

Structural aspects of visual similarity

STEPHEN E. PALMER

University of California, Berkeley, Berkeley, California 94720

The hypothesis that visual representations for lines and/or points are independent structural units was tested using similarity judgment and speeded discrimination for pairs of six-segment letter-like figures. The stimuli were constructed such that each of two comparison figures had five segments in common with a standard figure. One figure was similar to the standard in its higher order structure (connectedness and closedness properties), whereas the other differed. The results show that the figures with similar higher order structure were systematically judged more similar to the standard than the figures with different structure. The former were also more difficult to discriminate from standards than the latter, as indicated by both time and error measurements. These effects were less pronounced in sequential than in simultaneous comparisons.

One basic representational issue for any theory of perception and pattern recognition is the nature and size of its structural units. For example, the basic units of template theory are representations of points and of their defining characteristics: location and color (cf. Palmer, in press). Each pattern is represented as a subset of all possible points, and the similarity between any two patterns is some measure of the agreement between their point sets. Progressively more complex and sophisticated theories can be constructed by postulating representations for larger structural units and their defining characteristics. Different perceptual theories postulate representations of line segments (e.g., Rumelhart & Siple, 1974; Rumelhart, Note 1), angles (e.g., Gibson, 1969; Selfridge & Neisser, 1960), and still larger portions of figures up to and including the figure as a whole (Gibson, 1969; Palmer, 1975, 1977). This paper investigates the nature of structural units through perceptual similarity to determine whether certain simple classes of representations are tenable for theories of human perception.

The distinction of concern is between theories that represent only lines and/or points vs. those that represent larger structural units such as angles, subfigures, and so forth. In the former class are all versions of template theory and certain line-based feature theories. For example, Rumelhart's representational assumptions specify that letters are coded as sets of line representations (Rumelhart & Siple, 1974; Rumelhart, Note 1). These line representations are not grouped into angle representations or any other higher order units prior to their concatenation by association into letter representations. As Rumelhart & Siple (1974) point out, their representational assumptions correspond roughly to the characteristics of Hubel and Wiesel's (1962) "simple" line detectors.

Underlying such a representational system is the assumption that segments are independent structural units; that is, segments do not enter into any larger

units (except by association into letters), nor do any properties depend on more than one segment. One implication of this assumption is that in computing the similarity of two figures, the variable of primary interest is just the number of segments that are identical in the two figures except for the global position of the figures themselves; that is, the number of segments that overlap when the two figures are superimposed. Likewise, the time to perceive that two figures are not identical should depend largely on the number of common segments. Template theories make the independence assumption at the level of points. Here, even line representations do not exist. Because points are assumed to be independent, the previous statements about similarity and discrimination are obviously true for template theories as well as line-based feature theories.

The present experiments investigate the nature of perceptual units for simple straight-line figures by testing these implications of segment and point independence. Figure 1 illustrates the stimulus structure forming the basis for the test. Six-segment figures were generated and designated the "standards" (S). Two modified versions (H and L) were then constructed by changing just one segment. Thus, both H and L figures have five segments in common with their corresponding standard figure and one segment that is different. According to the segment-independence assumption, then, the H and L figures should be seen as equally similar to the standard.

The two comparison figures for each standard were constructed such that one had a structure similar to that of the standard (the "high-similarity" or H figure) and the other had a structure different from that of the standard (the "low-similarity" or L figure). Although the H and L figures were not generated according to a precise algorithm, certain types of transformations characterize their relationship to the standards. The H figure generally contains the same subjective parts

as the standard (according to the experimenter's intuitions), but with one part rotated or reflected with respect to the other parts. In no case was the H figure a complete rotation or reflection of the standard. In the L figures, structure was altered by connecting parts that were separate in the standard (or separating parts that were connected) and by closing parts that were open in the standard (or opening parts that were closed). In other words, the H figures preserve connectedness and closedness properties of the standard, whereas the L figures do not. These are global properties of configurations of segments that cannot, in principle, be defined in a point- or segment-based representation without violating the independence assumption.

