure were allowed to continue, some of the picture information from the other Information subset would have been displayed. This would cause interpretive difficulty for the Same information/Different information manipulation. Therefore, whenever a subject reached 50% and had not yet identified the picture, the clarification procedure was stopped, the fully clarified picture was displayed, and 50% was recorded as the dependent measure. A criterion p value of 0.05 was used for all statistical tests.

Results and Discussion

To facilitate comparisons between Experiments 1a and 1b, the data were analyzed in parallel fashion. First, there was a reliable effect of prior presentation. First occurrences (NEWS) requiring clarification required 35.7% clarification for identification, while OLD items required only 18.3%, $t(17) = 19.598$.

Table 2: Outcome of Experiment 1b. Entries are in percent clarified at correct identification.

New:	35.714
Old:	18.263
Subdividing Olds into...	
Previously Clarified:	16.437
Previously Viewed Clear:	21.916
Subdividing Previously Clarified into...	
Same Information:	14.141
Different Information:	18.733

Second occurrences necessarily occur later in the series than do first occurrences. This raises the possibility that the Old/New effect is, at least in part, a practice effect. Because subjects could be acquiring a general skill as the task progresses, they might become better able to perform the task later in the series. Since Old items occur later than New, this would give the impression of an advantage for Olds. In order to address this possibility, the effect of practice within an item type (News) was examined. From the data, it was possible to recover the performance on each individual trial in the sequence in which trials occurred. Performance on the first six News requiring clarification was compared to the last six News requiring clarification. There was no reliable difference in the amount of clarification necessary between early New items (mean serial position = 8.333) required 35.5% clarification, and late New items (mean serial position = 61.667) required 34.6% clarification, $t(17) = 0.448$. It therefore seems reasonable to conclude that the Old/New effect seen here is not a product of a general skill learned by the subjects. Rather, the difference reflects transfer on some sort: either category priming or episodic transfer.

Parallel to the analysis of Experiment 1a, the dependence of facilitation observed on the nature of the prior presentation was examined. The Previously Viewed Clear items required 21.9% clarification, while Previously Clarified items required 16.4%, $t(17) = 5.256$. This replicates the effect observed in Experiment 1a. In addition, the effect of matching the visual details between first and second presentations was examined. For half of the Previously Clarified items, the Same Information was presented on first and second presentation (Information subset A followed by A or B by B). For the other items, Different Information was presented on first and second presentations (A followed by B or B by A). Same Information items required 14.1% clarification, while Different Information items required 18.7% clarification, $t(17) = 4.412$. All of the effects observed in Experiment 1a are replicated here, further lending support to the notion of episodic transfer.

In the current experiment, there were a number of trials on which subjects had not yet identified the item by 50% clarification. As was mentioned above, if subjects had not yet identified the item by 50%, the item was presented in the clear and 50% was recorded as the dependent measure. This prevents observations above 50%. One potential problem with this practice is that the range of the dependent measure is restricted. If more of these truncations occur for one type of item than another, the size of the difference between the two item types will be affected. It is important to note that the number of truncations within an item type is positively correlated with mean performance for that item type. That is, New items had the most truncations, followed by previously viewed clear, and previously clarified had the fewest. The elimination of high outlier scores acts to reduce the mean value. This means that the differences reported here are actually conservative. If some way had been found to allow values of the dependent measure to exceed 50%, the mean differences would have been even greater than those reported here. Thus, the presence of truncations does not distort the pattern of results. Rather, it simply understates them.

The replication here of the pattern of results seen in Experiment 1a suggests that transfer effects are general at least across the 'coming around the corner' and the checkerboard masks. The 'coming around the corner' mask mimics an ecologically ubiquitous visual distortion, and identification routinely precedes full visual information. The demonstration that memory for prior encounters aids in identification under masking conditions mimicking these suggests that episodic transfer may be a very general phenomenon.

The masking used in the current experiment was designed to address the possibility that object components might be a privileged type of information, and that masking matched to components might have special consequences. The difference in the masking technique had no effect on the general pattern of results. This, of course, must be interpreted with some care. The mask used here was not matched to specific object components. Rather, it was made of a spatial frequency likely to be of the same order as object parts. Had the results of this experiment been notably different from those seen previously, it would have been a clear indicator for a new line of investigation. Because there was no indication that the mask acted differently than the 'coming around the corner' mask, it remains a topic for future investigation.

Chapter 3: Generalization in the picture identification task

Experiments 1a and 1b support the contention that performance of the clarification task depends upon the retrieval of prior encounters with an item. A minimal theoretical response to these findings would be a hybrid model of picture identification. In such a hybrid, memory for an item would be used to identify that item upon its second presentation. However, the prototype would be used to identify new items of that object category.

This is, in essence, the explanation Warren and Morton (1982) gave for their finding that prior presentation of an item had greater facilitative effects on the identification of that same item than on the identification of a different member of the category. In addition, the new category member was easier to identify than members of new categories. Their explanation was that item memory was used to identify previously seen items, and that the identification of new members of the category was easier than members of new categories because the knowledge structure subserving picture identification, a "pictogen" in their model, was primed by the prior item. That is, facilitative effects are a mixture or hybrid of item memory and category priming.

One explanation for the apparent episodic basis of picture identification seen in Experiments 1a and 1b is that for identical items, identification does depend upon item memory. However, by a hybrid account, this use of item memory is a special case that applies only to the identification of previously seen items. For the identification of all other category members, the prototype is used, and not item memory. One hazard of hybrid models is that they can be underspecified. With two alternative paths to item identification, it is necessary to specify the conditions under which each is used. In the absence of this specification, the model fails to do the one thing models must be required to do: predict performance. The Warren and Morton model can be termed a strong hybrid. In a strong hybrid, item memory is used only on previously seen, literally old items, while the prototype or pictogen is used for all other items.

The special expertise view of picture identification has two diagnostic phenomena: specificity and generalization. Experiments 1a and 1b document the relationship between the amount of transfer and the match in visual detail, or specificity. However, specificity is only half the story. The crucial component in the special expertise view is generalization: the process by which prior solutions are adapted to new problems. By the prototype view, identification is done by comparing the stimulus to prototypes, which are stable knowledge structures. The prior presentation of an item can prime the prototype, reducing the amount of stimulus information necessary for identification. Therefore, the amount of information necessary for identification depends upon the resemblance of the test item to the prototype and the amount of priming. By the episodic view, identification depends upon the adaptation of a prior solution to the current problem. This adaptation depends upon the similarity between the test stimulus and the prior item. Therefore, the amount of transfer depends not only on categorical relatedness between the two items, but also on their perceptual similarity. To this end, the effects of perceptual similarity on transfer performance was examined.

Experiment 2

Use of the picture identification task has largely been confined to memory issues. However, Jacoby, Baker, and Brooks (1989) used the task to investigate identification of new category members. There were two major findings. First, greater transfer was found if items had been clarified in training than if they had been viewed clear. This finding was replicated in Experiments 1a and 1b here. Their second finding was that performance on new members of the name categories seen in training was better than on those that had not been seen in training, replicating Warren and Morton (1982). Further, and more importantly, the amount of transfer to same name items was dependent upon how the training items had been encoded. Less specificity (or more transfer) was seen if the training items had been viewed clear, while greater specificity (or less transfer) was seen if the training items had been clarified. With performance on both old items and same name items dependent upon the same encoding manipulation, the result implies that both test item types are under control of the same processes and knowledge.

A hybrid prototype view can make an account of the results as well. In a strong hybrid, item memory is used to identify old items, while the prototype is used to identify new category members. A possible explanation of the Jacoby, Baker, and Brooks result would be that the encoding manipulation changed the balance of item memory and category priming. This implies that item memory and category priming are in a trading relationship. With greater item memory, there is less category priming and vice versa. As Jacoby, Baker, and Brooks (1989) point out, although increasing attention to visual detail might increase transfer to old items, this ought not reduce the amount of priming. Item memory and category priming are paths to identification relying upon separate and presumably independent information sources. The current experiment, and the one that follows are designed to more pointedly address a related weakness in the strong hybrid prototype model.

Within an object category, there is tremendous variability in the perceptual characteristics of items. This is true for man-made objects, such as chairs as well as natural categories, such as dogs and fish. In addition, because the stimuli are pictures, the perceptual characteristics of an item can be very diverse with different viewing conditions and object orientations. By the strong hybrid model, the similarity between the first and second examples of a category ought not have an effect on the magnitude of priming effects. Only the amount of category priming should have an effect. In contrast, if the second example of a category acts as a retrieval cue, then the more similar it is to the first example, the more successful the item retrieval, and the greater the magnitude of transfer. To this end, the effects of interitem similarity on transfer in the picture identification task were examined.

Method

Subjects: The subjects were 24 McMaster University students enrolled in an introductory psychology class. Subjects were run individually, and none had participated in the previous experiments.

Materials: Line drawings similar to those in the preceding experiments were used as stimuli. The 96 drawings represented 24 object categories, with four examples per category. Within a quartet of category members, perceptual similarity was varied. See Figure 3 for an example of a quartet.

Each quartet can be subdivided into two pairs. Between pairs, items were less perceptually similar than within pairs.

