Journal of Experimental Psychology: Human Learning and Memory
1976, Vol. 2, No. 4, 456–466

Structural Units and the Redintegrative Power of Picture Fragments

Gordon H. Bower
Stanford University

Arnold L. Glass
Rutgers University

We suppose that line drawings are perceived and represented in memory as a hierarchy of related parts and subparts, as dictated by Gestalt laws like common direction and spatial proximity. Therefore, a figure fragment comprising a natural part of an originally studied pattern should serve as a strong retrieval cue for redintegrating memory for the pattern, whereas an equally large fragment suggesting either no units or misleading units should lead to poorer recall. This was confirmed in an experiment in which subjects studied 33 nonsense line drawings; recall of each was tested with good, mediocre, or bad (misleading) fragments of the original patterns. Good cues had about five times more redintegrative power than bad cues. A second experiment testing multiple-choice recognition memory showed that subjects confused an originally studied pattern about four times as often with a structurally similar distractor as with a structurally dissimilar distractor (which had an equal-sized change.) Thus, memory cuing by fragments and memory confusions with slightly altered distractors indicate the significant constituents of a figure.

Our primary interest here is in the way simple line drawings are represented in memory. Our second interest is in validating a particular methodology which provides evidence about such representations. Our third aim is to propose and test a theoretical algorithm that mechanically segments and assigns a psychological representation to a set of simple (nonsense) line drawings. This is important for understanding the "parsing rules" for forms shown to the human visual system. The underlying assumption is that what is remembered about a line drawing is determined by the output of a perceptual parser and by whatever associations memory adds immediately to that output (e.g., verbal labeling of what the shape resembles).

Regarding the memory representation of line drawings, a number of theories differ according to the size and structure of the units in their representation of the drawing. At one extreme is the *holistic* or *holographic* hypothesis which supposes that memory for a visual shape consists of an analog copy or template of the line drawing. We may think of this as a photographic copy which can be retrieved or regenerated in imagery upon demand. At the opposite extreme is the *elementaristic* hypothesis which supposes that a line drawing is encoded as a collection of elemental, nonrelational features such as line segments, curves, and angles, and the directional orientation of these elements. Other hypotheses suppose that a line drawing is represented as a *structural hierarchy* of significant parts; the significant parts correspond to relations among visual features at a lower level of analysis, and the structural hierarchy describes the relationship among successively larger parts.

Several methods exist for collecting converging evidence regarding the internal representation of a picture. One method involves perceptual judgments of the similarity of two figures (e.g., Goldmeier, 1972; Palmer, 1974). Each hypothesis supposes that similarity judgments reflect the percentage overlap of identical "units" in the encodings assigned to the two figures: One tries to find

This research is supported by Grant MH-13950 to the first author from the National Institute of Mental Health. We thank Jeremy Rosenblum for helping with the first experiment.

Requests for reprints should be sent to Gordon H. Bower, Department of Psychology, Stanford University, Stanford, California 94305.

critically differentiating triples of stimuli, A, B, C, such that whether Figure A overlaps more with B than with C depends on what the perceptual units are; in this way, similarity judgments may help identify the units of the representation. Along with generalization errors, this direct scaling method is also frequently used to investigate the similarity of verbal materials (paraphrases, words, nonsense syllables). Another method to study figural units examines the time it takes a person to decide that a large pattern contains a particular fragment (or probe) as a subpart. For example, Reed (1974) reported that intuitively "good" subparts were found more quickly than "poor" subparts embedded in a pattern. Assuming that the normal analysis of the pattern breaks it into good subparts, a good probe will find an immediate match, whereas a poor probe will be "seen" only by breaking and recombining in nonintuitive ways pieces of normal subparts. This same point is at issue in studies of camouflaging and hidden (embedded) figures.

The new method tried in our study is memory cuing by fragments of an originally studied pattern: After studying a set of nonsense line drawings, the subject receives a series of cues, each comprising a fragment of one or another of the drawings, and is asked to recall any of the drawings the cue brings to mind. Interest centers on the redintegrative power of the various fragments or parts of a figure. The method was suggested, in passing, by Köhler (1947, p. 170), but he reported no evidence for it, nor were we aware of any.

Memory theories suppose that a cue reactivates memory of a pattern on the basis of some resonance or similarity between the codings of the cue and the original pattern. For example, in a template theory, a probe retrieves its corresponding pattern from memory according to the maximal cross-correlation between the probe and the images of the pictures in memory. In the elemental feature theory, (see, e.g., Bower, 1967), retrieval depends on selecting from memory that figure coding which overlaps maximally in features with the probe. The structural hierarchy theory supposes that a good part will match a subunit of the memory code of the pattern, causing redintegration of the whole pattern; on the other hand, a poor fragment will match either no pattern code in memory or will weakly match many pattern codes, so no particular figure would be strongly suggested. Our experiments attempt to validate this theory of memory cuing as a method for collecting convergent evidence regarding figural parsing.

