

Identifying objects from a haptic glance

ROBERTA L. KLATZKY

Carnegie Mellon University, Pittsburgh, Pennsylvania

and

SUSAN J. LEDERMAN

Queen's University, Kingston, Ontario, Canada

Subjects identified common objects under conditions of a "haptic glance," a brief haptic exposure that placed severe spatial and temporal constraints on stimulus processing. They received no advance cue, a superordinate-level name as cue, or a superordinate and basic-level name as cue. The objects varied in size relative to the fingertip and in the most diagnostic attribute, either texture or shape. The data suggest that object recognition can occur when global volumetric primitives cannot directly be extracted. Even with no cue, confusion errors resembled the target object and indicated extraction of material and local shape information, which was sufficient to provide accuracy above 20%. Performance improved with cuing, and the effect of exposure duration was observed primarily with minimal cuing, indicating compensatory effects of top-down processing.

Many everyday circumstances require that you reach for and identify an object when it is not in the visual field. While watching a computer screen, for example, you may reach for a coffee cup or mouse. Other examples occur when you find an audio tape in the car while driving, grasp and extract keys from pocket or purse, reach for the telephone receiver with your head turned, contact and press an alarm-clock button, or find the light switch beside your bed during the night. These situations often involve not only identifying the object of interest, but also discriminating it from others contacted in the workspace. We have commonly observed that these achievements can occur within the first moments of contact.

Rapid identification of an object from brief initial contact is likely to reflect two underlying types of processes. One is data extraction from the object; a brief contact may be sufficient to extract a number of features that are diagnostic of the object's identity and even the location of contact. Below, we will describe research on the nature of information available from brief contact that supports this assumption. The second mechanism is top-down processing: The situations we have described are those in which a person has an advance hypothesis about the objects to be encountered in the workspace, and having a hypothesis about an object's identity is known to facilitate identification from limited data. The interplay between data extraction and top-down hypotheses is at the

heart of interactive activation models of pattern identification (e.g., McClelland & Rumelhart, 1981).

The purpose of the present research was to investigate how, and how well, we perform object identification from initial contact. We wished to determine whether identification from limited haptic contact is in fact possible, and if so, at what levels and under what circumstances it occurs.

Our term for initial contact with an object is a *haptic glance*. We define a haptic glance as brief, spatially constrained contact that involves little or no movement of the fingers. Although we are interested in contact on the order of 200 msec, for reasons explained below, we will consider longer intervals in this paper. To understand how object recognition might be possible with such minimal contact, it is important to consider both what sensory stimulation is present and what kinds of object exploration a glance might instantiate. We will review this issue next.

Information Provided by Exploration With a Haptic Glance

In keeping with Loomis and Lederman (1986), we have adopted one of several traditional approaches used to differentiate sensory systems—specifically, one that focuses on differences in their underlying neural mechanisms. Accordingly, our definition of the haptic system stresses the use of somatosensory inputs from skin (cutaneous), as well as from muscles, tendons, and joints (kinesthetic). We argue that such inputs, which are typically, although not necessarily, obtained via purposive manual exploration, are combined to derive one's perceptions of many external properties of objects and their spatial layout. The exposure conditions of a haptic glance would be expected to provide both cutaneous and kinesthetic information, although kinesthetic cues would be limited by the absence of prehension.

We acknowledge support from the Information Technology Research Centre of Excellence of Ontario and the National Science and Engineering Research Council of Canada. We also thank Judith Loftus, Michelle Pinto, and Steve Ito for their help with data acquisition and analysis. Correspondence concerning this article should be addressed to R. L. Klatzky, Department of Psychology, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213-3890 (e-mail: klatzky@cmu.edu).

Our previous work has stressed the importance of purposive haptic exploration in object identification (see Klatzky & Lederman, 1993, and Lederman & Klatzky, 1993, for reviews). A starting point for this program was the identification of a set of *exploratory procedures* (EPs)—that is, relatively stereotyped patterns of contact and movement relative to objects that are associated with the extraction of particular object properties (Lederman & Klatzky, 1987). An EP is said to be associated with a property when it is (1) preferred for extracting that property under conditions of free exploration, and (2) optimal for extracting that property (in terms of accuracy, or in case of a tie, speed). When an EP is not optimal for extracting some property, it may still be sufficient to make reasonably fine discriminations.

We have demonstrated that an EP acts like a *gateway* to perception of an object property, in that if the appropriate EP is not performed, the associated property is minimally used for identification and discrimination tasks (e.g., Klatzky, Lederman, & Reed, 1989; Lederman, Klatzky, & Reed, 1993). Thus it is important to consider what EPs are executable during a haptic glance, and accordingly, what information about object properties might be available.

The principal EPs and associated properties are *lateral motion*, movement tangential to the object's surface associated with sensing of surface properties like texture; *pressure*, usually applied normal to a surface and associated with compliance; *static contact*, associated with sensing of thermal properties (e.g., apparent coolness); *contour following*, movement along edges necessary to extract precise shape; *enclosure*, molding to surfaces and edges associated with global shape and volume; and *unsupported holding*, used to sense weight. Although the EPs were initially defined for free hand movement, some are instantiated, at least to some extent, by brief contact with the fingertips.

If there is essentially no movement of the fingers, a haptic glance constitutes the EP of static contact, associated with thermal sensing. Deformation of the finger pad to accommodate surface discontinuities, and relative positioning of multiple fingers so as to maximize contact with a surface, constitute local versions of the enclosure EP. Structural properties that would be provided by this EP would include the shape, size, and orientation of tactual edges, signaled by discontinuities in pressure; in contrast, continuous pressure gradients on the fingertip (such as one would find on the surface of a sphere) and relative finger positioning would indicate surface shape and/or orientation relative to the plane of the hand. Sufficient pressure normal to the surface to maintain contact, if not overt force, could give rise to information about the compliance of the touched surface. Tangential movement of the skin across the object surface during the glance would constitute lateral motion or, if the glance fell on an edge, contour following, associated with texture and shape encoding, respectively.

In summary, a haptic glance could provide local information about an object's material (thermal properties, compliance, texture) and local geometry. Obviously it

would not directly provide information about an object's geometry outside the scale of the fingertip. The distribution of mass, which may tell the haptic system a great deal about an object's part structure (Solomon, Turvey, & Burton, 1989), would also be unavailable from a haptic glance.

We have recently verified (Lederman & Klatzky, 1995) that discriminations about a number of properties can be made with glance exposure, using a search task (Treisman & Gelade, 1980) with an apparatus designed to bring surfaces up to 1–6 outstretched fingers (Moore, Broekhoven, Lederman, & Ulug, 1991). A number of discriminations using this task produce very fast search rates (i.e., slopes of under 30 msec). Such rapid discriminations have been found for contrasts on the dimensions of roughness, compliance, apparent coolness (thermal conductivity), and local shape (e.g., edge/no edge; hole/no hole; deep/shallow indentation). Slopes under 30 msec mean that a search over all six fingers could take place within 200 msec (although time for finger adjustment, entering the intercept of the search function, may still be needed). Further, our pilot work suggested that 200 msec was a reasonable approximation to the phenomenon with which we began: a brief, purposive contact that is sufficient for identification.

Structure and Material as Cues for Haptic Object Identification

A haptic glance would appear to eliminate the most fundamental information needed for object identification: global structure, at least for objects outside the scale of the fingertip. Biederman and associates (Biederman, 1987; Biederman & Cooper, 1991, 1992; Hummel & Biederman, 1992) have provided substantial evidence for a model of visual recognition called *recognition-by-components*, in which objects are represented as sets of volumetric primitives, or *geons*. A description of relatively few volumetric primitives and their spatial relations is sufficient to recognize an object and is invariant over transformations of viewpoint. Indeed, Biederman and Ju (1988) showed that adding non-edge-based information such as color did not facilitate visual object recognition (but see Wurm, Legge, Isenberg, & Luebker, 1993, for color effects).

