COGNITIVE PSYCHOLOGY **10,** 178–202 (1978)

Interference in Memory for Pictorial Information

JOHN R. ANDERSON

Yale University

AND

REBECCA PAULSON

University of California, San Diego

Subjects studied either faces composed from visual features or verbal facts composed from concepts. Recognition times were increased for both faces and facts when they were composed of elements that occurred in multiple study items. In Experiment 1 the interfering effect of other study items was much larger for verbal facts than for faces. This difference was largely eliminated in Experiment 2 where care was taken to control the features by which the faces were encoded. Experiment 2 also showed that verbal information could interfere with pictorial information and vice versa. However, this cross-modality interference was much weaker than within-modality interference. The data are consistent with the ACT theory in which pictorial material and verbal material are stored together in an abstract propositional network. The subnode model (Anderson, *Language, memory, and thought*, Hillsdale, N. J.: Lawrence Erlbaum, 1976) can account for the greater within- than cross-modality interference.

There is a considerable controversy over the nature of the information representations in various cognitive tasks. Paivio (1971) proposed that there are at least two distinct modes of representation—a verbal mode and a visual or imagery mode. Anderson and Bower (1973) and Pylyshyn (1973) criticized this position and argued that all information is represented in a single, abstract propositional mode. The difference between verbal and visual information, they argued, is not a matter of the nature of the information representation, but rather a matter of the content that is encoded in that representation. More recently, Paivio (1976) and Kosslyn and Pomerantz (1977) have criticized the propositional position and reasserted the claim of a distant visual or imagery mode of representation.

This research was supported by Grant No. BNS76-00959 from the National Science Foundation. We would like to thank Lynne Reder, Robert Moyer, Brian Ross, Miriam Shustack, Edward Smith, and two anonymous reviewers for their comments on the manuscript.

Correspondence should be sent to John R. Anderson, Department of Psychology, Yale University, New Haven, CT 06820.

0010-0285/78/0102-0178$05.00/0

The hypothesis of a propositional representation is relatively rigorous and has been embodied in a number of computer simulation models (e.g., Anderson, 1976; Anderson & Bower, 1973; Baylor, 1971; Kintsch, 1974; Moran, 1973; Norman & Rumelhart, 1975). In contrast, the imagery hypothesis has been plagued by vagueness. However, recent work of Kosslyn and Shwartz (1977) represents a commendable effort to bring precision to the imagery hypothesis. Anderson and Bower (1973) and Pylyshyn (1973) criticized an imagery theory that was based on a metaphor with a photograph. Kosslyn and Pomerantz (1977) and Paivio (1976) have complained that this is a "straw man" position that few hold seriously. However, they do not provide a rigorous definition of what an imagery theory might be. Nonetheless, they argue that a number of recent phenomena require an imaginal representational format in distinction to a propositional format. These include studies of mental rotations (e.g., Cooper & Shepard, 1973; Shepard & Metzler, 1971), perceptual comparisons based on long-term memory information (Moyer, 1973), and effects of image size (Kosslyn, 1975).

An issue of particular concern to us is whether it is necessary to assume different long-term memory representations for verbal and visual material. Advocates of an imagery position have argued that an imagery representation tends to be used for visual material and a verbal or propositional one for verbal material. Propositional theorists propose an abstract propositional representation for all memory. Clearly, this issue of memory representations is a subset of the larger issue of whether imagery representations are necessary. Even if it could be established that a propositional representation is appropriate for all long-term memory phenomena, it is possible that an imagery representation will prove useful for other cognitive phenomena. This suggestion has been made by Anderson and Bower (1973).

It is clearly a step away from parsimony to propose two distinct memory representations. However, it is possible that this step could help gain parsimony in other areas. If it could be shown that memory for pictures and memory for verbal material obeyed different laws, then it might prove easier to account for these different laws with different representations. It would be particularly inelegant to claim that the same propositional representation obeyed one set of memory laws when it encoded information from a verbal source and a different set of laws when it encoded information from a pictorial source.

There have been a number of attempts to show that memory for pictorial information obeys different laws than memory for verbal information. It had been claimed that pictorial memory has greater capacity (e.g., Shepard, 1967; Haber, 1970; Standing, 1973). These are demonstrations that meaningful visual stimuli (e.g., faces) are much better remembered than meaningful verbal material. However, further analysis seems to indicate that, if there is better memory for pictures, it is not due to their pictorial quality per se. Material that is not familiar to the subject is not well retained (e.g., snowflakes—Goldstein & Chance, 1971). The same visual item is much better remembered when it is given a meaningful interpretation than when not interpreted (Wiseman & Neisser, 1971; Bower, Karlin, & Dueck, 1975). It has also been shown with verbal material that the memory for the material depends on the richness of interpretation that can be given to the material. Thus, it seems that both verbal and visual information are subject to a "depth of processing" principle (Craik & Lockhart, 1972; Craik & Tulving, 1975; Anderson, 1976) such that more deeply or richly processed material is better remembered. The fact that subjects display better memory for both visual and pictorial material only when they are meaningfully processed is consistent with the view that the long-term memory representations underlying both types of information are propositional—i.e., a representation which encodes the abstract meaning of the material rather than its physical characteristics.

Even if we do concede an abstract propositional representation to both pictures and linguistic material, this would leave unexplained the alleged superiority of pictorial material when meaningfully encoded over verbal material when meaningfully encoded. It may be that, even though pictures and linguistic material are coded in the same long-term memory format, pictorial material enjoys a more rapid rate of encoding. This may reflect an inherent superiority of the visual system or, perhaps, greater practice at encoding visual material. The latter possibility would be consistent with the observation of poor memory for unfamiliar pictorial material. However, there is another possible explanation: The experiments demonstrating superiority of visual material have involved an important recognition component. Either the task is recognition or the subject is required to recall to a pictorial stimulus. In such cases it is hard to equate the similarity among the pictorial stimuli with similarity among the linguistic stimuli. Clearly, recognition memory for pictures will be poor if sufficiently similar foils are used. Interestingly, a recent series of experiments (e.g., Mandler & Johnson, 1976; Mandler & Ritchey, 1977) indicate that subjects are best at detecting changes that seriously alter the interpretation of a picture and are not so good at detecting larger physical changes that do not change a picture's interpretation. This is also consistent with the hypothesis of a memory representation that encodes abstract meaning rather than physical characteristics.

The Fan Effect

These considerations suggest that a profitable paradigm for comparing pictorial versus linguistic information is one that focuses on retrieval properties of these two stimulus types and not properties that can be

attributed to acquisition. Accordingly, the experiments reported here are concerned with whether the same laws of interference hold for the two kinds of material. The way we have chosen to study this question involves what has been called the *fan effect* (Anderson, 1974; Anderson, 1976, chap. 8; Lewis & Anderson, 1976; Thorndyke & Bower, 1974). This is the phenomenon that has been most studied in a sentence recognition paradigm where subjects commit sentences to memory. The concepts composing a sentence can occur in varying numbers of other study sentences. The fan effect refers to the fact that subjects are slower to recognize sentences or reject unstudied sentences that are composed of concepts that have occurred in many study sentences.