In half of the stimulus triads, the H and L figures were equated for the location, length, and orientation of the changed segments. That is, in the "matched" triads, the same segment of the standard was changed in both comparison figures, and the corresponding segments were equally different from it in absolute distance, orientation, and length. Thus, the matched segments control for the possibility that partial propertywise matches on nonidentical segments might influence similarity in a segment-based system such as that of Rumelhart (Note 1). In the "unmatched" stimuli, no attempt was made to equate the H and L figures for any aspect of the changed segments.

If supersegmental structure such as connectedness and closedness is an important component of perceptual representation, then the standards should be perceived as more similar to the H figures than to the L figures. Experiment 1 investigates this possibility using subjective similarity ratings, and Experiment 2 does so using speeded discrimination for physical identity.

EXPERIMENT 1: SUBJECTIVE SIMILARITY RATINGS

Method

Stimuli. Twenty stimulus triads were constructed in the manner described previously (see Figure 1). The 40 stimulus pairs, 20 sets of S/H and S/L pairs, were drawn on two pages of a test booklet, each standard appearing once per page. The order of pages was reversed in half the booklets.

Procedure. Subjects were given written instructions to rate the similarity of each figure pair on a scale of one (very dissimilar) to five (very similar). They were given the booklet and told to look over the 40 pairs briefly to get a feel for the range of similarities present. They were then instructed to use the entire range of the response scale in making their ratings. Two sets of rating data were obtained from each subject, the second session being conducted at least 24 h after the first to minimize memory effects.

Design. The four substantive factors of the design were: comparison type (S/H or S/L), triad type (matched or unmatched), figures (the 20 triads), and subjects. All factors were combined orthogonally except triad types and figures, for which the 20 levels of figures were nested within the two levels of triad type. Page order was counterbalanced both within and between subjects.

Subjects. The eight subjects were five males and three females between the ages of 20 and 26 with normal vision. All but one

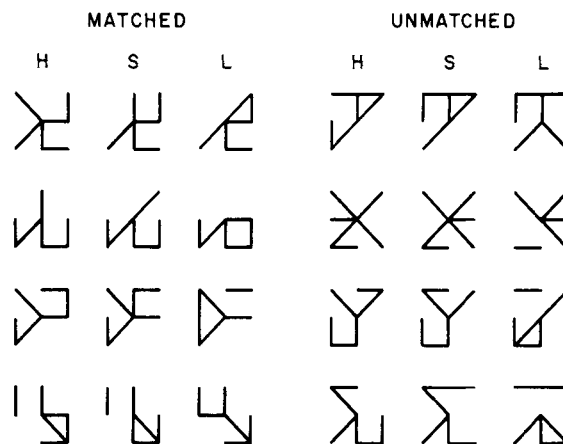


Figure 1. Examples of the stimuli. Each stimulus triad consists of a six-segment standard figure (S) and two comparison figures (H and L). The H figures preserve higher order structure (closedness and connectedness) of the standards, whereas the L figures do not. See text for matched-unmatched distinction.

were undergraduates at the University of California, San Diego, and all were paid for their participation.

Results and Discussion

Mean similarity ratings for the comparison and triad type factors are shown in Figure 2. These data clearly indicate that the H figures are perceived to be more similar to the standard figures than are the L figures. A within-subjects analysis of variance confirmed that this main effect of comparison type was highly significant both in the group data [$F(1,7) = 245.09$, $p < .01$] and for each individual subject [ranging from $F(1,20) = 83.53$ to $F(1,20) = 339.46$, $p < .01$ in each case]. The strength of this effect is further indicated by the fact that there were only three reversals (i.e., where the standard was judged more similar to the L