The emphasis on structural rather than material properties in object representations is also consistent with theories of Rosch (e.g., 1978), who suggested that structural descriptions are used to partition objects into basic-level categories. Tversky and Hemenway (1984) amended Rosch's theory to suggest a partitioning of categories by descriptions in terms of parts. Further, Lederman and Klatzky (1990) found that shape tended to discriminate basic-level categories even when subjects were told to think about identification by touch. Thus, this work suggests an ecological constraint on object recognition within vision or touch: If basic-level categories are distinguished on a structural basis, it is essential to extract structural information in order to successfully identify an object. A study by Klatzky, Loomis, Lederman, Wake, and Fujita (1993) showed that structural information was

sufficient, if not necessary, for object recognition by touch: When access to material properties of objects was minimized (by having subjects wear heavy gloves, using rigid objects without moving parts, and with no lifting permitted), identification was still highly accurate.

Under glance exposure, structural information about the object is not directly available beyond the local probe. This raises the question of whether objects beyond the scale of the fingertip can be identified at all, particularly in the absence of top-down hypotheses. If so, there are two possibilities: Either people can build an adequate volumetric description from very sparse spatial data, or the basis for recognition by a haptic glance is something other than a global structural description of the object. In the latter case, the description must be made in terms of material and/or more local structural features.

Work of Kappers, Koenderink, and Lichtenegger (1994) is relevant to the issue of building an adequate volumetric description from sparse data. They investigated the haptic perception of objects that were discriminable only by shape. Two parameters defined each object: a *shape* index, essentially determined by the relative curvatures of the two principal axes, and a *curvedness* index, reflecting the total curvatures. The objects were 20 cm in diameter, extending beyond the surface contacted by the hand. They were designed so that the shape and curvedness indices remained fairly constant over all points on the surface, and hence they were, in theory, locally discriminable. The shapes were also complex, the shape-index endpoints corresponding to spherical paraboloids and the objects near the midpoint being asymmetrical saddle shapes. After prefamiliarization, both visual and haptic, with the shape-index scale, subjects were asked to assign scale values to each of a set of objects. In general, there tended to be a regular relationship between the shape index and the subject's response; thus, subjects were able to discriminate and categorize the shapes (this varying to some extent with curvedness). However, the strategy for exploration was not a local, static sample; subjects instead actively scanned the object, usually with the whole hand. It is unknown whether subjects could have made the required discriminations on the basis of a local sample alone, but their clear preference was not to do so.

Even without the constraints of a haptic glance, material descriptions are likely to be important in tactual object recognition, particularly in comparison with their role in vision. In the same study that showed that structure was sufficient for object recognition, Klatzky et al. (1993) found that haptic object recognition was augmented when subjects used an open-fingered glove as opposed to one that sharply reduced access to material. The augmentation was greatest when the data allowed about the object were most sparse—namely, when exploration involved probing with a single extended finger.

Questions Addressed by the Present Study

The present study was intended to address several questions. The first, and most fundamental, was, how good

is object identification with a haptic glance? To answer, we determined baseline levels for identification with a haptic glance over a substantial corpus of objects. Although the constraints on the corpus of objects may limit the generality of the precise identification rates obtained, the data indicate performance well above chance. The second question was, what local information can be obtained from a haptic glance? The foregoing discussion indicates that material and shape cues should be available; converging evidence was sought from patterns of confusion errors.

The remaining issues can be couched in terms of the interactive activation approach to pattern recognition (e.g., McClelland & Rumelhart, 1981), in which object identification reflects the interplay between information from the stimulus (data-driven, or bottom-up processing) and the identifier's expectations or hypotheses (conceptually driven, or top-down). In these terms, we have severely limited data-driven processing. The limitations in global structural information, which is thought to be so important in visual object recognition, are likely to be particularly severe. Our third question was, then, what is the relative contribution of material and global structure to recognition under these conditions? To address this question, we used objects for which either texture or shape was known to be particularly diagnostic under conditions of free exploration. Differences in identification of these objects would suggest differences in the relative contributions of material and structural properties. This issue is not independent of the scale of the object, as discussed below.

Fourth, our manipulations allowed us to separate, at least to some degree, the spatial and temporal constraints inherent in the haptic glance and ask how important they are. Temporal constraints were manipulated by using two exposure durations. To examine spatial constraints, we varied whether the object itself was large or small relative to the area spanned by the fingers. Glance-imposed limitations on the spatial sample should be evident in an effect of object size for both texture- and shape-diagnostic objects, but in different directions: When texture is diagnostic and it is homogenous across the object, larger objects allow contact with a larger skin surface, which may facilitate identification by providing redundant information. But if shape is diagnostic, a small object that lies entirely within the finger contact area should be more identifiable than one having shape cues that extend beyond the fingertips—unless global shape can be determined from local cues, in which case we would expect no effect of object size.

A fifth question was, can top-down information compensate for the lack of perceptual data about the stimulus in this situation? To investigate the role of top-down information, we used three cuing conditions, varying in the amount of object-category information presented to the subject in advance of contact. In one condition, subjects were cued with a yes/no question that identified the object at the superordinate-level category and provided a basic-level name to be verified or disconfirmed (the

Sup+Bas/Y-N task). In the two other conditions, they were cued with either a superordinate category label or no cue and were to produce an object name (the Sup/Naming task and No-cue/Naming task, respectively). We predicted that more informative cues should reduce the need for stimulus information and thus alleviate the effects of temporal and spatial constraints. Alleviation of temporal constraints would be indicated by a decrease in the effects of exposure when stronger cues were present. To the extent that the cue compensated for an inability to encode global shape or, alternatively, augmented the ability to do so, it should particularly reduce errors on large objects for which shape is diagnostic; thus, one would expect the effect of cuing to vary both with diagnostic property and with object size. Finally, if texture and shape differ generally in informativeness under a haptic glance, these differences should also be reduced in the stronger cue conditions.

METHOD

Stimuli

The stimuli were a set of 44 objects, partitioned according to three variables: (1) the attribute that had been judged most useful in categorizing the objects at the basic level, texture or shape; (2) the size of the object, large or small; and (3) whether the object occurred as a positive or negative instance in the yes/no condition of the task. These variables are explained in more detail below.

The first variable is called the *most diagnostic attribute*, or MDA. The items were initially selected from a pool compiled by Lederman and Klatzky (1990). The constraints on the original pool were that each object be familiar and come from a superordinate category having two reasonably common basic-level categories, that one of those categories have two subordinate categories with an identifiable MDA, and that the MDAs of the subordinates vary over the pool. To determine norms for the diagnosticity of attributes, Lederman and Klatzky had subjects indicate, for each of a series of named objects, which tactually accessible attributes would be most diagnostic when the object was to be identified by touch. The attributes were identified from a closed list including texture and shape. Items were chosen for the present study from two subsets in the Lederman/Klatzky norms for categorization at the basic level: those objects for which texture was more diagnostic than shape, and those for which shape was more diagnostic than texture. In most cases, the dimension that was the more diagnostic between texture and shape was also the object's MDA from the norming data and will be referred to as the MDA here.¹ The particular objects that were physically presented were selected to be common exemplars of the named category (although the baseball bat was below regulation size, to avoid excessive weight on the contact sensor).

An object's size was defined according to the lengths of major axes of the planar surface exposed to the exploring hand. The item was classified as "small" if its smaller axis was 3.8 cm or less in length, which was intended to fall within the width spanned by the three middle fingers of a typical hand. It was called "large" if both major axes exceeded 3.8 cm. A part or feature of the object exposed to the hand might be smaller, however, along one axis, in which case the longer axis of the part was required to be 5 cm or more (e.g., the stem of the pipe or the rim on the bread pan). Placement of the object relative to the hand was different for small and large objects, as described below.

The third variable was whether the object was in a negative or positive set. In the yes/no version of the task, as in Lederman and Klatzky (1990), subjects were asked questions of the form, "Is this

[superordinate category name] a [basic category name]?" (E.g., Is this writing implement a pencil?) For a positive item, the basic category name given in the question was also the name of the presented object. For a negative item, a different object was presented than that named in the normed items, although it was in the same superordinate category.