The fan effect has been interpreted within the context of an associative network theory called ACT (Anderson, 1976). This theory assumes that each concept is represented by a node in memory and that each concept is connected to the propositions involving it by associated links. The more propositions involving a concept, the greater the fan out of the concept node. A particular probe is recognized by activation spreading from the concepts in the probe to activate the propositional structure encoding memory of the sentence queried in the probe. The speed of spread is inversely related to the fan out of the concept—the more sentences sharing a concept, the longer to activate any sentence. The negative judgments are determined by a waiting process which produces a rejection if the concepts in a probe do not result in an intersection of activation after a specified time. This waiting process is adjusted to reflect the average fan out of the concepts in the probe. The process waits longer for higher fan concepts because it takes longer to activate propositions involving these concepts.

This fan effect defined on reaction times is related to the more traditional demonstrations of stimulus specific interference defined on probability correct (see Anderson, 1976, for a discussion of this relationship). We focus on reaction time because it provides a more sensitive measure and a measure more amenable to treatment within the ACT theory.

The fan effect has so far only been demonstrated with verbal material. This experiment was conducted to determine whether the analog of the fan effect could be obtained with pictorial material. In this research the material used were faces composed from the IDENTI-KIT materials. This set of material allows one to systematically construct faces out of independent visual features. Examples of the faces so constructed are illustrated in Fig. 1. The faces used in Experiment 1 were constructed from four face features. The hair style was one feature, the chin style a second, the mouth–nose combination a third, and the eye–eyebrow combination a fourth.

A given face feature could appear in a large number of faces (nine), a medium number of faces (three), or in just one face. Each face was asso-

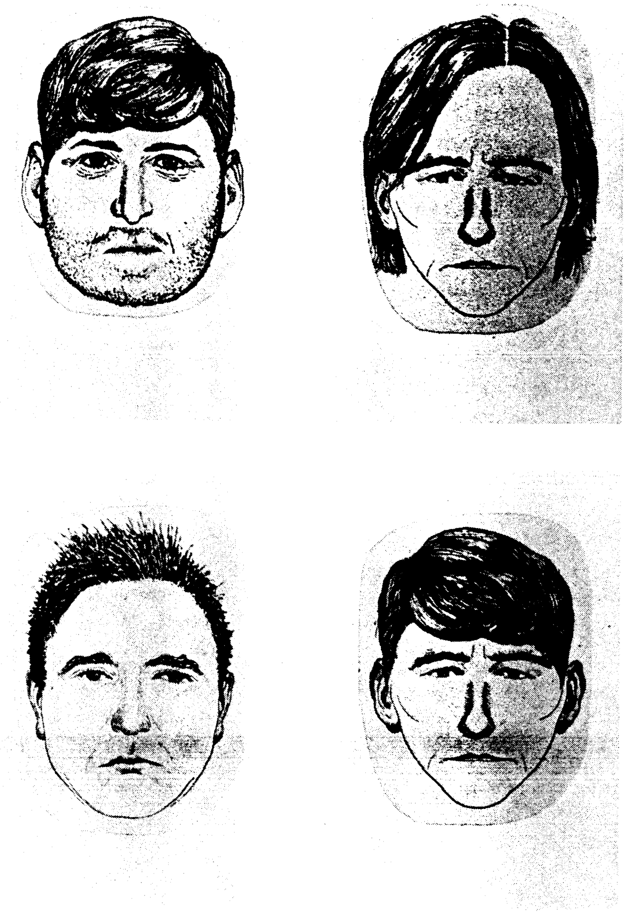

FIG. 1. Examples of experimental materials constructed with aid of the IDENTI-KIT.

ciated with a particular profession. Figure 2 shows a schematic network representation for part of the information. Each face is represented by a central node connected to an eye, mouth, hair, and chin feature as well as a profession. We refer to different exemplars of a particular attribute with numbers, e.g., chin 7 is different from chin 4. Figure 2 illustrates that some features like eyes 3, mouth 6, and chin 2 occur in a large number of faces, others like hair 3 occur in three faces, and others like chin 7 and eyes 1 in a single face.

These experiments were performed as a test of whether such an associative network representation, which has proven useful for verbal material,

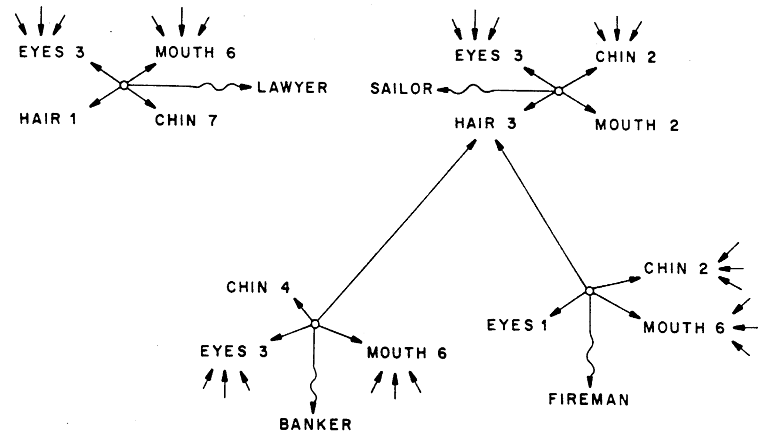

FIG. 2. A schematic representation of what a network encoding might be for the faces material used in the experiments.

would also prove useful for pictorial material. This would be evidence that both kinds of information have the same representation in memory and hence evidence for the propositional theory. Consider how ACT would recognize a face on the basis of such a representation. For instance, suppose it were presented with a face composed of eyes 3, chin 2, hair 3, and mouth 2. It is assumed that nodes in memory corresponding to these features would be activated by perceptual processes. The subject would be able to recognize the face when activation spread from these features to activate the encoding of the face in memory. The speed of this activation would depend on the fan out of the features or, equivalently, the number of faces associated with each feature. Thus, the ACT theory predicts subjects should be slower to recognize (or reject) faces composed of features which occurred in multiple study faces.

In addition to the question of whether pictorial and verbal material have the same representation, this experiment is of interest simply because it addresses the question of whether interfering effects can be obtained for facial material. Some research might be interpreted as indicating that faces are stored as gestalt units and would not be subject to interfering effects defined on subparts. For instance, Smith and Nielson (1970) found no effect of number of relevant features on same recognition times for schematic faces. There is evidence that configural information is important in face recognition and that reliance on such configural information increases with age (Carey & Diamond, 1977; Harmon, 1973; Rock, 1974). On the other hand, Bradshaw & Wallace (1971), using IDENTI-KIT material, found that time to reject a foil varies with the number of different features. They take this as evidence against a gestalt representation of the face. These experiments which have addressed the gestalt issue have all involved an immediate memory paradigm. It is possible that faces could have an immediate gestalt representation but a nongestalt representation in long-term memory. To display a fan effect would be evidence against a gestalt long-term memory representation for pictorial information. It would be particularly impressive to do this with faces, which seem to have such strong configural properties.