Although one could validate the cuing method by using intuitions about what are good versus poor subparts of a figure, we desired a more general formulation of intuitions. To this end, our experiments tested two rules for combining line segments or curves into larger structural units. The rules are the Gestalt principles of common direction and proximity in space (see Wertheimer, 1938) applied to combining line segments into structural units.

The rule of *common direction* states that line segments that are continuous in the same direction will be combined and encoded as a single structural unit. For instance, in Part A of Figure 1, this rule will conjoin Segment a to b rather than to c, forming the Line Unit ab, with c as another unit. The rule similarly combines Segments a with b, and c with d, in Parts B and C of Figure 1. The curve in Part D is a single unit because it is continuous in one direction, whereas Part E has a change of direction (inflection point) which causes the rule to break it into two different units (a and b).

The second rule is the rule of *minimal angle,* which is like a spatial proximity notion for combining line segments. When three or more segments intersect at a point, the two with the minimal angle between them are combined to form a single structural unit. For example, in Figure 2, Part A, Lines a and b form a smaller angle than do either a and c, or b and c; therefore, a and b form one structural unit, whereas c is a separate unit. Similarly, in Part B, Figure 2, the rule of minimal angle combines Segments d and e into one unit, and f and g into another unit. Part C of Figure 2 illustrates a conflict between how segments would be unitized by common direction (Unit hi) versus minimal angle (Angle ij). Our intuitions suggested that common direction should have priority

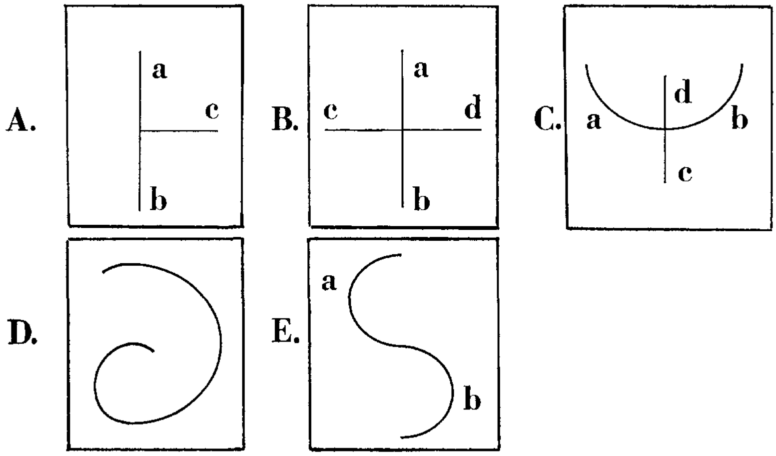

FIGURE 1. Illustrations of rule of common direction.

in resolving such conflicts, so Unit hi was assumed to be the first grouping derived.

A corollary to the minimal angle rule is this: If a line segment contacts other line segments at its two ends, it will be combined into a unit with that segment having the minimal angle with it. For instance, in Figure 3, Part A, Segment b will be combined into a unit with Segment a rather than c because Angle ab is less than Angle bc. Similarly, in Part B, Segments f and g are melded first, leaving Segment e to be melded with d. In Part C, Segment i makes equal angles with h and j, so the three parts are melded together as the jagged unit, hij.

Figure 4 illustrates how two entire figures are segmented according to the rules. In Figure 4, Part A, the rule of common direction applies to the four intersections labeled 1, and it melds the segments into the Line Units ab, cd, ef, and gh. By the law of minimal angle, Line ab is melded with Line cd (rather than gh), and Line ef is melded with Line gh (rather than cd). The remaining Segments j, k, m, and n are melded into the Units jk and mn because they touch. The subunits of Part A are summarized in Part B. Similarly, the stick figure, Part C, would be parsed by the rules into head, torso, and arms (by common direction) and legs (by minimal angle); these units are shown in Part D.

While other rules are probably required in order to analyze all possible two-dimensional figures into intuitively plausible constituents (see, e.g., Glass, 1975; Guzmán, 1969), the rules of common direction and minimal angle seem to be high priority principles for any analysis scheme. These rules are used for analyzing the patterns and cue fragments in the experiments below.

EXPERIMENT 1

The idea tested was that a fragment corresponding to a structural unit of a memorized pattern would be a strong retrieval cue for that pattern, whereas an equally overlapping fragment not corresponding to a structural unit would be a poor cue. For

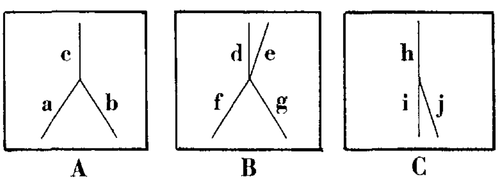

FIGURE 2. Illustrations of rule of minimal angle.