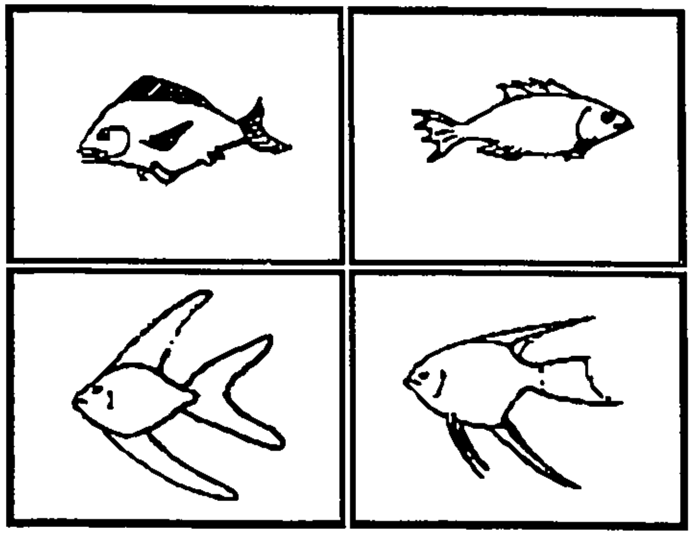

Figure 3: An example of a quartet used in Experiments 2, 4, and 5.

Design: The experimental design is a 2 by 3 within subjects design. Subjects clarified pictures in a 48 trial series, with each object category appearing twice. On average, there was a lag of 11.3 items between first and second presentations of a category, with the range from 7 to 15. The first example of the category could be presented intact (Viewed Clear) or Clarified. The second example of the category always required clarification, and could be the same category member (Ident), the perceptually similar category member (SN/Sim), or a perceptually dissimilar category member (SN/Diff). Six counterbalance orders were constructed to rotate items through conditions across subjects.

Procedure: At the outset of a trial, the prompt "Press a key..." was displayed. Once the subject pressed a key, the trial began. If the picture was to be presented clear (Viewed Clear), the intact picture was displayed for approximately five seconds. Subjects were required to name the item out loud, and were instructed to study the picture for as long as it remained on the screen. If the trial required clarification, an essentially blank screen was displayed. The experiment used the same checkerboard masking procedure as Experiment 1b. Every time the subject pressed a key, approximately 2% of the picture information was added to the display. One checkerboard mask was used. Also, there were no truncations in this experiment. Once a subject had completely clarified one set of squares, clarification of the complementary set began.

Subjects gave their hypotheses verbally, and the same naming criterion was used as that in Experiments 1a and 1b. If a subject's hypothesis was incorrect, he was instructed to continue clarifying. When subjects had correctly named the object, they were instructed to press the ESCAPE key on the computer. Doing so resulted in the intact picture then being displayed for five seconds.

The measure resulting from a trial was the number of key presses necessary for correct identification of the item. These were converted to a percent measure of clarification before being submitted to an analysis of variance. For all statistical tests, a criterion p value of 0.05 was used.

Results and Discussion

In the analysis of variance, the effect of Viewed Clear vs. Clarified on first presentation was marginally reliable, $F(1,23) = 4.20$, $MSe = 142.989$. Items required less clarification if the first category member had been Clarified (37.58%) than if it had been Viewed Clear (41.66%). The interaction between prior presentation (Viewed Clear vs. Clarified) and Test Picture Type (Identical, SN/Sim, SN/Diff, and New) was not reliable. The failure to get an interaction is a departure from the results reported by Jacoby, Baker, and Brooks (1989). However, the experimental design here is quite different than that used by Jacoby, Baker, and Brooks. They used a between subjects design, discrete training and test phases, interpolated recall, a longer retention interval, and different stimuli. Each of these changes individually is unlikely to cause trouble, but it is hard to anticipate the cumulative effects of all these changes together. The current experiment was not intended as a replication, so the lack of a significant interaction is of secondary concern. However, it would be of interest to determine which of these factors, either singly or jointly, is responsible for the difference in results.

Table 3: Outcome of Experiment 2. Entries are in percent clarified at correct identification.

	Encoding type	
	Viewed Clear	Clarified
Test picture type		
Identical	32.964	26.345
SN/Sim	42.253	38.129
SN/Diff	49.761	48.251

The effect of picture type was reliable. Identical items required the least clarification (29.7%), SN/Diff items required the most (49.0%), and SN/Sim items required an intermediate amount (40.2%), $\underline{F}(2,23) = 33.093$, MSe = 136.143. This effect is consistent with the hypothesis that transfer is graded across picture types. To further examine the source of the result, performance on SN/Sim items was compared to SN/Diff items. SN/Sim items required reliably less clarification than did SN/Diff, $\underline{t}(23) = 3.45$. This result clearly indicates that performance on new category members is not equal. The strong hybrid leads to the prediction that SN/Sim and SN/Diff items should require the same amount of clarification. The demonstration that performance differs between SN/Sim and SN/Diff items is clearly contrary to this position.

In a post hoc test, the effect of Viewed Clear versus Clarified on first presentation was reliable for Identical items. Items clarified on first presentation required 26.3% clarification upon second presentation, while those items previously seen Clear required 33.0% clarification, $t(23) = 3.248$. This replicates the findings of Experiments 1a and 1b. Because the visual characteristics of SN/Sim test items are different than those of the prior item, the episodic account implies that the amount of transfer should be less. In addition, the Viewed Clear versus Clarified manipulation should be less effective.

This implication is borne out by the data. For SN/Sim items, if the prior item had been viewed clear, 42.3% clarification was required for identification, while for those SN/Sim items whose prior item had been clarified, 38.1% clarification was required, $t(23) = 1.143$, $p = 0.265$. While the mean differences for SN/Sim items between Viewed Clear and Clarified are in the same direction as for Identical items, the mean difference has been reduced to the point where the difference is no longer reliable. This makes sense, as the Viewed Clear versus Clarified manipulation of similarity of visual detail can only be effective if the items presented on first and second presentations have the same details to present. This is true for Identical items, but not for SN/Sim items. Therefore, the manipulation ought to be, and is much less effective for SN/Sim items.

The advantage of SN/Sim items over SN/Diff items is consistent with the episodic account. Transfer depends upon retrieval of the prior item.

This retrieval depends upon the overlap in cues between encoding and retrieval. SN/Diff items were selected to be as different as possible. Presumably, a SN/Diff item has less in common with its prior item than would a SN/Sim. Therefore, transfer should be even less for SN/Diff items than it is for SN/Sim. This is precisely the pattern of results observed.

There is an alternative account for the different levels of facilitation on SN/Sim and SN/Diff items. While the members of each quartet of pictures are all called by the same name, this does not assure that a single abstract knowledge structure is used to identify them all. One cannot equate name with category. For example, both the instrument used to hit a baseball and winged mammals are named "bat". However, it is absurd to suggest that a common knowledge structure is used to identify both types of object, despite the common label. It is therefore not forced that all items in a quartet are identified by reference to one knowledge structure despite a common label. This substantially complicates the prototype account, but complication per sé is not evidence against a position. Occam's razor suggests only that when two equally effective accounts for a phenomenon are proposed, the simpler account is to be preferred. It does not imply that complicated explanations are in error.

With the provision that there can be multiple prototypes for a single name category, the prototype account can be made to account for the effects demonstrated here. The identification of items within a perceptually similar pair, then, depends upon one prototype. However, if two same name items are perceptually different, an upright piano and a grand piano for example, they may be represented by two different prototypes. Presentation of one upright might prime only the prototype for upright pianos and not the prototype for grand pianos. Therefore, there will be priming for a SN/Sim item (other uprights), but perhaps no priming for SN/Diff items (grand pianos). Because category priming operates for SN/Sim items, but occasionally not for SN/Diff items, performance on SN/Diff items should be worse. While this explanation complicates the prototype account, it does fit the results of the experiment here.

Experiment 3

In Experiment 2, the performance on new members of a category is demonstrated to depend upon the similarity between the two category members. As discussed in Experiment 2, allowing an unspecified number of prototypes per category would account for the effects of similarity. However, this makes the prototype account much more difficult to test directly. It is possible, however, to test the explanatory sufficiency of priming. Priming is the temporary increased activation of a knowledge structure, allowing the knowledge structure to identify an item either faster, or on the basis of less information. If priming is to be used as a way of increasing the flexibility of a stable knowledge structure, then it is, of necessity, a temporary change in activation.

In contrast to the prototype account, the special expertise view suggests that transfer should be observed as long as the original encoding event can be retrieved. Retrieval depends upon cuing conditions, not the time since the encoding event. By this view, there is no necessary time line to transfer effects.

By the hybrid view, different mechanisms are involved in facilitative effects of prior experience on Identical and Same Name items. Item memory underlies facilitation on Identicals, while category priming underlies facilitation on new category members. There is no necessary reason for facilitation on Identical items to decrease with delay. However, facilitation on Same Name items must decrease across a delay. To fail to do so would be a fundamental violation of the basic premise of the prototype account: that the knowledge structures used in the identification of pictures are general to the category and stable. This is a necessary implication of the structure of prototypes and pictogens.

To examine whether different mechanisms are responsible for the facilitation on Identical and Same Name items, the durability of the effects of prior exposure were examined over a 24 hour delay. Therefore, there should be an interaction between delay and old or new pictures within a category. As was suggested earlier, performance on SN/Sim and SN/Diff items might rely upon different knowledge structures. The use of perceptually disjunct categories might be a special case in which a single name might correspond to a number of different knowledge structures. To allow for this possibility, it is necessary to examine the effects of delay on performance within items sets for which it is not reasonable to suggest multiple prototypes.

Method

Subjects: The subjects were 18 McMaster University students enrolled in an introductory Psychology class. None had participated in the preceding experiments, and were run individually.

Design: The experiment is a two by three within subjects design. The experiment was conducted in three phases. There was a learning phase and two test phases. The first factor is time of test. Subjects were tested both immediately after the learning phase, and again after a 24 hour delay. The second factor is test picture type. A picture in a test phase could be an Identical repetition of a picture seen in the learning phase, of the Same Name category as one in the learning phase, or of a New name category. Six counterbalance orders were constructed that rotated items through the six cells of the design across subjects.