EXPERIMENT 1

Experiment 1 was performed to test whether the fan effect could be obtained with pictorial material. A second order purpose was to see whether the size of the fan effect was comparable for verbal and pictorial material. To assess this, a control set of material was constructed involving four adjectives describing a person. In all ways the materials composed from the four adjectives were treated like the materials composed from the four face features.

It would be useful to spell out the logic of this experiment: If we find that the same laws of interference apply to pictorial material as to verbal material, this will be taken as evidence for an abstract propositional representation in long-term memory. Of course, a dual code model could also predict these phenomena. One would simply postulate identical laws of interference for the imaginal code as for the verbal code. However, it would be clearly a significant step away from parsimony to do so. This point has been made in general terms in Anderson (1978). In that paper it was shown formally that it is not possible to discriminate between representational contrasts, such as propositional versus dual code, solely in terms of their ability to account for data. There is a way of explaining any empirical phenomenon within either representational system. However, one explanation may prove more parsimonious than the other. If similar interfering effects can be found for pictorial and verbal stimuli, this would be one piece of evidence that the propositional theory is more parsimonious.

Method

Materials. The face features were created through use of the IDENTI-KIT.[1] The four features chosen were hair style, chin (combination of beard and chin elements in the IDENTI-KIT), mouth (mouth and nose), and eyes (eyes and eyebrows). Features were chosen to be maximally different to aid discrimination. There were seven values chosen for

[1] The IDENTI-KIT Company (1230 East Warner Avenue, Santa Ana, Ca. 92705) produces a kit used by law enforcement agencies to translate eyewitness descriptions into composite portraits of suspects.

each feature attribute. One value was designated the null value which would appear with nine faces, one or two values were designated the 3-fan valves, and four or five were designated 1-fan. The features were combined into faces so that each face consisted of two null values and two non-null values. Altogether there were eighteen faces: in six faces both non-null values were 1-fan. In six faces one feature was 1-fan and the other 3-fan. In six faces both features were 3-fan. The assignment of feature values to conditions was random and was rerandomized for each pair of subjects. Adjective quartets were constructed from 28 adjectives in the same way faces were constructed from the 28 feature values.

In addition to the 18 faces that subjects studied, 18 foil faces were constructed from the same features. Two of the features in each foil face were the null values and two were non-null values. The foil faces were so constructed that six had two non-null values that were both 1-fan, six had a 1-fan and a 3-fan value, and six had both values be 3-fan. Eighteen foil adjective descriptions were constructed in the same manner.

The faces were like those illustrated in Fig. 1 and were presented by slides. The four adjectives were presented on slides vertically arranged. Also names of professions were typed onto slides. The subjects' task would be to learn to pair the profession names with the faces or the adjective descriptions.

Subjects. Thirty-two subjects were recruited from the general undergraduate population at the University of Michigan. Sixteen subjects were run in the faces condition and sixteen subjects were run in the adjective condition. Subjects were paid $2.00 an hour for their participation in the experiment. The subjects in the faces condition spent about 2 hr in the experiment; whereas, subjects in the adjective condition required about 3 hr. The difference in time reflects the fact that subjects were much slower learning the material in the adjective condition.

Procedure. The procedures were in all respects identical for the adjective and face material. All material was presented to subjects by two slide projectors controlled by the IBM 1800 at the Human Performance Center. The IBM 1800 was programmed to present the material in a different random order to each subject. The experiment was run in three phases. In the first phase, subjects were required to learn the profession associated with each of the faces or adjective descriptions. The face or description was presented via one slide projector and the subject tried to recall the name. If he recalled or gave up, the experimenter pressed a button which caused the second slide projector to present the name. Subjects continued studying and recalling the professions until they were able to recall all 18 to the faces or the adjective descriptions. The purpose of this phase of the experiment was to have the subject commit the faces or adjectives to memory.

We were afraid that subjects might have associated the profession to only a subset of the face features or a subset of the adjectives. Therefore, in the next phase the real faces were mixed in with foils. Subjects had to recall the names of the real faces and respond to the foils and the answer "nobody." The foils were constructed to contain various subsets of features of the study sentences. True to our fears subjects usually committed many errors when they made this transition. Subjects were kept in this phase until they could recall the names of the 18 targets and could correctly reject the 18 foils.

The third phase of the experiment involved collecting subject reaction times to recognize targets and reject foils. Subjects were presented with slides containing the faces or adjectives with no mention of professional label. They had to press one of two buttons to indicate whether this was one of the people they had studied and learned a label to, or whether this was a foil combination. Subjects were allowed to choose which hand (right or left) to assign to response (yes or no). Ten blocks of trials were constructed, each consisting of the 18 targets and the 18 foils. Reaction time was measured from the presentation of the slide on the screen to the subject's pressing of one of two buttons, indicating acceptance or rejection of the slide. After the subject pressed the button the other projector presented the profession or the word "nobody" providing the subject with feedback as to the accuracy of his response.

Results

Mean reaction times were computed for each subject × block × response type (target vs foil) × fan combination. Mean reaction times (collapsing over block) are presented in Table 1. Trials on which subjects made errors were excluded in these calculations. Because of the large differences in mean reaction times, data from the face material and the adjective material were subject to separate analyses of variance. These reaction times were subject to a 4-way analysis of variance with subjects as a random effect. Note that since materials varied with subjects, any statistical tests involving subjects as the error term will also involve error variance due to materials. This is one of the ways recommended by Clark (1973) for dealing with the materials-as-fixed-effects problem.

There were highly significant effects ($p < .01$) of all variables for both faces and adjectives. The factor of blocks did not interact with the other factors. Therefore, we have collapsed over blocks in Table 1. Also note from Table 1 that the error rates correlate strongly with reaction times within the face data ($r = .87$) and within the adjective data ($r = .98$) although not that well between the two sets of data. Because of the correlation of RT with error rates we will just focus on the mean reaction times reported in Table 1. The standard error of the means in Table 1 is 17.2 msec for the faces data and 48.8 msec for the adjective data.

TABLE 1

REACTION TIMES (IN MSEC) AND ERROR RATES FROM EXPERIMENT 1

	Faces	
	Yes	No
1–1	877 (.082)	972 (.088)
1–3	1002 (.101)	1023 (.085)
3–3	1104 (.130)	1122 (.146)
	Adjectives	
	Yes	No
1–1	1991 (.085)	2029 (.085)
1–3	2186 (.100)	2267 (.101)
3–3	2824 (.229)	3001 (.226)

There are a number of effects that are quite apparent in the data of Table 1. First, reaction time increases with fan both for faces and for adjectives. All 32 subjects showed a fan effect. Second, subjects took significantly longer to reject foils than to accept targets. Third, subjects took over twice as long to respond to the adjectives as the faces. This is another manifestation of the greater general difficulty that subjects had with the adjective material. They also had taken much longer to learn the material. The effects of response and fan are larger for the adjectives than faces. Subjects are 45 msec slower to foils than targets for faces but 99 msec slower to adjective foils than adjective targets. Comparing the extremes of the fan effect, 1–1 vs 3–3, the size of the effect is 189 msec for faces versus 903 msec for adjectives. As this contrast is particularly interesting we performed a test of its statistical difference. It is quite significant [$t(300) = 14.38; p < .001$].