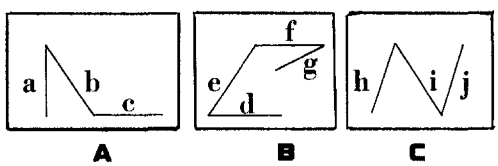

FIGURE 3. Illustrations of corollary to rule of minimal angle.

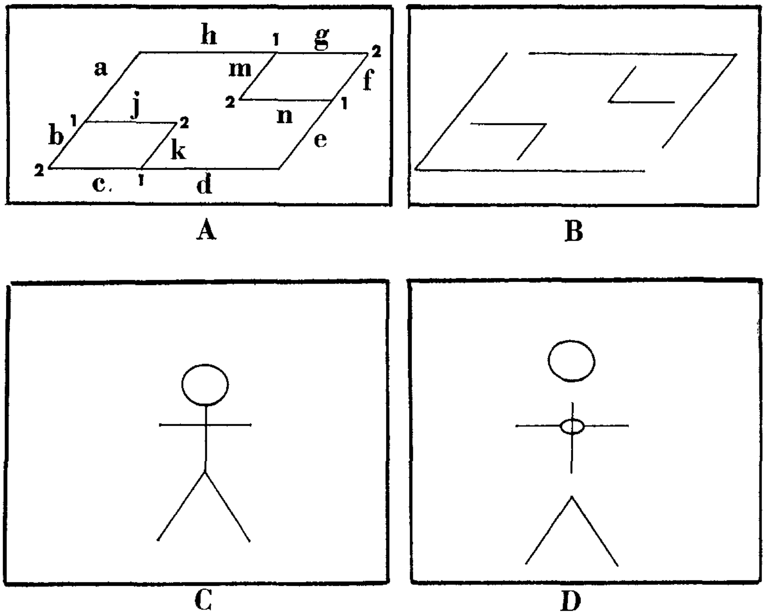

FIGURE 4. Illustrations of parsing of complex patterns (A and C) into parts (B and D).

each studied pattern we devised three types of fragments, which we call good, mediocre, and bad cues. This describes the relation of the fragment to the subparts of the theoretical parsing of the to-be-remembered figures.

The good cue for a figure was comprised of one or more salient subparts as dictated by our intuitions and parsing rules. Good cues are shown just to the right of three illustrative patterns in Figure 5. The parsing of the good cue for Pattern A of Figure 5 would contain two of the units of the original, namely, abcd and jk (see Figure 4, Part A). Similarly, the good cue for the stick figure uses two units (torso and arms) plus a strong cue for a third (head). Pattern C is described by the rules as "four vertical lines cut by a diagonal," and the good cue clearly contains vertical lines as units distinct from the diagonal unit.

Mediocre cues (see Figure 5) overlap as much (in terms of common line lengths) with the original pattern as do good or bad cues. Yet mediocre cues should be poorer retrieval cues because they contain not only fewer constituents but also fragments of constituents. For instance, the rules analyze the mediocre cue for Pattern A as Line cd connecting to Line e (so the cde "corner" matches the representation of the original pattern); but the a and b lines of the mediocre cue are left as dangling structural units which mismatch those assigned to the original pattern. In the mediocre cue of Pattern B, Line f is only part of a former angle, and connectivities of parts are not indicated. The mediocre cue for Pattern C connects parts of two structural units differently than they are connected in the original pattern. For instance, in the original pattern the diagonal line criss-crosses four lines, but in the mediocre cue, Lines g, h, i, and j merely touch the diagonal, k, without crossing it.

Bad cues (see Figure 5) are misleading in the sense that the parsing rules assign to them structural units that do not correspond in any way to the units in the representation of the original pattern. For example, the two jagged lines in the bad cue for Pattern A each comprise a structural unit. By the rule of minimal angle, Segments m and n form a unit as do Segments o and p, although neither of these units occur in the parse of the original figure. Similarly, in the bad cue for Pattern B, Segments q and r form a single unit by the rule of minimal angle, al-

