Journal of Experimental Psychology:
Human Perception and Performance
2001, Vol. 27, No. 2, 290–302

0096-1523/01/$5.00 DOI: 10.1037//0096-1523.27.2.290

Interaction of Scale and Time During Object Identification

Thomas Sanocki
University of South Florida

On each trial a target object and a fragment of the target (or a control stimulus) were presented briefly enough to be integrated together. The stimuli were masked, and identification accuracy was measured. The fragments were large or small in size scale, and were presented early in processing (fragment before target) or late in processing (fragment after target). When presented early, large-scale fragments tended to facilitate identification more than small-scale fragments, but when presented late, small-scale fragments facilitated more than large-scale fragments. Facilitation effects from common feature fragments supported the idea of a spatiotemporal dependency, in which the efficiency of processing a piece of information depends on other pieces of information that have been processed. This is a strong type of global-to-local processing and can be interpreted within a structural description framework.

Because object recognition is a process that is defined over time, the nature of its time course—its microgenesis—is of central importance. Perhaps the most parsimonious view is that pieces of critical information are extracted and processed independently of each other during microgenesis (e.g., Estes, 1978; Navon, 1991; Parker, Lishman, & Hughes, 1996; Sanocki, 1990; Shibuya & Bundesen, 1988; Townsend, 1981). An alternative hypothesis is that there are dependencies—that the interpretation of one piece of information is affected by other pieces or configurations of information that have been processed. Beginning with the assumption that objects are represented at more than one level of scale, a number of researchers have suggested that processing of features at a smaller, finer, or more local level of scale may be affected by prior processing of features at a more global level (e.g., Broadbent, 1977; Navon, 1977; Neisser, 1967, 1976; Norman & Bobrow, 1976; Palmer, 1975; Sanocki, 1993). For example, large-scale features could provide a framework within which smaller scale features could be integrated, and this framework could reduce computational complexity (Sanocki, 1993, 1997, 1999; see also Broadbent, 1977; Navon, 1977; Norman & Bobrow, 1976; Palmer, 1975). This is an interaction between levels of scale in which the efficiency with which a piece of information is processed varies, depending on its spatiotemporal context. The purpose of the present article is to develop this idea and report results consistent with it. Scale was defined in terms of the size of features.

This research was supported by Grant DBS-9213246 from the National Science Foundation, and by a Research and Creative Scholarship Award and a Sabbatical Award from the University of South Florida. I thank Mike Braunstein, Marc Carter, Pierre Jolicoeur, Glyn Humphreys, Dennis Parker, and Chris Woodard for helpful comments; and Beth Chapman, Matthew Guida, Chris Guy, Mark Menke, Chuck Sarchione, Darcy Welsh, and Carole Young for their help in conducting the experiments.

Correspondence concerning this article should be addressed to Thomas Sanocki, Department of Psychology, University of South Florida, BEH 339, Tampa, Florida 33620-8200. Electronic mail may be sent to Sanocki@chuma.cas.usf.edu.

Relevant Prior Research

The experiments were conducted with an immediate priming method, and previous experiments involving primes at differing levels of scale are most relevant. Parker et al. (1996) used primes in a sequential same–different matching task with faces and objects. Scale was defined in terms of spatial frequencies. The primes were high-pass (high spatial frequency) or low-pass (low spatial frequency) components of target images and were presented immediately before the second (target) stimulus. Parker et al. were testing for a possible coarse-to-fine sequence in identification, and they assumed that if coarse information is of higher processing priority during initial processing, then the coarser low-pass primes should be more effective than high-pass primes. Contrary to predictions, Parker et al. found that high-pass primes were somewhat more effective than low-pass primes in speeding the same–different judgments. Parker et al. concluded that high-frequency information is more ecologically important than low-frequency information throughout the time course of processing.

There is a major problem with the interpretation of the Parker et al. (1996) experiments, however. The high- and low-pass primes were not equated on the utility of the information they contained: High-pass primes may have contained more of the information that is necessary for the same–different judgments than low-pass primes. Distinctive information might be critical in this situation, and high-pass primes may provide more of it. Indeed, Parker et al. argued that the high-pass information was more ecologically valid, and when presented alone the high-pass primes were responded to more quickly than low-pass primes (Experiment 4). Therefore, advantages for high- over low-pass primes could be attributed to utility of the information provided by the primes. Low-pass primes may have facilitated subsequent processing in a coarse-to-fine interaction, but the effect could have been less than the effect of the increased utility of high-pass prime information.

In the domain of scene perception, Schyns and Oliva (1994) obtained evidence consistent with the idea of scale-based changes in processing over time. They used hybrid scenes obtained by overlaying high-pass components of one scene over the low-pass components of another scene. The main finding was that observers were more likely to identify the hybrid as the low-pass scene when processing time was short (30 ms) but as the high-pass scene when processing time was long (100 ms). This result is consistent with the idea that lower spatial frequency information is relatively important early in processing and that higher spatial frequency information subsequently becomes important. The results do not necessarily implicate dependencies between levels of scale, however, because processing at the two levels could proceed independently. Lower spatial frequency information could be processed with a fast time course, being predominant early in processing, whereas higher spatial frequency information could be processed with a slower time course, becoming predominant later in processing. A number of previous results are consistent with an early predominance of lower spatial frequency information (e.g., Krueger & Chignell, 1985; Lupker, 1979; Townsend, Hu, & Kadlec, 1988).

In more recent work, Oliva and Schyns (1997) used the hybrid spatial frequency stimuli in a priming paradigm. They obtained evidence that identification does indeed proceed at both scales independently, and that high-frequency information can be dominant in some cases. However, the high- and low-frequency primes were not equated on utility of information, so advantages for high-frequency primes can be attributed to differences in the usefulness of the information they provided. In addition, the hybrid stimuli present conflicting information to the visual system—conflict between the interpretations (scene categorizations) of higher and lower frequency components—which may alter observers' strategies. Because the levels conflict, facilitory interactions between levels are likely to be discouraged.

In summary, when scale has been defined in terms of spatial frequency, no clear evidence of interactions between levels of scale has been found. The strongest results (Schyns & Oliva, 1994) can be explained in terms of independent processing at two levels of scale with differing time courses. The idea of a dependency between levels is stronger than this—it requires that processing at one level effect processing at the other level, increasing its efficiency under the appropriate conditions.

One possible reason for the lack of evidence for dependencies may be that scale should be defined in a manner other than spatial frequency. Although there is extensive evidence that the visual system selectively processes spatial frequency information (e.g., DeValois & DeValois, 1990; Graham, 1989), spatial frequency is not a reliable basis for parsing scenes and objects into functionally relevant components such as objects and parts because a global analysis represents an entire, uninterpreted image (Palmer, 1999; Pinker, 1984). There are other methods for defining scale that may be important in high-level vision. Many approaches to object perception assume that high-level vision decomposes objects hierarchically into levels of scale defined by structural properties rather than by spatial frequency (e.g., Biederman, 1987; Hummel & Biederman, 1992; Marr & Nishihara, 1978; Oden, 1977; Palmer, 1977; Quinlan, 1991).

Various properties have been proposed as the basis for structural decomposition. In some models, structural units are segmented at points of deep concavities (Hoffman & Richards, 1984), and the units themselves (parts) are subsequently defined by properties such as symmetry, closure, and volumetric edge relations (e.g., Biederman, 1987). In such approaches, the whole object is the largest scale unit, followed by parts (which may subdivide into large and small size; see Biederman, 1987), followed by properties of parts such as edges, vertices, and symmetries. A similar but nonidentical decomposition is based on size scale: Objects are first decomposed into their largest structures (outline or overall shape) and then into smaller components and markings (see Sanocki, 1993). An important benefit of the size-scale decomposition is that larger structures can provide frames for smaller structures, helping to localize and resolve them (Sanocki, 1997, 1999).