Discussion

The question this experiment was designed to investigate was whether a fan effect could be obtained for pictorial material. We have clearly established such an effect. This result provides evidence for an abstract propositional representation. However, there are a couple other results in this experiment which suggest some important after-the-fact questions. One of the striking effects is the much greater difficulty subjects experienced with the verbal than the pictorial material. This was manifested both in the slower learning of the material and the two to three times slower recognition times in Table 1. We think this can be interpreted in the same way the earlier results on superior memory for pictures can be interpreted. That is, these differences can be attributed to the subject's greater facility at encoding facial than verbal material. The mean recognition time for the faces was just over a second. This is about the time it would take the subjects to read four unrelated words. Thus, there is no way subjects could read, encode, recognize, and respond to the four adjectives as fast as they did to the faces. It is uncertain whether the encoding advantage reflects an inherent superiority for faces or whether it reflects greater practice at encoding faces. Another difference between faces and the adjective quartets may be amount of preexperimental interference. The features making up the faces were relatively unique, but the words were familiar and would have many preexperimental associations.

The overall differences in learning rate and reaction time are not a problem for a propositional theory which concerns itself with long-term retention and not encoding. On the other hand, it is potentially a problem for propositional theories like ACT that there was such a large difference in the size of the fan effect for faces versus for adjectives. There are a number of possible explanations for this effect. One is to relate the difference in fan effect to the overall difference in mean reaction time. One might propose that the sources of increased time interacted with the fan factor producing particularly large effects for the verbal material. It should be pointed out that in an analysis of fan effects from many experiments with verbal materials, Anderson (1976) noted a strong relationship between mean reaction time and the size of a fan effect. In light of these past results, this explanation seems plausible.

A second explanation lies in the possibility that subjects did not always encode the face in terms of the four features that we manipulated. Suppose, for instance, that subjects treated mouth and chin as a single feature. Then a pairing a 3-fan mouth and 3-fan chin would be functionally for the subject a 1-fan feature that only occurred in a single face.[2] This also is an explanation we find quite plausible. Experiment 2 will serve, in part, as a test of this explanation.

Either of these two explanations leaves the basic conclusion of a network representation for pictorial material unperturbed. A third explanation, however, assumes that subjects typically represent pictorial material in an imagery medium which can be rapidly accessed and which does not show a fan effect. The reason a small fan effect was obtained for pictorial material is because occasionally subjects described the pictures to themselves verbally. These verbal encodings were represented in network form. According to this hypothesis, the fan effect for faces was produced by a portion of trials on which subjects verbally encoded the material. All sixteen subjects displayed a fan effect with the faces material. So, if the

[2] We would like to thank Lynne Reder for suggesting this interpretation of the data.

fan effect reflects verbal encoding only occurring sometimes, all subjects must have engaged to some degree in verbal encodings.

This hypothesis can be tested by one analysis of the data. This hypothesis predicts that subjects would display long reaction times when they had to retrieve verbal encodings of the pictures. This means that short reaction times should reflect pure imagery processing and should not display a fan effect. Therefore, we looked at the shortest reaction times for each subject in each of the six conditions. The means (computed over subjects and yes–no) of these minimum reaction times were 547 msec in the 1–1 condition and 573 msec in the 3–3 condition. Thus, the fan effect is considerably reduced. However, this reduction would be expected under most assumptions about the reaction time distributions. The remaining difference is of marginal statistical significance [$t(15) = 1.58\, p < .10$, one tailed]. The second experiment will present data which will be difficult (but not impossible) to explain according to the verbal encoding hypothesis.

EXPERIMENT 2

One purpose of the second experiment is to evaluate the second explanation offered of the reduced fan effect for pictorial materials. This explanation is that the features in terms of which subjects encoded the faces were not always the same as the four features we manipulated. In this experiment we will take a number of steps to maximize the probability that one of the features subjects use in encoding faces is hair style. Hair style will be the only feature involved in the fan manipulation.

A second purpose of this experiment is to determine if the verbal information will interfere with pictorial material and vice versa. The first experiment found interference among pictorial materials and among verbal materials. If this experiment finds interference across these information sources this will be evidence that the two kinds of information are stored together in memory. However, if no cross interference is found, this will be evidence that, while the two types of information may obey the same storage laws, they are stored in separate locations. Either way, there would be further evidence on the central question of this paper—whether there are any functional differences in the long-term storage of verbal versus pictorial material. We were able to investigate this issue of cross interference by having subjects learn both verbal and pictorial information about the hair styles.

Method

The design of this experiment and the materials are somewhat more complicated than Experiment 1. To facilitate comprehension of the detailed experimental specifications to follow, it would be useful to provide a

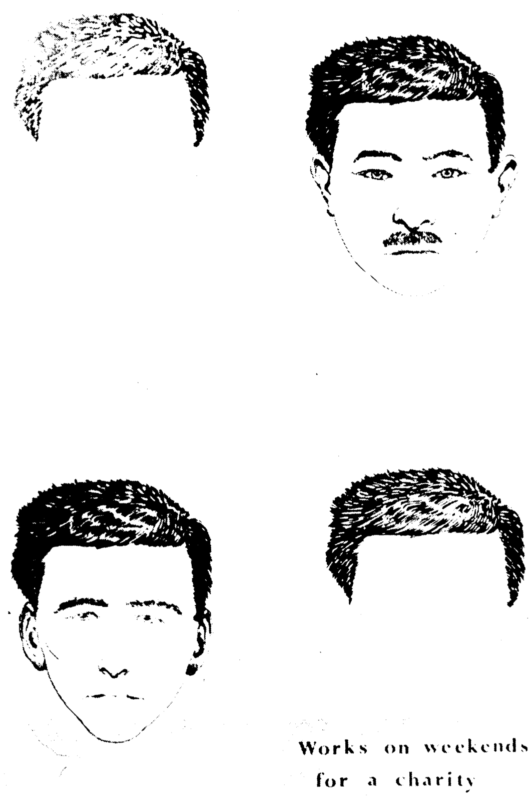

FIG. 3. Examples of the materials presented to subjects in Experiment 2.

brief overview of the experiment, highlighting significant features of the subject's history through the experiment. A subject started out the experiment with a familiarization phase in which he simply saw hair styles as in the upper left quadrant of Fig. 3. After this he studied faces and combinations of hair style plus predicate. A single subject might study the three illustrated in the other quadrants of Fig. 3. The subject learned to associate professions to each of these. So he might learn that the face in the upper right quadrant was *the lawyer*, the face in the lower left, *the doctor*, and, as for the lower right quadrant, the hair style belonged to and the predicate was true of *the banker*. After learning these professional associations, subjects then proceeded to a speeded recognition phase of the experiment. Here they say these faces or hair style plus predicate combinations mixed

in with distractors that recombined hair styles, face features, and predicates. We will principally be interested in the speed with which they made these judgments.