The meaning of the MDA variable is slightly different for positive and negative items. The MDAs from the Lederman/Klatzky norms refer to the target item that was named in yes/no questions. In the case of positive items, this is also the item that was presented, so the same MDA applies. However, the negative items differed from the named item. The present negative items were selected, as much as possible, so that in the yes/no task, the MDA of the named item was a dimension that differentiated it from the presented item. That is, when texture was the MDA for a named object, the presented (negative) item was discriminable in texture but similar in (local) shape. The converse was true for the shape MDA. For example, when "crayon" was the named item and texture was the MDA, the presented item was a magic marker, which is similar in shape to a crayon but discriminable in texture. Selection of the negative items in this way does not guarantee that the associated MDA (from the Lederman/Klatzky norms) is most diagnostic for free identification, but it does indicate that the dimension is diagnostic for discrimination.

Within each of the diagnostic dimension categories (texture and shape), there were 12 positive and 10 negative items, and within each of those sets, half were classified as small and half as large. The Appendix lists the items and indicates their diagnostic attribute, size, and positive/negative status.

Procedure

Subjects identified objects under three cuing conditions, with cuing a between-subjects manipulation. In the Sup+Bas/Y-N task, subjects were cued just as in the earlier study of Lederman and Klatzky, with a question that named the superordinate category and basic-level object (e.g., is this container top a cork?). The subject indicated whether the basic level name was correct. In the Sup/Naming task, the superordinate-category name corresponding to the yes/no question was provided (It's a container top: What is it?), and in the No-cue/Naming task, no information was given in advance (What is it?).

On each of the series of trials, the experimenter first placed an object on a piezoelectric sensor on the table top. Some objects were placed on an intervening base that was scaled to compensate for the height of the object, so that the total height of the base (if any) and object, relative to the sensor, was approximately 15 cm for all items. The height of the object was therefore not a cue to its identity.

The blindfolded subject placed his/her arm at the top of a vertical metal bar, approximately at shoulder height and to the right of the object. A comfortable bend in the elbow, adjusted freely by the subject, resulted in the object's being directly below the hand. The experimenter sighted down through the subject's hand and adjusted the object's position so that the finger(s) would contact it as described below. The experimenter then verbally gave the subject no cue, a superordinate category cue, or a yes/no question naming the superordinate category and presenting a putative basic-level category, as was appropriate to the condition. When the experimenter pressed a key on the computer, a tone sounded, indicating that the subject should begin moving his/her arm down the metal bar at a comfortable speed. After descending approximately 20 cm, the arm touched an elastic band placed around the bar at a height just above the top of the object to indicate that the object was near, and the fingers then contacted the object. The subject was instructed to gently touch the objects with the fingertips only (as many as made contact with the object) and to avoid applying strong pressure or moving the fingers on the object's surface. At contact, the sensor sent a signal to the computer, and at 100 msec less than the target exposure duration, either 200 or 3,000 msec, a 100-msec tone sig-

naled that the arm should be withdrawn upward. The subject was instructed to lift the arm off the object at the onset of the tone.

The experimenter monitored the subject's hand position throughout the trial, including the liftoff from the object. If the hand moved during descent so that it missed the object, which rarely occurred, the trial was repeated.

The absolute precision allowed by manipulating target exposure in this way is obviously limited, but there is little doubt that the two conditions, which differed by more than an order of magnitude in signaled duration, differed substantially in actual contact duration. The signal to withdraw anticipated the desired end of contact by 100 msec. If the period from that signal to the onset of arm withdrawal occurred in the range of an auditory reaction time—about 150 msec (Woodworth & Schlosberg, 1954)—subjects would be leaving the stimulus within 50 msec of the intended end of contact. Of course, moving the mass of the arm would take somewhat longer than the usual digital response, extending the period slightly. However, because subjects took part in only one duration condition, they became used to the rhythm of the upward and downward arm motion and hence may have initiated withdrawal in considerably less than a simple RT. Monitoring by the experimenter ensured that the subject did not visibly extend the contact duration.

To avoid biasing against large, shape-diagnostic objects by placing critical features outside the contact area, we tried to maximize access to critical features at contact. Therefore, positioning of the objects relative to the hand was done so as to satisfy two constraints: (1) Keep the functional size of the object—the size exposed to the fingertips—consistent with its designated category, large or small. (2) Provide the maximum information possible. Under the first rule, the objects classified as small, but with a longer axis exceeding 3.8 cm, were oriented so that the longer axis was parallel to the arm. In contrast, large objects having an axis of the exposed surface that was less than 3.8 cm long were oriented with the long axis across the fingertips. A small object having two axes that fell within 3.8 cm was also oriented with the longer axis across the fingertips, so as to maximize the available information. If there was a feature that was particularly informative with respect to identifying the object or, in the yes/no task, to disconfirming a negative item (e.g., the pouring lip on a pitcher), the object was placed so that the feature was centered below the fingers. As examples of these rules, the chalk (small) was oriented so that its long axis was aligned with the fingers, but the rim of the bread pan (large) and hook (small) were oriented so that the long axis was across the fingers; the file was placed so that the ridged portion (designated as a discriminating feature) was just below the fingertips. Details about the placement of individual objects are available from the authors.

After retracting the hand, the subject responded and gave a confidence rating on a 1–5 scale. In the yes/no condition, the response was yes or no, and in the no-cue or superordinate cue condition, the subject gave a name or a “don't know” response. He or she was informed that a “don't-know” response was equivalent to zero confidence. After a don't-know response in the superordinate and no-cue conditions, subjects were invited to indicate what they knew about the object. The experimenter then asked them to indicate whether they had knowledge about additional attributes, if not mentioned: its material, shape, size, weight, temperature, texture, and hardness.

Subjects and Design

The subjects were college students who took part in the study to satisfy a course assignment. Cue conditions (Sup+Bas/Y=N task, Sup/Naming task, and No-cue/Naming task) and target exposure durations (200 and 3,000 msec) were varied between subjects, for a total of six groups of 12 subjects each. Characteristics of the items (size, MDA, positive/negative) were within-subjects variables. The 44 items were randomly ordered for presentation to each subject. Four practice items preceded the stimulus set proper.

Control Conditions

Two control conditions were performed to determine the objects' identifiability under free exploration and the guessing rate given the superordinate cue. In the first control, 10 subjects freely explored the objects and attempted to generate the name of each. If they were incorrect or could not provide a name, the superordinate cue was given, and if that failed to elicit the response, the yes/no question was asked. Overall, 93% of the responses were correct, given no cue whatsoever. Only 2 of 440 responses were incorrect given the yes/no cue, in both cases for corduroy (which was rather soft). Each object was scored for free identification for purposes of correlations reported below. If all subjects were correct with no cue, the score was 4; if any subject required a superordinate cue but not more, it was 3; if any required a yes/no cue, it was 2; and if any was incorrect with the yes/no cue, it was 1.

The second control condition was used to estimate a guessing rate for the Sup/Naming task. The 22 subjects were given the set of 28 superordinate category names and were asked to generate, for each name, a designated number of objects. That number was set individually for each superordinate category, corresponding to the number of exemplars of that category in the stimulus set (from 1 to 4 objects per category). Thus, each subject generated 44 object names, the same number as in the stimulus set. The mean number of responses that matched stimulus items was 10.82 ($SD = 1.64$), or 25%. This constitutes a lenient estimate of the guessing rate for the superordinate cue, because it does not require that multiple items within a superordinate category be guessed in the correct order. If the proportion of correct guesses for a superordinate category is discounted by multiplying it by the probability that a guessed item is in the correct position in the set of items from that category, the guessing rate drops to 16%. However, this estimate may be too low, because experimental subjects can use previously identified objects from the category when generating a new guess. Thus the guessing rate given a superordinate-category cue should be estimated as within the range of 16%–25%. Since our interest is in demonstrating that subjects who were exposed to the objects with a haptic glance performed above chance, performance with a superordinate-category cue is conservatively required to be above 25%.