Materials: The stimuli in the experiment were 54 pairs of line drawings. These line drawings were constructed and prepared the same way and with the same equipment as in Experiments 1a, 1b and 2. Each pair was made of two examples of an object category. The two members of a category shared structural characteristics and name, but were clearly different examples of the object category. Some of the SN/Sim pairs used in Experiment 2 were used as stimuli here.

Procedure: The experiment was subject paced. In the learning phase, subjects clarified a series of 36 pictures, each from a different name category. The clarification procedure used was the checkerboard masking used in Experiment 2. Subjects paced both the intertrial interval and the rate of clarification. When a subject offered a hypothesis, the correctness of the hypothesis was judged by the same criteria as in the preceding experiments.

After completion of the 36 training pictures, the subject was given a brief two to five minute rest. During this time, instructions for the first test phase were given. Subjects were informed that some of the items in the upcoming phase were ones they had seen before, and that some were new. The procedure in the test phase was the same as that in the learning phase. Subjects then proceeded to the first test phase. After completion of the first test phase, the subjects left, and were instructed to return at the same time the following day. On the second day, the same instructions were given for the second test phase as for the first. Subjects then performed the second test phase.

The datum resulting from a trial was the number of key presses required for identification. The number of key presses was converted into the percentage clarified, and submitted to an analysis of variance. In all statistical tests, a criterion $\underline{p}$ value of 0.05 was used.

Results and Discussion

In the analysis, an effect of test session was found. Subjects required reliably more clarification to identify pictures in the second test session than in the first. In the immediate test phase, subjects required 27.9% clarification, while in the second test phase, they required 31.4%, $\underline{F}(1,17) = 10.1$, MSe = 31.647. This may well be attributable to different demand characteristics in the two test phases. The first session, containing a training series of 36 clarifications and 27 test trials, was scheduled for a one hour time slot, whereas the second session, containing only 27 test trials, was similarly scheduled for a one hour time slot. The pressure for rapid performance in the first session may have served to raise the subjects' motivation relative to the second session. While reliable, the effect does not appear important. First, the effect is small. Using an $\underline{\omega}^2$ statistic (Keppel, 1982) to estimate effect size, the effect of test session accounts for 2.76% of the variance. Secondly, this effect is not a memory phenomenon. Performance on New items increased across the two test sessions as much as did performance on Olds. Clearly, effects of retention interval cannot be used to explain the change in performance across test sessions, as the effects are as large for New items as Old.

Table 4: Outcome of Experiment 3. Effects of delay on picture identification. Entries are in percent clarified at correct identification.

	Time of Test	
	Immediate	Delayed
Test picture type		
Identical	18.206	23.194
SN/Sim	28.565	29.860
New	37.037	41.074

More germane to the current issue, there was a reliable effect of picture type. Those items that were identical repetitions of training items (Identicals) required 20.7% clarification, new examples of name categories seen in the learning phase (SN/Sim) required 29.2%, and entirely new pictures of new name categories (News) required 39.1% clarification, $\underline{F}(2,34) = 89.8$, MSe = 33.839. This replicates the effect seen in Experiment 2.

To more clearly determine the source of the effect, two post-hoc analyses were performed. Reliably less clarification was needed for Identicals than for Same Name pictures, $\underline{t}(17) = 7.11$. Also, Same Name pictures required reliably less clarification than New pictures, $\underline{t}(17) = 7.96$. Thus, prior presentation of an item aids not only in the subsequent identification of that item, but also aids in the identification of items of the same name category.

Crucial to the issue of priming is the interaction of test picture type with time of test. A hybrid view would demand that for the Same Name pictures, facilitation should dissipate with time. There was no interaction of picture type with test session, $\underline{F}(2,34) = 1.018$. Such robustness of facilitative effects of a prior viewing is clearly at variance with any priming account that relies upon activation of essentially stable knowledge structures such as prototypes. While it is sometimes hazardous to argue from a lack of effect, we failed to detect an interaction despite considerable experimental and statistical power. For example, the effect of test session was reliable even though it accounts for less than 3% of the observed variance. Thus, this lack of interaction should be taken seriously.

An alternative analysis consistent with Snodgrass's (1987, 1990) relative savings analysis of transfer in the picture identification task is to assess each subject's performance on Identical and Same Name items relative to performance on New items. One supporting argument for this analysis is that factors other than facilitation, such as general skill or motivation, are controlled and therefore less likely to mask or magnify other effects. A way of approximating Snodgrass' relative savings measure is to divide the subject's performance on Identical and Same Name items by the subject's performance on New items for that test session. This procedure yields a proportion measure.

Using this treatment of the dependent measure, there is a main effect of picture type. Identicals required less clarification (0.552) than did Same Name items (0.763), $F(1,17) = 48.461$. In contrast to the previous analysis, there was a reliable interaction between test picture type and time of test, $F(1,17) = 5.766$. The interaction in this analysis is summarized in Table 5. What is important to note is that the interaction between picture type and delay is backwards from that predicted by a priming account. Performance on Same Name items is relatively better after the 24 hour delay than it was on the immediate test (although in absolute terms, performance was slightly worse). Performance on Old items, on the other hand, is slightly worse after the delay than on the immediate test in both absolute and relative measures.

Table 5: Results of Experiment 3 using adjusted savings measure. Entries are the proportion of clarification relative to that required for New pictures in that test session.

	Time of Test	
	Immediate	Delayed
Test picture type		
Identical	.513	.591
SN/Sim	.783	.744

The pattern of results across the preceding series of experiments is clear. Subjects use item memory in performing the picture identification task for both old and new stimuli. The amount of observed facilitation on old items is dependent upon the match in visual detail between first and second presentations. Two observations argue against a simple version of the strong hybrid model. First, the amount of facilitation depends upon the similarity between the training item and the test item: facilitation is greater for SN/Sim items than SN/Diff. Second, the facilitation seen on SN/Sim items is of the same durability as that seen on Identical items, and far too durable to attribute to priming.

A minimal theoretical response to these observations is what can be termed a weak hybrid. A weak hybrid is one in which item memory is used for any item type: identical, same name similar, and same name different, though in the last case, one would expect item memory to be of little effect given the perceptual dissimilarity between the memory item and the target. This weak hybrid model suggests that category general knowledge is represented and used by subjects, but that abstract knowledge enjoys no absolute privileged status in the identification of pictures and objects. The weak hybrid model reduces to the episodic view with the provision that subjects also have, and can use, knowledge of properties general to a category. Prior solutions are adapted to new problems, and use of prototypes or algorithmic solutions are then limited to the case where no prior solution can be adapted, or knowledge of what is generally true is specifically required.

Chapter 4: The role of intent in special expertise

Having advanced the position that performance of the picture identification task depends upon special expertise, it is time to investigate some aspects of how item memory might be used in the task. The results of experiments in this chapter are not primarily useful for selecting between the two views contrasted in preceding chapters. Rather, they function primarily to provide constraint on how the picture identification task is performed.

One implication of the episodic view is that transfer effects might depend upon a retrieval based strategy. By this view, subjects could actively attempt to recall prior items, and use the resulting item as a guide in interpreting the visual array. This volitional use of memory is not a necessary attribute of the episodic view, but rather is one possible way in which episodes could be used to aid picture identification. For example, case based reasoning is a recent application of the special expertise view to problem solving. In case based reasoning, new problems are solved by retrieving prior solutions, and using those prior solutions as aids in analyzing the current problem. Explicit, volitional retrieval is not a necessary part of case based reasoning, but it is often discussed as if it were.

The contrast to the volitional use of prior items is the automatic and unaware intrusion of a prior item in the processing of the test item. Logan (1989) suggests that obligatory retrieval is an important component of automaticity. Automaticity is characterized by impenetrability to strategy or intent, as the Stroop demonstration vividly illustrates. In a similar vein, cognition is replete with examples of the use of prior knowledge in perception, or top-down effects. Yet, these top down effects are rarely accompanied by awareness of the source of the effect. The strategy of the next experiment is to determine whether the item specific transfer performance observed in the previous experiments is selectively dependent on conditions that are often thought to allow deliberate, volitional retrieval.

Experiment 4

It seems clear that there are at least two ways in which a prior episode can be used in the clarification task. One possibility is that subjects adopt a time consuming strategy of deliberate recall, and then use the recalled items as aids in interpreting and identifying the masked picture. Another way in which item memory can affect performance in the task is through a more automatic intrusion of a prior episode into the processing of the test item. This sort of intrusion is unlikely to be strategic or accompanied by any sort of awareness. The importance of distinguishing between these two modes of retrieval has been demonstrated in a series of experiments on recognition memory by Jacoby (1990). The current experiment investigates the extent to which episodic effects in the picture identification task are dependent upon a time consuming, deliberate retrieval strategy. A common conception of volitional recall is that it requires time and effort. In the current experiment, some subjects are given unlimited time between steps of clarification, while for others, clarification proceeds automatically at a fast rate. If episodic facilitation results from a time consuming retrieval strategy, then the amount of episodic facilitation should differ between these conditions. One does not have to believe that volitional recall is necessarily time consuming to be interested in this contrast. However, if the difference between automatic and self paced clarification produced dramatically different results, then it would be important for our interpretation of the processes underlying this kind of picture identification.

A second intent behind the experiment is to examine the possibility that multiple kinds of information are used in the clarification task. Item memory is used in the task in the identification of Identical and SN/Sim items, as evidenced by the preceding experiments. While subjects have item memory, they also have a good deal of analytic knowledge about object categories. That is, they know that birds have wings and fish have fins. There is no reason to suspect that this type of information is not used in the clarification task. Rather, it would be astonishing if it were not used. The picture type for which this type of information would be most useful is SN/Diff. The prior presentation of a category member might serve to make that category more available to the subject. Hirshman, Snodgrass, Mindes and Feenan (1990) documented this in an experiment in which a list of category names presented in an initial study phase facilitated later identification of pictures in those categories. With the increased availability of the category, subjects evidently were able to use their analytic knowledge of object categories to help them identify test pictures. In the current experiment, this would lead to an advantage of SN/Diff items over News.