Material. We selected 18 hair styles from the IDENTI-KIT, six partial faces (minus hair style) constructed from IDENTI-KIT material, and six predicates. To associate a partial face with a hair style, the face and hair style were just presented combined to compose a complete face. To associate a predicate with a hair style, the hair style was presented disembodied from any face with the predicate written below it. The six predicates used were:

beats his children every night
speaks well in several languages
works on weekends for charity
does crossword puzzles on Sunday
never pads his expense account
plants a large garden each spring

In all ways the materials combining predicate and hair style was treated identically to the material combining partial face and hair style.

Table 2 provides a summary of how these materials were used to construct the stimuli for the experiment. Two different hair styles were assigned to each of eight conditions. These eight conditions were created by factorially combining the conditions of 0, 1, or 2 partial faces associated with the hair style and the conditions of 0, 1, or 2 predicates associated with the same hair style with the exception of the 0–0 condition. These numbers of hair styles, faces, and predicates are multiplied by two in

TABLE 2

SUMMARY OF ASSIGNMENT OF MATERIAL TO CONDITIONS FOR A SUBJECT IN EXPERIMENT 2

		Number of partial faces learned to hair styles		
		0	1	2
Number of predicates learned to hair styles	0		2 hair styles 2 partial faces 0 predicates	2 hair styles 4 partial faces 0 predicates
	1	2 hair styles 0 partial faces 2 predicates	2 hair styles 2 partial faces 2 predicates	2 hair styles 4 partial faces 2 predicates
	2	2 hair styles 0 partial faces 4 predicates	2 hair styles 2 partial faces 4 predicates	2 hair styles 4 partial faces 4 predicates

Table 2 because there were two instances of each condition. The 0–0 condition could not be realized because it involves presenting nothing to the subject for study. Therefore, there is nothing to test. The hair style illustrated in Fig. 3 represents an instance of a 2–1 condition. That is, the subject learns to associate two partial faces and one predicate to the hair style. Altogether, the subjects learned four hair styles in the 0-face (there is no 0–0 condition), six in the 1-face, and six in the 2-face. Thus, they learned $(4 \times 0) + (6 \times 1) + (6 \times 2) = 18$ hair style-face combinations. To achieve this, each of the six partial faces were associated to three hair styles. Similarly, they learned 18 hair style–predicate combinations and each predicate was associated to three hair styles.

Corresponding to the target items foil items were created by re-pairing hair styles with partial faces and predicates. The same number of foil partial faces and predicates were associated with a hair style as the number of target faces and predicates. So, if a subject learned to associate one face and two predicates to a hair style, they would encounter three foil items using the hair style—one involving a face and two involving predicates. A different combination of hair styles and faces and of hair styles and predicates was used for each pair of subjects (as there were only 15 subjects a new combination was used with the fifteenth subject). The hair style–face or hair style–predicate combinations were photographed and presented by slide. Slides of the hair style alone were also constructed for the familiarization phase of the experiment.

Subjects. Fifteen subjects were recruited from the undergraduate population at the University of Michigan. They were paid $2.00 an hour for participation in an experiment that lasted 3 hr.

Procedure. The experiment began with a familiarization phase in which an attempt was made to make the hair styles sufficiently familiar that the subject would be likely to use these as one feature in encoding the faces. Subjects were shown the 18 hair styles alone and asked to rate how much they liked each. Subjects assigned numbers on a scale from 1–7 where 1 meant strongly dislike and 7 meant strongly like. This rating phase continued until subjects were able to assign the same numbers to the hair styles twice in a row. Subjects were only told after their first rating that they would have to achieve a consistency in ratings before they would leave this familiarization phase.

After this familiarization, the experiment proceeded much as it had for Experiment 1. Subjects first learned to associate professions to the 18 face–hair style combinations and to the 18 predicate–hair style combinations. After doing this to the criterion of one perfect recall, they entered a phase where the 36 targets were mixed in with 36 distractors. Subjects had to continue to recall the professions to the targets and identify the distractors as "nobody." They continued in this phase until they had correctly responded to all 72 stimuli in a single sequence of trials. Then they went into a phase of the experiment where they had to make speeded recognition judgments of the 72 items (36 targets and 36 foils). Reaction times were again collected for these judgments. Subjects went through 10 blocks of 72 such trials. As in Experiment 1 the presentation of material was controlled by computer and order of presentation was randomized for each subject.

Results

The data can be classified according to the factors of *dimension*: whether the subject was judging the hair style with a face or with a predicate; *response*: whether it was a target, requiring a yes response, or distractor, requiring a no response; *within fan*: the number of items paired with the hair style within the tested dimension (this could be 1 or 2); and *cross fan*: the number of items involving the hair style on the other dimension (this could be 0, 1, or 2). The mean reaction times were computed for each subject, excluding errors. Table 3 presents the reaction time data and error rates for the various conditions defined by the four factors. There is only a modest correlation ($r = .52$) between reaction time and error rate. However, all of the main effects defined on reaction times are mirrored in error rates. The reaction time data were subjected to an analysis of variance using only subjects' correct times. The factors in this analysis of variance were dimension, response, within fan, cross fan, and subjects. The standard deviation of the reaction times in Table 3 is 50.1 msec. This is based on the overall condition × subjects interaction and has 345 degrees of freedom. This is the term that will be used in the statistical tests unless otherwise stated. The standard error of a condition in this experiment is larger than Experiment 1 because there were fewer observations per condition. Since there are more conditions, many of the contrasts (collapsing across conditions) are as reliable as in Experiment 1. All main effects in this analysis were significant, the least significant being cross fan [$F(2,345) = 5.37$; $p < .01$]. Subjects take longer to reject distractors, longer to reject or accept predicates, longer the more facts there are on the within dimension, and longer the more facts there are on the cross dimension. There were three significant interactions—within × dimension [$F(1,345) = 4.23$; $p < .05$], within × cross [$F(2,345) = 4.39$; $p < .05$], and within × cross × dimension [$F(2,345) = 4.43$; $p < .05$]. We will discuss these interactions after considering a further analysis of the data.