RESULTS

The principal dependent variables were accuracy and confidence. To assess accuracy, the Sup and No-cue/Naming tasks, in which object names were freely given, required a set of scoring rules. A given name was scored correct if it was accurate at the basic level or below; a qualifier below the basic level was not necessary (e.g., “ball” was sufficient for a tennis ball). The name of a material was not sufficient unless the object itself was a material sample, in which case guesses about the object were accepted (e.g., “corduroy” or “corduroy pants” was accepted for a square of corduroy). An incorrect subordinate name that was in the correct basic-level category and consistent with the explored sample (e.g., “pump” for a man's shoe) was accepted, but one that was grossly inconsistent (e.g., “pool ball” for “tennis ball”) was not. Adjectival qualifiers in isolation were not accepted (e.g., “hard”), but qualifier–noun combinations were accepted as long as the correct qualifier or a close alternative (e.g., “felt ball” for a tennis ball) was given and the noun was correct at the basic level. The two authors and a third experimenter collectively agreed on any ambiguous items, following these rules.²

Overall, performance was well above chance. Accuracy was 74% in the Sup+Bas/Y-N task (relative to 50% chance), 47% in the Sup/Naming task (relative to the upper chance estimate from the control condition, 25%), and 21% in the No-cue/Naming task. Texture- and shape-diagnostic objects were recognized at similar levels when averaged within each cuing condition and overall (about 50% on the average for both sets of objects). Table 1 reports the means and standard deviations over subjects for each of the principal experimental conditions.

Because there are very different guessing levels for the different cues, analyses of variance (ANOVAs) of accuracy were conducted separately within each cuing condition, on the factors of exposure time and three item characteristics: MDA, positive/negative, and size. An ANOVA of confidence was conducted across the three cuing conditions. ANOVAs with subjects and items as the units of observation were conducted. In the subject ANOVAs, exposure was a between-observations variable and the item variables were within observations; in the item analyses, exposure was a within-observations variable and the item variables were between observations. To avoid idiosyncratic effects of items or subjects, we report only effects that were significant in both analyses. In the following reports, the *F* tests and *p* values are given for the subject and item analyses in succession (labeled *F*₁ and *F*₂, respectively).

To preview the results, there was an effect of exposure time on accuracy in the No-cue/Naming task, but this was qualified in the Sup+Bas/Y-N task and absent in the Sup/Naming task. There was an interaction between size and MDA in the analyses of both accuracy and confidence, in both the no-cue and superordinate-cue conditions. The form of this interaction was as predicted: An increase in size facilitated identification when texture was the MDA and impeded it when shape was the MDA.

This interaction was qualified or absent in the Sup+Bas/Y-N condition, however.

Accuracy: No-Cue/Naming Task

In the No-cue task, there was a significant effect of exposure [*F*₁(1,22) = 14.77, *p* < .001, and *F*₂(1,36) = 27.02, *p* < .0001]. Exposure also interacted with size [*F*₁(1,22) = 10.50, *p* < .01, and *F*₂(1,36) = 4.32, *p* < .05]. The proportion correct increased from .15 to .31 as exposure increased from 200 to 3,000 msec for the small objects, and it increased from .15 to .22 for the large. The effect was significant within both levels of size [*t*₁(22) = 4.27, *p* < .001, and *t*₂(21) = 5.22, *p* < .0001, for the small objects, and *t*₁(22) = 2.38, *p* < .05, and *t*₂(21) = 2.94, *p* < .01 for the large objects]. There was also an interaction between size and MDA [*F*₁(1,22) = 33.02, *p* < .001, and *F*₂(1,36) = 4.22, *p* < .05], shown in the left panel of Figure 1. The form of the interaction is that described above.

Accuracy: Sup/Naming Task

The only significant effect was an interaction between size and MDA, similar to that found for the No-cue condition [*F*₁(1,22) = 35.92, *p* < .001, and *F*₂(1,36) = 6.33, *p* < .025]. This is shown in the right panel of Figure 1. There was no hint of an exposure effect (means = 46% and 48% for the 200-msec and 3,000-msec conditions).

Accuracy: Sup+Bas/Y-N Task

This condition produced a four-way interaction among size, MDA, positive/negative, and exposure [*F*₁(1,22) = 7.02, *p* < .025, and *F*₂(1,36) = 4.63, *p* < .05]. If one examines the size × MDA interaction at each combination of the other variables, no condition produces the pattern described above and shown in Figure 1. There were also two-way interactions involving exposure × positive/negative [*F*₁(1,22) = 6.34, *p* < .025,

Table 1
Mean Accuracy and Confidence by Exposure, Cue Condition, and Item Type

Cue Condition	200-msec Exposure								3,000-msec Exposure							
	Texture				Shape				Texture				Shape			
	Small		Large		Small		Large		Small		Large		Small		Large	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Accuracy																
NC pos	.11	0.11	.24	0.16	.21	0.14	.07	0.09	.32	0.24	.26	0.13	.38	0.16	.21	0.16
NC neg	.08	0.10	.27	0.10	.20	0.15	.02	0.06	.15	0.17	.28	0.16	.40	0.21	.12	0.20
SUP pos	.32	0.18	.50	0.18	.65	0.15	.43	0.26	.29	0.20	.46	0.23	.81	0.17	.47	0.23
SUP neg	.42	0.25	.43	0.14	.60	0.24	.33	0.23	.42	0.20	.50	0.23	.63	0.17	.30	0.23
YN pos	.55	0.25	.82	0.13	.64	0.24	.72	0.13	.83	0.16	.78	0.15	.85	0.13	.78	0.16
YN neg	.70	0.18	.73	0.18	.83	0.24	.77	0.28	.65	0.15	.75	0.12	.85	0.15	.55	0.26
Confidence																
NC pos	1.83	0.41	3.01	0.74	2.14	1.17	2.17	0.80	2.34	0.45	2.96	0.67	2.52	1.16	2.14	0.83
NC neg	2.18	0.82	2.32	1.59	2.10	0.61	1.70	1.14	2.01	0.55	2.32	1.25	2.55	1.25	1.78	0.73
SUP pos	2.78	1.11	3.28	0.60	3.24	0.93	2.81	0.64	2.64	0.79	2.90	0.78	3.51	0.80	2.85	0.25
SUP neg	2.83	0.77	3.37	1.03	3.48	1.06	2.60	1.27	2.48	0.72	3.22	1.34	3.38	1.30	2.60	0.96
YN pos	3.97	0.38	4.10	0.40	3.80	0.55	3.80	0.55	3.88	0.48	3.88	0.61	3.80	0.65	3.60	0.30
YN neg	3.62	0.46	4.25	0.51	4.15	0.40	3.58	0.35	3.63	0.24	4.17	0.47	4.10	0.38	3.32	0.31

Note—NC, No-Cue/Naming task; SUP, Sup-cue/Naming task; YN, Sup+Bas/Y-N task. Most diagnostic attribute is either texture or shape; size is either small or large; pos, positive; neg, negative.

