Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval times

STEPHEN MICHAEL KOSSLYN
The Johns Hopkins University, Baltimore, Maryland 21218

This paper provides support for the notion that images, once formed, are a distinct form of internal representation, processed differently than other forms of internal representation. In two experiments, subjects decided as quickly as possible whether or not named animals had given properties. When imagery was not used, people verified properties more quickly in accordance with how strongly associated the property was with the animal in question. When images of the whole animal were consulted, in contrast, subjects were faster in accordance with increasing size of the property, and not with increasing association strength. However, if subjects imaged only the local region where a property ought to be found, and did not consult an image of the whole animal, the size of a property no longer influenced verification time. These results and their implications for the debate over imagery vs. propositional representation were discussed.

It long has been believed that visual images serve as repositories of information (see Paivio, 1971). Recently, this notion has been reexamined by theorists who wish to decompose the construct of imagery into more primitive constructs. This effort has been directed towards demonstrating that a single primitive form of representation, the abstract "proposition," is sufficient to account for all information representation (see Anderson & Bower, 1973; Pylyshyn, 1973). This enterprise would not seem immune nor exempt from the traditional critiques of reductionism (see Putnam, 1973, for a recent treatment). That is, it is entirely possible that emergent properties result when primitives are combined. A television picture, for example, can be described in terms of whether or not each cell in a matrix is "on;" this sort of description, however, fails to convey what picture is depicted by the gestalt, the total configuration, of the cells. Thus, even if images are constructed or derived from some sort of underlying "propositional" (in the sense of a computer scientist) basis, once an image is constructed, it conceivably then may serve as a qualitatively distinct form of internal representation (see Kosslyn, 1975). The two experiments of the present paper attempt to demonstrate that different principles determine ease of retrieving information from images and ease of retrieving information from (presumably) propositional representations.

People require more time to verify that an animal has a given property (e.g., an ear) if they consult a subjectively small image of the animal compared to when they inspect an image that seems subjectively larger (see Kosslyn, 1975). This result was obtained when subjects adjusted size directly, and when size was manipulated indirectly (e.g., by having a target animal, like a rabbit, imaged adjacent to a correctly proportioned image of an elephant vs. a fly). Kosslyn reasoned that recognition of properties of an image was easier given some minimal size for some of the same reasons that properties of larger objects are easier to discern in perception, and that manipulating an image's size had effects because the sizes of its component properties were altered: the parts of smaller images themselves are smaller, and thus more difficult to classify. A clear implication of this supposition is that size of properties per se should influence ease of using an image to retrieve them: larger properties of imaged objects should be "easier to see" (if you will) than smaller properties.

Size of a property is a very different sort of variable from those usually given consideration when studying speed of accessing knowledge. Experiments in the area of "semantic memory" consistently have shown that the more highly related, associated, or produced in response to a noun a predicate is, the more quickly their relation can be verified (in the case of "true" predicates—see Conrad, 1972; Loftus & Suppes, 1972; Schaeffer & Wallace, 1969; Smith, Shoben, & Rips, 1974; Wilkins, 1971). For example, if "stripes" was rated more highly associated with zebra than with "back," the statement "A zebra has stripes" would be verified more quickly than "A zebra has a back."

Table 1
Animals and Properties Used in Experiment I

Animal	Set 1 True Properties		Animal	Set 2 True Properties	
	High Association/ Low Area	Low Association/ High Area		High Association/ Low Area	Low Association/ High Area
Mouse	Whiskers	Back	Shark	Teeth	Tail
Alligator	Teeth	Tail	Monkey	Hands	Shoulders
Bear	Claws	Legs	Rat	Eyes	Fur
Bee	Stinger	Wings	Iguana	Tongue	Legs
Cobra	Fangs	Tail	Panda	Paws	Chest
Owl	Eyes	Chest	Billygoat	Beard	Back
Lion	Teeth	Legs	Raccoon	Mouth	Stomach
Frog	Tongue	Back	Rabbit	Nose	Back
Cat	Claws	Head	Horse	Hooves	Stomach
Duck	Feet	Feathers	Octopus	Suckers	Skin
Rattlesnake	Rattle	Belly	Beaver	Teeth	Back