The principal focus of this experiment is on the effects of within vs cross fan and how this might interact with whether subjects were judging predicates or faces. Figure 4 is an attempt to highlight the data relevant to these issues. Here we have the reaction times plotted separately for predicate and face judgments, as a function of the number of facts on the within and cross dimensions. All data are collapsed over response (yes vs no). The cross plot is also collapsed over the within factor and the

TABLE 3

REACTION TIMES (IN MSEC) AND ERROR RATES FROM EXPERIMENT 2[a]

		Recognition of faces							
		Yes Number of predicates				No Number of predicates			
		0	1	2	Mean	0	1	2	Mean
Number of faces	1	1158 (.023)	1099 (.033)	1135 (.027)	1130 (.028)	1152 (.013)	1239 (.053)	1330 (.087)	1241 (.051)
	2	1221 (.057)	1308 (.098)	1270 (.063)	1266 (.073)	1386 (.057)	1432 (.055)	1379 (.050)	1399 (.054)
Mean		1190 (.035)	1203 (.066)	1202 (.045)	1199 (.050)	1270 (.035)	1336 (.054)	1355 (.069)	1320 (.053)
		Recognition of predicates							
		Yes Number of faces				No Number of faces			
		0	1	2	Mean	0	1	2	Mean
Number of predicates	1	1251 (.020)	1443 (.043)	1399 (.043)	1364 (.035)	1368 (.023)	1627 (.117)	1638 (.063)	1545 (.068)
	2	1563 (.115)	1621 (.127)	1560 (.113)	1582 (.118)	1838 (.067)	1742 (.067)	1788 (.067)	1790 (.067)
Mean		1407 (.068)	1532 (.085)	1480 (.078)	1473 (.077)	1603 (.045)	1685 (.092)	1713 (.065)	1667 (.067)

[a] The numbers reported in this table are not identical to those in earlier drafts of this paper. We found a systematic error in the analysis program which had scored the last forty trials as errors.

within plot collapsed over the cross. Here it is clear that the effect of the within dimension is much larger than the cross. The effect of number of within facts for predicate judgment (231 msec) is significantly larger than the same effect for face judgment (147 msec) as reflected by the within × dimension interaction. However, the attempt in this experiment to assure that subjects encode the faces in terms of the hair feature has greatly reduced the difference in the size of fan effect for verbal material (predicates) relative to that for pictorial material (faces). The difference in the size of the fan effects in Experiment 1 had been in the order of 5 to 1. The remaining difference is only on the order of 1.5 to 1. The reduction in the difference supports the hypothesis that subjects in Experiment 1 produced a smaller fan effect for faces because of our failure to control the

features by which they encoded the faces. The remaining difference between the two fan effects may reflect occasional configural encodings that did not use hair style as a feature. This experiment is also consistent with an attempt to relate differences in the fan effect of differences in overall recognition time. Note that the recognition times are much closer for pictorial and verbal material in this experiment than Experiment 1.

One further feature to note about the data in Fig. 4 is that the increase with fan on the cross dimension is occurring from 0 to 1 fan and not from 1 to 2 fan. Averaging the two cross curves together there is a 72 msec increase in going from 0 to 1 propositions and a −2 msec change from 1 to 2 propositions. A contrast comparing the size of these two differences is marginally significant [$t(345) = 1.71; p < .05$; one tailed]. The model we will present predicts that the increase from 0 to 1 will be greater than the increase from 1 to 2. Therefore, we will define the fan effect for the cross dimension as the average of 1 and 2 fan minus the 0 fan.

Now that we have the cross dimension fan effect defined, we can consider the two interactions involving this factor. The interaction between within fan and cross fan reflected the fact that the cross fan effect was much larger when the within fan was 1 than when it was 2. When the within fan was 1 the size of the cross fan was 131 msec; when the within fan was 2 the cross fan effect was 11 msec. The model we will present expects this cross-by-within interaction. This interaction is modulated by the existence of a three-way interaction involving the factors of dimension, within fan, and cross fan. The message of this three-way interaction is that the above two-way interaction is more true for the predicate material than the face material. For the predicate material, the cross fan effect was 218 msec

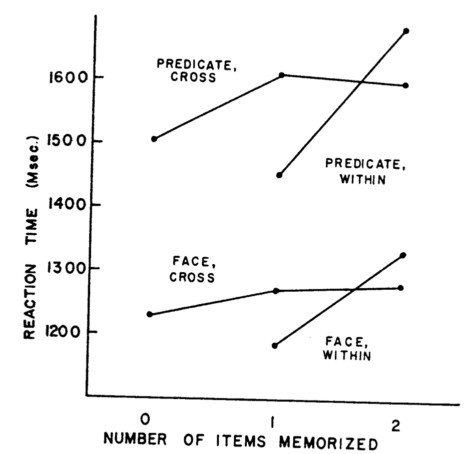

FIG. 4. The effects of within and cross fan for recognition judgments of faces and hair style–predicate combinations.

when the within fan was 1 and −23 msec when the within fan was 2. For the face material, the cross fan effect was 46 msec when the within fan was 1 and 44 msec when the within fan was 2. This three-way interaction is not predicted by the model we will present. We suspect this might be one of the spurious .05 effects that one must be wary of in an experiment that has 15 main effects and interactions.

Figure 4 collapses over the response, yes vs no. There were no interactions with this factor, but we thought it worthwhile to consider how reliable the cross fan effect is over this factor. A test for the cross fan effect is significant for both responses [$t(345) = 1.81$ for positives; $p < .05$; one tailed, and $t(345) = 2.79$ for negatives; $p < .005$].

Given the rather small size of the cross fan effect, one might be suspicious of these tests which rely on an error variance pooled over all conditions and on averaging a number of conditions together. These tests have a possible danger because of nonindependence of error variance across the averaged conditions in a within-subject design. More seriously, there is a potential problem in that variances might not be identical across conditions. Therefore, to deal with these potential problems and to get a further measure of the reliability of the cross fan effect, we calculated the cross fan effects for each subject for recognition of faces and separately for recognition of predicates. We choose to break the cross fan down in this way because of the importance of establishing the same effects for face judgments as predicate judgments. Even though the between dimension × cross fan interaction was not significant [$F(2,345) = .83$], the between fan effect did appear weaker for faces than predicates. Of course, any such division of the data will result in a reduction of the power of the statistical tests.

Twelve of the fifteen subjects showed a cross fan effect for predicates. That is, they were interfered with by face fan in predicate judgments as indicated by a positive cross fan contrast. A t test for whether the mean of 15 subject contrasts was different from zero was significant [$t(14) = 3.19$; $p < .005$]. Ten of the 15 subjects displayed interference of predicate fan on face judgments. The test for their contrasts did not achieve conventional standards of significance [$t(14) = 1.31$; $p < .15$]. Given the uncertainty of this contrast defined on reaction time, we also looked at the cross fan contrast defined on percent correct. Twelve out of 15 subjects displayed the predicted effect of face fan on predicate judgments with one tie. A test of whether their contrasts were different from zero was highly significant [$t(14) = 3.37$; $p < .005$]. Finally, for percentage correct, 11 of the 15 subjects showed the effect of predicate fan on face fan. The test for the significance of these contrasts was positive [$t(14) = 2.60$; $p < .025$]. So there clearly is an interfering effect of face fan on predicate judgment. Also combining the results of reaction time and percentage correct, there seems pretty clearly to be an effect of predicate fan on face judgment.

Discussion

The data of Experiment 2 are consistent with the hypothesis that verbal and pictorial information are stored together in memory. There is an effect of fan involving both pictorial material and verbal material. The size of these effects are fairly similar in contrast to Experiment 1. This can be attributed to better control over how the subjects encode the faces. Also, there are interfering effects of pictorial material on verbal and of verbal material on pictorial. These results are all predicted by ACT which is one embodiment of a theory that makes no distinction in its representation of verbal vs pictorial information.