If subjects are using their knowledge of category general features, then items that have more of these features should be identified more easily than those that have fewer. If this knowledge of the features that are typical for a category are the main basis of identification, then typical items should always be easier to identify that atypical items even if similar instances have been seen in the same context. This constant advantage for typical items need not follow if subjects use memory for a prior item or prior occurrence to identify a test item. By the results of Experiments 1a, 1b, and 2, it is clear that facilitation depends upon the match in visual detail between the two encounters with an item. As there is no necessary relationship between the information available in the perceptual display and the categorical predictiveness of that information, and since episodic transfer depends upon the match of visual detail, episodic transfer may be independent of typicality effects. That is, the effects of typicality may be attenuated or perhaps eliminated on Identical items.

Method - Typicality Ratings

Subjects: Subjects were 12 students enrolled in an introductory psychology class at McMaster University. Subjects were run in small groups, numbering one to four. None of the subjects had participated in the previous experiments.

Materials: 25 quartets like those used in Experiment 2 were prepared. Each quartet was printed onto a separate page in the organization illustrated in Figure 3. The quartets were then photocopied onto transparencies. Response sheets were prepared for the subjects. On the response sheet there was a list of the names of the 25 object categories. One random sequence of categories was created, and that sequence was fixed for all subjects. Next to each name was four spaces for responses. Across the top of the response sheet was a numerical scale to be used in rating the typicality of each picture. The scale was from 1 (Very typical) through 7 (Very atypical). Also, a dictionary definition of typicality and instructions on how to perform the task were printed across the top of the page.

Procedure: After seating the subjects, the experimenter gave the definition of typicality and checked to make sure all the subjects understood what was meant by the term. They were then given the rating instructions. The quartets were displayed using an overhead projector. The quartets were displayed until all subjects had given a numerical rating to all four members.

To be sure that picture type differences do not result from category availability or other factors other than prior presentation, it is necessary to rotate quartets, and therefore categories, through the cells of the design. In order to test for the effects of typicality, then, it is necessary to determine which quartets contain variation in typicality. To accomplish this, the average typicality ratings within each perceptually similar pair of same name items was computed, yielding two scores per quartet. Next, the quartets were ranked according to the size of the difference between the two scores.

Method - Experiment 4

Subjects: Subjects were 48 students enrolled in an introductory psychology course at McMaster University. Subjects received partial course credit for their participation, and were run individually. None of the subjects had participated in the previous experiments.

Materials: Stimuli used in the experiment were 20 quartets of category members as used in Experiment 2. In addition, there were 12 filler items. Each of these fillers was from a different name category. These fillers were taken from the same sources as the other pictures used here.

Design: The experimental design was a two by four mixed design. The first factor, test method, was a between subjects factor. Subjects either clarified the picture at their own pace (Self Paced; N = 32) or the display clarified automatically and rapidly, and they were to stop the clarification when identification was possible (Automatic; N = 16). The second factor is test picture type. The experiment was conducted in two discrete phases: a learning phase and a test phase. One member from each of the critical item quartets was presented during the test phase. That picture was an Identical repetition of the member presented in training, the visually similar to a training item(SN/Sim), or visually different from a training item training (SN/Diff). If no member of the quartet was presented during the learning phase, then that test item was New. All filler items appeared in the training and test phases. Thus, fillers in the test phase were Identical. There were five Identical, five SN/Sim, five SN/Diff, and five New pictures in the test set, in addition to 12 filler items. Data were not analysed for the filler items.

Counterbalances were constructed such that each quartet member appeared in each cell of the design equally often across subjects. A pseudo-random sequence of item types was created, and then fixed for all subjects such that the distribution of test item types was approximately uniform throughout the series, but with fillers concentrated at the beginning of the test phase. Quartets were assigned such that one of the targets in each picture type were relatively typical, and one was relatively atypical.

Procedure: Pictures were displayed by an Apple //e computer on a 14 inch monitor. During the learning phase, subjects named aloud each picture as it was displayed on the computer screen, and were instructed to study the picture for as long as it remained on the screen. Presentation duration during the learning phase was about seven seconds. The intertrial interval in the first phase was subject paced. After completing the learning phase, the subjects were given the test phase instructions.

Preceding every test trial, the computer displayed the prompt "Press a key for next trial". When the subject pressed the key, a random dot pattern was displayed on the screen. Every time the subject in the Self Paced condition pressed the space bar of the computer, some of the random dots were removed, and picture information was displayed in its place, thereby increasing the signal to noise ratio. This process is documented in Vokey, Baker, Hayman, and Jacoby (1986). The subjects were instructed to continue pressing the space bar until they thought they might know the identity of the picture. In the Automatic clarification condition, the subject held down a key, causing continuous and quite rapid clarification. The rate of clarification was approximately 5% per second. Once the subject released the key, the clarification stopped, and the subject was required to make the identification. After successful naming of the picture, the picture was displayed fully clarified for about two seconds before proceeding to the next trial. The dependent measure was the amount of clarification upon correct naming of the picture.

In the Self Paced group, subjects were urged to offer any hypotheses they might have as to the identity of the picture. If a subject guessed the name incorrectly, they were informed that they were wrong and were allowed to continue the trial. This was not possible in the Automatic Clarification group. Here, subjects began a trial, and once they offered a hypothesis and stopped the clarification, but offered an incorrect name, they were told they were wrong and the trial discarded. The same naming criterion was used here as in the previous experiments.

Results and Discussion

The data from this experiment were subjected to two separate analyses. The first ignores the typicality manipulation and is concerned with the issue of analysis and the inspection time manipulation. The second analysis uses only the data for the Self Paced group and is concerned with the effects of typicality on the picture identification task.

Replicating the results of Experiments 2 and 3 and demonstrating the robustness of the effect over a variety of masking conditions, there was a reliable effect of test picture type. Identical pictures required the least clarification (37.3%), and New pictures required the most (48.2%). Same Name pictures required an intermediate amount of clarification, with SN/Sim items requiring less (40.0%) than did SN/Diff (43.7%), $\underline{F}(3,138)=33.102$. The results of this analysis are summarized in Table 6.

Table 6: Outcome of Experiment 4. Entries are in percent clarified at correct identification.

	Test procedure	
	Automatic	Self Paced
Test picture type		
Identical	39.027	35.552
SN/Sim	43.140	37.206
SN/Diff	46.488	41.074
New	51.390	44.647

There was also a reliable effect of test procedure. Subjects in the Self Paced group required less clarification (39.6%) than did the Automatic Clarification group (45.0%), $\underline{F}(1,46)=26.412$. The source of this effect seems to be the degree of certainty required before offering a hypothesis. The Self Paced group could afford to offer a response at a much lower level of certainty than could subjects in the Automatic Clarification group. If a subject in the Self Paced group offered an incorrect hypothesis, they were told so and allowed to continue the trial. The Automatic Clarification group were not allowed to correct their errors with subsequent responses. It therefore seems likely that these subjects required greater certainty, and therefore more information, before responding.

More important to the issue addressed here is the interaction between testing procedure and picture type. If the picture type effect seen in Experiment 2 and 3 and replicated in the Self Paced group here, depends upon a time consuming retrieval of prior items, the picture type effect should be diminished in size for the Automatic clarification group. There is no evidence here to suggest that the picture type effect is smaller in the Automatic clarification group than in the Self Paced group, $\underline{F}(3,138)<1$. From Table 6, it is apparent that the manipulation of Automatic versus Self Paced had no detectable effect on the picture type effect. Thus, there is no evidence to suggest that the picture type effect relies upon a selectively slower mode of retrieval of prior items.

In the analysis focussing on typicality effects, there was also an effect of picture type. Identical items required 33.0% clarification, SN/Sim items required 34.8%, SN/Diff items required 41.2%, and News required 43.1%, $\underline{F}(3,93) = 11.779$, MSe = 130.613. More relevant to the issue at hand, there was also a reliable effect of target typicality. Typical test items required 35.1% clarification, while atypical test items required 40.9%, $\underline{F}(1,31) = 16.381$, MSe = 130.168. This finding is consistent with the well documented effects of typicality. In the picture identification task, atypical items require more clarification than do typical items. While the picture identification paradigm is relatively new, and the effects of typicality on performance of the task are of a different form that those seen in verification and categorization tasks, the effects seen here are quite consistent with those seen in other paradigms. In both verification and categorization tasks, performance on atypical items tends to be both slower and less accurate. The finding that atypical items require more clarification converges with this general pattern of results.

More relevant to the issue of episodic transfer is the interaction between typicality and picture type. As mentioned earlier, the visual details emerging from the mask need not be related to categorical relatedness. Since performance on Identical and to a lesser extent SN/Sim items appears to rely upon item memory and item memory reflects the visual details present, typicality effects may be attenuated for these item types. By such a view, typicality should interact with picture type. However, the interaction did not approach significance, $F(3,93) = 1.085$, $MSe = 262.694$.

There are three reasons for caution in interpreting the lack of interaction. First, the data analyzed in this case is but a small subset of that collected. There is one observation per subject per cell of the design. The small number of observations very seriously reduces the statistical power of the analysis. This problem can be addressed in future experiments by simply increasing the number of quartets.

The second reason for caution is not as easy to address. First, note that differences in typicality are within quartet differences. That is, if a target is designated as typical, this means that its rated typicality is greater than members of the perceptually different same name pair. Typicality was computed this way because the experimental design required that quartets be rotated through cells of the design. However, there is no assurance that the typicality ratings for one quartet are in any way comparable to the typicality ratings for another. For example, one quartet may contain two very typical items and two moderately typical ones. Another quartet might contain two moderately atypical items and two very atypical ones. The differences of typicality within the quartets might be of equal magnitude, but there may be great variability in the typicality of typical targets. All that can be assumed is that the distribution of typicality among typical targets has a higher mean than the distribution for atypical targets. However, the amount of overlap in the two distributions is a matter of speculation. If this overlap is substantial, the typicality manipulation might be rather ineffectual.