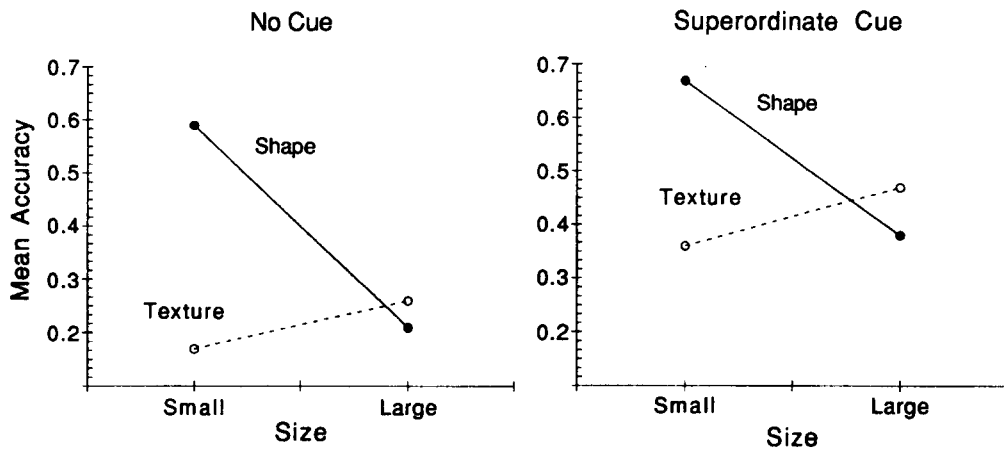


Figure 1. Effects of object size and MDA (texture, shape) on accuracy in the No-cue/Naming and Sup/Naming tasks.

and $F_2(1,36) = 10.03, p < .01$], and exposure \times size [$F_1(1,22) = 9.66, p < .01$, and $F_2(1,36) = 7.49, p < .01$]. Tests of the exposure effect at each level of positive/negative revealed that the positive items showed a significant exposure effect, the proportion correct increasing from .68 at 200 msec to .81 at 3,000 msec [$t_1(22) = 2.90, p < .01$, and $t_2(23) = 2.53, p < .05$], whereas the negatives showed no effect. The small objects showed an exposure effect, the proportion correct increasing from .68 to .80 [$t_1(22) = 2.87, p < .05$, and $t_2(21) = 2.81, p < .05$], but the large showed no effect. Although the three-way interaction among size, positive/negative, and exposure was not significant, examination of Table 1 reveals further that the effect of exposure is primarily on the items that are both small and positive.

Confidence

The ANOVA across cue conditions, along with the factors exposure, MDA, size, and positive/negative, showed only two significant effects across both subjects and items. One was the effect of cue [$F_1(2,66) = 60.29$ and $F_2(2,72) = 87.45, ps < .0001$], and the other was an interaction between MDA and size [$F_1(1,66) = 32.78, p < .0001$, and $F_2(1,36) = 14.76, p < .05$]. The effect of cue was due to a clear decline in confidence as the cue provided less information (mean = 3.86, 3.00, and 2.26 for the Sup+Bas/Y-N task, Sup/Naming task, and No-cue/Naming task, respectively). The MDA \times size interaction was of the same crossover type found previously: The confidence for texture items increased with size (from 2.85 to 3.32), whereas the confidence for shape items decreased (from 3.23 to 2.74). Although higher order interactions of size and MDA with other factors were not significant, the form of the interaction was somewhat variable across cue conditions. The analysis of the Sup/Naming task revealed the same type of MDA \times size interaction [$F_1(1,22) = 44.31, p < .001$, and $F_2(1,36) = 5.09, p < .05$]. The MDA \times size interaction was the only effect evidenced in the No-cue/Naming task, although it was marginal in the item analysis [$F_1(1,22) = 30.63, p <$

.001, and $F_2(1,36) = 3.08, p < .10$]. The upper panels of Figure 2 show these effects. The Sup+Bas/Y-N task produced an overall MDA \times size interaction as well [$F_1(1,22) = 25.29, p < .001$, and $F_2(1,36) = 8.44, p < .01$]; however, this was mitigated by a three-way interaction among MDA, size, and positive/negative [$F_1(1,22) = 25.75, p < .001$, and $F_2(1,36) = 4.83, p < .05$], shown in the bottom panels of Figure 2. The three-way interaction reflects the fact that the crossover interaction of MDA and size found in the other cue conditions was evident in the yes/no condition only when the positive items were considered.

Factor Analysis

Table 2 shows the results of a principal-components factor analysis performed on all variables in the study, with varimax rotation. Items were the units of observation in the input. The analysis accounted for 75% of the variance with three factors. The first received high loadings on all of the measures from the no-cue condition and the confidence ratings in the Sup/Naming task, with more moderate loadings from Sup-cue accuracy measures. The second had high loadings on all measures from the Sup+Bas/Y-N task and no others, and the third had high loadings from the free-exploration control and the Sup-cue accuracy measures, as well as somewhat lower loadings from the Sup-cue confidence measures. This analysis indicates that the ordering of the items with respect to identifiability is different in the Sup+Bas/Y-N task from the others.

Error Analysis

Subjects' errors in the Sup-cue and No-cue/Naming tasks took the form of confusion errors, in which the subject produced an object name that was not the same as that of the object presented. There were 579 error responses, which were analyzed in detail to determine the nature of the confusion, as follows. For each presented object, the proportion of error responses differing from it with respect to each of four dimensions was scored.

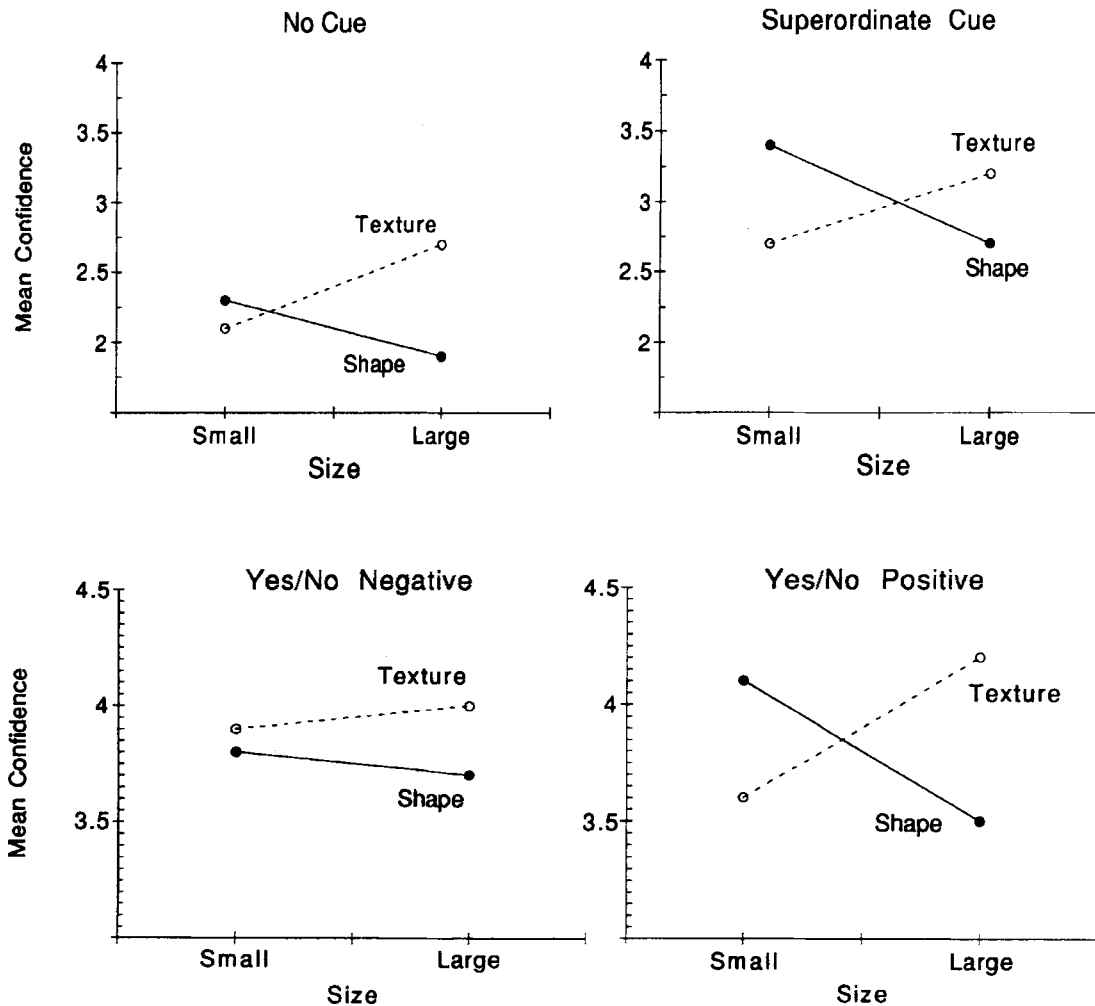


Figure 2. Effects of object size and MDA on confidence in each task.