There were also a number of interesting results about which there were no clear predictions in the ACT theory, one way or the other. One was the fact that the size of the fan effect on the cross dimension was much smaller than the size on the within dimension. The second nonpredicted result was the greater fan effect for the cross dimension with only one within fact than with two. The third nonpredicted result was that the cross fan effect was totally between zero propositions and one. There was no difference between one and two. These results can be explained by an extension of a proposal put forth in Anderson (1976) to explain the partial separation of experimental knowledge about a concept from real world knowledge about that concept. The suggestion made there was that subjects set up a subnode of the concept to contain experimental information. Similarly, we propose for this situation that the subjects set up two subnodes for the hair style concept, one to represent the hair style when it is part of a face and one to represent the hair style when it is paired with a predicate.

Figure 5 illustrates, quite schematically, the memory representation that we have in mind. There is a node in memory which represents the hair style. Connected to it are two subnodes for the hair styles in the two contexts. It is these subnodes that are involved in representing the faces and the predicates. The subject should be able to tell, after a very cursory check of the probe, whether he should use the face subnode or the predicate subnode. He would retrieve the appropriate subnode by a process of spreading activation from the main node. Once he has retrieved

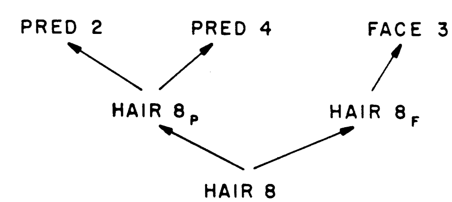

FIG. 5. A schematic network structure illustrating how a node for a hair style can be connected to two subnodes.

the subnode, he could focus on it and spread activation from that node to achieve the encoding of the face or predicate.

This representation, together with the assumptions of the ACT activation theory, predict almost all the complexities of the fan effects in the experiment. First, there should be an effect of number of within items. This variable has its effect on the speed at which activation spreads from the subnode to activate the correct face encoding or predicate encoding. Second, this model predicts the effect of fan on the cross dimension. The difference between no cross items and any cross items is very important because just one cross item will cause an additional path to be built from the concept to a subnode. This interfering path will take activation away from the correct path and so slow retrieval of the correct subconcept. The difference between one and two cross items should be much less important because in both cases there will be the interfering path to a subnode. Subjects might be somewhat slower in the two cross fan case because the interfering path will be somewhat stronger due to more frequent repetition. As predicted, the observed effect was larger for an increase from zero to one proposition than from one to two propositions on the cross dimension. This model also predicts a lesser effect of cross than within fan. Cross fan only affects activation of the link to the subconcept; whereas within fan affects activation of the more complex structure encoding face or predicate (This complexity is not illustrated in the schematic structure of Fig. 5). In the ACT theory the effect of fan is multiplied in reaction time by the complexity or size of the network structure to be activated.

This theory also predicts the interaction between within and cross fan. When the fan is high on the within dimension, the link leading to the correct subnode will be stronger because of greater frequency of exposure. Therefore, the interference will be less from the link to the subnode for the other dimension.

We feel that this subnode notion is an important addition to the ACT activation theory. It can be used to explain not just the data from this experiment, but also why people can retrieve relatively rapidly real world information they possess about concepts, even when they know many facts about the concepts. We would propose that information about a familiar concept like Kissinger is organized, perhaps hierarchically, into a subnode structure and that the subject has only to search facts attached to one subnode. Note, however, that there is an important prerequisite for a subnode structure to be effective. It must be possible for the subject to have some ready basis for selecting the correct subnode from which to focus search. In Experiment 2, this basis was provided: The probe contained either a face or a predicate. Similarly, the subnodes for familiar concepts must be readily selected by the context within which one is asked about a fact. For instance, we might suppose that facts about Kissinger are divided into subconcepts concerned with physical appearance, role in Middle East, personality, his role in Africa, his role as fashioner of detente, etc. The subnode appropriate to answering a particular question could be selected on the basis of certain key words in the question.

One might wonder how the subnode structure in Fig. 5 differs from assuming that there are two kinds of representations—one verbal and one pictorial. The important differences are: (1) Similar representational formalisms could be used to minimize interference within purely verbal or purely visual material. That is, the subnode hypothesis is more general than just a means of separating visual and pictorial information. (2) It predicts similar laws of interference for pictorial and visual material (3) It predicts interference, if reduced, between pictorial and verbal material.

GENERAL DISCUSSION

One is naturally tempted to make comparisons between the two experiments in terms of the size of the within-fan effects for verbal and pictorial information. Table 4 presents the data relevant to making that comparison. Note that for both faces and verbal materials there is a reduction in the size of the fan effect from Experiment 1 to Experiment 2. This reduction is predicted by the ACT theory. Both the faces and the verbal material consisted of four elements in Experiment 1. In the high fan case the fan of these elements were 9, 9, 3, and 3 and in the low fan case 9, 9, 1, and 1. ACT basically predicts that the rate of activation will be a function of the inverse of the sum of the inverse of the fans (Anderson, 1976; pp. 265; Eq. 5). So:

$$\text{Hi-fan-rate-Expt-1} = \alpha/(1/9 + 1/9 + 1/3 + 1/3) = 1.125\alpha \quad (1)$$

$$\text{Lo-fan-rate-Expt-1} = \alpha/(1/9 + 1/9 + 1/1 + 1/1) = .450\alpha \quad (2)$$

where α reflects a time scale factor. In Experiment 2 we can also consider the faces as consisting of four features, only one of which was involved in the fan manipulation. The predicates contained basically three concepts each. So the hair style plus predicate consists of four features. The fans of three features were 3 and the fan of the hair style was 1 or 2 depending on whether it was in the hi-fan and lo-fan case. So our predicated rates are:

$$\text{Hi-fan-rate-Expt-2} = \alpha/(1/3 + 1/3 + 1/3 + 1/2) = .667\alpha \quad (3)$$

$$\text{Lo-fan-rate-Expt-2} = \alpha/(1/3 + 1/3 + 1/3 + 1/1) = .500\alpha \quad (4)$$

Thus, the difference in activation rates in Experiment 1 is $.675\alpha$ whereas it is $.167\alpha$ in Experiment 2. This predicts that the ratio of fan effects between the experiments should be 4.04–1. The observed ratio for the verbal material was 3.91–1 which is quite close. The ratio for the face material, 1.29–1, is way off. However, the fan effect in Experiment 1 was hypothesized to have been reduced because subjects were encoding combinations

TABLE 4

Comparison of Within-Fan Effects for Verbal and Face Material in Experiments 1 and 2—Reaction Times and Error Rates (in Parentheses)

	Faces		Verbal	
	Expt 1	Expt 2	Expt 1	Expt 2
Lo-Fan	924 (.085)	1186 (.040)	2010 (.085)	1455 (.052)
Hi-Fan	1113 (.138)	1333 (.064)	2913 (.228)	1686 (.093)

of our independent cues configurally. That is, it is a case of E and S disagreeing about what the features were.