The third reason for caution in interpreting the interaction has to do with the effects of target typicality on naming. In a verification task, subjects are presented with sentences such as "Penguins are birds" or "Robins are birds." The range of typicality in such an experimental preparation can be quite large. However, in the picture identification task, the task is to name the picture. Pictures of penguins tend to be called penguins, and robins tend to be called robins. If subjects call different members of a quartet by different names, there is no reason to suspect that the subjects are consulting a single knowledge structure in identifying all quartet members. To combat this, the members of a quartet were constructed such that subjects called all members by the same name. Unusual or especially distinctive quartet members tend to be called by different names that the rest of the quartet. During the process of stimulus development, such items were modified or replaced until all quartet members were called by the same name. The net effect of this process may have been to reduce the variability of item typicality within quartets. The range of typicality within a quartet, then, might be rather small. This, taken with the possibility that differences in typicality between quartets may be large, suggest that the typicality manipulation may have been rather weak. The manipulation was not completely ineffectual, as is evidenced by the reliable main effect of typicality. However, the manipulation might not be sufficiently strong to show the effects in an interaction.

While the current analysis lacks statistical power and perhaps an adequate typicality manipulation, the results here suggest that coordination between the picture identification task and traditional paradigms is in order. The main effect of typicality seen here is in line with typicality effects in other paradigms. In a sense, this is to be expected, as traditional paradigms are ultimately aimed at the same set of concerns as the picture identification task. There are independent reasons to suspect that the picture identification task has attributes to recommend it for some purposes. However, if the task is to be of general utility, it will be necessary to coordinate the picture identification task with other paradigms.

Experiment 5

Bruner and Potter (1964) have reported results suggesting that inappropriate hypotheses on the part of the subject can impede identification. They had subjects identify color slides which were presented blurred, and came into focus over time. If the slide started off very blurred, subjects required a higher level of focussing to finally identify the picture than if it started less blurred. More interestingly, if the picture came into focus slowly, subjects required a higher level of focussing to identify the picture than if it came into focus quickly. Apparently, having a long time to inspect the picture acted to the detriment of the subject. Bruner and Potter reported that subjects persevered with inappropriate hypotheses that had been developed early in the trial, and sometimes misinterpreted an aspect of the picture on the perceptual level. For example, a concavity would be interpreted as a convexity. These misinterpretations interfered with identification, and were more probable with long inspection times.

In Experiment 4, the times to identification on a substantial proportion of trials were unusually slow, which suggests that the same type of process might be going on in the present study as that in the Bruner and Potter study. A base line against which to evaluate how slow these identifications were is provided by a pilot study not reported here. The learning phase in this study involved naming pictures masked by a 50% clarified random dot mask such as that used in Experiment 4. While no formal data was collected, the experimenter noted that these subjects almost never failed to identify the picture despite the fact that, since it was the learning phase, all of the items were new. However, in the test phase of Experiment 4, about 30% of the New test items required more than 50 % clarification. Judged against this criterion, subjects "ought" to have been able to identify the test pictures in Experiment 4 at or before 50% clarification. Those trials requiring more than 50% clarification can be thought of as delayed identifications, paralleling the delayed identifications seen in the Bruner and Potter study.

Following the logic of Bruner and Potter, one might have expected that the self paced condition in Experiment 4 would produce more delayed identifications that would the automatic clarification condition. This should have been especially true for the SN/Diff and New pictures for which people would be less likely to have a correct hypothesis available. However, there was no significant interaction between picture type and clarification condition. Both conditions in Experiment 4 produced delayed identifications relative to the 50% criterion, but the possibility of extra time afforded by the self paced condition did not differentially affect the two picture types. Admittedly, this lack of interaction could be the result of any number of extrinsic factors, from lack of statistical power to instructional variables. However, there seemed to be a number of fundamental problems to refining the power of the contrast between the self paced and the automatic clarification procedures. The most fundamental of these is that to speed up the rate of clarification in the automatic clarification condition to make the manipulation stronger would change the nature of the dependent measure. At some rate of automatic clarification, the amount of clarification at identification would be more a measure of the decision and reaction time of the subjects than a measure of the amount of information necessary for identification. On the other hand, encouraging subjects to be more cautious in producing hypotheses in the self paced condition would at some point produce a trade off between time taken and information required for identification.

In the current experiment, a different approach was taken to controlling the effect of perseveration on a hypothesis in the clarification task. If delayed identifications are due to perseveration on a current hypothesis, then they might be reduced or eliminated if that hypothesis were interrupted and the next step in clarification for that picture were given after intervening, unrelated activity. To produce this, a discrete trial procedure was adopted. On the initial trial of this procedure, a picture was presented at a low level of clarification. If it was not identified, then another picture chosen at random from the pool of remaining items. On some later trial, the original picture would reappear at a higher level of clarification. Thus, between each successive step of clarification for a given picture, other items appeared which should serve to interrupt any ongoing hypothesis. If the delayed identifications are due to perseveration of hypotheses, then this procedure should reduce or eliminate them.

An additional manipulation was introduced in this experiment to further investigate the influence of deliberate analytic activity. For one group, each clarification step was on the screen only one second, while for the other, it remained on for as long as the subject wished. If perseveration on inappropriate hypotheses serves to reduce the effectiveness of deliberate analysis, thereby masking the contribution of this kind of processing to identification performance, then reducing these perseverations ought to better reveal it. Thus, the manipulation of inspection time should interact with picture type, with the advantage of Sn/Diff items over News being greater in the Unlimited time condition than in the Limited time condition.

Method

Subjects: Thirty two McMaster University students enrolled in an introductory psychology course served as subjects for partial course credit. Subjects were run individually, and none had participated in the preceding experiments.

Design: The experimental was a 2 by 4 mixed design, with one within subjects and one grouping variable. The experiment was arranged in two phases: learning and test. In the learning phase, subjects named a series of pictures, each from a unique name category. The within subjects factor was test picture type. Test pictures could be Identical. Same Name Similar (SN/Sim), Same Name Different (SN/Diff), or New name. The grouping factor was test phase presentation duration. In the test phase, half the subjects were allowed unlimited inspection time on each trial, while the rest were allowed one second per trial.

Materials: The stimuli used were the same 24 quartets of line drawings used in the other experiments reported here. In addition, 20 filler items were used, all of different name categories. The stimuli were displayed on an Apple //e computer equipped with a monochrome monitor.

Procedure: The experiment was done in two phases: a learning and a test phase. In the learning phase, the prompt "Press a key..." was displayed. Once the subject pressed a key on the computer, the line drawing was presented for five seconds. Subjects were required to name the picture out loud, and were instructed to study the picture for as long as it remained on the screen. The 20 fillers and one member from 18 of the quartets were presented in the learning phase.

After the 38 learning trials, subjects were given the test phase instructions. Subjects were informed that each trial was a different viewing of a different picture. At the outset of a test trial, the prompt "Press a key..." was displayed in the center of the screen. When the subject did so, a picture was displayed masked with a random dot mask. In the unlimited inspection time group, subjects were allowed to inspect the picture as long as they wished. They were asked to identify the picture with as few key presses as possible. If subjects offered a hypothesis, they were given immediate feedback by the experimenter. The same naming criterion was used here as in the other experiments. If the hypothesis was correct, the subject was instructed to press the ESCAPE key of the computer. Doing so presented the picture fully clarified for five seconds. If the subject had no hypothesis, or one that was incorrect, the subject was instructed to press the space bar of the computer to go on to the next trial. In the limited inspection time group, the masked picture was displayed for one second. The screen was then blanked, and a response was required of the subject. The subject either offered a hypothesis or stated that they had no hypothesis. If the hypothesis was correct, the subject pressed the ESCAPE key, presenting the fully clarified picture for five seconds. Otherwise, the subjects pressed the space bar to go on to the next trial.

The sequence of test pictures was partially random. The computer program designated a pool of ten items. For each of these items, the computer recorded the level of masking on each presentation, starting from 14% for critical items (members of a quartet), and 18% for fillers. The program randomly selected one of the pool of ten, and presented it at a masking level approximately 3% greater than its prior presentation (if any). If the item was not identified on that trial, it was returned to the pool. If an item was correctly identified, its current level of clarification was recorded and the item was removed from the pool, and another item was added to the pool starting with a masking level of 14% or 18% depending on item type. The net effect of this was a series of trials with obvious variability in masking level from trial to trial.

The fillers were clustered with eight of the first pool of ten items being fillers, eight clustered at the end of the test list, and the remainder being randomly distributed through the rest of the list. The clustering at the start of the list was done to allow subjects some practice with the task before encountering many critical items. At the end of the list, the last few items to be identified are very likely to be presented repeatedly in succession as the pool is decreasing in size. The clustering of fillers at the end was to assure that when this happened, only fillers would not yet be identified. Any deviation from the discrete trial format at the end of the list would affect only fillers, which are not analysed.

The datum resulting from a trial was the amount of clarification at identification. These values have been converted to a percentage measure, and submitted to an analysis of variance.

Results and Discussion

There was a reliable effect of inspection time. Those subjects allowed unlimited inspection time required 33.3% clarification for correct identification, while those subject in the limited inspection time group required 38.2%, $\underline{F}(1,30) = 16.278$. Clearly, allowing subjects unlimited time to scrutinize the test picture allowed them to identify the picture with less clarification.