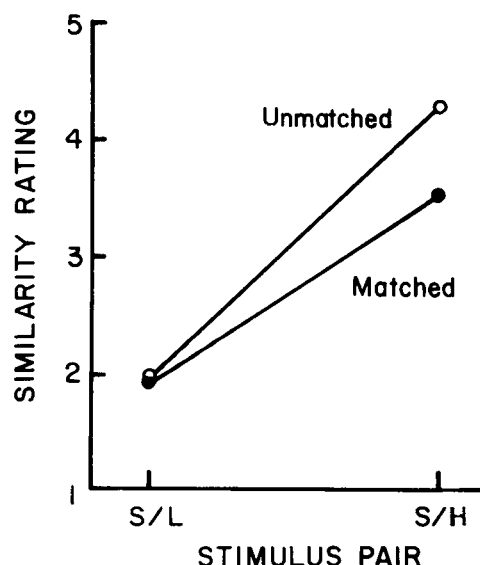


Figure 2. Mean similarity ratings for S/H and S/L pairs of both matched and unmatched triads.

figure than to the H figure) among the 160 mean judgments in the subjects by figures design.

Figure 2 also shows that the unmatched H and L figures were judged to be generally more similar to their standards than the matched H and L figures were to theirs [$F(1,7) = 56.44$, $p < .01$]. The triad-type factor also interacted with the comparison-type factor [$F(1,7) = 37.74$, $p < .01$]. The similarity difference between S/H and S/L comparisons was greater for the unmatched triads than for the matched ones. Because the strongest evidence against the segment-independence assumption was provided by the matched triads, separate analyses were performed for just these stimuli. A strong main effect of comparisons was still evident both in the group data [$F(1,7) = 117.99$, $p < .01$] and in each individual subject's data [ranging from $F(1,20) = 24.20$ to $F(1,20) = 104.04$, $p < .01$ in each case]. This result provides clear evidence against the assumption that segment representations are independent of each other. By inference, it supports the hypothesis that high-order part structure must be present in the representations of figures and that aspects of this structure are important for making similarity judgments.

EXPERIMENT 2: DISCRIMINATION REACTION TIME

Because conscious judgments may be unreliable indicators of normally unconscious perceptual processing, the stimuli from Experiment 1 were used in a speeded discrimination task. Subjects had to determine whether two figures were identical. Both latency and accuracy of S/H and S/L "different" responses are measures of discrimination difficulty and can be further interpreted as converging evidence about perceptual similarity. On the basis of the results of Experiment 1, it is expected that S/H discriminations will take longer and/or be less accurate than S/L discriminations. Whether similarity will be manifest in time or error measures (or both) will depend on the nature of speed-accuracy tradeoffs adopted by subjects.

Most current theories of same-different visual comparison posit two underlying processes (e.g., Bamber, 1969; Cooper, 1976; Nickerson, 1972). One is a rapid, holistic automatic comparison, sometimes discussed as a "template-matching" process (Reed, 1973). The second is a slower, sequential comparison of individual aspects. In order to account for established patterns of reaction times (RTs) for simultaneous and sequential comparisons (e.g., Egeth, 1966; Nickerson, 1967), it is usually assumed that the holistic, template-like process dominates performance in sequential comparisons, whereas the componential, feature-like process dominates performance in simultaneous comparisons (Bamber, 1969; Reed, 1973). To investigate the possibility that sequential comparisons rely more heavily on template-like representations, both

simultaneous and sequential presentation conditions were used. If sequential comparisons take place solely on template representations, then there should be no differences between S/H and S/L conditions for the sequential task. Naturally, if performance in both sequential and simultaneous conditions is some probabilistic mixture of both types of processing, then there should still be some effect in sequential comparisons, but it should be smaller than in simultaneous comparisons.

Method

Stimuli. The stimuli were identical to those used in Experiment 1. They were presented on a fast-decay (90% in .63 msec) cathode-ray tube oscilloscope, appearing light green on a black background. The figures measured .75 x .75 in., subtending approximately 1.5 deg of visual angle when viewed from the subject's position. In simultaneous presentation, the two figures were separated by .5 in. horizontally, so that the display subtended about 4 deg of visual angle in this direction.