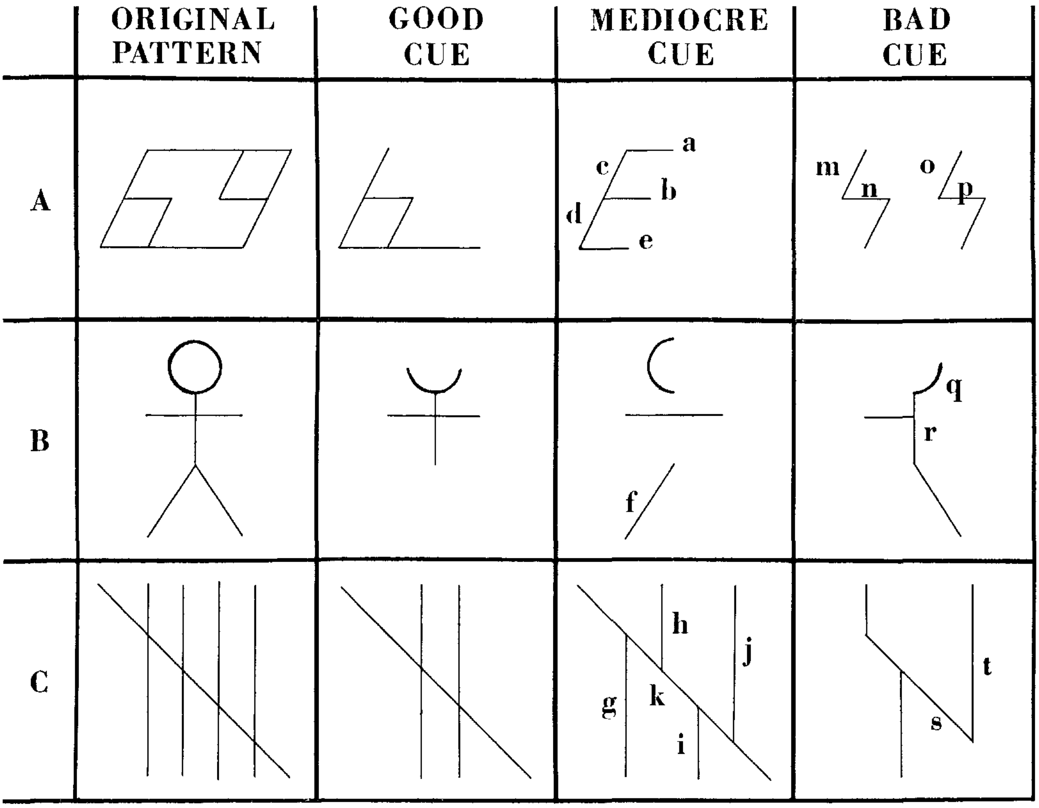

FIGURE 5. Illustrations of good, mediocre, and bad cues for three memorized figures, Experiments 1 and 2.

though these are fragments of different parts in the original. Finally, in the bad cue for Pattern C, Segments s and t are melded into a single unit by the rule of minimal angle, although they are fragments of different units in the original.

To summarize, we expect a good cue to redintegrate memory of its studied pattern better than a mediocre cue, which will be better than a bad (misleading) cue. The cues were constructed a priori according to our intuitions and the formal codings assigned by the rules of common direction and minimal angle.

Method

Materials. A good, mediocre, and bad cue, like those shown in Figure 5, was constructed for each of 27 patterns. The 27 patterns were two-dimensional, unshaded, line drawings seen as a "single good figure" and were generated at random, with the proviso that each have intuitively good and bad cues. An attempt was made to control for the amount of the original pattern reproduced in each cue as measured by the length of its line segments. While this could not be tightly controlled for individual patterns, over the entire item set the amount of pattern reproduced in each type of cue did not vary substantially. Each cue was appropriate for only one pattern. The patterns and cues were drawn with a black felt-tip pen on white cardboard 2.3 × 3.0 cm. A tracing procedure was used so that each cue appeared exactly the same and in exactly the same position when presented alone as it did as part of the original pattern.

Procedure. Twenty-four high school and college students, who were paid for their participation, were tested in groups of 5. The subjects were told at the outset that they would have to recall the pictures. First, each of 33 patterns was manually presented one at a time for 7 sec. The set of 33 patterns consisted of the 27 test patterns followed by 6 filler patterns. Next a subject was given as much time as necessary to draw each of the patterns that could be freely recalled. Finally, for 81 trials a subject was presented with a cue and had

7 sec to recall the pattern it came from. If a subject indicated that he or she could recall the pattern, as much time as necessary was given to draw it.

The cues were presented in three blocks of 27 trials, one cue per pattern in each block. Each block consisted of 9 good, 9 mediocre, and 9 bad cues. Over the entire series, each subject thus saw every cue for every pattern. The order of the cues was randomized within each block, and the three blocks of cues were completely counterbalanced across the six groups of subjects.

Results and Discussion

No subject ever recalled the wrong pattern to a cue, and no recall drawing was so distorted or incomplete that there was doubt about what pattern had been recalled. Two scorers graded the recalls and had nearly perfect reliability.

Free recall averaged 46% for the critical patterns. The cued recalls across the three tests were partitioned according to the number of prior recalls of a given target pattern (counting free recall). These data are summarized in Table 1. Inspection of Table 1 reveals that good cues had much greater redintegrative power than mediocre cues which, in turn, were better than bad cues. A second finding was that a pattern was more likely to be recalled to a given cue the more often it had been recalled previously. This latter effect is shown by the consistent increase in recall in Table 1 with number of prior recalls (down the columns) regardless of the type of test cue under consideration (across rows).