There was also a clear test picture type effect. Identical pictures required 31.7% clarification, SN/Sim items required 35.3%, SN/Diff required 37.4%, and New items required 38.5%, $\underline{F}(3,90) = 12.372$. This picture type effect replicates the pattern of effects seen in Experiments 2, 3, and 4.

Of greater interest regarding the contribution of deliberate analytic activity to picture identification is the interaction between picture type and amount of inspection time. The interaction did not approach significance,

$\underline{F}(3,90) < 1.0$. The results of this experiment are summarized in Table 7. The hypothesis behind the experiment was that the facilitation of SN/Diff items over News reflected deliberate analysis of the test display, since the categories of SN/Diff items occurred in the training phase while the categories for New items did not. The perceptual dissimilarity between training items and SN/Diff test items makes episodic retrieval an unlikely source of facilitation. However, in both Experiments 4 and 5, manipulations of available time had no differential effects on different picture types as had been conjectured. Further, there is no evidence to suggest that the form of the picture type effect differs across Experiments 4 and 5 despite the substantial change in method.

Table 7: Outcome of Experiment 5. Entries are in percent clarified at correct identification.

	Presentation duration	
Test picture type	Unlimited	Limited
Identical	29.181	34.213
SN/Sim	32.549	38.071
SN/Diff	35.454	39.409
New	35.945	41.080

The discrete trial procedure did have one notable impact on task performance. In this experiment, there were fewer delayed identifications. In Experiment 4, approximately 30% of New trials required more than 50% clarification, while in this experiment, approximately 10% of New trials required more than 50% clarification. By changing the task from continuous clarification to discrete trials, it was hoped that errors in the subjects' analysis of the test picture might be easier to overcome, since they have multiple opportunities to initiate analysis on a given test picture. The decreased number of delayed identifications observed here is consistent with this suggestion.

There is one important point to note here. Overall, the amount of clarification necessary for identification is somewhat lower in this experiment than in the preceding ones (35.7% in Experiment 5, 42.3% in Experiment 4). The 50% criterion is an arbitrary one, and is located some distance above the mean performance. Any manipulation that serves to shift the entire distribution of scores lower down the axis will reduce the number of observations falling above a criterion that is fixed to the axis, as that criterion will be further away from the mean of the distribution. In effect, the 50% criterion is further above the mean performance in Experiment 5 than it is above the mean for Experiment 4. This fact alone would lead to a decrease in the number of observations above the 50% criterion. Therefore, the decrease in the number of delayed identifications can be taken as no more than suggestive. However, processing differences between the discrete versus continuous trial procedures remains one direction for future research into the processes underlying performance of the clarification task.

Experiment 6

In Experiment 6, a different approach was taken to finding the source of interfering hypotheses. In all preceding experiments, the stimuli were constructed such that retrieval of a perceptually similar item would lead to retrieval of the correct object category for that test item. However, in object categories, there is tremendous variability in how objects can look. There is no reason, therefore, that two perceptually similar items need be of the same object category. By contriving stimuli such that the retrieved item is of the wrong object category, it might be possible to demonstrate episode based interference. If a subject were to encounter a test item perceptually similar to a training item, but of a different name category, the subject's ability to identify the test item might be inhibited. This might produce delayed identifications for object categories relative to those categories for which there were no interfering items. For example, for the materials in Figure 4, having seen the chest of drawers might interfere with the subject's ability to identify the target piano.

Method

Subjects: The subjects were 12 McMaster university students enrolled in an introductory psychology class, and participated for partial course credit. Subjects were run individually, and none had participated in the preceding experiments.

Materials: For the experiment, 16 quartets of items were constructed. Two of the pictures in each quartet were perceptually similar same name items. One of these pictures was designated the target picture. One member of the quartet was a perceptually dissimilar same name item. The remaining quartet member was of a different object category. The target picture was duplicated and modified such that it became a member of a different object category. See Figure 4 for an example of a quartet. For descriptive ease, these different name, perceptually similar items will be referred to as "false friends."

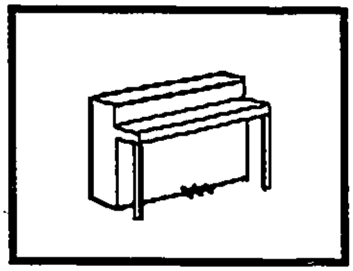

Same Name similar training item

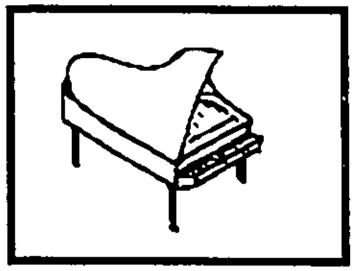

Same Name different training item

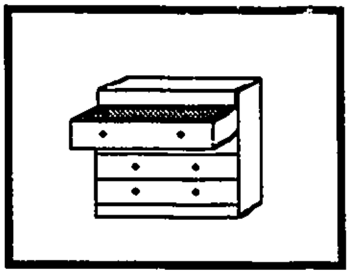

Different Name similar training item

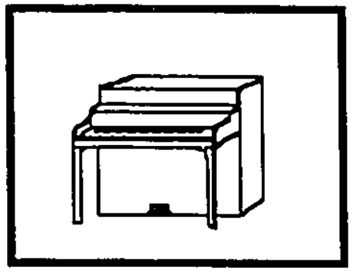

Target item

Figure 4: An example of a quartet used in Experiment 6. The target item appeared in the test phase. In the training phase, either the Same Name similar or Same Name different training item was presented. For half the target items, the Different Name similar item (False Friend) also occurred in the training phase.

Design: The experiment was a two by two within subjects design. The experiment was run in two phases. In the first phase, subjects were presented with 24 pictures. From four of the quartets, the item perceptually similar to the target picture was presented. From four other quartets, the same name different picture was presented. From the other eight quartets, the different name "false friend" item plus a same name item (half similar and half dissimilar) were presented. In the second phase, only the target pictures were presented. Thus, half the target pictures were similar to a same name training item and half dissimilar. Orthogonal to this, half were similar to a false friend item. Counterbalances were constructed such that each target picture occurred equally often in each of the four cells of the design.

Procedure: At the outset of a training trial, the prompt "Press a key..." was presented on the screen. Once the subject pressed the space bar of the computer, the picture was presented on the computer screen for approximately seven seconds. Subjects were required to name the picture aloud, and were instructed to study the picture for as long as it remained on the screen.

After the learning phase, the subjects were given instructions for the test phase. At the outset of a trial, the "Press a key..." prompt was displayed. Once subjects did so, a random dot pattern was displayed. Now every time the subjects pressed the space bar, random dots were removed, and replaced with pixels from the target picture. Subjects continued pressing the space bar until they thought they might know the identity of the picture. Subjects verbally stated their hypotheses. The correctness of subjects' hypotheses was judged using the same naming criterion as in the preceding experiments. Once a subject correctly named the picture, they were instructed to press the ESCAPE key on the computer. Doing so presented the picture fully clarified for approximately three seconds. The picture was then removed, and the next trial began. The sequence of items was random and unique for each subject.

The datum resulting from each trial was the number of key presses necessary for correct identification. These values were converted into the percentage clarified, and submitted to an analysis of variance. In addition, a record was kept of the hypotheses subjects offered.

Results and Discussion

Replicating the previous experiments, there was a reliable effect of picture type. Target items similar to a same name training picture (SN/Sim) required 33.6% clarification while those dissimilar to a training item (SN/Diff) required 38.5%, $\underline{F}(1,11) = 30.848$.

Of greater interest is the effect of the presence of false friends. Those items for which a false friend had occurred in the learning phase required 35.9% clarification, while those for which no false friend had occurred required 36.1% clarification, $\underline{F}(1,11) < 1.0$. Thus, there is no evidence that the presence of a false friend slowed identification performance. Further, there was no interaction between the two variables, $\underline{F}(1,11) < 1.0$.

Table 8: Outcome of Experiment 6. Entries are in percent clarified at correct identification.

	"False Friend"	
Test picture type	Present	Absent
SN/Sim	33.455	33.749
SN/Diff	38.431	38.482

Not all subjects are equally skilled at the picture identification task, and not all targets are equally easy to identify. If it were the case that between target item variability were greater than between subject variability, an item analysis would be a more sensitive test. The analysis was repeated with items as the replicate rather than subjects. The pattern of results is identical. There is a reliable effect of picture type,

$F(1,15) = 17.053$. There was no reliable effect of the presence or absence of false friends, $F(1,15) < 1.0$. Also, there was no interaction, $F(1,15) < 1.0$.

One possible explanation of the lack of effect is that the presence of different name perceptually similar items simply had no effect at all. To test this possibility, the frequency with which the name categories of the different name items was examined. Unfortunately, rather few false hypotheses of any type were offered. Subjects offered the name of a different category, perceptually similar item eight times. Of these eight, the hypothesis occurred on a perceptually different, same name target item six times. With frequencies this small, it is hazardous to draw any firm conclusions. However, the occurrence of these "false friend" name categories does at least suggest that the name categories of the false friends were available to the subjects, and that these name categories were (perhaps) more likely to occur to subject on perceptually dissimilar test trials than on similar test trials. However, the availability of the "false friends" did not affect identification performance.

The failure to demonstrate an effect of false friends is puzzling. The central premise of Experiment 5 was that delayed identifications might be a result of deliberate analytic activity. The Limited versus Unlimited inspection time conditions ought to have been powerful manipulations of such deliberate analytic activity and resulted in an interaction with picture type. The ineffectiveness of these manipulations suggests that the source of delayed identifications may not be deliberate analytic activity. In Experiment 6, a different source of interfering hypotheses was investigated. If the obviously robust picture type effect results from prior presentations associating a particular shape with a category, then supplying false friends should have led to interference.