Research on Selective Interference

There might appear to be a contradiction between these experiments and the many experiments that have been reported on selective interference (e.g., Segal & Fuzella, 1970; Brooks, 1968; Salthouse, 1975). For instance, Segal and Fuzella (1970) showed that subjects were poorer at detecting visual stimuli when imaging pictures and poorer at detecting auditory stimuli when imaging sounds. It is the case that almost all such selective interference tasks have involved perceptual or immediate memory tasks. Therefore, it is conceivable that the mental representations used in these tasks are quite different than the ones subjects use to retain the information in our experiments which extended over hours. That is, there may be a difference between immediate and long-term representation of information.

However, the selective interference results are explainable within a propositional framework as has been noted by some imagery theorists (e.g., Kosslyn & Pomerantz, 1977). That is, propositional representations that come from the same modality would tend to overlap more in terms of their component features than would representations from different modalities. For instance, representations from the same modality might share such relations as *above*, *left-of*, features like *pointed*, *horizontal*, etc. Since the propositions describing the to-be-detected element would be similar to those involving the interfering material, they would be difficult to distinguish from each other, and they would be confused one for the other.

Implications for Theories of Face Memory

While the principal concern of this research is whether pictorial and verbal information are similarly represented in memory, this research is

also relevant to the issue of whether faces have Gestalt representations or representations that are feature composites. Clearly, the composite position is favored by the result of interference in face memory on the basis of feature overlap. However, the outcome is not a total victory for the composite position. If our interpretation is correct of why the fan effect was less for pictorial than verbal material in Experiment 1, then it is clear that there are configural principles at work. Our interpretation of this result was that often subjects did not use our features of hair, eyes, mouth, and chin as the features with which to encode in the face. Rather than using these physically separated elements, they used features that were configural combinations of these elements. So, while a face representation may be a composite of features, these features are often quite configural. This outcome is not inconsistent with the ACT theory which uses nonconfigural representations. ACT only requires that the nodes and their links behave as independent elements. It does not deny that there might be configural perceptual processes which map input into these nodes. So, as long as these configural effects are occurring in perception, they are irrelevant to evaluating a theory of memory representation.

REFERENCES

Anderson, J. R. Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 1974, **5**, 451–474.

Anderson, J. R. *Language, memory, and thought*. Hillsdale, N. J.: Lawrence Erlbaum, 1976.

Anderson, J. R. The status of arguments concerning representations for mental imagery, *Psychological Review*, 1978, **85**, in press.

Anderson, J. R., & Bower, G. H. *Human associative memory*. Washington, D. C.: V. H. Winston, 1973.

Baylor, G. W. A treatise on the mind's eye. Technical report. Institute of Psychology, University of Montreal, Canada, July, 1971.

Bower, G. H., Karlin, M. B., & Dueck, A. Comprehension and memory for pictures. *Memory and Cognition*, 1975, **3**, 216–220.

Bradshaw, J. L., & Wallace, G. Models for the processing and identification of faces. *Perception and Psychophysics*, 1971, **9**, 443–448.

Brooks, L. R. Spatial and verbal components of the act of recall. *Canadian Journal of Psychology*, 1968, **22**, 349–368.

Carey, S., & Diamond, R. From piecemeal to configurational representation of faces. *Science*, 1977, **195**, 312–314.

Clark, H. H. The language-as-fixed-effect fallacy: A critique of language statistics in psychological Research. *Journal of Verbal Learning and Verbal Behavior*, 1973, **12**, 335–359.

Cooper, L. A., & Shepard, R. N. Chronometric studies of the rotation of mental images. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.

Craik, F. I. M., & Lockhart, R. S. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 1972, **11**, 671–684.

Craik, F. I. M., & Tulving, E. Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, 1975, **104**, 268–294.

Goldstein, A. G., & Chance, J. E. Visual recognition memory for complex configurations. *Perception and Psychophysics*, 1971, **9**, 237–241.

Haber, R. N. How we remember what we see. *Scientific American*, 1970, **222,** 104–112.

Harmon, L. D. The recognition of faces. *Scientific American*, 1973, **229,** 70–83.

Kintsch, W. *The representation of meaning in memory*. Hillsdale, N. J.: Lawrence Erlbaum, 1974.

Kosslyn, S. M. Information representation in visual images. *Cognitive Psychology*, 1975, **7,** 341–370.

Kosslyn, S. M., & Pomerantz, J. R. Imagery, propositions, and the form of internal representations. *Cognitive Psychology*, 1977, **8,** 52–76.

Kosslyn, S. M., & Shwartz, S. M. A simulation of visual imagery. *Cognitive Science*, 1977, **1,** 265–295.

Lewis, C. H., & Anderson, J. R. Interference with real world knowledge. *Cognitive Psychology*, 1976, **7,** 311–335.

Mandler, J. M., & Johnson, N. S. Some of the thousand words a picture is worth. *Journal of Experimental Psychology: Human Learning and Memory*, 1976, **2,** 529–540.

Mandler, J. M., & Ritchey, G. H. Long-term memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, 1977, **3,** 386–396.

Moran, T. P. The symbolic imagery hypothesis: A production system model. Unpublished Ph.D. Dissertation, Carnegie–Mellon University, 1973.

Moyer, R. S. Comparing objects in memory: Evidence suggesting an internal psychophysics. *Perception and Psychophysics*, 1973, **13,** 180–184.

Norman, D. A., Rumelhart, D. E. and the LNR Research Group. *Explorations in Cognition*. San Francisco: Freeman, 1975.

Paivio, A. *Imagery and verbal processes*. New York: Holt, Rinehart & Winston, 1971.

Paivio, A. Images, propositions, and knowledge. In J. M. Nicholas (Ed.), *Images, perception, and knowledge*. The Western Ontario Series in the Philosophy of Science. Dordrecht: Reidel, 1976.

Pylyshyn, Z. What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, 1973, **80,** 1–24.

Rock, I. The perception of disoriented faces. *Scientific American*, 1974, **230,** 78–86.

Salthouse, T. A. Simultaneous processing of verbal and spatial information. *Memory and Cognition*, 1975, **3,** 221–225.

Segal, S. J. & Fusella, V. Influence of imaged pictures and sounds on detection of auditory and visual signals. *Journal of Experimental Psychology*, 1970, **83,** 458–464.

Shepard, R. N. Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior*, 1967, **6,** 156–163.

Shepard, R. N., & Metzler, J. Mental rotation of three-dimensional objects. *Science*, 1971, **171,** 701–703.

Smith, E. E., & Nielson, G. D. Representations and retrieval processes in short-term memory: Recognition and recall of faces. *Journal of Experimental Psychology*, 1970, **85,** 397–405.

Standing, L. Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology*, 1973, **25,** 207–222.

Thorndyke, P. W., & Bower, G. H. Storage and retrieval processes in sentence memory. *Cognitive Psychology*, 1974, **5,** 515–543.

Wiseman, S., & Neisser, U. Perceptual organization as a determinant of visual recognition memory. Paper presented at meetings of the Eastern Psychological Association, Spring, 1971.

(Accepted November 16, 1977)