The figures were paired as follows. For each stimulus triad, there were four "same" and four "different" trials. The "same" trials consisted of two S/S comparisons, one H/H comparison, and one L/L comparison. The "different" trials consisted of two S/H comparisons and two S/L comparisons. Thus, there were 160 trials per session, 80 "same" and 80 "different" comparisons.

Procedure. Subjects were given instructions regarding the nature of the task. They were instructed to respond as quickly as possible while keeping their overall error rate at or below 5%. Before the first session began, subjects were given 40 practice trials. If any subject had remaining questions, they were answered at this point.

In the simultaneous condition, a trial consisted of the following sequence. A ready signal was presented for 500 msec. This signal was a pair of 3 by 3 dot matrices (the endpoints of the 16 possible line segments) in the positions where the figures would appear. The stimulus pair was then presented. The subject terminated the display by making a response. Because subjects usually knew when they had made a mistake, no specific feedback was given. After a 2-sec delay, the next trial began.

In sequential presentation, a trial consisted of a ready signal followed by successive presentations of the members of the stimulus pair. The ready signal was a single 3 by 3 dot matrix in the position where the figures would appear. The first figure remained on the screen for 1 sec. After a 500-msec delay, the second (test) figure appeared. The presentation of the test stimulus was terminated by the subject's response. After a 2-sec delay, the next trial began.

Each session began with 20 practice trials using stimuli different from the experimental material, but similarly constructed. The first 80 trials were followed by a mandatory 1-min rest period. At the end of each session, there was a delay of approximately 5 min before the next session began. All subjects were tested individually in eight sessions, four per day for 2 days. In half the sessions, the stimulus pairs were presented simultaneously (side by side). Position was balanced for the S/H and S/L trials. In the other half of the sessions, the stimulus pairs were presented sequentially. Presentation order was also balanced for S/H and S/L pairs.

For each subject, the order of the presentation conditions was balanced within days (blocks of four sessions). The order was either simultaneous, sequential, sequential, simultaneous or sequential, simultaneous, simultaneous, sequential. These orders were balanced between subjects. For all four sessions in the same day, one hand was used for "same" responses and the

other hand for "different" responses. Responses were switched to opposite hands on successive days. The order of hand assignments was reversed for half the subjects.

Design. The seven substantive factors of the design were: presentation mode (simultaneous and sequential), days (first and second), sessions (first and second in a given day), stimulus pairs (S/S, S/H, and S/L), type of triad (matched and unmatched), figures (the 20 triads), and subjects. All factors were combined orthogonally except type and figures, for which the 20 levels of figures were nested within the two levels of type. There were two observations per cell.

Subjects. The subjects were the same as in Experiment 1.

Results and Discussion

Reaction times (RTs) were measured to the nearest millisecond and were statistically analyzed as follows. Using correct responses only, cell means were computed for all substantive factors other than figures. "Wild" observations were defined as reaction times more than 5 SDs above or 3 SDs below the corresponding cell mean. (Asymmetrical cut-offs were used because the RT distributions were positively skewed.) Less than .5% of the data typically fell into the "wild" category. Cell means were then recomputed, discarding both incorrect and "wild" responses. These missing data were then estimated with the mean of the appropriate cell in the design. Statistical analyses were performed after applying a logarithmic transformation to homogenize variances.

"Different" responses: S/H vs. S/L. The main purpose of the experiment was to determine whether discriminating the standards from H figures takes more processing time than discriminating them from the L figures. The results for these two conditions are shown in Figure 3. Subjects were able to respond significantly faster to the S/L pairs than to the S/H pairs in both the simultaneous [$F(1,7) = 306.02$, $p < .01$] and the sequential [$F(1,7) = 34.69$, $p < .01$] presentation conditions. This effect was present for each subject individually in both simultaneous [ranging from $F(1,160) = 85.98$ to $F(1,160) = 172.61$, $p < .01$] and sequential [ranging from $F(1,160) = 12.94$ to $F(1,160) = 27.07$, $p < .01$] presentations. Performance was also more accurate on the S/L pairs than on the S/H pairs in both simultaneous [$F(1,7) = 35.51$, $p < .01$] and sequential [$F(1,7) = 5.78$, $p < .05$] conditions.