In an effort to distinguish between retrieval of property information from images and retrieval when imagery is not used, the first experiment to be reported used items such as "cat claws" and "cat head," where the smaller property also was more strongly associated. When imagery is not used, association strength—not size—should facilitate evaluation time. In contrast, when imagery is used, we would not necessarily expect an abstract relationship like association strength (which has been taken to reflect the ordering of propositions in a list, e.g., Anderson & Bower, 1973) to affect time to "see" a property. The previous results of Kosslyn (1975) would lead us to predict, however, that size should influence the ease of detecting a property on an image: smaller properties should take more time to evaluate than larger properties. Since we are interested in how information is represented, we are only concerned with "true" evaluation times, where the subject actually locates a representation of the sought information; "false" responses may tell us more about the processes that operate on representations than the representations themselves, and thus will not be addressed here.

EXPERIMENT I

This experiment consists of three separate, but related, parts. These parts differ primarily in the precise nature of the imagery instructions, and are designed to discover whether different ways of using imagery are distinguishable from each other, as well as whether or not these ways of using imagery are distinguishable from situations where imagery is not used. The first part is an attempt to determine if the size of a property per se is related to the ease with which it can be recognized on an existing image. Subjects in this condition image an animal and later use this image to decide whether or not a presented property is appropriate. The second part is concerned

with the ease with which a property can be generated on an image of a whole animal. Rather than probe an existing image, subjects in this part are asked to construct an image of the probed property—on the whole animal in question—at the time of probe. The third part is an extension of this last task; again, the subjects are asked to generate an image of a property at the time of probe. In this procedure, however, the property is not supposed to occur on an image of the whole animal; instead, the subjects are asked to image only the relevant section of the animal itself.

Method

Three groups of subjects participated in a task that required judging the appropriateness of properties for given animals. All subjects participated in two blocks of trials: Instructions for the first block made no mention of imagery, and merely urged the subject to make his judgments as quickly and accurately as possible. Instructions for the second block, in contrast, requested the subject to evaluate properties by reference to visual images of the animals; the three groups of subjects differed primarily in the particulars of these imagery instructions, as described above.

Materials. To gather materials for the reaction time experiment, a separate group of subjects first was asked to participate in a ratings task. Twenty-two animal names, each paired with five "true" properties, were printed on a page. The subjects rated on a standard 7-point scale "how strongly associated" each property was with the animal. Four separate rating sessions, where animals were paired with new properties, were necessary before enough items were found; each item used in the experiment was rated by 45 Stanford undergraduates. At the conclusion of this rating procedure, each animal finally was paired with two "true" properties, one being more highly associated but smaller in area than the other. In addition, two "false" (uncharacteristic) properties were selected for each animal (e.g., for mouse, "stripes" and "stinger"). The animals and "true" properties are listed in Table 1.

The animals were divided randomly into two sets (each containing 11 animals). The order of the 44 animal-property pairs in each set was randomized (except that no animal or property could be probed twice within three items). Two tape recordings were made, each containing both sets of items (the sets were recorded separately, one following the other). On the first recording, the animal names were followed 5 sec later by one of the four properties used with the animal; 10 sec after this a new animal was presented. On the second tape, the animal and property names

were read one after the other, with no pause between the two; 10 sec after each pair a new pair was presented. Both tapes were constructed such that a clock was started upon presentation of the property name (see Kosslyn, 1975); the clock was stopped when the subject depressed either of two response buttons.

Procedure. The three separate parts of the experiment were conducted consecutively. All three groups received two blocks of trials; the first block never was preceded by imagery instructions, while the second always was. Since interactions, not main effects, are primarily of interest, this seemed an acceptable procedure; it also seemed likely that if the imagery block preceded the non-instruction block, some subjects would adopt an imagery strategy during that block as well. Half of the subjects in each group received one of the two sets of items in the no-imagery block, and half received the other. This assured that items occurred equally often in both imagery and no-imagery conditions.