Some partial evidence of dependencies between levels of size scale was provided by Sanocki (1993, Experiment 4). In the experiment, observers identified brief stimulus presentations as one of six objects. The presentations consisted of a target and either a fragment of the target or a control stimulus. The fragments were either large or small in scale (relatively small, localized interior features or imprecise, large, exterior features). Temporal relations were manipulated by presenting the fragments either before or after the target. When the fragments were presented early in processing (fragment–target order), large fragments were more effective than small fragments, whereas the opposite occurred when the fragments were presented late in processing (target–fragment order). This result is consistent with the idea that large-scale information is important early in processing and that small-scale details are important later in processing. Because each type of fragment was more effective within one spatiotemporal context, the interaction cannot be attributed to differences in the utility of the information the fragments provided. The usefulness of the information depended on when the fragments were presented.

However, there were limitations with this experiment. First, the results could be explained by independent processing at large and small scales with differing time courses. (This type of explanation is explicated and contrasted in more detail after Experiment 1 is reported.) Second, there were differences in baseline performance rates between the fragment–target and target–fragment orders, which could influence the magnitude of the effects. Third, the large-scale fragments were imprecise—that is, they were made by duplicating lines from the targets and then separating them by small amounts around the original lines (a manipulation motivated by the idea of coarse-to-fine processing). Imprecision could interact with the time course of processing. That is, imprecise fragments might be tolerated in the fragment–target order but could interfere with higher resolution late processing in the target–fragment order. A fourth problem was that the results were limited to particular masking conditions. One purpose of the present experiments was to explore the robustness of these effects.

The Present Experiments

As shown in Figure 1, each trial began with a fixation point and (in Experiments 1 and 2) a mask. Then a fragment and a target were presented successively in the same parafoveal position, followed by the mask and then a forced-choice identification response. (Figure 1 illustrates a large-scale fragment followed by a target.) The stimuli were presented briefly enough to be integrated together (Eriksen & Collins, 1967). The forced choice involved several similar alternatives, and accuracy was measured. I used the two possible fragment–target orders to manipulate temporal relations. There were large size-scale fragments, small size-scale fragments, and a control stimulus that provided no identity-relevant information. None of the stimuli were imprecise. The main concern was the extent to which the fragments increased identification

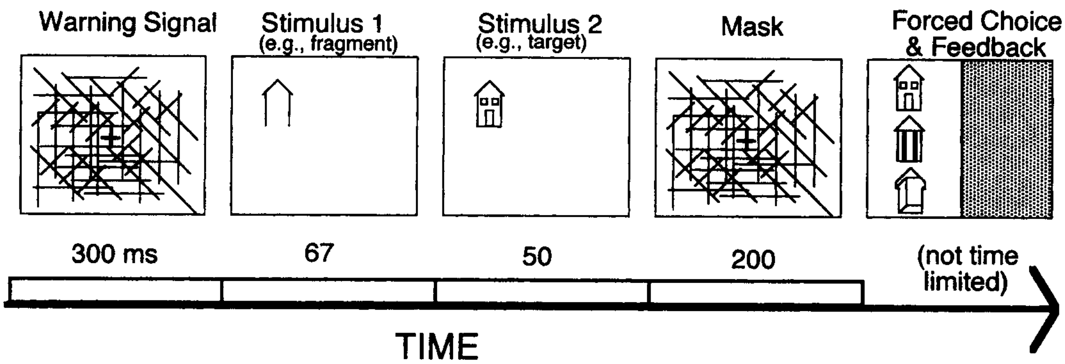

Figure 1. Events on each trial of the experiments. The durations of Stimulus 1 and Stimulus 2 were adjusted during the experiment. The times shown are means from Experiment 1, fragment–target order with a short fragment duration.

accuracy relative to the control stimulus. Partial support for the idea of a dependency during object identification would be provided if the relative effectiveness of the fragments varied with their temporal order. Specifically, large-scale fragments should be relatively more effective in the fragment–target order (early in processing), whereas small-scale fragments should become more effective in the fragment–target order (late in processing). Such effects could not be attributed to utility of information in the fragments because the relative effectiveness would depend on temporal order. A further prediction that distinguishes the idea of a dependency from the hypothesis of independent processing with differing time scales is presented in Experiment 2.

Two other variables were manipulated in the experiments. First, the large-scale fragments were designed to emphasize 3D volumetric information or 2D image information. This manipulation can be relevant to differing approaches to defining the large-scale components of object recognition. One type of approach emphasizes the 3D, volumetric nature of structural units (e.g., Biederman, 1987), whereas another approach emphasizes the 2D, view-, or image-based nature of structural units (e.g., Bulthoff, Edelman, & Tarr, 1995). If the structural units of object identification are volumetric in nature, then fragments that emphasize properties crucial for a 3D, part-based interpretation (high-depth fragments) should be more effective early in processing (fragment–target order) than fragments that do not. In contrast, if identification is based on matching of 2D images or viewpoints (e.g., Bulthoff et al., 1995), then 3D information should not be crucial, and identification could be facilitated as strongly by low-depth fragments as by high-depth fragments early in processing. The other variable that was manipulated was the duration of the fragments within the two orders. The purpose of this manipulation was to examine the possibility of quantitative differences in how quickly large- and small-scale information is processed within early and late processing.

Experiment 1

Method

Stimuli. The six targets are shown in the top row of Figure 2. The targets were divided into two similar shape sets (buildings and vehicles) of three objects each. On a given trial, any of the targets was equally likely to occur, but the forced-choice identification alternatives were restricted to the three objects in the target's similar shape set.

Two large-scale and two small-scale fragments were generated from each target. The two small-scale fragments for a given object contained mutually exclusive, small, localized portions of interior segments that were unique to that object (Figure 2, rows 4 and 5). The longest dimension of a small-scale fragment was always less than 55% of the corresponding target object's longest dimension. The large-scale fragments were the same size as the target and consisted mainly of large-scale segments. There were two types of large-scale fragments (rows 2 and 3). Following Biederman (1987), high-depth large-scale fragments contained the vertices and segments necessary to define the largest 3D parts of the target. Low-depth large-scale fragments contained about the same total line length, but the segments did not define complete 3D solids. The average total line lengths were 168 mm in the high-depth large-scale fragments and 159 mm in the low-depth large-scale fragments. To validate the assumed differences in depth for these fragments, an independent sample of 36 observers rated the fragments for amount of depth. The means on a 6-point scale (6 = *high depth*) were 5.2 and 1.9, respectively ($p < .001$).

A small cross was used as the control stimulus. The neutrality of this stimulus is supported by previous experiments in which it produced performance very similar to two other control stimuli that differed substantially in size, shape, and total line length (Sanocki, 1993). These results imply that size is not crucial for the control stimulus.

At the viewing distance of approximately 70 cm, the largest dimensions of the targets were 2.9° (houses) and 3.4° (vehicles). The lines of the stimuli were approximately 1.3 min in width. A mask was used that consisted of haphazardly arranged horizontal, vertical, and oblique lines approximately 2.6 min thick, with a fixation cross in the center.