As in the rating data, an interaction was present for the RT data such that the difference in performance was greater for the unmatched than the matched pairs. Separate analyses for only the matched pairs confirmed significant RT differences between S/L and S/H conditions for both simultaneous [$F(1,7) = 242.50$, $p < .01$] and sequential [$F(1,7) = 21.14$, $p < .01$] conditions. Corresponding analyses of error rates also indicate significant effects for just the matched pairs in simultaneous [$F(1,7) = 42.94$, $p < .01$] and sequential [$F(1,7) = 6.62$, $p < .05$] presentation. Thus, the basic prediction was confirmed: Standards are

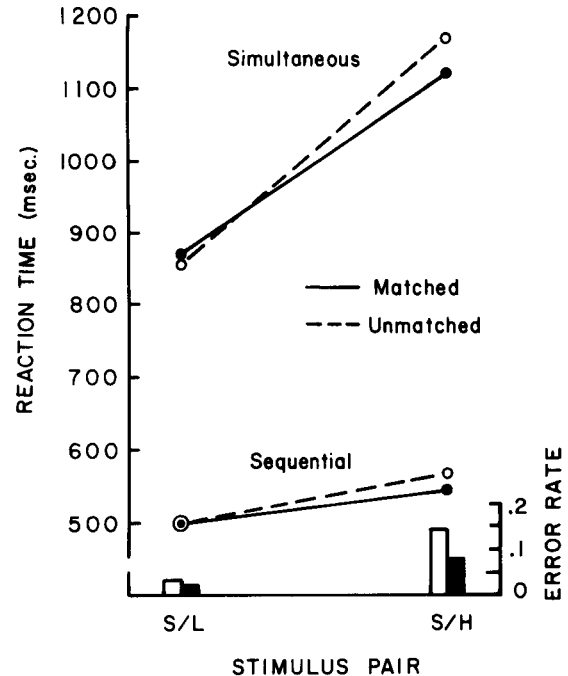


Figure 3. Mean reaction time (in milliseconds) for S/H and S/L pairs in simultaneous and sequential presentation conditions. Histograms at the bottom show overall error rates for simultaneous (light) and sequential (dark) presentation in these conditions.

generally more similar to the H figures than to the L figures as indicated by discrimination performance.

Although significant differences between S/H and S/L comparisons were present for both presentation modes, their relative magnitudes are noteworthy. The RT difference in simultaneous presentation was nearly 300 msec, whereas the difference in sequential presentation was only about 50 msec. This difference was highly significant [$F(1,7) = 65.92$, $p < .01$] and was present for each individual subject. One interpretation is that although the segment representations are not independent in either simultaneous or sequential comparisons, they are functionally much more so in a sequential match than in a simultaneous match. This would be the case, for example, if some proportion of matches were performed on a template representation and the rest on some higher level representation in which lower level parts were combined into larger structural units. Obviously, this is just one of many ways to account for the difference.

Correlations were employed to assess the correspondence between the rating data from Experiment 1 and the latency data from the present study. Significant correlations were found for each pair of tasks: ratings vs. RT in simultaneous presentation [$r = .77$, $t(38) = 7.46$, $p < .01$], ratings vs. RT in sequential presentation [$r = .76$, $t(38) = 7.22$, $p < .01$], and RTs in simultaneous vs. sequential presentation [$r = .84$, $t(38) = 9.44$, $p < .01$]. However, these correlations were

due almost entirely to gross differences between S/H and S/L pairs. After removing this main effect, the only significant correlation was that between the two RT measures [$r = .49, t(38) = 3.45, p < .01$]. The somewhat disappointing correspondence between the two types of measures may be due to differential importance of different variables in the two tasks or the presence of sufficient variability in one or both to make only gross differences reliable.