The first group of subjects participated in a task designed to investigate the effects of using an existent image during property verification. This group, the Pause/Whole group, received the tape where a 5-sec pause separated the presentations of the animal and the property names. The no-imagery instructions directed the subject to "think about the properties of the whole animal, not just part of it" as soon as he heard the animal name. In this first block, the subject never was told explicitly *not* to use imagery, but simply was not instructed to do so. The instructions emphasized responding as quickly as possible, by depressing one of two buttons (hand of response being counterbalanced over subjects), while keeping errors to a minimum. Ten practice trials, using nontest animals, preceded the actual test of items.

The second block of trials (using the set of items not used in the no-imagery block) was preceded by imagery instructions. These instructions emphasized making a visual image of a typical instance of the entire animal upon hearing its name. When the property was presented, the subject was to evaluate its appropriateness by examining his image of the animal. As soon as he found the property, he was to respond "true." If he "looked" at the appropriate area or part of the animal and did not find the property, the subject was to respond "false" by depressing the appropriate button. It was emphasized that all responses in this block should be based on the image, even if the subject intuitively "knew" the answer more quickly than he could "see" the property. Each subject was given extensive training (see Kosslyn, 1974, 1975) and was probed repeatedly during 10 practice trials (which preceded the test trials and used nontest animals) as to his mental processes. After responding to the 44 test items, the subject was asked to estimate the percentage of the time he followed instructions, and was questioned in order to discover whether he had inferred the hypotheses and intentionally responded in a particular way (see Kosslyn, 1975).

The second group of subjects participated in a task designed to investigate the effects of generating and using an image of the whole animal at the time of verification. This group, the No-Pause/Whole group, received the second tape, where properties were recorded immediately after the animal names. The instructions in the no-imagery block were the same as those of the Pause/Whole group, except that the subject was not told to think about the animal upon hearing its name. The imagery instructions now directed the subject to look for the property on an image made only after hearing both members of the animal-property pair. That is, as soon as he heard both stimulus words, the subject was to image the animal and look immediately for the relevant property. It was underscored that the image a subject used should always be one of the entire animal, not simply the isolated part relevant for that particular probe. In all other respects, the instructions and procedure for the No-Pause/Whole group were identical to those given to the Pause/Whole group.

The final group of subjects participated in a task designed to investigate the effects of generating and using an image of only the part of an animal relevant to the probed property during verification. This group, the No-Pause/Part group, also received animal and property names contiguously (on Tape 2). The

instructions and procedure used in the no-imagery condition for this group were absolutely identical to those given the No-Pause/Whole group. The only change in the imagery instructions involved replacing the request to use an image of the *whole* animal during verification. These subjects were asked instead to image only the relevant *part* of the animal. They were told explicitly to try to image not the whole animal, but rather simply the section of the animal's body where the named property ought to be located. In all other respects, the procedure and instructions were the same as those used with the No-Pause/Whole group.

Subjects. Sixteen different people were assigned, in an unsystematic fashion, to each of the three groups. These people were Stanford undergraduates fulfilling an introductory psychology course requirement. All subjects claimed (when queried afterwards) to have followed the imagery instructions at least 70% of the time during the second block. Furthermore, no subject fathomed the purposes or hypotheses of the experiment. Thus, data from all subjects were included in the analysis.

Results

An analysis of variance was performed on the data. Because we wished to generalize over both subjects *and* items, the quasi-F statistic, denoted F', was necessitated (Winer, 1971). Following each statistical value will be the mean square error for that comparison over subjects (represented SMS_e) and items (IMS_e). Only correct responses were analyzed, and the overall error rate was 6.44%.