An extreme alternative hypothesis is that episodic transfer is not operating on a category level at all. Rather, similar prior presentations might help to perceptually organize a new stimulus. If so, a subject might be less likely to integrate fragments of the target picture with adventitious patterns in the mask. Additionally, the subject might find it easier to organize the lines in the target into basic shapes independently of category identification. However, this hypothesis also is not supported by the results of Experiment 6. By definition, SN/Diff items are items that do not receive facilitation from perceptually similar items within their own category, but these items were not identified more rapidly when the perceptually similar false friends were provided. This suggests that perceptual similarity in and of itself is not sufficient to produce facilitation.

A more moderate possibility is that the effect of similarity to prior items in the same category depends on a combination of perceptual and conceptual cuing. If so, then the lack of effect of false friends might result from two antagonistic effects. The false friend aids performance by facilitating perceptual organization of the target, but interferes with performance by making the wrong object category available. This suggests that there might be a strong effect of prior stimuli that are perceptual as well as conceptual false friends: stimuli that encourage the subject to organize the stimuli incorrectly as well as not making the appropriate category contextually available.

Another additional research project could directly address why the false friends had no effect on perseveration despite some evidence that they were a source of incorrect hypotheses. One possibility is that when the subject discovers that the false friend name category is incorrect for the target item, there is already sufficient clarification to allow perception of features that cue the correct category. Perseveration might be limited to cases in which perceptual organization is incorrect or in which, despite correct perceptual organization, there are no highly diagnostic cues for the correct category. If, using the example of Figure 4, the pedals and the keys of the target upright piano were faint, then perseveration might occur even after discovering that the false hypothesis was questionable.

Summary

The common denominator to the experiments in Chapter 4 is that factors that could reasonably be expected to have important consequences for identification performance were manipulated. In Experiments 4 and 5, factors that ought to have strong consequences for deliberate analytic activity were manipulated, but no differential effects on different picture types were detected. However, in Experiment 4 evidence was found that perseveration on inappropriate hypotheses was occurring, implying deliberate analytic activity. The conditions of Experiment 4 were such that the contributions of deliberate analysis might have been masked or attenuated. The procedure used in Experiment 5 was designed to reduce perseveration on inappropriate hypotheses, thereby making it more likely that the effects of deliberate analysis would be seen. However, the experiment still failed to differentially affect the contributions of deliberate analysis.

All of this could mean no more than that deliberate hypothesis testing makes no substantial contribution to performance of the picture identification task. However, the outcome of Experiment 6 was not dependent on an effective manipulation of deliberate analytic behavior. In this experiment, false friends were provided that could have affected unreflective first impressions as well as deliberately pursued hypotheses. However, the false friends manipulation did not interact with the picture type effect either. The source of the observed delayed identifications may lie more at a level of initial perceptual organization than at the level of assigning already organized partial stimuli to object categories.

Chapter 5: Discussion

A summary of issues and findings

A prominent view in psychology is that picture and object identification typically rely upon general semantic or categorical knowledge, and that episodic knowledge is an unusual route to identification. This position has been characterized in this thesis as the robustness position to capture the idea that the knowledge structures upon which identification depends are robust in the sense that they are invariant across category members, situations, and viewing conditions. This view is implicit in the prototype view of concept learning, and an explicit goal in Biederman's theory of object identification. The contrasting position is that of special expertise. In the special expertise view, specific prior experiences are used as the knowledge base for identification. This process is implicit in episodic theories of concept learning, and bears strong kinship with Newell and Simon's view of expertise in that solutions of prior perceptual problems are adapted to new problems.

In most of the discussion of this thesis, the robustness and special expertise positions have been cast in rather extreme terms: reliance solely on abstract knowledge structures and complete reliance on specific prior processing episodes. Undoubtedly, aspects of each positions are important to object identification. That people can describe an average dog despite never having seen "the average dog" is beyond dispute. This kind of knowledge can only be described as an abstraction. What is disputed here is whether this kind of knowledge is the primary means by which pictures and objects are identified, and whether special expertise is reserved for special situations.

According to the robust view of object identification, performance on a given item or category should be stable. That is, the amount of visual information necessary for identification should not change appreciably after a single prior experience, since the knowledge base and identification procedure for naturally occurring categories have been well developed. An analogy can be made with the "running mean" calculation,

$$M_n = M_{n-1} + (X - M_{n-1})/N,$$

where M_n is the new mean, M_{n-1} is the previous mean, X is the new value to be included in the mean, and N is the number of observations. Once a significant number of values have been included in the mean, the additive term drops to (essentially) zero. By analogy, the mean is the abstract knowledge structure, X is the new item, and the additive term is the new learning resulting from the new item. Thus, the knowledge base or identification procedure upon which picture identification depends should be subject to very little revision once the category has been well learned. One resource that some robustness positions have for responding to individual events is priming. In its explanatory sense, priming is a transient change in the activation or availability of a stable knowledge structure, which thereby reduces the information or time necessary for identification of members of that category. Priming must be transient or it violates the goal of having identification rely mainly upon stable and robust knowledge.

The special expertise position is based upon episodic representation. Past experiences are retrieved and used in the process of identification. There are two phenomena diagnostic of special expertise: durability and transfer that is sensitive to similarity to prior episodes. Because special expertise is the adaptation of prior episodes to new perceptual arrays, the effects of prior experience ought to be detectable for as long as the prior episode can be retrieved, and episodic memory is very long lasting. In Experiment 3, effects of prior presentation were found after a 24 hour delay: a delay far too long to be explained by priming. Further, no evidence was found indicating that the effects of the prior experience had diminished over the delay.

The second important property of the special expertise position is transfer dependent upon specific similarity. In Experiments 1a and 1b, matching the visual details across two presentations of an item resulted in greater facilitation than if the details mismatched. In Experiments 2, 3, 4, and 5, greater facilitation was found on Identical items than on Same Name similar items. Further, in Experiments 2, 4, 5, and 6, greater facilitation was found on Same Name Similar items than on Same Name Differents. These results demonstrate that performance depends critically on how similar the prior encoding event is to the current test situation in either the visual details presented or similarity between nonidentical training and test items.

The effects of specific similarity are consistent across a variety of procedures and task demands: both explicit and implicit. Interleaved training/test series and discrete phase experiments both show the same basic pattern. Allowing or discouraging deliberate analysis of the test display failed to affect the pattern of results, suggesting that the pattern indicative of special expertise does not result from a special strategy adopted by the subjects. Requiring concurrent memory judgments and therefore explicitly orienting subjects to episodic memory did not affect performance. These are all null results, and there are important limitations on the evidential value of null results. However, across a considerable variety of experimental preparations, these null results provide convergent evidence that the episodic basis of performance in the picture clarification task is not a consequence of a special strategy adopted by subjects under restrictive task demands.

In addition, the materials used in the experiments are conventional renderings of common object categories. As a result, the set of object categories to which an as of yet unidentified item could belong is very large. This stands in contrast with the typical two alternative forced choice categorization tasks used in concept learning experiments. Normally, subjects in a categorization experiment are required to decide if an item belongs to category A or category B. When identifying objects under normal "real world" conditions, people very rarely have only two alternatives to decide between. A further consideration is that pictures represent a kind of perceptual variability, both within category and within item, that is a common property of real world identification problem. Artificial materials typically fail to represent this kind of variability or restrict it greatly. Thus, clarification of pictures is at least more like identifying real world objects than is the identification of artificial materials that severely restrict perceptual variability.

Finally, the experimental methods used here do at least mimic some conditions that occur in the real world. Ecological validity was one of the design considerations for the experimental preparations used here. The 'coming around the corner' masking technique mimics a common identification problem. The random dot masking used in Experiments 4, 5, and 6 are similar to defocussing and viewing through snow or fog. While the checkerboard masking technique is somewhat contrived, it does relate to extant theories of object identification. The stability of results across this range of masks suggests that this paradigm is tapping processes which may be general in dealing with distortion and occlusion in the world.

In summary, the use of prior episodes in the identification of pictures, or special expertise, is a phenomenon of sufficient breadth to be interesting in understanding normal object identification. It is demonstrated with kinds of distortion plausible in the world, across a variety of task demands, and with drawings of naturally occurring object categories. The evidence for special expertise cannot be attributed reasonably to a few highly specific conditions of distortion, restrictive task demands, or artificial stimuli. Thus, on the basis of the work reported here, special expertise would seem a normal, rather than exceptional, component of picture and object identification.

Implications for theories of concept learning

The work reported here speaks to the larger issue of concept learning. There has been a great deal of recent work documenting that item knowledge is a means by which stimuli can be classified: more than would be wise to review here (see Brooks, 1987). The experiments reported here extend that work. The difference between identifying an object as a dog and categorizing it as one is largely one of task description. The methods used here are directly concerned with identification, but the results are relevant to issues in categorization.

One tradition in concept learning research typically involves the presentation of a series of artificial stimuli that the subject categorizes. Subsequently, subjects are presented with stimuli that are especially diagnostic of what the subject might have learned, and classification accuracy is measured. The researcher can then infer what the subject has learned and how that knowledge is represented from diagnostic patterns of error. This kind of inference cannot be made with the materials used here since subjects (virtually) never fail to correctly identify an item. However, in the current experiments, it is clear that subjects use prior episodes to identify pictures. While the current work does not address learning directly, it does address what has been learned from the subject's innumerable experiences with members of naturally occurring object categories. Thus, the demonstrations of special expertise here lend support to the position that conceptual knowledge is distributed across prior episodes rather than primarily by abstract summaries of that prior experience.

There is a further issue to consider. In concept learning research, heavy emphasis has been placed on subjects' ability to correctly categorize items. In some important sense, category members are equivalent, and the ability to categorize cats as cats and dogs as dogs is a valuable and useful skill. However, in real world situations, it is as important to discriminate the virtuous Lassie from the malevolent Cujo as it is to categorize both as dogs. The ability to discriminate between category members is a vital skill. Such discriminations are only possible if item specific characteristics represented and subjects exploit those characteristics.