Procedure and design. The experiment was controlled by a Macintosh SE microcomputer (Apple Computer, Cupertino, CA) with a 9 in, 60 Hz monitor. As shown in Figure 1, each trial began with that mask containing the central fixation cross (300 ms), which was replaced by the target and the fragment (or control stimulus). The target and fragment appeared successively, in either of the two possible orders. The stimuli were then replaced by the mask (for 200 ms), and then a 100-ms blank stimulus followed by the forced-choice display. The stimuli were transferred to the screen within 20 ms, by the same transfer routine, involving identical machine operations, in all conditions (the transfer routine was based on Rensink, 1990). The target, fragment, and control stimulus were centered 4.3° from fixation on the same randomly chosen corner of an imaginary square. Position of the fragment relative to the rest of the target was preserved. Observers were instructed to begin each trial by focusing on the fixation cross and were required to respond on each trial. No mention was made of the fragments. As noted, the forced-choice alternatives were

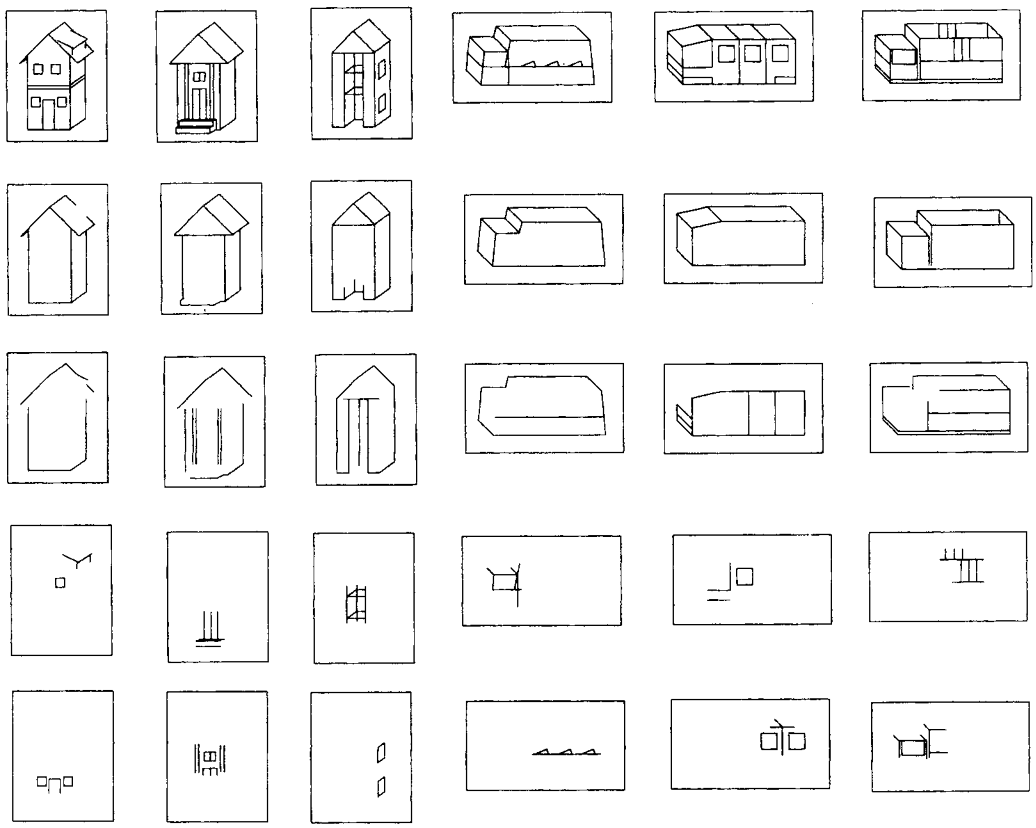

Figure 2. Target objects and primes in Experiment 1. Targets are in the top row, followed by high-depth and low-depth large fragments, and then two rows of small fragments. The enclosing rectangles were not presented to participants.

limited to the three objects from the target's similar shape set. The position of the three alternatives was constant across trials. Observers responded by clicking a mouse in the region corresponding to the object they thought occurred. Auditory feedback was given after the observer's response to signal the correctness of the response.

The target duration is termed the *base duration.* It was initially 67 ms and was adjusted for each observer to maintain an intermediate level of performance throughout the experiment. There were two fragment durations for each observer: the base duration and the base duration plus a fragment increment. For one group of observers, the fragment increment was 17 ms, and for the other it was 34 ms. In addition, on every trial 17 ms were added to the duration of the first stimulus to offset the effects of the forward mask. In the previous experiment (Sanocki, 1993, Experiment 4), performance on control-stimulus trials was higher for control–target sequences than for target–control sequences (Sanocki, 1993), suggesting that the forward mask decreased the effects of the first stimulus by delaying its perception. The additional time for the first stimulus in the present experiment was assumed to absorb most of those effects, making the effective stimulus durations of the first and second stimuli approximately equal.

Trials were organized in blocks of 144. There were three test blocks in the experiment, preceded by an additional 60 practice trials. After practice and after each block of trials, the base duration was adjusted. If level of accuracy for the preceding block was outside of limits of 45% and 75%, the duration was increased or decreased 17 ms, within fixed limits of 34 and 100 ms. The average mean base duration was 50 ms.

Within each block, each combination of target object and stimulus condition occurred once in a randomly determined order. There were 24 stimulus conditions, defined by the 2 stimulus orders, 6 fragment types, and the 2 fragment increments. The fragment types were control 1, control 2 (identical to control 1), high-depth large scale, low-depth large scale, local 1, and local 2. In the main analyses, the data were collapsed within the control-stimulus, large-scale, and small-scale fragment conditions. There were 72 trials in each of the main six conditions (defined by stimulus order and large vs. small vs. control fragment). Observers were given short breaks every 36 trials within a block.

Observers. Twenty-six observers served in each fragment increment group. They were undergraduates enrolled in introductory psychology courses at the University of South Florida and received course credit for their participation. Five additional observers were replaced because of low performance levels at the end of the experiment—either their ending target duration was at the upper limit, or their performance level in the last two blocks was close to chance levels (<45%).

Results

Scale and time. The main concern was the relative amount of facilitation for large-scale and small-scale fragments during the two temporal orders. Facilitation was defined as the difference in accuracy between a fragment condition and the corresponding control condition. These effects are shown in Figure 3a. Accuracy in control conditions is reported in the figure along the abscissas. Planned comparisons were conducted with analysis of variance (ANOVA), and an alpha level of .05 was used.

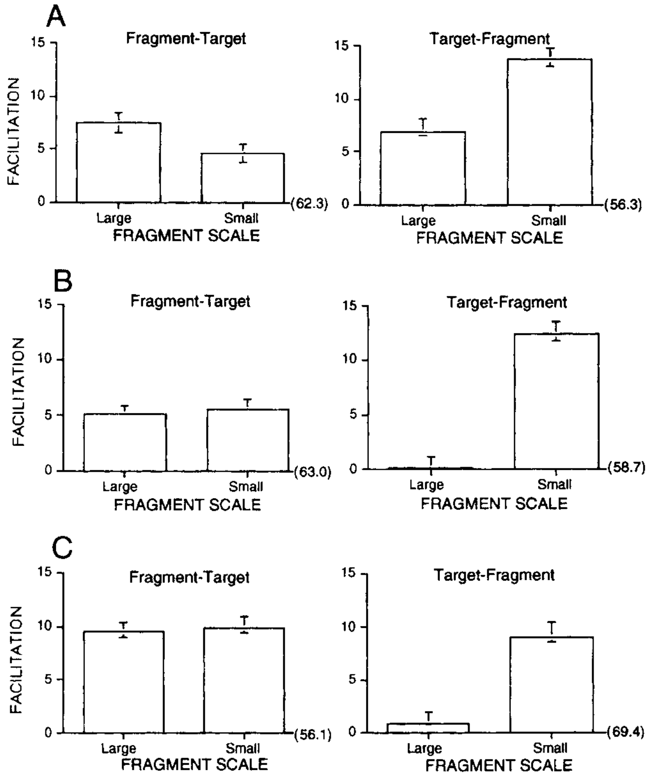

Figure 3. Facilitation effects (percentage correct with fragment minus control) and standard errors in each stimulus order for Experiments 1, 2, and 3 (A, B, and C, respectively). Control percentages are reported along the *x* axes.

Planned comparisons indicated that each of the four facilitation effects shown in Figure 3a was reliable ($ps < .001$). More important, an overall qualitative change was obtained: Large-scale fragments were more effective than small-scale fragments in early processing, whereas small-scale fragments were more effective late in processing. Interactions involving scale were assessed with the large- and small-scale fragment data only (the control data were omitted). The interaction between scale and time was reliable, $F(1, 50) = 53.10, p < .001$, for the interaction of large versus small fragment and stimulus order. The interaction between scale and order was similar across the two fragment duration groups, $F(1, 50) < 1$ for the three-way interaction of scale, order, and group. Performance with large-scale fragments was higher than with small-scale fragments in the fragment–target order, $F(1, 50) = 10.16$, $p < .001$; the opposite occurred in the target–fragment order, $F(1, 50) = 42.38, p < .001$.