In general, when subjects were not asked to use imagery, they evaluated high-association properties more quickly, even though these properties also were smaller than the low-association properties; in contrast, when subjects used imagery, they now assessed the low-association/large-area properties more quickly [for the interaction between property type and imagery instructions, $F'(1,53) = 9.20$, $p < .01$, $SMS_e = .052$, $IMS_e = .329$]. The effects of instructions on time to evaluate the two types of properties also was evident in analyses of variance done for each group separately: for the interaction between property type and instructions in the data of the Pause/Whole group, $F'(1,35) = 5.66$, $p < .05$, $SMS_e = .035$, $IMS_e = .189$; for the No-Pause/Whole group, $F'(1,37) = 7.95$, $p < .01$, $SMS_e = .065$, $IMS_e = .128$; for the No-Pause/Part group, $F'(1,35) = 4.34$, $p < .05$, $SMS_e = .057$, $IMS_e = .152$. These effects of instructions were about the same in the three groups; no interaction occurred between property type, instructions, and group, $F' < 1$. It should be noted, however, that contrasts performed on the imagery data revealed significant effects of size only for the first two groups; there was no effect of property type in the imaged items from the No-Pause/Part group, $F' < 1$. These results are visible in Figure 1. As is evident in the figure, imagery instructions slowed evaluation times down relative to the no-imagery condition, $F'(1,52) = 103.01$, $p < .001$, $SMS_e = 1.564$, $IMS_e = .256$. There were no differences between overall group times, $F' < 1$, nor was any other effect or interaction in the "true" data significant.
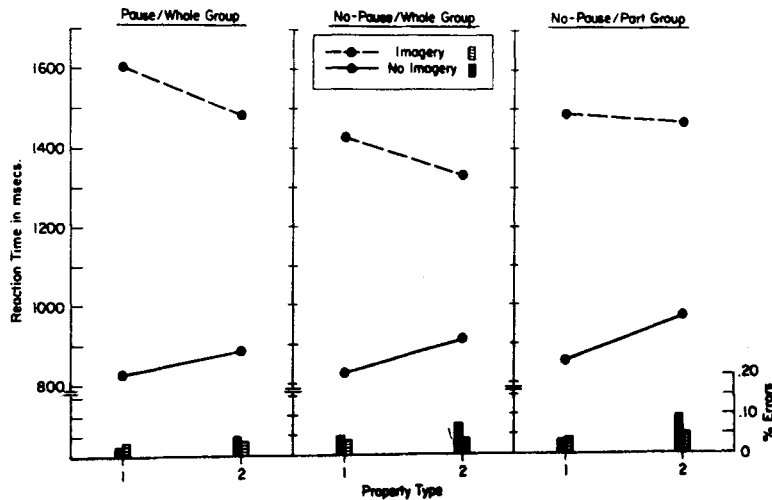
Figure 1. The time required to verify two types of properties in three different conditions when imagery was or was not used. Type 1 properties are high association/small area and Type 2 properties are low association/large area.

Finally, errors (depicted at the bottom of Figure 1) generally were positively correlated with evaluation times.

Discussion

When people used images of the whole animal, they could verify the larger property more quickly (e.g., head vs. claws for a cat), even though the smaller property was more highly associated with the animal in question. In contrast, when no imagery instructions were administered, subjects were faster not with the larger characteristic, but with the more highly associated, smaller one. Interestingly, neither association strength nor size dictated time to construct an isolated part of an image; discussion of this result will be deferred until after the next experiment.

The fact that imagery instructions resulted in longer overall evaluation times is also of some interest. In the present situation, however, instructions and presentation order are confounded. In these sorts of experiments, however, the effects of practice (i.e., lower reaction times), not of fatigue, usually occur with later items (e.g., Conrad, 1972). Thus, it seems unlikely that the huge differences in overall evaluation times due to instructions are attributable merely to order effects. The present data, however, do not allow us to discriminate between three explanations of this finding: First, in the present experiment, subjects were urged repeatedly to be certain to use imagery. These instructions may have caused the subjects to be extra conservative in determining when they had adequately consulted their images. Second, imagery may follow propositional search, and thus take more time to assess. Jorgenson and Kintsch (1973) rejected this explanation because imagery instructions did not significantly increase verification times in their experiment (although a trend in this direction was present). Their failure to obtain a significant difference may have been due to the precise nature of the imagery instructions and their inclusion of

difficult-to-image items (e.g., "Calendar projects movies"), which may have discouraged serious use of imagery. The third explanation posits that propositional search and imaging occur in parallel, with imagery simply requiring more time. That is, perhaps imaging involves additional processes, like retrieving from the semantic representation of the property procedures (in Winograd's, Note 1, sense) used to classify parts of the image. Such an account of the longer latencies obtained when imagery is used presupposes, of course, that the operations used in the nonimaginal case—but not used with imagery—consume less time than those used in imaginal verification, but not otherwise.