One way in which people could correctly categorize objects and correctly discriminate between category members is to use different knowledge bases for the two tasks: abstract knowledge to identify the category and episodic knowledge to identify the specific item. First, a person could identify an object as a dog, and then determine which dog it is. This is implicit in current abstractive views of concept learning, and an explicit design feature of Biederman's theory of object identification. Such a view presumes that category identification is prior to item identification.

An equally plausible hypothesis is that some part of item identification precedes categorization. The assignment of an item to a category is subsumed by item identification. To identify a beast as Lassie includes the assignment of the beast to the category "dog" provided that you previously knew Lassie to be a dog. Such a process is not restricted to literally familiar items. Experiments 2, 3, 4, 5, and 6 document transfer graded according to similarity between training and test items, implicating the retrieval of the training item in identifying the test item. Applied to a categorization task, this retrieval of a prior item includes the retrieval of sufficient information for categorization, as the item had been assigned to a category as part of the training procedure: it was named.

One attraction of this view is its flexibility. It is clear that people have item knowledge, as they are able to discriminate between category members. When required to identify an item, people retrieve the relevant information from episodic memory. When task demands change and the item is to be assigned to a category, people could retrieve information from episodic memory and report a different part of that information. Item identification and categorization could differ not in knowledge base, but rather in what is reported from retrieved episodic knowledge.

Kahneman and Miller (1986) present a strong argument that abstraction is not an automatic component of experience. They suggest that when judgment is required, prior episodes are retrieved, and an abstraction is formed. The required judgment is then made with reference to this temporary abstraction. By this view, abstraction is post hoc in that it is performed when a judgment is necessary, and not as an automatic consequence of experience. A similar position has been developed by Medin and Ross (1989). They argue that the process of abstraction is not autonomous, meaning that it is similarly post hoc. In parallel with Kahneman and Miller's arguments, they argue that episodic knowledge is recruited and abstraction performed when required. It is not the obligatory consequence of experience that abstractive accounts of concept learning would have it.

While the work reported here extends the case that episodic memory is used in the identification of pictures and furthers the argument that categorization often relies upon an episodic knowledge base, there are two significant puzzles. Typicality effects have long been taken as evidence of the existence of prototypes. In Experiment 4, typicality influenced ease of identification. This finding converges with other effects seen in category verification and categorization tasks. A concept closely related to (but not identical with) typicality is frequency of instantiation. Typical items tend to occur more frequently than do atypical items. The finding that typical items are easier to identify that atypical ones could be explained if frequency of instantiation is taken into account. Typical new items would be more likely to have a similar item in memory than would an atypical new item. The prior presentation of the item or a similar one ought to remove this disadvantage of atypical items. Typicality effects were not detectably smaller after prior presentation, suggesting that this conjecture is incorrect or incomplete. If the picture identification task is to be of general utility in concept learning research, it will be necessary to coordinate the findings using the task with findings using other paradigms. This requires that an account for the persistence of typicality effects after prior presentation be found. The use of prior episodes in picture identification is not yet well enough understood to make such an account.

The second major puzzle is that under a variety of task demands and experimental preparations, the size of the difference in performance on SN/Diff and New items has remained essentially constant. It is quite unlikely that this difference reflects episodic retrieval, as SN/Diff items were constructed to be as different as possible. Since a SN/Diff test item presents a different set of perceptual cues than the training item, retrieval of that training item is an unlikely basis of performance. If episodic retrieval is an unlikely explanation, then category level knowledge would seem the best explanation. If category level knowledge is the basis for the advantage of SN/Diff items over News, then one would expect manipulations such as inspection time or clarification technique (Automatic versus Self Paced) to affect the size of the advantage. There is no evidence of such an interaction. Clearly, the expectation that category level transfer should yield to these manipulations was in error.

There are two kinds of explanation for the lack of effect. First, and most obvious, is that the manipulations were too weak to be effective, but other manipulations could have been used that would have been effective. This possibility hard is to address until those more effective manipulations are found if they do exist. On the basis of the evidence available, category level transfer does not appear to be affected by manipulations that ought to change subjects' reliance on analytic strategies in picture identification.

The second possibility is that categorical transfer in picture identification operates at a level that is impenetrable to the subject's choice of strategy. Category priming is one mechanism of this type. Experiment 3 was a sensitive test that could have detected the dissipation of priming effects over a 24 hour delay. The experiment was sensitive enough to detect an effect accounting for less than three percent of the observed variance. That this experiment did not detect a dissipation of priming effects should be taken as strong evidence that priming did not contribute to the facilitation seen on same name test items.

The robustness of the picture type effect across this variety of manipulations is a puzzle. From the outset, it has been acknowledged that performance of the picture identification task depends upon both episodic and categorical information. Experiments 4 and 5 sought to manipulate the contribution of categorical information, but in neither case were picture types differentially affected. Experiment 6 sought to manipulate the utility of episodic information, but the picture type effect was similarly unaffected. Future research will require manipulation of the combination of these information types, such as the use of perceptual and conceptual false friends, or directly manipulating the salience of highly diagnostic features. One goal of future research using the picture identification task is to more fully investigate the coordination of item level and category level information.

BIBLIOGRAPHY

Bartram, D. J. (1974). The role of visual and semantic codes in object naming, **Cognitive Psychology, 6**, 325-356.

Biederman, I. (1987). Aspects and extensions of a theory of human image understanding. In Z. Pylyshyn (Ed.), **Computational Processes in Human Vision: An interdisciplinary perspective.** New York: Ablex.

Biederman, I. and Ju, G. (1988). Surface versus edge-based determinants of visual recognition, **Cognitive Psychology, 20**, 38-64.

Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch and B. Lloyd (Eds.), **Cognition and Categorization**, Hillsdale, NJ: Erlbaum.

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.** Cambridge, England: Cambridge University Press.

Bruner, J. S. and Potter, M. C. (1964). Interference in visual recognition, **Science, 144**, 424-425.

Gollin, E. S. (1960). Developmental studies of visual recognition of incomplete objects, **Perceptual and Motor Skills, 11**, 289-298.

Hirshman, E., Snodgrass, J. G., Mindes, J., and Feenan, K. (1990). Conceptual priming in fragment completion, **Journal of Experimental Psychology: Learning, Memory, and Cognition, 16**, 634-647.

Jacoby, L. L. (1983a). Perceptual enhancement: Persistent effects of an experience, **Journal of Experimental Psychology: Learning, Memory, and Cognition, 9**, 21-38.

Jacoby, L. L. (1983b). Remembering the data: Analyzing interactive processes in reading, **Journal of Verbal Learning and Verbal Behavior, 22**, 485-508.

Jacoby, L. L. (1990). Separating consciously controlled from automatic processes of memory: Effects of dividing attention at test. Manuscript submitted for publication.

Jacoby, L. L., Baker, J. G., and Brooks, L. R. (1989). Episodic effects on picture identification: Implications for theories of concept learning and theories of memory, **Journal of Experimental Psychology: Learning, Memory, and Cognition, 15**, 275-281.

Jacoby, L. L. and Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception, and concept learning. In G. H. Bower (Ed.), **The Psychology of Learning and Motivation: Advances in research and theory, Volume 18**. New York: Academic Press.

Kahneman, D. and Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives, **Psychological Review, 93**, 136-153.

Keppel, G. (1982). **Design and Analysis: A researcher's handbook**, (pp.87-92). Englewood Cliffs, N. J.: Prentice-Hall Inc.

Legge, G. E., Cohen, M. A., and Stromeyer, C. F. (1978). Spatial frequency masking with briefly pulsed patterns, **Perception, 7**, 161-166.

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

Logan, G. D. (1989). Automaticity and cognitive control. In J. S. Uleman and J. A. Baugh (Eds.), **Unintended Thought**. New York: Guilford Press.

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? **Cognitive Psychology, 22**, 1-35.

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

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

Medin, D. L. and Smith, E. E. (1984). Concepts and concept formation, **Annual Review of Psychology, 35**, 113-138.

Mitchell, D. B. and Brown, A. S. (1988). Persistent repetition priming in picture naming and its dissociation from recognition memory, **Journal of Experimental Psychology: Learning, Memory, and Cognition, 14**, 213-222.

Nelson, D. L., Reed, V. S. and Walling, J. R. (1976). The pictorial superiority effect, **Journal of Experimental Psychology: Human Learning and Memory, 2**, 523-528.

Newell, A. and Simon, H.A. (1972). **Human Problem Solving.** Englewood Cliffs, New Jersey : Prentice Hall.

Parkin, A. J. and Streete, S. (1988). Implicit and explicit memory in young children and adults, **British Journal of Psychology, 79**, 361-369.

Snodgrass, J. G. (1987). Sources of learning in the picture fragment completion task, Paper presented at the meetings of the Psychonomic Society, Seattle, Washington, November, 1987.

Snodgrass, J. G. (1990). Sources of learning in the picture fragment completion task. In S. Lewandowsky, J. C. Dunn, and K. Kirsner (Eds.), **Implicit Memory: Theoretical issues.** Hillsdale, N. J.: Lawrence Erlbaum.

Tulving, E. and Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory, **Psychological Review, 80**, 352-373.

Tversky, B. and Hemenway, K. (1984). Objects, parts, and categories, **Journal of Experimental Psychology: General, 113**, 169-193.

Vokey, J. R., Baker, J. G., Hayman, G., and Jacoby, L. L. (1986). Perceptual identification of visually degraded stimuli, **Behavior Research Methods, Instruments, and Computers, 18**, 1-9.

Warren, C. and Morton, J. (1982). The effects of priming on picture recognition, **British Journal of Psychology, 73**, 117-129.