Role of 3D shape information. The effects of high- and low-depth large-scale fragments early in processing are relevant to different approaches for defining the large-scale structural units in object recognition. The data for the two types of large-scale fragment were analyzed separately within each temporal order, and the means and facilitation effects are shown in Table 1. In the fragment–target order, performance levels were the same for high- and low-depth fragments, $F(1, 50) < 1$. This (initial) result does not provide unique support for either volumetric or view-based approaches.

When presented late in processing (target–fragment order), there was a marginal advantage (6.1% in accuracy) for the low-depth over high-depth fragments, $F(1, 50) = 2.78$, $p = .10$. The facilitation effect for the low-depth fragments was reliable ($p < .01$), but the effect for high-depth fragments was not. It is possible that some large scale features of the low-depth fragments reinforced relations involving distinctive features provided by the initial target. Alternatively, there may have been more distinctive information in the low-depth fragments than in the high-depth fragments, with the effects of the information being strongest during high-resolution late processing.

Effects of fragment duration. The duration of the fragments was varied in order to examine possible differences in the rates with which large- and small-scale information was processed during early and late processing. The amounts of facilitation caused by the large- and small-scale fragments increased with the duration of the fragments, $F(1, 50) = 4.37$, $p = .04$, for the interaction of fragment duration and control stimulus versus large- and small-scale fragments. The facilitation effects and associated mean accuracy levels are reported in Table 2. These data are collapsed over fragment duration group because there were no interactions involving group and fragment increment.

An omnibus ANOVA was conducted with the large- and small-scale fragment data to examine whether the fragment-duration effects varied as a function of scale or order. They did not; none of the interactions involving fragment duration approached reliability. As can be seen in Table 2, the facilitation effects and their increases with duration (duration differences in the table) were generally small and constant across scale and order, with one

Table 1
Results for Large-Scale High-Depth Versus Low-Depth Fragments in Each Experiment

Order and depth	Fragment	Control	Facilitation
	Experiment 1		
Fragment–target order			
High-depth fragment	73.1	62.3	10.8
Low-depth fragment	73.1	62.3	10.8
Target–fragment order			
High-depth fragment	59.9	56.3	3.6
Low-depth fragment	66.0	56.3	9.7
	Experiment 2		
Fragment–target order			
High-depth fragment	69.4	63.0	6.4
Low-depth fragment	67.1	63.0	4.1
Target–fragment order			
High-depth fragment	58.1	58.7	−0.6
Low-depth fragment	59.6	58.7	0.9
	Experiment 3		
Fragment–target order			
High-depth fragment	63.2	56.1	7.1
Low-depth fragment	68.0	56.1	11.9
Target–fragment order			
High-depth fragment	66.7	69.4	−2.7
Low-depth fragment	73.8	69.4	4.4

Note. Facilitation is the difference between large-scale fragment and the corresponding control stimulus.

Table 2
Effects of Fragment Duration in Each Condition of Each Experiment

Order and scale	Fragment	Control	Facilitation	Duration difference
	Experiment 1			
Fragment–target order				
Short large-scale fragment	67.5	60.6	6.9	
Long large-scale fragment	74.0	64.0	10.0	3.0
Short small-scale fragment	65.0	60.6	4.4	
Long small-scale fragment	70.9	64.0	6.8	2.5
Target–fragment order				
Short large-scale fragment	62.8	57.2	5.7	
Long large-scale fragment	63.5	55.4	8.1	2.4
Short small-scale fragment	67.8	57.2	10.6	
Long small-scale fragment	72.3	55.4	16.9	6.4
	Experiment 2			
Fragment–target order				
Short large-scale fragment	66.7	61.1	5.6	
Long large-scale fragment	69.6	64.9	4.7	−0.9
Short small-scale fragment	66.2	61.1	5.1	
Long small-scale fragment	71.0	64.9	6.1	1.0
Target–fragment order				
Short large-scale fragment	60.7	59.3	1.4	
Long large-scale fragment	56.9	58.1	−1.2	−2.6
Short small-scale fragment	68.3	59.3	9.0	
Long small-scale fragment	73.9	58.1	15.8	6.8
	Experiment 3			
Fragment–target order				
Short large-scale fragment	66.7	56.9	9.7	
Medium large-scale fragment	65.5	54.5	11.0	
Long large-scale fragment	64.8	56.9	7.8	−1.9
Short small-scale fragment	65.8	56.9	8.9	
Medium small-scale fragment	64.2	54.5	9.7	
Long small-scale fragment	68.1	56.9	11.2	2.3
Target–fragment order				
Short large-scale fragment	69.4	69.8	−0.4	
Medium large-scale fragment	70.7	70.3	0.4	
Long large-scale fragment	70.7	68.2	2.5	2.9
Short small-scale fragment	76.9	69.8	7.1	
Medium small-scale fragment	80.0	70.3	9.7	
Long small-scale fragment	78.6	68.2	10.4	3.3

Note. Facilitation is the difference between accuracy with fragments and accuracy with the corresponding control stimulus. Duration effect is the difference in facilitation between the longest and shortest duration.

exception: The rather large increase in facilitation for small-scale fragments late in processing (target–fragment order). However, this difference did not produce any reliable interactions.

Discussion

The main result was that an interaction between scale and time was observed—large size scale fragments were more effective than small size scale fragments when presented in the fragment–target order, whereas small-scale fragments were more effective when presented in the target–fragment order. The results cannot be attributed to utility of information in the fragments because each fragment was more effective in one spatiotemporal context—utility was time-dependent.

Although an interaction between scale and time was obtained, the data for the large-scale fragments pose a potential problem for the idea of a global-to-local dependency. One would expect that large-scale fragments would have less of an effect in later processing than in early processing. However, the effects in the two stimulus orders are generally similar for those fragments (see Figure 3a). An explanation of this is that the large-scale fragments may have contained some smaller scale distinctive features that were useful in later processing (although less than the small scale fragments). In Experiment 2 the large fragments contained no distinctive features, and therefore this explanation can be tested. The prediction is that the facilitation effects in later processing should be greatly reduced. Also, one complication of the present experiment is that the baselines provided by the control conditions were somewhat different in the two stimulus orders. The control level was somewhat higher in the target–fragment order (see the percentages listed along the abscissa of Figure 3a), leaving less room for facilitory effects of fragments. The level of baseline performance can influence the size of the facilitation effects, meaning that absolute comparisons between orders should be conducted with caution.

The effects of depth in the large-scale fragments were mixed, and the effects of fragment duration were generally small. These effects are discussed after additional data are provided in Experiments 2 and 3.

Of primary interest in this article is the scale by time interaction. One interpretation of the interaction is that it is caused by a spatiotemporal dependency, in which the processing and representation of one type of information affects the processing of another type of information. In particular, a global-to-local dependency can be characterized as follows (Sanocki, 1993, 1997, 1999). The general ideas are that early-arriving information is used to build an initial representational framework and that this framework aids the interpretation of later arriving information. Early large-scale information provides a large framework that can subsequently integrate small-scale information across the entire object. In later processing, a framework is typically available, and small-scale information is easily interpreted, providing details useful for discriminating between objects. Within this conceptualization, two specific types of dependency are produced in the present experimental design. In the fragment–target order, the effectiveness with which the details of the late target are processed depends on the usefulness of the framework provided by the fragment. The large-scale fragments provide the largest, most effective framework. Small-scale fragments provide a smaller framework that is partially useful, as well as some distinctive features that are processed somewhat inefficiently. The control stimulus provides no useful framework (or distinctive features). Thus, the differences in frameworks produce the ordering in accuracy between the three conditions. In the target–fragment order, an effective framework is provided by the target during the first interval. During the late fragment interval, the small-scale fragments provide highly distinctive information that is efficiently processed because a large-scale framework is available. The large-scale fragments provide less distinctive information and therefore have a smaller impact on performance. The control stimulus provides no distinctive information.