The dimensions were as follows: rigid versus compliant, rough versus smooth, continuous surface versus 3-D edge (a bump, step, or hole within the interior of the region touched), and material (metal, glass or porcelain, plastic, paper, cloth, or "other"). For example, when the target object was a piece of corduroy and the subject responded "silk," it was a roughness mismatch; when the subject said "paper napkin," it was a material mismatch. A continuity mismatch occurred, for example, if the object was a stick of gum and the response was "paper clip."

The target object and each erroneous response were classified on each of the four dimensions (one object was judged unscorable on compliance and two on roughness), and a mismatch was counted on a given dimension if the dimensional classification of the two objects differed. Note that the same error response (i.e., erroneously named object) could contribute to mismatches on multiple categories; for example, it might mismatch the presented object on both roughness and compliance. Further details of scoring can be obtained from the authors. Two independent scorers achieved agreement of

91% on a subset of the responses if they were allowed to agree in advance about the classifications of the presented object on each of the four dimensions, and 87% if they did the scoring entirely independently.

Table 3 indicates the average proportion of the total error responses given to a presented object that mismatched along each dimension, by MDA and by size of the presented object. (The averages are over objects.) The majority of the error responses, though differing in name from the presented object, matched that object in all of the scored dimensions. Material mismatches were not uncommon, occurring on about 1/4 of the errors, but other categories of mismatch occurred in fewer than 10% of the cases.

There were clear default tendencies, or biases, in the properties of objects that subjects reported. These are indicated by Table 4, which shows the direction of the errors on each dimension by MDA. For example, the top left entry is the proportion of erroneous responses to rigid, texture-MDA objects that were compliance mismatches; these are responses that named a compliant object. The table clearly shows, regardless of MDA, that

Table 2
Results of Factor Analysis, Showing Loading of
Each Measure on Three Factors

Condition	Factor		
	1	2	3
No-Cue 200 Conf	0.83*	0.07	-0.09
No-Cue 3,000 Conf	0.90*	0.16	0.00
No-Cue 200 Acc	0.74*	0.40	0.32
No-Cue 3,000 Acc	0.74*	0.42	0.26
Sup 200 Conf	0.73*	0.14	0.35
Sup 3,000 Conf	0.72*	-0.04	0.50
Sup 200 Acc	0.43	0.32	0.71*
Sup 3,000 Acc	0.44	0.20	0.81*
Yes-No 200 Conf	0.21	0.77*	0.21
Yes-No 3,000 Conf	0.49	0.72*	0.15
Yes-No 200 Acc	0.03	0.74*	0.24
Yes-No 3,000 Acc	0.09	0.82*	0.10
Free	-0.16	0.29	0.77*
Variance accounted	33.3%	22.7%	18.9%

Note—The measures are labeled by cue condition, exposure duration, and whether the measure is confidence or accuracy. “Free” refers to the item rating from the free identification task. *Highest loading for variable.

there was a tendency to err in the direction of naming rigid, smooth, and continuous objects; these values appear to act as defaults. For example, 41% of the errors on texture-MDA, rough objects named objects judged to be smooth, and fewer than 1% of errors on texture-MDA, smooth objects named objects judged to be rough.

The default tendencies shown in Table 4 account for a pattern evident in Table 3. Compliance and roughness mismatches occurred more often with items for which texture was the MDA, whereas mismatches in continuity occurred more often with items for which shape was the MDA. This pattern reflects not only response biases, as seen in Table 4, but also asymmetries in the distribution of object properties across MDA. Specifically, shape-diagnostic objects were primarily rigid and frequently had a discontinuity in the form of a protrusion or hole, whereas texture-diagnostic objects were more often compliant and tended to have continuous surfaces. The errors in Table 3 are high in just those conditions where there are more objects with the nondefault value. For example, there were more opportunities to misperceive rough surfaces as smooth within the texture MDA, since the shape MDA had no rough objects.

Default values might reflect bottom-up or top-down biases; that is, they might correspond to the stimulus properties that are perceived most easily or the properties that are most common in the domain of real-world objects. The present “default values” of smooth, rigid, and continuous do not directly conform to predictions based on either of these processes, however. They are not identical to the dimensional values that produced “pop-out” in our haptic search task (Lederman & Klatzky, 1995), which included rough surfaces among smooth ones, edges raised from the finger plane among flat surfaces and flat surfaces among edges, and rigid surfaces among compliant surfaces and compliant among rigid.

Thus, being available with early perceptual pickup (i.e., popout) appears neither to preclude nor to guarantee a stimulus feature’s being a default value in the present paradigm. Although we do not have distributions of property values in real-world common objects, other data suggest that the present defaults are not the property values most available in memory. In previous norming data (unpublished), we asked subjects to generate as many objects as occurred to them having high, intermediate, or low values on a number of dimensions, including roughness, compliance, and shape complexity. Subjects generated about equal numbers of rough and smooth objects (46 and 50) and about equal numbers of simple and complex objects (56 and 58), but nearly twice as many rigid items as compliant ones (49 vs. 29). Thus the present defaults were not the obvious result of memory-based, top-down processing.

We also examined the pattern of dimensional mismatches as a function of cue condition and exposure. (These analyses used number of mismatches per object; because each object occurred in all cue and exposure conditions, normalization to control for differences in total error responses per object was not necessary.) Mismatches were greater for the No-cue condition than for the Superordinate-cue condition for all dimensions, as indicated in an ANOVA over objects with the factors of dimension (four levels), cue, and exposure, the only significant effects being that of cue [$F_2(1,43) = 17.28$], and dimension [$F_2(3,129) = 27.47$, $ps < .001$]. On the average, there were 1.45 errors with mismatches per object with no cue and .49 with a superordinate cue. The material mismatches were identical (2.7 per object) at the two exposures. Mismatches in compliance, roughness, and continuity all tended to decrease with exposure (averaging .59 and .41 per object at 200 and 3,000 msec, over the three dimensions), but not significantly; in an ANOVA over objects with the factors of dimension (three levels only, excluding material), cue, and exposure, the effect of exposure was $F_2(1,43) = 1.91$, $p = .174$, nor did the exposure \times cue interaction reach significance [$F_2(1,43) = 2.09$, $p = .155$].

DISCUSSION

The introduction to this paper raised a fundamental issue with regard to early haptic processing and object recognition: If basic level categories are largely differentiated on the basis of global object structure, what can

Table 3
Proportion of Total Erroneous Responses to a Target Object
That Mismatched on Each of Four Dimensions,
by Most Diagnostic Attribute (MDA) and Size

Dimension	MDA		Size	
	Texture	Shape	Small	Large
Compliance	.069	.025	.034	.060
Roughness	.088	.003	.048	.043
Continuity	.028	.086	.067	.047
Material	.254	.264	.253	.266

Table 4
Direction of the Errors on Each Dimension

Dimension	MDA	
	Texture	Shape
Rigid → compliant	.041	.028
Compliant → rigid	.154	.125
Rough → smooth	.410	NA
Smooth → rough	.006	.007
Continuous → 3-D edge	.020	.017
3-D edge → continuous	.333	.177

Note— $X \rightarrow Y$ indicates target objects having value X on a dimension, and for which the subject erroneously reported an object having value Y . The entry gives the proportion of the total error responses to objects with value X , and the given most diagnostic attribute (MDA) that named objects having value Y .

we know about an object's identity after initial local contact? One possibility is that we can know little, not enough to recognize the object. Alternatively, the data extracted under glance exposure may be sufficient, together with top-down expectations, to allow identification to occur. In that case, one can ask what the basis is for object identification by a haptic glance and under what circumstances it may occur. We now revisit the questions raised in the introduction.

How Good Is Object Identification With a Haptic Glance?