Presentation modes. The most obvious difference between performance in the two presentation modes was that subjects responded more quickly and accurately when the stimulus pairs were presented sequentially than when they were presented simultaneously [$F(1,7) = 582.58, p < .01$]. This may be partially attributed to a methodological artifact. In the simultaneous condition, the time required to process both figures was measured; in the sequential condition, only the time to process the second figure was measured. There are other differences in the data, however, that suggest that this is not the only factor operating. One difference lies in the size of the similarity effect mentioned earlier: Comparison type had much larger effects in simultaneous presentation. Another lies in accuracy data: Subjects made almost twice as many errors in simultaneous as in sequential presentation. Two other differences are discussed in the following sections.

"Same" vs. "different" responses. The relative speeds of "same" and "different" responses were not the same in simultaneous and sequential presentation. In the simultaneous condition, "same" responses were intermediate between the two types of "different" responses. They were slower than S/L responses [$F(1,7) = 89.75, p < .01$] and faster than S/H responses [$F(1,7) = 10.35, p < .05$]. In the sequential condition, however, "same" responses were faster than either type of "different" responses. "Same" responses were clearly faster than S/H responses [$F(1,7) = 23.94, p < .01$] and faster than S/L responses [$F(1,7) = 7.51, p < .05$]. This pattern of RTs for simultaneous and sequential presentations is similar to results reported by Egeth (1966) and Nickerson (1967) using multi-dimensional stimuli such as large red triangles and small green squares.

Practice effects. Because there were two sessions per day for 2 days, there were two types of practice effects in the present results. The effects of days are presented in Figure 4A, and those of sessions in Figure 4B. For sequential presentation, there was a significant sessions effect of 50 msec across comparison conditions [$F(1,7) = 7.47, p < .05$]. For simultaneous presentation, the sessions effect was also significant [$F(1,7) = 10.87, p < .05$], being somewhat greater in the S/H and S/S comparisons. These effects could be due to either familiarity with the motor responses (because response-hand assignments were the same for

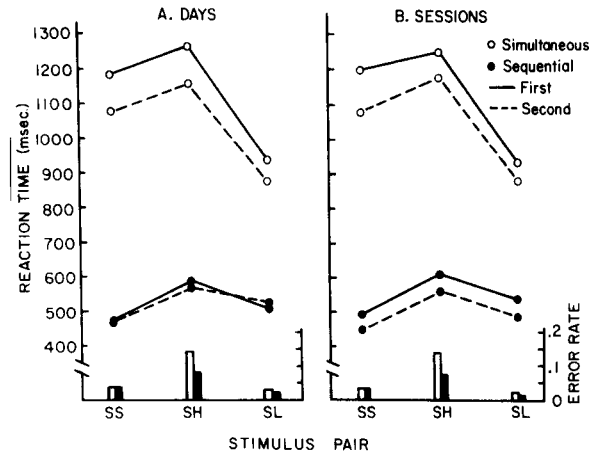


Figure 4. Mean reaction time (in milliseconds) for "same" (S/S) and "different" (S/H and S/L) trials in simultaneous and sequential presentation. Part A shows the effects of first and second days of testing; Part B shows the effects of first and second sessions within days.

all sessions in a day) or to learning. Practice effects due to days, however, were present in the simultaneous condition [$F(1,7) = 5.61, p < .05$], but not in the sequential condition ($F < 1$). These practice effects cannot be due to motor response facilitation, but must be attributed to processing changes. This result is consistent with subjective reports. As subjects became familiar with the figures, they claimed they learned where to look for the differences between the figures, especially for the S/H pairs.