At least two kinds of alternative interpretations of the crucial interaction can be considered. In one explanation, the differential effectiveness of the large- and small-scale fragments is caused not by the fragments themselves but by their differential susceptibility to masking. This idea is examined in Experiment 3. A second explanation is that the interaction arises from a simple difference in the extraction rates over time for distinctive features at large and small scales. According to this *extraction-rate* explanation, the perceptual system extracts and processes distinctive features independently at the two levels of scale. Large-scale features are extracted at a high rate during initial processing, whereas small-scale features are extracted at a low rate initially and a higher rate later in processing. In the case of spatial frequency information, there is evidence that coarse (lower spatial frequency) information is processed faster than fine information (e.g., Breitmeyer, 1975). In the present case, the extraction of distinctive large-scale features could occur at a high rate relatively early, during the first stimulus interval. The extraction process could slow during the second stimulus interval, with the result that large-scale fragments would be more effective during the first interval than during the second. In contrast, extraction of small-scale features may begin slowly (during the first stimulus interval) and perhaps not increase to high rates until the second interval. Consequently, small-scale fragments would be more effective during the second interval than during the first.

Experiment 2

In the dependency hypothesis, a crucial effect of early large-scale information is that it can provide a framework for interpreting subsequent information. One way to distinguish this idea from the extraction-rate explanation is to use large-scale fragments that consist only of features common to the forced-choice alternatives. Such fragments can provide an appropriate framework for interpreting subsequent information, but they do not provide distinctive information that contributes directly to the forced-choice discrimination. Therefore, if common feature fragments facilitate identification, the effect can be attributed to the framework that they provide (see Sanocki, 1991a, 1993, 1997, 1999). Because large-scale common feature fragments do not provide distinctive features, the extraction-rate explanation predicts no effect of the fragments and consequently no Scale × Time interaction. (A statistical interaction would arise if there were differences over time for small-scale features but no effect at all for large-scale features, but this type of interaction would be trivial.) Thus, the main purpose of Experiment 1 was to use large-scale common feature fragments in a test for the Scale × Time interaction.

A second purpose of Experiment 2 was to further examine the importance of 3D volumetric and 2D view-based information early in processing. The large-scale fragments again varied in whether 3D or 2D features were emphasized. However, because the fragments consisted only of common features, there were no distinctive features that could contribute directly to performance.

Method

The target objects were modified somewhat to give objects within each similar shape set the same 3D shape. For each similar shape set, two large-scale fragments were created from segments common to the targets in the set. The stimuli are shown in Figure 4. For each shape set, one fragment was high depth and one was low depth. Following Biederman (1987), high-depth fragments were created by deleting contour mid segments while leaving all vertices, which are necessary for recovery of parts. Low-depth fragments were limited to exterior segments (silhouette) and thus did not contain vertices necessary for part recovery. The average total line lengths for high- and low-depth fragments were 114 mm and 106 mm, respectively. For each target object, two small-scale fragments were created from mutually exclusive sets of small scale internal distinctive segments, as in Experiment 1. The procedure was the same as before except that 30 undergraduates participated in each duration group. One additional observer was replaced because of a low level of performance at the end of the experiment. The average base duration of the stimuli was 54.4 ms.

Results and Discussion

Scale and time. The main result was an interaction of scale and time, $F(1, 58) = 65.56, p < .001$, for the interaction of large versus small Fragment × Stimulus Order. This interaction was constant across the two fragment duration groups, $F(1, 58) < 1$. The facilitation effects are shown in Figure 3b. Early in processing (fragment–target order), both large- and small-scale fragments caused facilitation effects ($ps < .001$), and the levels of perfor-

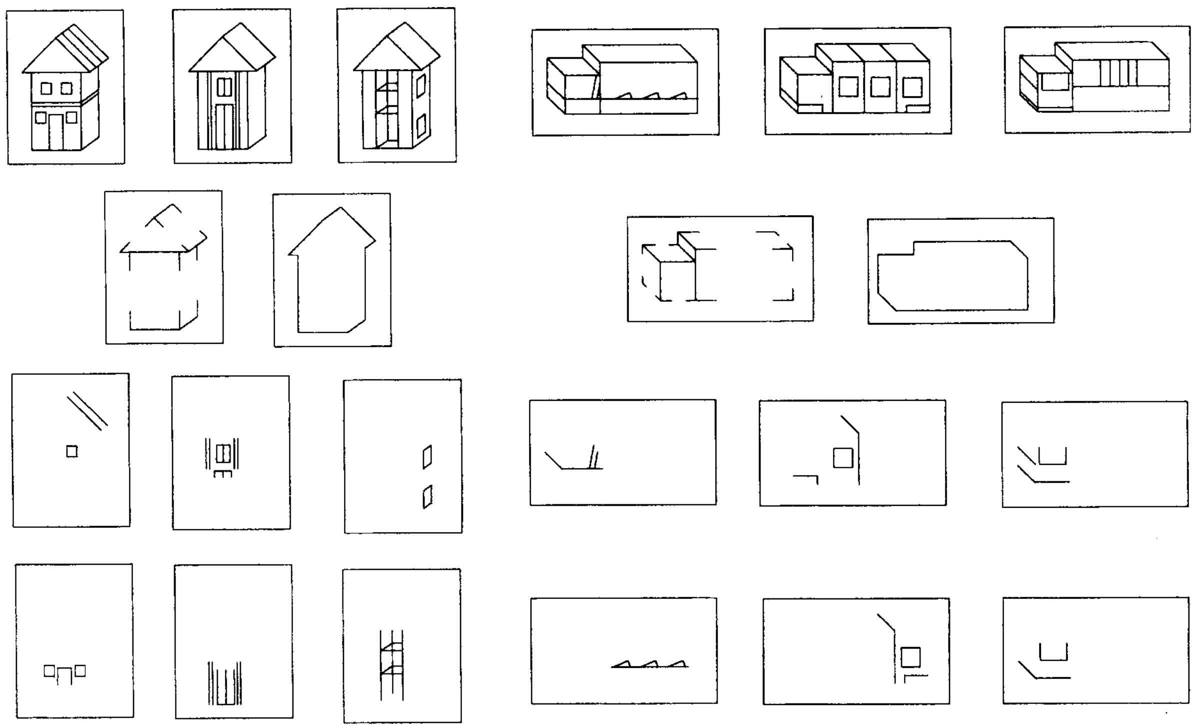

Figure 4. Target objects and fragments in Experiment 2. The targets are in the top row, followed by the high- and low-depth large fragments in the next row, and the small fragments in the bottom two rows. The enclosing rectangles were not presented to participants.

mance were similar for the two fragment types, $F(1, 58) < 1$. In later processing, only small-scale fragments caused facilitation ($p < .001$), and performance was higher with small-scale fragments than with large-scale fragments, $F(1, 58) = 139.80$, $p < .001$. Thus, the relative effectiveness of large- and small-scale fragments changed over time. Because the large-scale fragments provided no distinctive features, the interaction cannot be explained in terms of temporal differences in the rates with which large- and small-scale distinctive features were extracted.

The facilitory effect of early large-scale fragments can be explained by assuming that they provide a framework that aids processing of distinctive information in the subsequently presented targets. The lack of late facilitation for large-scale fragments suggests that the framework was not helpful after the target was presented. Presumably, the target would have already provided a framework, and the large-scale fragments could not increase performance by providing distinctive information because they contained no distinctive features.

The overall level of performance with early large-scale fragments was similar to that with early small-scale fragments. The similarity in performance for the two fragment types can be explained by the fact that small-scale fragments contained distinctive features that contributed directly in performance, whereas the large-scale fragments did not. The distinctive features could boost performance levels for small-scale fragments relative to common feature fragments. The distinctive features of small-scale fragments were even more effective in later processing, where they caused a very large advantage relative to large-scale fragments. Note that the facilitory effects for large-scale fragments decreased markedly over time. This result is consistent with the argument that the lack of such a decrease in Experiment 1 was due to the distinctive features in those large-scale fragments.