TABLE 1

PROPORTION CORRECT RECALL, EXPERIMENT 1

Number of prior recalls	Cue type: Good	Mediocre	Bad
Trial 1			
0	65	34	08
1	86	41	22
Weighted average	77	38	14
Trial 2			
0	51	16	05
1	84	39	29
2	96	60	32
Weighted average	73	39	24
Trial 3			
0	51	04	08
1	90	35	22
2	96	56	32
3	95	88	38
Weighted average	78	43	28

The reliability of the difference between the mediocre and bad cues was assessed by a sign test across items and subjects. The mediocre cue was more effective for 22 of the items, the bad cue for 4, and there was one tie, for a $p < .001$. The mediocre cue was more effective for 29 subjects out of 30, with one tie, $p < .001$. The difference between the good and mediocre cues was also tested across items and subjects. For 26 of 27 items the good cue was superior, with one tie, $p < .001$. For all 30 subjects, the good cue was superior to the mediocre cue. Hence, a good cue was a significantly stronger cue for recall than a mediocre cue, which was significantly stronger for recall than a bad cue.

The results support the hypothesis that patterns are encoded in memory in the structural units specified by the rules of common direction and minimal angle. Even after a pattern had been recalled three times previously, its percentage recall to the bad cue (38%) was still less than the initial probability of free recall (46%). It is as if the cue could not make contact with the representation of the pattern in memory unless the structural units of the cue corresponded closely to those in the memory trace of the figure.

EXPERIMENT 2

One explanation of the better retrieval capability of the good cue supposes that a good cue is more likely than a bad cue to elicit the same verbal label as that elicited by the original pattern. This hypothesis, of course, by-passes the critical question of what stimulus factors might cause the good cue to be labeled similarly to the original pattern. In any event, to lay this criticism to rest, we collected verbal associations to the good cues and the original patterns to check how often the members of a pair were described by the same label. Pilot tests with five subjects giving associations to the bad cues and the whole patterns were abandoned since no one ever gave the same label to the bad cue and its whole pattern. However, the good cue associations seemed worth collecting.

Method

Twenty-five of the 27 whole patterns and their good cues from Experiment 1 were shown to 10 subjects individually. The subjects were faculty and staff at Rutgers University and were run individually. They first gave a verbal label (name or descriptive phrase) for the good cues then for the original (whole) patterns.

Results

Of interest is the percentage of good-cue-whole-pattern pairs which elicited similar verbal labels. An extremely lenient criterion of "same label" was adopted, namely, that the phrases elicited by the cue and the whole pattern have at least one content word (or synonym or paraphrase) in common. By that criterion, only 36 or 14% of the 250 pairs revealed any degree of overlap. It is quite likely that this 14% overlap figure is an overestimate of the true overlap, given the premises of the verbal explanation. In particular, many subjects volunteered that they "saw" the good-cue fragment when they later judged the whole pattern, so in the ambiguous labeling task there was probably some bias to give the whole pattern the same label they had given to the good cue (e.g., "one loop," then "two loops"). In any event, this low overlap percentage can not begin to account for the nearly 90% retrieval effectiveness of the good cues in Experiment 1. That is, perceptual redintegration from a good cue that succeeds 90% of the time could not be driven by a verbal mediation process that succeeds probably less than 14% of the time. The verbal mediation hypothesis would appear to be putting the cart before the horse.

EXPERIMENT 3

If a pattern is encoded as a set of structural units, then these structural units characterize the *equivalence class* of the memorized pattern. Thus, variation in details of a pattern that alter its structural units will be noticed more readily than variations that do not alter its structural units. Therefore, two patterns containing the same structural units should be more difficult to discriminate in a forced-choice recognition memory task than patterns with differing structural descriptions. As one illustration, Patterns A and B in Figure 6, which differ only in the length of Segment a, would be encoded as two straight lines by the rule of common direction, since the lengths of the lines are irrelevant to this rule. Hence Patterns A and B should be difficult to discriminate. On the other hand, since Segments c and d in Pattern C are not connected, the rule of common direction does not apply, so they would not be formed into a single structural unit. Hence, people should discriminate Pattern A from Pattern C more easily than from Pattern B.

In Experiment 3, the rules of common direction and minimal angle were again used to segment patterns into structural units consisting of lines, angles, and curves. Our distractors and our predictions then were generated from this theoretical analysis. We needed one further rule, namely, lines without endpoints such as those in Figure 7, Part A were considered to be different structural units than those with endpoints (Part B). This distinction between closed and unclosed structural units corresponds closely to the Gestalt principles of good figure and closure, though they are not stated in quite the same way.