The failure to find any difference between the three groups in overall time to assess the imaged items is slightly disconcerting. A simple, straightforward conceptualization of this situation would surely lead us to expect the imaged items in the Pause/Whole group to have been evaluated more quickly than the imaged items in the No-Pause groups. In the former case, a subject already had much of the necessary construction done prior to probing, while in the latter case, he must both construct and evaluate the image at the time of probe. Perhaps temporal uncertainty (see Welford, 1968) is operating to increase times in the Pause situation. Also, it is entirely possible that the subjects in the Pause group used different strategies, with different time-consuming components, than those used by subjects in the No-Pause groups (see Kosslyn, 1974, for an elaboration of this idea).

EXPERIMENT II

The set of items used in the first experiment were highly selected and quite unusual; most often, association strength and size are positively correlated (e.g., as in the lion's mane, a tiger's stripes, etc.). One could argue that our results are somehow a consequence of the peculiarity of our item set. Thus,

in the following experiment, items were not preselected with size in mind at all. Instead, a group of people provided post hoc ratings on relative size and association strength of properties for given animals, as well as rating the familiarity of the animals themselves. The mean of each variable, along with "node distance" (in the sense of Collins & Quillian, 1969), was calculated for each item used previously in a verification task; these values then were included in a regression analysis of the evaluation times measured in the verification task. This analysis was an attempt to discover which variables were most important in determining time to verify properties in four conditions, roughly equivalent to the no-imagery and imagery conditions in the Pause/Whole and No-Pause/Part groups of Experiment I.

Method

Four groups of people determined as quickly as possible whether or not named animals had given properties. They heard statements like "A lion has a mane," where the predicate followed immediately after the verb or was separated by a brief pause. Within each of these conditions, one group verified the statement via consulting an image, whereas the instructions for another group made no mention of using imagery.

Materials. The data analyzed here are a subset of those collected by Kosslyn and Nelson (unpublished manuscript, reported in part in Nelson & Kosslyn, 1975). That experiment utilized 54 true and 54 false items, half of which were statements incorporating a "can" relation (e.g., "A mouse can see"), and half incorporating a "has" relation (e.g., "A lion has a mane"). Only the "has" relations were considered in the present analysis, it being unclear how to assess the "relative size" of the parts of an image involved in depicting a "can" relation. The items were selected to represent a wide distribution along two dimensions: association strength and "node distance" (i.e., level in a nested hierarchy where the property would be stored most efficiently; see Collins & Quillian, 1969; Conrad, 1972). Size was not a factor considered when selecting the items.

The same items, in the same random order (except that no entry could occur twice within four consecutive items), were recorded on two tapes. On one tape, the predicate followed 5 sec after the noun, a brief pause being inserted just prior to the property word (e.g., "A cow has . . . horns"); on the other tape, the sentence was read in a relatively rapid manner, without any pause. As in Experiment I, a clock was started upon presentation of the property name and was stopped when the subject responded by depressing either of two buttons. Sentences were separated by 10 sec.

Ratings. Twenty-six Johns Hopkins students (who were about the same age and from similar backgrounds as the students who participated in the verification task to be described below) rated the 27 true "has" animal-property pairs (three properties for each of the nine different animals). They rated the pairs in terms of: (1) familiarity of the animal, (2) strength of the association between the property and animal, and (3) relative size of the property on an image of the animal. All ratings were done on a standard 7-point scale. A mean rating was obtained for each item on each of the three dimensions. In addition, each pair was categorized in terms of the "node distance" reflected in the relation. That is, the properties were ranked in terms of the level of hierarchy at which they would be most efficiently stored according to Collins and Quillian's (1969) scheme. These four variables were chosen because they seemed relatively separable and distinct, and because they had been shown to affect evaluation times in other experiments. Variables like "conjoint frequency" and "production frequency" were not included because they seemed related to association strength (but

were more difficult to assess); similar considerations led to the inclusion of measures of familiarity, but not "frequency," of the animal.