Role of 3D shape information. Facilitation effects for the high- and low-depth large-scale fragments are shown in Table 1. The early facilitation effect (fragment–target order) for high-depth fragments was 6.4%, which was 1.6 times as large as for low-depth fragments, $F(1, 58) = 4.70$, $p < .05$, for the comparison between high- and low-depth fragments. Both facilitation effects were reliable when considered individually ($ps < .001$). In later processing, neither type of large-scale fragment was effective, nor were there differences between the fragments, $F(1, 58) = 1.25$, $p > .20$.

The advantage for early high-depth fragments suggests that 3D information is especially important early in processing. This importance could be attributed to the information about volumes provided by the primes or perhaps to the ease of decomposing those primes into parts. However, the facilitation for low-depth fragments indicates that complete 3D information is not necessary for facilitation. Interpreted within the global-to-local conceptualization, these results suggest that an effective framework can be constructed from outline or 3D information, although 3D information may be more important.

Effects of fragment duration. In the present experiment, there were differences in the effects of fragment duration as a function of scale and time. In the omnibus ANOVA, on the large- and small-scale fragment data, fragment duration interacted with order, with fragment scale, and with order and scale, Duration × Order: $F(1, 58) = 5.70, p = .02$; Duration × Scale: $F(1, 58) = 17.40, p < .001$; and Duration × Order × Scale: $F(1, 58) = 12.23, p < .001$. As can be seen in Table 2, the apparent reason for these interactions is a large increase in facilitation in one condition—small-scale fragments in the target–fragment order. The increase in facilitation with fragment duration in this condition was reliable, $F(1, 58) = 9.53, p < .01$. This result suggests that small-scale information was processed very efficiently late in processing.

Experiment 3

The results of Experiment 2 distinguish the idea of a dependency from the idea of independent processing at two levels with differing time courses. The main purpose of Experiment 3 was to test the dependency hypothesis under different masking conditions. Large-scale fragments with common features were used, as in Experiment 2. According to a masking explanation of the previous results, the crucial interaction is caused by differential susceptibility of the fragments to masking. In particular, the interaction could be caused by small-scale fragments being especially susceptible to the initial (forward) mask. This susceptibility could reduce their effectiveness during the first stimulus interval relative to large-scale fragments and relative to small-scale fragments in the second stimulus interval.[1,2]

In order to test this idea, the initial mask was removed in Experiment 3. If differential susceptibility to masking is the cause of the Scale × Time interaction, then the interaction should be markedly reduced.

Also, there was some concern that differences in the relative levels of accuracy in the two stimulus orders may have contributed to the patterns of effects in the previous experiments. The baseline levels of performance provided by the control-stimulus conditions (reported along the abscissas in Figure 3) were 4% to 6% lower in the target–fragment order than in the fragment–target order. This leaves more room for improvement in that order, possibly increasing the size of facilitation effects and fragment-duration effects. To offset the baseline differences, the durations in target–fragment order were increased slightly relative to the fragment–target order.

Finally, to provide a more powerful within-participant assessment of fragment-duration effects, three levels of fragment duration were used with each participant. The levels of duration covered the same range as the pairs of groups in the previous experiments.

Method

The stimulus design was the same as in Experiment 2. Because the absence of a forward mask increased the perceptibility of the targets, the targets were modified by reducing the amount of distinctive features in each. This procedure was done to keep the target durations (which were determined by accuracy levels) similar to that of the previous experiments. The stimuli are shown in Figure 5. The large-scale high- and low-depth fragments were the same as in Experiment 2. Also as previously, two small-scale fragments were generated from each target (although the details differed from Experiment 2).

The experimental design was similar to the previous experiments with the exception that three rather than two fragment durations were used, and there was only one group of participants. There were two blocks of 216 trials, each consisting of every combination of the six targets and 36 stimulus conditions (3 Fragment Durations × 2 Orders × 6 Fragment Types). The procedure was the same as before except for the timing of the stimuli. Because there was no forward mask, no time was added to the first stimulus. An increment was added to the durations of both stimuli in the target–fragment order. This increment, as well as the base duration, was adjusted every 36 trials. The increment for the target–fragment order was initially 34 ms and was adjusted by 17 ms if the difference between the two orders (target–fragment − fragment–target) was outside of 0% and 12% limits. The increment averaged 5.7 ms during the experiment. The base duration of the stimuli averaged 71.2 ms. Twenty-four undergraduates participated.

Results and Discussion

Scale and time. The main result was an interaction of scale and time, $F(1, 23) = 10.84, p < .01$, for the interaction of large versus small Scale × Stimulus order. The facilitation effects are shown in Figure 3c. In the fragment–target order, both large- and small-scale fragments caused facilitation effects ($ps < .001$), and the levels of performance were similar for the two fragment types, $F(1, 23) < 1$. In later processing, only small-scale fragments caused facilitation ($p < .001$), and performance was higher with small-scale fragments than with large-scale fragments, $F(1, 23) = 34.21, p < .001$. Thus, the relative effectiveness of the primes changed over time, replicating the Scale × Time interaction without a forward mask. The early facilitation effects were similar for large-scale common feature fragments and small-scale fragments, as in Experiment 2. This result indicates that the large- and small-scale fragments were not differentially affected by the mask in Experiment 2. The results are inconsistent with the explanation of the Scale × Time interaction in terms of differential masking effects. The obtained interaction is consistent with the idea of a global-to-local dependency in object identification.

Compared with Experiment 2, the magnitudes of the early facilitation effects were greater in the present experiment for both large- and small-scale fragments. In fact, the facilitation effects for small-scale fragments were slightly larger in fragment–target order than in the target–fragment order. However, the level of accuracy in the control condition was 13% lower in the fragment–target order than in the target–fragment order (as intended). The lower baseline can increase the effects of early fragments because there was much more room beneath the ceiling to facilitate performance in that order. Note that the absolute accuracy level for small scale primes was 66.0% in the fragment–target order and 78.5% in the target–fragment order. Also, note that the lack of a premask in this experiment should increase the impact of the first stimulus, including the early fragments. Increased impact should result in larger facilitation effects.

Role of 3D shape information. Facilitation effects for the high- and low-depth large-scale fragments are shown in Table 1. Performance was higher with low-depth fragments than with high-depth fragments. In the fragment–target order, low-depth frag-

[1] I thank Pierre Jolicoeur for suggesting the possible importance of masking.

[2] A second logical possibility is that the large-scale fragments are more susceptible to the second (backward) mask and that the effects of large-scale fragments would be as large as those of the small-scale fragments if a backward mask was not used. This possibility seems unlikely because the large-scale fragments consisted only of common features.