Since the critical prediction concerned recognition-memory errors, additional data were collected to help determine the source of errors. A possible source of errors is the initial encoding of a figure. For example, if a subject did not notice the break between Segments c and d in Part C of Figure 6, he would surely confuse it wtih Part A of Figure 6. To check initial encoding, therefore, each subject was asked to draw each pattern immediately after he or she had seen it. If he

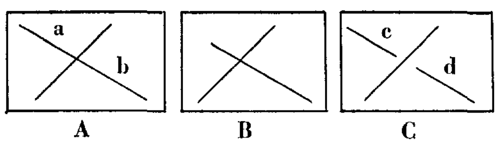

FIGURE 6. An original pattern (A), a distractor that doesn't have different structural units (B), and a distractor that does have different units (C), Experiment 3.

A.

B.

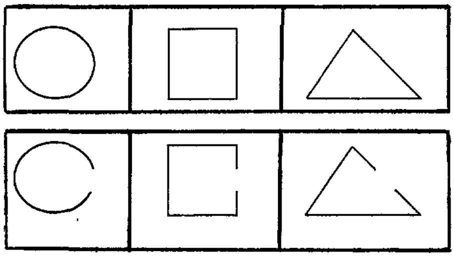

FIGURE 7. Illustration of rule of closure. (Patterns in Part A form different structural unit than their counterparts in Part B.)

made an error in his immediate drawing corresponding to a discrepancy in a distractor that he was to see later, then that test would not be counted. As a further check on initial encoding, subjects were asked during recognition testing whether they had remembered the pattern using a perceptual or a verbal code. A tendency to code patterns verbally might lead to the same kinds of errors as predicted by the structural hypothesis. For example, Part A and B of Figure 6 are both more likely than Part C to be called an X, and hence would be more likely to be confused if verbal labels had been used to code them into memory. Finally, subjects were asked to indicate during the recognition test what the difference was between the two test patterns, that is, the original pattern and its distractor. It was hoped that this response might indicate how the two figures were encoded.

Method

Materials. The materials consisted of a study set of eight patterns along with four test sets, each containing eight pairs of test patterns and distractors. The eight patterns were selected from those used in Experiment 1 on the condition that distractors of interest could be generated easily for them. Two of the patterns from the study set are shown to the left in Figure 8. Four alterations on each pattern generated four distractors per pattern, as illustrated in Figure 8. The first two alterations *removed* part of a line from the pattern; the last two alterations *added* or extended a line in the pattern. For each pattern the four alterations added or removed a line of exactly the same length. The first and third alterations generated distractors that, according to the parsing rules, contained the same structural units as the unaltered pattern. To illustrate, the original Pattern A in Figure 8 is segmented by the rules into two unclosed angular units, abc, connected to a third angular unit, de. Distractor A-1 is obtained by deleting Line c, but A-1 still contains similar unclosed angular units to those of the original. Distractor A-3 is obtained by extending Line e, but without changing the structural units. On the other hand, Distractor A-2 is obtained by deleting Line e, which destroys a former structural unit (Angle ed). Similarly, Distractor A-4 is obtained by extending Line c, producing a *closed* structural unit at the top which mismatches the former, unclosed, unit in pattern A.

As a further illustration, the S pattern in Figure 8 is analyzed by the rules as two connected unclosed curves (f and g). Distractor B-1 is derived by shortening Curve f, B-3 by extending Curve g, but neither change alters the structural description of the pattern. On the other hand, Distractor B-2 is derived by shortening Curve g, with the result that the lower unit disappears; Distractor B-4 is derived by extending Curve f, with the result that a completely different structural unit (a closed curve) is formed at the top. So by adding or deleting a small piece of line, we created either a theoretically similar or a very different distractor, depending on how the change was made.

The original patterns and distractors were drawn in ink on graph paper, and then xerox copies were made. The xeroxed patterns and distractors varied in length from 2.5 cm to 3.2 cm. A xerox of each study pattern was scotch-taped to a 3 × 5-in. (7.62-cm × 12.70-cm) card. A test pair was constructed by scotch-taping a pattern and one of its distractors side by side on a 3 × 5-in. (7.62 × 12.70-cm) card. The original pattern appeared equally often on the left and right sides of the test card. Four different test sets of eight cards were constructed, containing a different distractor for each original pattern. Each test set contained two of each of the four kinds of distractors.

Procedure. Twenty graduate and undergraduate psychology students participated in the experiment. Each subject was told that the line drawings would have to be reproduced from memory. Each of the original patterns was shown to the subject, one at a time, and he had as much time as he wanted to look at each drawing. When he indicated he was sure he could reproduce the drawing from memory, it was removed and he tried to reproduce it. No feedback was given regarding his drawing. No subject took less than 2 or more than 4 min to study the set of eight drawings. The order of presentation of the patterns was counter-balanced across subjects.