Procedure. Four groups of subjects were tested: The Pause/No-Image group heard the tape with the pause inserted prior to the property word, and was instructed simply to evaluate the sentences as quickly as possible; these instructions were basically the same as the no-imagery instructions administered to the Pause/Whole group of Experiment I. The Pause/Image group also received the tape with the pause, and received imagery instructions of the sort given subjects in the Pause/Whole group of the first experiment. The No-Pause/No-Image group heard the stimulus tape where statements were recorded without a pause after the verb, and received instructions identical to the no-imagery instructions given subjects in the No-Pause groups of Experiment I. Finally, the No-Pause/Image group was instructed to base all decisions upon consultation of images of the stated relations; these instructions did not require the subject to image the entire animal, but rather requested that he wait until the property word occurred, and then image that characteristic on the animal. In essence, these instructions were equivalent to those of the No-Pause/Part group of Experiment I.

All subjects first evaluated eight practice items (using nontest animals and properties) prior to hearing the actual test items. Details of the procedure were the same as those employed with the corresponding group in Experiment I.

Subjects. For each group, data from 15 different Stanford undergraduates (who were paid volunteers randomly assigned to groups) were analyzed. These subjects all had error rates under 10%. This constraint seemed desirable because the data were analyzed item by item, and Item by Subject interactions seemed less likely to occur if all subjects correctly evaluated as many items as possible.

Results

Four mean evaluation times were obtained for each animal-property pair, one per group (i.e., instructional condition); these means were obtained by averaging over the 15 subjects' times to verify the pair. Only correct trials for the true "has" statements were analyzed, and the overall error rate for the subjects included in the analysis was 7.6%; errors were correlated positively with evaluation times in each group. Four separate step-wise multiple regression analyses were conducted, one for data from each instructional condition. The mean evaluation times were treated as the dependent variable, and the means for each of the three variables described earlier, plus node distance, were entered for each item as independent variables.

Unlike many regression analyses, the purpose of the present analysis was not to obtain a multiple correlation. Nor were we particularly interested in effects due to all of the variables; familiarity and node distance were included by way of avoiding possible confoundings between these variables and the remaining two. Instead, we were interested in the *order* that association strength and relative size were entered into the equation for each set of data. The variables were entered into the equation in accordance with how much variance they accounted for: more important variables were entered sooner. In addition, we wanted to know whether significant amounts of

Table 2
Results of the Regression Analyses of Data from Experiment II

No Imagery		Imagery	
Variable	F(1,25)	Variable	F(1,25)
Pause			
Node Distance	9.33**	Node Distance	5.71*
Association Strength	11.52**	Relative Size	8.21**
Familiarity	4.35*	Familiarity	4.19
Relative Size	< 1	Association Strength	1.04
No Pause			
Familiarity	5.88*	Association Strength	6.56*
Node Distance	3.80	Node Distance	2.76
Association Strength	4.88*	Familiarity	3.31
Relative Size	< 1	Relative Size	< 1

*p < .05   **p < .01

variance were accounted for by the addition of each of these variables in the different instructional conditions.

The results of the regression analyses are presented in Table 2. The variables are listed in the order they were entered into the equation, and an F ratio is listed to the right of each variable. The F values were calculated by dividing the variance accounted for by inclusion of a variable by the residual variance; this allowed us to estimate the significance of the variance accounted for by each independent variable. Significant effects of association strength indicate that evaluation times decrease as association strength increases. Similarly, significant effects of the relative size of a property indicate that larger properties were verified more quickly. In addition, it should be noted that times *decreased* with increasing node distance, precisely the opposite of Collins and Quillian's (1969) findings. This result is replicated (using nonregression techniques) and discussed in Nelson and Kosslyn (1975).

Pause/No-Imagery group. Association strength accounts for more of the variance in evaluation times for this group than does relative size: not only was association strength entered sooner, but it also accounts for significant amounts of variance (p < .001); relative size, on the other hand, was of no consequence at all (p > .25).

Pause/Imagery group. When subjects were asked to consult an available image of the whole animal, relative size now was the more important variable. Not only did it enter the equation before association strength, but it accounts for significant amounts of variance (p < .01); in contrast, association strength was of no importance here (>.25).

The results from the two Pause groups dovetail nicely with the results from the Pause/Whole group of Experiment I. Size was important when images were consulted; when images were not inspected, however, association strength was the more potent factor in determining ease of verification.