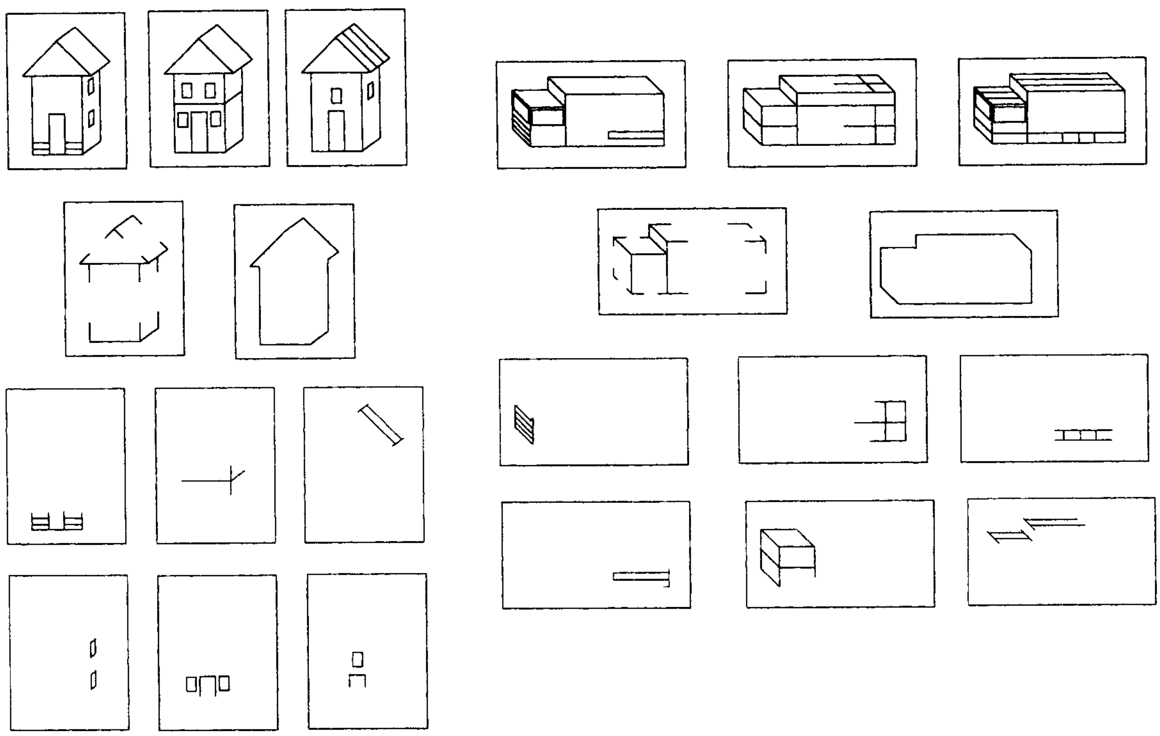

Figure 5. Target objects and primes in Experiment 3. The targets are in the top row, followed by the high- and low-depth large fragments in the next row, and the small fragments in the bottom two rows. The enclosing rectangles were not presented to participants.

ments caused an 11.9% facilitation effect, which was 1.7 times as much facilitation as high-depth fragments, $F(1, 23) = 7.73$, $p = .01$, for the comparison of low- vs. high-depth fragments. Both facilitation effects were reliable when considered individually ($ps < .001$), however. The finding of facilitation for both fragment types is consistent with the idea that both types of information can provide an effective early framework. However, the advantage for low-depth fragments contrasts with the results for the same stimuli in Experiment 2—in that experiment, early priming effects were greater for high-depth fragments.

The reversal in the early depth effects between Experiments 2 and 3 is somewhat surprising. Several differences between Experiments 2 and 3 may have contributed to the reversal of the depth effect. The most likely explanation appears to be the forward mask. Structural information is helpful in overcoming the effects of a randomly patterned mask (McClelland, 1978), and it is possible that the mask interfered most with the fragments that provided the less complete structure. The high-depth fragments provided a complete 3D structure—that is, the role of each segment of the fragment in a structural unit (3D part) was clear within the fragment itself. In contrast, the low-depth fragments provided a more ambiguous structure because the outline segments did not form simple 2D or 3D structural units; their role within a simple unit was not clear without the rest of the object (see Figure 4). Therefore, if structural completeness makes a fragment less susceptible to masking, the high-depth fragments would have been less affected by the mask of Experiment 2 than low-depth fragments.

Other differences between the experiments (including the relative timing and details of the stimulus sets) are also possible causes of the changes in the depth effect. However, there is no obvious reason why these factors would be important. Moreover, the similarity between experiments in the crucial Scale × Time interaction seems to argue that the timing and stimulus differences were not much of an influence on the nature of processing.

The effectiveness of both high- and low-depth fragments early in processing across the three experiments implies that the initial structural units of object recognition may be flexible in nature, involving either high-depth (volumetric or part-based) or low-depth information. The exact level of effectiveness of the differing types of fragments appears to vary with stimulus conditions. Future research on this issue should include comparisons across different stimulus conditions.

In the target–fragment order, there also was an advantage for low-depth fragments over high-depth fragments, $F(1, 23) = 11.45$, $p < .01$. Low-depth fragments caused a reliable facilitation effect ($p = .01$), but high-depth fragments did not. This facilitation effect late in processing for the low-depth fragments is rather surprising because the fragment contained no distinctive features. The same fragments did not cause facilitation in Experiment 2. The low-depth fragments in Experiment 1, which contained distinctive features, did cause facilitation late in processing. It is possible that in the present experiment some outline features reinforced relations involving distinctive features provided by the initial target and that a similar effect occurred in Experiment 1. The effect is not predicted by the global-to-local conceptualization, nor by the alternative explanations that are based on feature extraction rates or masking.

Effects of fragment duration. The effects of fragment duration were negligible in the present experiment. As can be seen in Table 2, the largest increase in facilitation with fragment duration was only 3%—the effect for late small-scale fragments. This effect did not approach reliability, $F(1, 23) < 1$ for the interaction of small scale versus control stimulus and short versus long duration. There also was no effect of fragment duration in omnibus ANOVAs on all of the data or on the large- and small-scale fragment data. Thus, there were no effects of fragment duration that were reliable across experiments.

General Discussion

The main result of this research was the interaction of size scale and time—the relative effectiveness of large size-scale and small size-scale fragments changed with processing order. In general, large-scale information was more effective early in processing, whereas small-scale information was more effective late in processing. The interaction was obtained across three experiments differing in a number of details. The results of Experiments 2 and 3 were inconsistent with an explanation in terms of independent processing at two levels of scale because large-scale common-feature fragments facilitated processing even though they contained no distinctive features. Experiment 3 was inconsistent with an explanation in terms of masking because the crucial interaction was replicated without an initial mask.

The results are important in two general ways. First, they provide positive evidence of the importance of large-scale information in early identification processing. Size was an effective definition of scale. The positive results obtained with the structural manipulation of scale contrast with those obtained with spatial frequency manipulations (Oliva & Schyns, 1997; Parker et al., 1996). Second, the results are consistent with the idea of a global-to-local dependency during identification, in which initial information provides a framework that aids the processing of subsequent, detailed information (e.g., Sanocki, 1991a, 1993, 1997, 1999). This framework produces changes in the efficiency with which information is processed, depending on the spatiotemporal context.

The microgenesis of processing is known to vary with the density of pattern elements (e.g., Kimchi, 1998), and it can vary with the stimulus conditions. For example, the exact effects of fragments in the present experiments differed with several factors, and the effects of depth in large scale primes varied with the masking conditions between Experiments 2 and 3. Accordingly, further experimentation examining the generality of the present results is in order. Nevertheless, the present results provide important positive evidence of the existence of scale interactions and global-to-local dependencies.

Common Versus Distinctive Features

Common feature fragments have been crucial in this research because they cause framework-based dependencies without directly influencing performance through distinctiveness. That is, common features can influence performance by influencing the processing of other information, but they cannot influence performance by providing distinctive information directly.

Of course, identification performance will ultimately depend on the processing of distinctive information. This information may be primarily high frequency in nature (Oliva & Schyns, 1997; Parker et al., 1996). The present claim is that such information is not processed independently of other information in the stimulus; the efficiency with which it is processed can vary depending on structural information that has been processed previously.

Constructing Structural Descriptions

A key idea in the global-to-local conceptualization is that initial information can provide a framework for localizing and integrating subsequent information. Integration is important because object identification involves a number of different types of relations. These include relations between major parts (e.g., Hummel & Biederman, 1992; Quinlan, 1991) and relations between elementary features such as line or edge segments (e.g., Enns & Prinzmetal, 1984; Pomerantz & Pristach, 1989; Sanocki, 1991b, 1997, 1999; Sanocki & Oden, 1991; Treisman & Paterson, 1984). Relations can be encoded within structural descriptions (e.g., Oden, 1977; Palmer, 1975).