Theoretical implications. The present results are consistent with the general concept of dual process theories of same-different comparisons in a number of respects. First, the pattern of "same" and "different" RTs is essentially the same as that found in multi-dimensional matching tasks (e.g., Egeth, 1966; Nickerson, 1967). That is, "different" RTs increased monotonically with similarity in both simultaneous and sequential presentation, but "same" RTs were faster than either S/H or S/L "different" RTs in sequential presentation, whereas "same" RTs were intermediate to these conditions in simultaneous presentation. These are just the sort of RT patterns expected if sequential comparisons are dominated by a rapid holistic process, while simultaneous comparisons are dominated by a slower aspect-by-aspect process. Second, there is a substantial practice effect over days for simultaneous comparisons, but none for sequential ones. This is consistent with the hypothesis that the holistic process is relatively "automatic" and strategy free, whereas the slower serial process is controlled by more flexible perceptual strategies that can be influenced by practice with the particular stimulus set. Third, the similarity effect for "different" responses is much larger for the simultaneous condition than for the sequential one. This finding is consistent with the notion that a template comparison (in

which the independence assumption holds for segments and/or points) may be involved in the automatic holistic process. However, there are many alternative explanations that do not require template representations. For example, a parallel comparison of high-order features could constitute the "holistic" process. Such an operation would be relatively insensitive to similarity differences in its temporal characteristics, but would still reflect them in its error characteristics. In fact, Cooper (1976) has recently found evidence that some subjects (her Type 1 subjects) perform sequential same-different comparisons such that their speed for "different" judgments is unaffected by similarity, whereas their errors increase monotonically with similarity. The parallel comparison process Cooper proposes does not depend on template representational assumptions, but would fit the present data as one of the two underlying processes. (Note that the presence of similarity effects in RTs for all eight of the present subjects indicates that none would be classified as "Type 1" subjects in Cooper's dichotomy.)

GENERAL DISCUSSION

The present results indicate that perceptual representation includes higher level structural units than simple points or line segments. The large differences between the S/H and S/L conditions in both similarity ratings and discrimination performance are inconsistent with the segment-independent assumption. In other words, the present studies show that organization of points and lines into larger structures with emergent properties must be an important determinant of perceptual similarity.

Further evidence for this conclusion was obtained in an informal follow-up experiment. Ten different subjects were asked to sort 20 randomly generated figures of the type used in the present experiments into any number of piles according to their impressions of similarity. The number of times two figures were sorted into the same pile was used as an index of similarity for these figures. The resulting proximity matrix was fit to a hierarchical data model using a hierarchical clustering program (Johnson, 1967). Although not all the binary divisions were easily interpretable, the first two can be clearly described as "is (or is not) closed" and "is (or is not) connected." Neither of these aspects is consistent with the independence assumption for line segments. Closed subfigures are structural units that, by definition, must contain at least three segments, and connectedness is a property of entire figures that cannot, in principle, be computed by any local perceptual operator (Minsky & Papert, 1969).

The present results reinforce other recent findings that point toward the importance of large structural units in perception. Using stimuli similar to the present

figures, Palmer (1977) found evidence against segment independence in four different tasks. Results showed that people's performance in (1) dividing figures into natural parts, (2) rating the "goodness" or "naturalness" of such parts within figures, (3) speeded search for parts within figures, and (4) synthesizing parts into whole figures was inconsistent with the segment-independence assumption. The main thrust of these results is that certain groups of segments within figures form "good" parts that are more cohesive and unitary than other "bad" groups of segments. Reed (1974) has also found evidence for "good" parts in perceptual memory using a part-probe task similar to that employed by Palmer (1977). In a similar vein, Bower and Glass (1976) showed that "good" parts were more effective in cuing recall of the figures than were "bad" parts. They also found that structurally similar figures were more likely to be confused in memory with presented figures than were structurally dissimilar figures. Taken as a whole, these results strongly suggest that "good" parts are explicitly coded as structural units in perceptual representation. If this view is approximately correct, the independence assumption cannot hold for segments or points, because the parts are multisegment subfigures in each of the above cases.