No-Pause/No-Imagery group. Association strength

again was entered into the equation prior to relative size, and accounts for a significant portion of the variance (p < .05), whereas relative size did not (p > .25).

No-Pause/Imagery group. Interestingly, association strength was entered into the equation before relative size and accounts for significant amounts of variance (p < .05); relative size of properties seemed to have no effects on evaluation times in this condition (p > .25).

The results from the No-Pause groups reflect those obtained from the No-Pause/Part group of Experiment I. In neither experiment did relative size affect evaluation time in this condition. The only disparity between the results from the corresponding conditions in the two experiments was in the imagery condition; association strength had significant effects here, whereas it had no effect in the No-Pause/Part imagery condition in Experiment I. This may be a consequence of the particular items used in either experiment, and/or of the fact that the present items occurred along with "can" statements.

GENERAL DISCUSSION

Both experiments demonstrated that if one consults an available image for a property, smaller properties are more difficult to "see." When imagery is not used, however, more associated properties are evaluated more quickly. Thus, once images are constructed, they attain status as a representational medium with special properties and characteristics of their own.

If an image of the whole object is not consulted at the time of the probe, size does not systematically affect verification time. This is not surprising given the findings of Kosslyn (1975, Experiment 5). In this experiment, subjects were able to image relatively small or large objects at any of four different subjective sizes (furthermore, the actual relative size of an imaged object did not affect the ease with which an image of it could be generated). Thus, the subjective size of parts of an animal imaged in isolation need not be related to actual differences in relative size of the parts, and differences in time to "see" the parts internally need not arise. Only in the context of an image of the whole object are its various parts necessarily large or small, and thus more or less easily discerned.

The findings of the present paper, then, provide justification for utilizing imagery as an explanatory contruct in psychology. Only after an image is constructed and consulted does size affect processing time. Further, when previously constructed images are accessed, association strength—which reflects ease of accessing nonimaginal representations—no longer affects processing time. These results are difficult to understand if the image is an epiphenomenon of no consequence, a mere incidental

by-product of more basic abstract processes, as some (e.g., Pylyshyn, 1973) have argued.

REFERENCE NOTES

1. Winograd, T. *Frame representations and the declarative/procedural controversy.* Paper presented at the Carbonell Memorial Conference, Pajaro Dunes, California, May 1974.

REFERENCES

ANDERSON, J. R., & BOWER, G. H. *Human associative memory.* New York: Wiley, 1973.

COLLINS, A. M., & QUILLIAN, R. M. Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior,* 1969, **8,** 240-247.

CONRAD, C. Cognitive economy in semantic memory. *Journal of Experimental Psychology,* 1972, **92,** 149-154.

JORGENSEN, C. C., & KINTSCH, W. The role of imagery in the evaluation of sentences. *Cognitive Psychology,* 1973, **4,** 110-116.

KOSSLYN, S. M. Constructing visual images. PhD dissertation, Stanford University, 1974.

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

LOFTUS, E. F., & SUPPES, P. Structural variables that determine the speed of retrieving words from long-term memory. *Journal of Verbal Learning and Verbal Behavior,* 1972, **11,** 770-777.

NELSON, K. E., & KOSSLYN, S. M. Semantic retrieval in children and adults. *Developmental Psychology,* 1975, **11,** 807-813.

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

PUTNAM, H. Reductionism and the nature of psychology. *Cognition,* 1973, **2,** 131-146.

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

SCHAEFFER, B., & WALLACE, R. Semantic similarity and the comparison of word meanings. *Journal of Experimental Psychology,* 1969, **82,** 343-346.

SMITH, E. E., SHOBEN, E. J., & RIPS, L. J. Structure and process in semantic memory: A feature model for semantic decisions. *Psychological Review,* 1974, **81,** 214-241.

WELFORD, A. T. *Fundamentals of skill.* London: Methuen, 1968.

WILKINS, A. J. Categorization time and category size. *Journal of Verbal Learning and Verbal Behavior,* 1971, **10,** 382-385.

WINER, B. J. *Statistical principles in experimental design.* New York: McGraw-Hill, 1971.