There is evidence that functionally important relations are encoded between large- and small-scale structures that are proximal and connected (Palmer, 1977; Sanocki, 1997) and between structures that share an axis or reference frame (Sanocki, 1999). Large-scale fragments may provide more effective frameworks than small-scale fragments because they are involved in local relations across the object. When distinctive information is limited to one local region, early large-scale fragments are effective if they are proximal and connected to the distinctive information but not if they are low on proximity and connectedness (Sanocki, 1997). These later results were also inconsistent with explanations that emphasize externality versus internality rather than size scale (i.e., outside-in processing sequences; see, e.g., Earhard & Walker, 1985; McClelland & Miller, 1979). In the experiments, ineffective fragments were as external as the effective, structurally related fragments. The conclusion was that proximity and connectedness were more important than externality.

An additional advantage of large-scale structures may be that they have more redundancy than small-scale structures because their large size means they activate more receptors. This difference could also contribute to the speed with which they are processed.

Given that large-scale fragments can serve as frames for integrating information, one could ask if small-scale fragments could serve the same function. The present argument is that they might but not as effectively because of their small size. Sanocki (1997) examined a relatively small prime that was also highly structurally related; it had a small and unreliable facilitation effect.

Internal Structure Versus Attention-Shifting

Some early proposals for global-to-local processing emphasized the importance of global information for guiding the allocation of limited capacity attentional resources (e.g., Broadbent, 1977; Navon, 1977; Palmer, 1975). The shifting of attention does not appear to be related to the effects of internal object structure examined in the present situation. This conclusion is supported by results of other experiments that used the fragment–target order. The crucial conditions involved fragments that could attract attention to the location of the target but were not highly structurally related to the critical information in the targets (as measured by proximity and connectedness). Such fragments included several control stimuli (Sanocki, 1993, and the present control stimulus) and several fragments containing target information low in structural relatedness (Sanocki, 1997). The fragments did not facilitate identification when the temporal conditions were similar to the present ones, whereas highly structurally related fragments did. The ineffectiveness of the unrelated fragments implies that the shifting of attention to the target's location is not crucial in the present situation. Also, shifting attention to locations near the critical features was not crucial if there was no structural relation between fragment and target (Sanocki, 1997). Thus, structural relatedness was important but the guiding of attention was not. However, when processing time was lengthened by increasing stimulus durations, there were large facilitory effects for all of the fragments that could attract attention to the location of the target as a whole (Sanocki, 1997; Sanocki & Sellers, in press). Therefore, the properties that guide shifts of attention may be distinct from those involved in the construction of a structural representation.

References

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–147.

Breitmeyer, B. G. (1975). Simple reaction time as a measure of the temporal response properties of transient and sustained channels. *Vision Research, 15,* 1411–1412.

Broadbent, D. E. (1977). The hidden preattentive process. *American Psychologist, 32,* 109–118.

Bulthoff, H. H., Edelman, S., & Tarr, M. J. (1995). How are three-dimensional objects represented in the brain? *Cerebral Cortex, 5,* 247–260.

De Valois, R. L., & De Valois, K. K. (1990). *Spatial vision.* New York: Oxford University Press.

Earhard, B., & Walker, H. (1985). An "outside-in" processing strategy in the perception of form. *Perception & Psychophysics, 38,* 249–260.

Enns, J. T., & Prinzmetal, W. (1984). The role of redundancy in the object-line effect. *Perception & Psychophysics, 35,* 22–34.

Eriksen, C. W., & Collins, J. F. (1967). Some temporal characteristics of visual pattern perception. *Journal of Experimental Psychology, 74,* 476–484.

Estes, W. K. (1978). Perceptual processing in letter recognition and reading. In E. C. Carterette & M. P. Friedman (Eds.), *Handbook of perception* (Vol. IX, pp. 163–220). New York: Academic Press.

Graham, N. V. S. (1989). *Visual pattern analyzers.* New York: Oxford University Press.

Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition, 18,* 65–96.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99,* 480–509.

Kimchi, R. (1998). Dominance of configural properties in visual form perception. *Psychonomic Bulletin & Review, 5,* 135–139.

Krueger, L. E., & Chignell, M. H. (1985). Same-different judgments under high speed stress: Missing feature principle predominates in early processing. *Perception & Psychophysics, 38,* 188–193.

Lupker, S. J. (1979). On the nature of perceptual information during letter perception. *Perception & Psychophysics, 25,* 303–312.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organisation of three-dimensional shapes. *Proceedings of the Royal Society of London, B200,* 269–294.

McClelland, J. L. (1978). Perception and masking of wholes and parts. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 210–233.

McClelland, J. L., & Miller, J. (1979). Structural factors in figure perception. *Perception & Psychophysics, 26,* 221–229.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9,* 353–383.

Navon, D. (1991). Testing a queue hypothesis for the processing of global and local information. *Journal of Experimental Psychology: General, 120,* 173–189.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

Neisser, U. (1976). *Cognition and reality.* San Francisco: Freeman.

Norman, D. A., & Bobrow, D. G. (1976). On the role of active memory processes in perception and cognition. In C. N. Cofer (Ed.), *The structure of human memory* (pp. 114–132). San Francisco: Freeman.

Oden, G. C. (1977). Integration of fuzzy local information. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 565–575.

Oliva, A., & Schyns, P. G. (1997). Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology, 34,* 72–107.

Palmer, S. E. (1975). Visual perception and world knowledge: Notes on a model of sensory-cognitive interaction. In D. A. Norman, D. E. Rumelhart, & the LNR Research Group (Eds.), *Explorations in cognition* (pp. 279–307). San Francisco: Freeman.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology, 9,* 441–474.

Palmer, S. E. (1999). *Vision science.* Cambridge, MA: MIT Press.

Parker, D. M., Lishman, J. R., & Hughes, J. (1996). Role of coarse and fine spatial information in face and object processing. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 1448–1466.

Pinker, S. (1984). Visual cognition: An introduction. *Cognition, 18,* 1–63.

Pomerantz, J. R., & Pristach, E. A. (1989). Emergent features, attention, and perceptual glue in visual form perception. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 635–649.

Quinlan, P. T. (1991). Differing approaches to two-dimensional shape recognition. *Psychological Bulletin, 109,* 224–241.

Rensink, R. A. (1990). Toolbox-based routines for Macintosh timing and display. *Behavior Research Methods, Instruments, & Computers, 22,* 105–117.

Sanocki, T. (1990). *Effects of font- and letter-specific experience on letter recognition.* Unpublished manuscript.

Sanocki, T. (1991a). Effects of early common features on form recognition. *Perception & Psychophysics, 50,* 490–497.

Sanocki, T. (1991b). Intra- and inter-pattern relations in letter recognition. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 924–941.

Sanocki, T. (1993). Time course of object recognition: Evidence for a global-to-local contingency. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 878–898.

Sanocki, T. (1997). Structural contingencies and shifts of object-based attention during object recognition. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 780–807.

Sanocki, T. (1999). Constructing structural descriptions. *Visual Cognition, 6,* 299–318.

Sanocki, T., & Oden, G. C. (1991). Perceptual adjustments on representations of familiar patterns: Change over time and relational features. *Perception & Psychophysics, 50,* 28–44.

Sanocki, T., & Sellers, E. (in press). Shifting resources to recognize a forming object: Dependencies involving object properties. *Visual Cognition.*

Schyns, P. G., & Oliva, A. (1994). From blobs to boundary edges: Evidence for time and spatial scale dependent scene recognition. *Psychological Science, 5,* 195–200.

Shibuya, H., & Bundesen, C. (1988). Visual selection from multi-element displays: Measuring and modeling effects of exposure duration. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 591–600.

Townsend, J. T. (1981). Some characteristics of visual whole report behavior. *Acta Psychologica, 47,* 149–173.

Townsend, J. T., Hu, G. G., & Kadlec, H. (1988). Feature sensitivity, bias, and interdependencies as a function of energy and payoffs. *Perception & Psychophysics, 43,* 575–592.

Treisman, A., & Paterson, R. (1984). Emergent features, attention, and object perception. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 12–31.

Received October 5, 1998
Revision received February 23, 2000
Accepted May 31, 2000 ■