Approximately 10 days later, each subject returned and was given a forced-choice recognition task on a test set involving eight pairs of figures. As each pair was presented, the subject was asked to indicate on the score sheet (a) which of the figures he had seen before, (b) how confident he was of his decision, (5 = certain, down to 1 = guess), (c) whether he felt he had encoded the original pattern perceptually or verbally, (d) how the two figures (the distractor and the original) differed

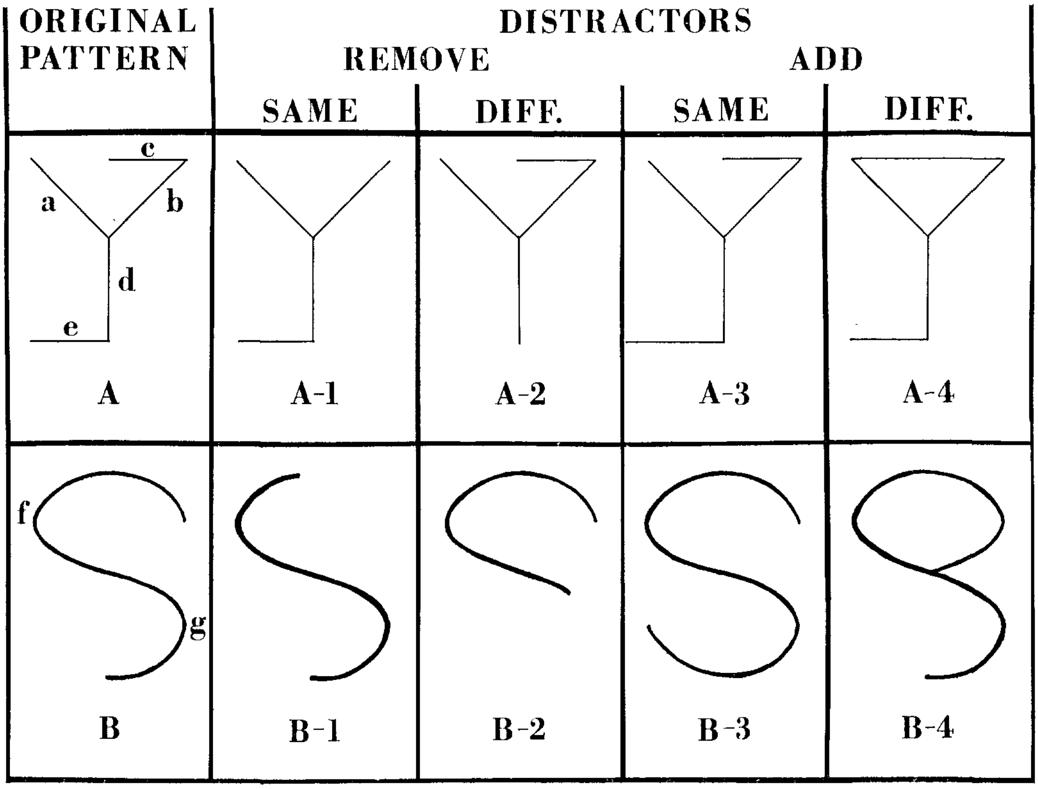

FIGURE 8. Two original patterns and their four distractors, Experiment 3.

from one another. Each of the test sets was presented to five subjects. There were only eight observations per subject. Each test set was always presented in the same order. All subjects completed their test within 5 min.

Results and Discussion

Nineteen errors were made in the immediate reproductions of the patterns. None of these errors were the same as discrepancies used in later distractors, so the data from all subjects on all tests were counted. Sixteen of the reproduction errors were followed by a correct response in the recognition task.

Turning to the main result, distractors that contained the same structural units as their corresponding patterns attracted 16 recognition responses (errors), whereas distractors that contained different structural units attracted only 4 errors, $\chi^2 = 6.85$, $p < .01$. Also, 19 out of 20 subjects gave higher confidence ratings to decisions about pairs that contained structurally different distractors. This difference was significant by a sign test, $p < .001$.

Five subjects indicated that they had always encoded the patterns perceptually, and one subject indicated he always encoded the patterns verbally. The remaining 14 subjects indicated perceptual encoding 61 times, verbal coding 31 times, both perceptual and verbal coding 18 times, and failed to respond twice.

The results confirm the hypothesis that distractors containing the same structural units as the original pattern would be confused more often with it. Subjects made significantly more errors and were significantly less confident on pairs with structurally similar distractors. Such differential errors occurred despite controlling the physical amount of change in the lines added or removed from the target pattern.

The tendency to choose structurally similar distractors might have been aided by the tendency of many subjects to use verbal encoding with some of the studied patterns. Structurally similar distractors would tend to call forth the same label as the original pattern (for example, an X or S). To the extent that the original was encoded and retained solely in terms of a verbal label, all test figures calling forth the same label would be indistinguishable.