All of this recent work can be seen as elaboration and extension of earlier research in the Gestalt tradition (e.g., Gottschaldt, 1926/1967; Wertheimer, 1923/1967). Of particular importance to the present topic is Goldmeier's (1936, 1972) classic treatise on visual similarity. Although his methods and stimuli were quite different from those employed here, Goldmeier reached a conclusion similar to the present one: Perceived similarity cannot be accounted for by the identity of independent components. Components enter into configural relationships with other components, resulting in larger structural units whose importance supercedes that of its constituents. The precise nature of these perceptual units, however, is problematic and remains an area of vigorous research.

REFERENCE NOTE

1. Rumelhart, D. E. *A multi-component theory of confusion among briefly exposed alphabetic characters* (Tech. Rep. 22). San Diego: University of California, Center for Human Information Processing, 1971.

REFERENCES

- BAMBER, D. Reaction times and error rates for "same"- "different" judgments of multidimensional stimuli. *Perception & Psychophysics*, 1969, 6, 169-174.
- BOWER, G. H., & GLASS, A. Structural units and the redintegrative power of picture fragments. *Journal of Experimental Psychology: Human Learning and Memory*, 1976, 2, 456-466.
- COOPER, L. A. Individual differences in visual comparison processes. *Perception & Psychophysics*, 1976, 19, 433-444.

- EGETH, H. W. Parallel vs. serial processes in multidimensional stimulus discrimination. *Perception & Psychophysics*, 1966, **1**, 245-252.
- GIBSON, E. J. *Principles of perceptual learning and development*. New York: Appleton, 1969.
- GOLDMEIER, E. Über Ähnlichkeit bei gesehenen Figuren. *Psychologische Forschung*, 1936, **21**, 146-208.
- GOLDMEIER, E. Similarity in visually perceived forms. *Psychological Issues*, 1972, **8**(Whole No. 29).
- GOTTSCHALDT, K. Über den Einfluss der Erfahrung die Wahrnehmung von Figuren. *Psychologische Forschung*, 1926, **8**, 261-318. In W. D. Ellis (Ed.), *A source book of Gestalt psychology*. New York: Humanities Press, 1967.
- HUBEL, D. H., & WIESEL, T. N. Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *Journal of Physiology*, 1962, **160**, 106-154.
- JOHNSON, S. C. Hierarchical clustering schemes. *Psychometrika*, 1967, **32**, 241-254.
- MINSKY, M., & PAPERT, S. *Perceptrons*. Cambridge, Mass: MIT Press, 1969.
- NICKERSON, R. S. "Same"- "different" response times with multi-attribute stimulus differences. *Perceptual and Motor Skills*, 1967, **24**, 543-554.
- NICKERSON, R. S. Binary-classification reaction times: A review of some studies of human information processing capabilities. *Psychonomic Monograph Supplements*, 1972, **4**(Whole No. 65), 275-318.
- PALMER, S. E. Visual perception and world knowledge: Notes on a model of sensory-cognitive interaction. In D. A. Norman, D. E. Rumelhart, & LNR Research Group (Eds.), *Explorations in cognition*. San Francisco: W. H. Freeman, 1975.
- PALMER, S. E. Hierarchical structure in perceptual representation. *Cognitive Psychology*, 1977, **9**, 441-474.
- PALMER, S. E. Fundamental aspects of cognitive representation. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, N.J.: Lawrence Erlbaum, in press.
- REED, S. K. *Psychological processes in pattern recognition*. New York: Academic Press, 1973.
- REED, S. K. Structural descriptions and the limitation of visual images. *Memory & Cognition*, 1974, **2**, 329-336.
- RUMELHART, D. E., & SIPLE, P. A. Process of recognizing tachistoscopically presented words. *Psychological Review*, 1974, **81**, 99-118.
- SELFRIDGE, O., & NEISSER, U. Pattern recognition by machine. *Scientific American*, 1960, **203**, 60-68.
- WERTHEIMER, M. Untersuchungen zur Lehre von der Gestalt. *Psychologische Forschung*, 1923, **4**, 301-350. In W. D. Ellis (Ed.), *A source book of Gestalt psychology*. New York: Humanities Press, 1967.

(Received for publication October 31, 1977;
accepted December 5, 1977.)