The present experiment indicated that identification of objects with a haptic glance can achieve a level of performance well above chance. A cue specifying the superordinate and basic-level category improved confidence relative to a superordinate-category cue, which in turn produced higher confidence than did no cue at all. Of course, the cue also increased absolute accuracy by virtue of increasing the guessing rate, but accuracy relative to guessing did not dramatically improve with a more informative cue: Accuracy was about 25% above chance in both the condition with no cue and that with the most informative cue, the basic-level name. In the absence of an advance cue, the highest accuracy—39%—was obtained with small items that could be identified by shape, under the longer exposure. In this case, the object fit into the "aperture" provided by the fingers. None of the cuing conditions approached the 93% accuracy achieved with free exploration (and no cue).

The present rates of object identification with spatially and temporally limited exposure are based on objects that were selected to be reasonably familiar and prototypical; however, other constraints applied, owing to our interest in diagnostic attributes and size. It is quite possible that objects selected specifically within other constraints, such as to maximize prior frequency of exposure, name familiarity, degree of association with the superordinate category name, or prototypicality of form, would yield higher rates. Those selected to be low on these attributes would presumably be less identifiable. The generality of the particular rates obtained here remains open to further research.

What Local Information Is Provided by Glance Exposure?

When data are sufficiently limited, an object's local geometry and material properties can be ambiguously perceived or misperceived. The analysis of confusion errors—responses naming objects other than the target—was directed at determining the extent of a subset of such misperceptions. It indicated that most erroneous responses matched the presented objects with respect to material and the three dimensions that were evaluated: roughness, compliance, and 3-D local shape variation. The mismatch levels for all coded categories were sufficiently low to indicate that a haptic glance allows substantial encoding of local structure and material properties. Thus this analysis confirms the conclusions of Lederman and Klatzky (1995), who used a speeded haptic search task.

When misperception of an object's material properties occurs, it presumably contributes to errors in judging the material from which it is made, producing confusions (e.g., between wood and glass). Mismatches between presented and reported objects with respect to material were not uncommon, constituting about 25% of error responses. A crude model might propose that material confusions occur whenever at least one material property, such as surface roughness, is misidentified. In that case, if we had scored all the dimensions critical to knowing an object's material, one would expect the sum of error rates on individual dimensions to meet or exceed the overall error rate for material (since an object producing an error on two dimensions contributes to the error rate for both, but would produce only one material error). In contrast, there is a considerably higher error rate for material than for the sum of the error rates on the material dimensions that we scored (roughness and compliance). At least by the model just described, this indicates that we have scored only a subset of the dimensions contributing to the identification of material. Indeed, this is not surprising, because our selection of which material dimensions to score was rather pragmatically based on those that were likely to give rise to reasonable reliability of scoring. For example, we did not include thermal conductivity, a dimension that is no doubt very important to material identification (see Katz, 1925/1989).

What Is the Relative Contribution of Structure Versus Material to Recognition Under Glance Conditions?

There was little overall difference between objects that had been identified as having texture as a diagnostic feature and those having shape. There was, however, an anticipated interaction between the object's most diagnostic attribute (MDA) and its size; the interaction was found in the majority of conditions across accuracy and confidence. The form of this interaction results from a difference in the implications of object size for texture-diagnostic and shape-diagnostic objects. For texture-diagnostic objects with homogeneous texture, a size in-

crease means that redundant information is available to the fingertips, which is found to improve performance. But when shape is diagnostic, a larger object means that critical shape cues are likely to extend past the fingertips, and accordingly, performance suffers.

Recognition of large, shape-diagnostic objects was not, however, at chance. Identification of such objects might reflect any of several mechanisms. The present data do not favor any one of these possibilities, and they are not mutually exclusive; different mechanisms might be used on different objects.

One possibility is that contact with a local shape sample is sufficient for one to infer the global shape of part or all of the object. For example, local contact with the rim of a bowl might provide enough information about width and radius of curvature for a person to infer its global shape (accuracy for bowl = 50% with no cue; 54% with superordinate cue). If an object has a small number of geons, *and* if they can be inferred from the local contact, object recognition might proceed much as it is described by recognition-by-components theory (Biederman, 1987; Biederman & Cooper, 1991, 1992; Hummel & Biederman, 1992).

A question for further research is whether local samples of a large (relative to the finger), but regular, shape such as a sphere are sufficient for someone to infer its global structure, providing an avenue for object recognition by means of volumetric primitives, even under severely constrained exposure conditions. Work of Kappers et al. (1994) described in the introduction casts some doubt on the ability to generate a complex global shape from a local, but predictive, sample.

The structure of an object and its division into parts will strongly influence the effectiveness of a local sample. Recognition of an object with a relatively large number of volumetric components, at least some of which were not accessible at the contact point, would depend on whether any geons that could be identified were sufficiently diagnostic of the object's identity. In the present study, this possibility was maximized by placement of the objects so that distinctive local features were contacted. In this case a local shape cue, such as the lip of a pitcher, could be sufficient to trigger identification (accuracy for pitcher = 8% with no cue; 13% with superordinate cue).

Another mechanism for identifying shape-diagnostic objects with a haptic glance is the encoding of material in compensation for loss of shape cues; for example, the distinctively waxy surface of a milk carton might be sufficient for recognition (accuracy for milk carton = 8% with no cue; 21% with superordinate cue).

How Important Are the Constraints of Space and Time in Glance Exposure?

The haptic-glance exposure condition is designed to produce a data-limited situation with respect to both time and space. Exposure time was manipulated here to investigate temporal limitations; object size (the scaling of the object to the fingertip) was manipulated to inves-

tigate spatial limitations. Recall that the interaction between object size and MDA reflects the fact that larger objects, which extend beyond the fingertip, lead to lower performance when shape is diagnostic and higher performance when texture is diagnostic, and thus it expresses effects of spatial limitations.

Exposure effects were found for accuracy in the No-cue/Naming task and for positive items (particularly small positives) in the Sup+Bas/Y-N task. The effect on the No-cue task was anticipated, under the assumption that longer exposures would compensate for the lack of advance information about the object category by providing more stimulus data. The effect in the Y/N task was unexpected. The fact that it was found primarily for small, positive items suggests that placement of the stimulus within the fingertip aperture allowed it to be confirmed, more so with longer exposures.

The effect of size on shape-diagnostic objects was observed in both accuracy and confidence measures, and across all tasks in at least one measure. The size \times MDA interaction also remained present at the longer exposure duration, suggesting that relaxation of temporal constraints did not compensate for spatial limitations.

Spatial constraints appear to produce somewhat more robust effects than do temporal constraints in these data. Limitations in the temporal constraints are indicated by the findings that in the factor analysis, conditions differing in exposure duration nevertheless loaded on a common factor, and that increased exposure duration did not significantly reduce the tendency for confusion errors to mismatch the target object on various dimensions. The limitations on exposure effects occurred despite the fact that exposure was manipulated over a 15:1 range, and the low exposure value was virtually as fast as could be implemented, given our apparatus. It should also be noted that even larger maximum exposure times, out to 5 sec, produced no greater exposure effect in pilot work.

Can Top-Down Processing Compensate for Data Limitations?

Haptic-glance presentation conditions produce a pattern-recognition context in which top-down processing, as initiated by advance cues, potentially plays an extremely important role. Overall, cuing increased confidence. (It also increased accuracy, but this effect is confounded with variations in guessing levels across the cues.) Our results indicate that as the amount of cuing in the situation increases, the impact of temporal limitations on haptic perceptual processing becomes less severe, as is indicated by the finding that the effects of exposure on accuracy were consistent across objects only in the minimally cued (No-cue/Naming) task. There were also indications that spatial limitations were alleviated by increased cuing, in that the size \times MDA interaction that reflects spatial constraints was absent for accuracy in the maximally cued (Sup+Bas/Y-N) task.

Cuing reduced the tendency for confusion-error responses in the naming tasks to mismatch the target on various material and structural dimensions. The ten-

gency for errors to match targets on the scored dimensions was greater for the superordinate cue than with no cue, which suggests that the superordinate cue allowed more information to be encoded from the stimulus even when full identification was not possible.