To check whether perceptual or verbal encoding influenced the results, recognition was calculated separately for items reported as coded in the two different ways. The percentage correct recognition proved to be practically identical—88% for patterns coded perceptually and 89% for patterns coded verbally. Thus, by this assessment, the type of subjective coding reported during learning did not influence later recognition memory. Moreover, subjects who used verbal coding subjectively reported "recognizing" the correct pattern during the test at a "visual level" before the old label they had used for it recurred to them. This suggests that the visual recognition is primary, and the overlap of elicited verbal labels is derivative and of secondary importance in this experiment. Similarity of verbal labeling for structurally similar patterns (as in our *same* distractors, Figure 8) does not detract from but rather strengthens the structural similarity hypothesis. If one asks why two of the figures elicited the same label, it is because the structural descriptions assigned to them by the perceptual system (simulated by our rules) were very similar. Thus, similarity of labeling can be viewed as another response indicator of similarity of structural descriptions, and is logically on a par with the recognition-memory judgments themselves rather than constituting an explanation of them.

General Discussion

In recent years, theories about the organization of memory have focused increasingly on the structural units involved in memory representations (see e.g., Fodor, Bever, & Garrett, 1974; Norman & Rumelhart, 1975; Tulving & Bower, 1974). However, up till now the types of units studied have been linguistic—phonemes, syllables, words, grammatical constituents, semantic primitives, and so on. The results of our experiments suggest that structural units in visual pattern memory also deserve attention. Such a line of research could form the basis for a "grammar" of visual shapes to complement the linguist's grammar of language.

The results suggest the utility of the retrieval cuing method for studying the structural units involved in memory representations of simple drawings. Fragments that comprise significant parts of a figure have the power to redintegrate the entire pattern from memory. Indeed, precisely this outcome was conjectured long ago by Köhler (1947). Moreover, the significant units of the figures studied seem to be those suggested by the Gestalt rules of common direction, proximity (minimal angle), and closure.

Regarding the demonstrated difference in effectiveness of good versus bad cues in Experiment 1, we have encountered two objections from colleagues. A first objection says that the result is a foregone conclusion because the good cue provides "more information" about the target pattern than does the bad cue. The problem with this objection is the slippery word "information": There is no *atheoretical* way to compute the information in a shape (or one of its fragments), whereas there are a few *theoretical* specifications of the information in a shape (see, e.g., Attneave & Arnoult, 1956) that would not declare all our good cues to have more information than our poor cues. When pressed, the more-information argument seems to be little more than an intuitive restatement of our structural units hypothesis, which is preferable because its rules go beyond intuitions.

The second objection is that the better retrieval capability of the good cue may be due to verbal labeling, that is to say, a good cue is more likely than a bad cue to elicit the same verbal label that the original pattern elicits. This alternative is implausible in light of the results of Experiment 2; the good cue elicited the same label as the original pattern much less frequently (14%) than the cue redintegrated memory of the pattern (90%), so it is unlikely that redintegration is driven by common verbal encoding. Also, in Experiment 3, subjects who reported verbally labeling the original patterns did not differ from nonlabelers in the frequency of later recognition confusion errors with structurally similar distractors. The results are consistent with the interpretation that the structural similarity of two figures is the cause of recognition confusion errors as well as common verbal labeling of the figures.

In the absence of strong alternatives, we tentatively accept the structural hierarchy notion, the Gestalt rules for figural segmentation, and the idea that natural subparts of a pattern comprise powerful retrieval cues for recalling that pattern.

REFERENCES

Attneave, F., & Arnoult, M. D. The quantitative study of shape and pattern perception. *Psychological Bulletin,* 1956, *53,* 452–471.

Bower, G. H. The multicomponent theory of the memory trace. In K. W. Spence & J. T. Spence (Eds.), *The Psychology of Learning and Motivation* (Vol. 1). New York: Academic Press, 1967.

Fodor, J. A., Bever, T. G., & Garrett, M. F. *The psychology of language.* New York: McGraw-Hill, 1974.

Glass, A. *Outline of a theory of pattern recognition.* Unpublished doctoral dissertation, Stanford University, 1975.

Goldmeier, E. Similarity in visually perceived forms. *Psychological Issues, Monograph,* 1972, Vol. 8(1).

Guzmán, A. Decomposition of a visual scene into three-dimensional bodies. In A. Grasselli (Ed.), *Automatic interpretation and classification of images.* New York: Academic Press, 1969.

Köhler, W. *Gestalt Psychology.* New York: Liveright, 1947.

Norman, D. A., & Rumelhart, D. E. *Explorations in cognition.* San Francisco: W. H. Freeman, 1975.

Palmer, S. E. *Structural aspects of perceptual organization.* Unpublished doctoral dissertation, University of California at San Diego, 1974.

Reed, S. K. Structural descriptions and the limitations of visual images. *Memory & Cognition,* 1974, *2,* 329–336.

Tulving, E., & Bower, G. H. The logic of memory representations. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation* (Vol. 8). New York: Academic Press, 1974.

Wertheimer, M. Laws of organization in perceptual forms. In W. D. Ellis (Ed.), *A Source Book of Gestalt Psychology.* New York: The Humanities Press, 1938.

(Received October 20, 1975; revision received January 20, 1976)