The effects of cuing can also be seen in the factor analysis, which assesses indications of different underlying processes under the various cuing conditions. The cuing conditions differ in the advance information provided by the cue, which is highest in the Sup+Bas/Y-N task and lowest in the No-cue/Naming task, and in the task itself, which in the Sup+Bas/Y-N task involves discrimination between the labeled and presented object and in the other conditions is open-ended identification. The analysis suggested that the maximally cued (Sup+Bas/Y-N) and minimally cued (No-cue/Naming) tasks were distinct from one another and the free-exploration control, with the Sup/Naming task appearing to share sources of variance with both the No-cue and free-exploration conditions. The free-exploration and Sup-cue conditions are similar in the potential for a substantial top-down component. Top-down influences are clearly made possible by the superordinate cue, and although the free-exploration control has no advance cue, we have found (Klatzky & Lederman, 1992) that the initial grasp of an object permitted during this condition provides enough information to form hypotheses and, potentially, to identify the object. Because the grasp functions like an initial cue in this sense, it is not surprising that the free identification measure and measures from the superordinate-cue conditions have high loadings on a common factor. The evidence for shared processes between the Sup-cue and No-cue/Naming tasks suggests that the two conditions are similarly driven by data limitations.

Conclusion

We began this paper by suggesting that people can commonly identify an object in their workspace after brief contact, without foveal vision. Our work helps to explain how such data-limited identification can occur. It indicates that when one to three fingers make contact with an object for approximately 200 msec, without active exploration but with hypotheses about what objects are likely to be present, individuals can determine, with substantial accuracy, what object they have encountered.

This work was motivated in part by an additional observation: Not only can people identify an object in their workspace without vision, but often, they can quickly go on to assume an appropriate manipulatory posture. Contact with the rim of a cup, for example, may be sufficient to orient the hand to grasp and lift it. This study demonstrates the availability of information from contact that could be used to initiate manipulation. Brief contact provides cues to local shape and/or material, which may be useful not only for object identification, but for determining current hand position in an object-centered representation. If that representation also provides information about the appropriate position for manipulation (along with other parameters such as the range of useful

grasping and lifting forces), the hand and arm movements necessary to arrive at the appropriate position for manipulation from the current position may be computed without ever viewing the object. How the information from a haptic glance is translated into effective manipulation is a focus for our future research.

REFERENCES

- BIEDERMAN, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.
- BIEDERMAN, I., & COOPER, E. E. (1991). Evidence for complete translational and reflectional invariance in visual object priming. *Perception*, **20**, 585-593.
- BIEDERMAN, I., & COOPER, E. (1992). Size invariance in human shape recognition. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 121-133.
- BIEDERMAN, I., & JU, G. (1988). Surface versus edge-based determinants of visual recognition. *Cognitive Psychology*, **20**, 38-64.
- HUMMEL, J. E., & BIEDERMAN, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, **99**, 480-517.
- KAPPERS, A. M. L., KOENDERINK, J. J., & LICHTENEGGER, I. (1994). Haptic identification of curved surfaces. *Perception & Psychophysics*, **56**, 53-61.
- KATZ, D. (1989). *The world of touch* (L. Krueger, Trans.). Hillsdale, NJ: Erlbaum. (original work published 1925)
- KLATZKY, R. L., & LEDERMAN, S. J. (1992). Stages of manual exploration in haptic object identification. *Perception & Psychophysics*, **52**, 661-670.
- KLATZKY, R. L., & LEDERMAN, S. J. (1993). Toward a computational model of constraint-driven exploration and haptic object identification. *Perception*, **22**, 597-621.
- KLATZKY, R. L., LEDERMAN, S. J., & REED, C. L. (1989). Haptic integration of object properties: Texture, hardness, and planar contour. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 45-57.
- KLATZKY, R. L., LOOMIS, J. M., LEDERMAN, S. J., WAKE, H., & FUJITA, N. (1993). Haptic identification of objects and their depictions. *Perception & Psychophysics*, **54**, 170-178.
- LEDERMAN, S. J., & KLATZKY, R. L. (1987). Hand movements: A window into haptic object recognition. *Cognitive Psychology*, **19**, 342-368.
- LEDERMAN, S. J., & KLATZKY, R. L. (1990). Haptic object classification: Knowledge driven exploration. *Cognitive Psychology*, **22**, 421-459.
- LEDERMAN, S. J., & KLATZKY, R. L. (1993). Extracting object properties through haptic exploration. *Acta Psychologica*, **84**, 29-40.
- LEDERMAN, S. J., & KLATZKY, R. L. (1995). *Early perceptual features for haptic object processing*. Manuscript submitted for publication.
- LEDERMAN, S. J., KLATZKY, R. L., & REED, C. L. (1993). Constraints on haptic integration of spatially shared object dimensions. *Perception*, **22**, 723-743.
- LOOMIS, J. M., & LEDERMAN, S. J. (1986). Tactual perception. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of human perception and performance* (chapter 31, pp. 1-41). New York: Wiley.
- MCCLELLAND, J. L., & RUMELHART, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, **88**, 375-407.
- MOORE, T., BROEKHOVEN, M., LEDERMAN, S. J., & ULUG, S. (1991). Q'Hand: A fully automated apparatus for studying haptic processing of spatially distributed inputs. *Behavior Research Methods, Instruments, & Computers*, **23**, 27-35.
- ROSCHE, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.
- SOLOMON, H. Y., TURVEY, M. T., & BURTON, G. (1989). Perceiving extents of rods by wielding: Haptic diagonalization and decomposition of the inertia tensor. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 58-68.
- TREISMAN, A. M., & GELADE, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, **12**, 97-136.

IVERSKY, B., & HEMENWAY, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, **113**, 169-193.
 WOODWORTH, R. S., & SCHLOSSBERG, H. (1954). *Experimental psychology* (rev. ed.). New York: Holt.
 WURM, L. H., LEGGE, G. E., ISENBERG, L. M., & LUEBKER, A. (1993). Color improves object recognition in normal and low vision. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 899-911.

NOTES

1. The MDA was originally identified for the object as a whole. We cannot guarantee that it remained the most diagnostic feature when the part that was contacted by the subject is considered, but the point of contact, selected for discriminability, generally retained the fea-

ture's diagnosticity. Two possible exceptions are the zipper (MDA = texture), where the fingers were placed over the tab and may not have contacted the teeth, and the envelope (MDA = shape), where the subject contacted the flap, which may not have been a discriminable contour.

2. In two cases involving negative items in the Sup+Bas/Y-N task, the contrast between the named and presented item was below the basic level as defined by common name (shoe: sneaker vs. oxford, and ball: baseball vs. tennis ball). By allowing subjects to name items at the basic level in the Sup and No-cue/Naming tasks, the same level of contrast was not enforced. However, if a subject chose to name an item at the subordinate level, and the name was clearly inconsistent with the presented object (e.g., saying "sneaker" for oxford), it was counted as an error. Failures to detect confusions at the subordinate level, when scoring the naming tasks, are undoubtedly few.

APPENDIX

Items Used in the Experiment by Most Diagnostic Attribute (MDA), Size, and Positive/Negative Status

For negative items, the first entry is the one named during yes/no trials and the second entry is the object that was actually presented in all conditions.

MDA = Texture	MDA = Shape
Size = Small	Size = Small
Positive	Positive
zipper	button
cork	key
chalk	hook
raisins	paper clip
file	clothes hanger
scotch tape (sticky side up)	padlock
Negative (named/presented)	Negative (named/presented)
crayon/magic marker	battery/electrical plug
cereal/rice	fork/spoon
eraser/post-it (sticky side up)	pen/pencil
noodle/stick of chewing gum	chopstick/plastic knife
	staple/binder clip
Size = Large	Size = Large
Positive	Positive
sandpaper	bowl
paper	milk carton
dish towel	sandal
sweater	pitcher
graham cracker	cupcake tin
corduroy	plate
Negative (named/presented)	Negative (named/presented)
can/plastic container	file card/envelope
baseball/tennis ball	pie plate/bread pan
bread/rice cake	wallet/coin purse
soap/deoderant	ashtray/pipe
toilet paper roll/washcloth	paddle/baseball bat
sneaker/oxford shoe	