CHAPTER 16

An Object Perception Approach to Static and Kinetic Subjective Contours

Philip J. Kellman and Martha G. Loukides

Subjective contours are among the most fascinating of perceptual phenomena. Although subjective contours and figures have often been considered isolated curiosities in perception, they have received considerable attention, both empirical and theoretical. In this chapter, we suggest that this attention has not been misallocated, and that subjective contours are not a peripheral curiosity. Specifically, we claim that subjective contours are examples of a process of unit formation that is pervasive in ordinary visual perception. This process underlies both subjective figure perception as well as a more common ability: the perception of the unity and boundaries of partly occluded objects.

The organization of this chapter parallels the development of our own thinking about subjective contours and object perception. In the first section (Direct Tests of Brightness and Configural Factors), we report empirical results on the persistent theoretical questions of the causal status of brightness and configural factors in subjective contours. These results, along with others, have led us to dispense with brightness factors in explaining subjective contours. In the second section (Subjective Figure Perception Across Space and Time: Kinetic Subjective Contours and Related Phenomena), we focus on the time dimension in subjective contour perception, introducing some new motion phenomena and noting older phenomena that a general theory should encompass. In the third section (Prospects for an Object Perception Theory: The Identity of Subjective Figure Perception and Ordinary Perception of Partly Occluded Objects), we present the fundamental claims of a new object perception theory, asserting that an identical unit formation process underlies subjective figures and perception of partly occluded objects in ordinary perception. In the fourth section (A Discontinuity Theory of Unit Formation in Subjective Figures and Partly Occluded Objects), we sketch a computational model of this unit formation process. Finally, in the last section (Prospects and Problems for a Discontinuity Theory), we note some prospective empirical tests of the theory, some relations to other theories, and some unresolved issues.

Direct Tests of Brightness and Configural Factors

Much debate about subjective contours has concerned the causal status of processes of brightness perception as opposed to configural factors (see Parks, 1984, for a review). By *configural* we mean factors having to do with the shape and/or arrangement of inducing elements: we will not distinguish for now among theories emphasizing simplicity, probability, local depth cues, and so on. A general strategy of research has been to try to produce subjective contours from brightness factors alone (Jory & Day, 1979; Kennedy, 1978a), or to produce subjective contours from configural effects in the absence of facilitating brightness factors (Halpern, 1981; Kanizsa, 1979; Kellman & Cohen, 1982; Parks, 1979; Prazdny, 1983). Suggestive findings on both sides of the issue have emerged but have not generally been regarded

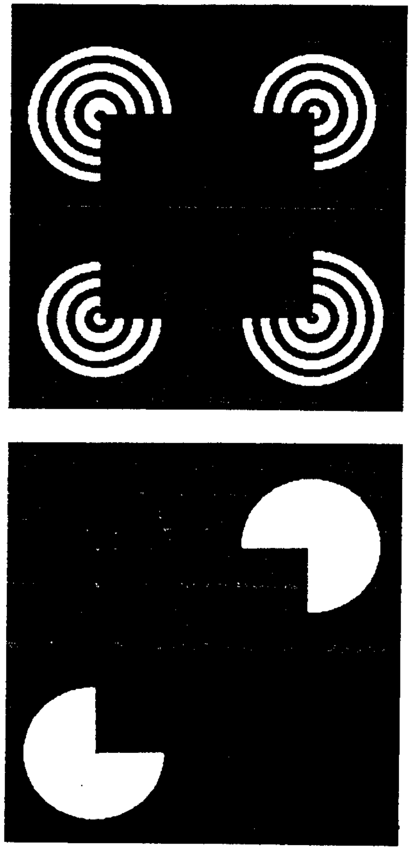

FIGURE 16.1. Subjective figures in physically and perceptually homogeneous space. (See text.)

as conclusive. One reason is that brightness and configuration are confounded in ordinary displays: inducing elements always have shape and normally differ in brightness or color from a background. Post hoc arguments for subtle configural or brightness effects can usually be made for any display in which subjective figures occur. Contributing to this explanatory latitude has been the imprecise specification of both configural (Coren, 1972; Kanizsa, 1974; Gregory, 1972) and brightness (Day & Kasperczyk, 1983b; Jory & Day, 1979) factors.

Configural Factors: A Direct Test

Our early research on this topic (Kellman & Cohen, 1982) attempted to test configural factors in the absence of induced brightness effects, using displays of the sort shown in Figure 16.1. The display in Figure 16.1b was also developed for similar reasons by Prazdny (1983). Using black and white inducing elements on a medium gray background, we sought to test subjective figure perception in the absence of any brightness differences across the subjective contours. Because individuals might have different luminance averaging characteristics, we tested three different proportions of black to white in the striped inducer displays (Figure 16.1a). Results were generally clear: naive subjects always reported robust subjective figures in these displays, and for every subject, at least one display was reported to have no figure–surround brightness difference. However, subjects often did report brightness differences, even with inducers containing equal proportions of black and white. Moreover, it was hard to rule out the possibility of subtle, unreportable brightness differences that might have supported contour perception.

Recently, we completed a more compelling test of contour perception in the absence of perceived brightness differences. Using displays of the sort shown in Figure 16.1a and 16.1b generated on a high-resolution computer system (Apollo DN660) with 256 gray scale values, we gave subjects precise control of the brightness characteristics of the displays. Against a constant, medium gray background subjects could simultaneously raise or lower the intensities of the black and white areas of the inducers to produce or eliminate brightness differences between the central area and the surround. For example, when these intensities were raised to their highest values, their space-averaged luminance far exceeded the luminance of the surround, inducing a darker central area. Similarly, by lowering the intensities of the black and white areas, the central area could be made to appear lighter than the surround. The black and white areas could be simultaneously raised or lowered through 100 equal brightness steps; between the extreme high and low values the brightness of the central area varied by about 3 Munsell steps. Thus, each step in subjects' ad-

justments varied the brightness of the central area by a tiny amount (about 0.03 Munsell steps on average). Subjects were shown two keys on a computer keyboard that varied the brightness up or down by single steps and two others for coarse adjustment that changed the brightness up and down by 10 steps per keypress.

We familiarized subjects with the subjective contour phenomenon using pictorial displays in which subjective contours are known to occur and others in which they are known not to occur; the latter procedure may be important in avoiding guessing tendencies (Kellman & Cohen, 1984). A pretest checked subjects' individual matching tendencies. Subjects were shown a rectangle composed of alternating black and white stripes, each one pixel wide. From the subjects' viewing distance, individual stripes were not visible and the rectangle appeared as a homogeneous gray. Subjects were instructed to adjust the color of the rectangle to make it match the background and disappear. All subjects succeeded, indicating that each could match the gray surround by equal amounts of black and white somewhere in the available range of intensities.

The subjective figure displays were presented initially with brightness settings making the central square appear somewhat darker than the background. Subjects were told to adjust the brightness of the display through the entire available range and to (1) try to make the edges of the central square go away, and (2) try to make any brightness difference between the central area and the surround go away.

Results were unequivocal. Every subject was able to eliminate all brightness differences between the central area and the surround. However, subjects were in general unable to eliminate clear subjective contours at any point in the adjustment range,[1] despite the fact that every subject was able to adjust the central square from appearing darker to appearing lighter than the surround by very small steps. Subjects were asked to rate contour strength for the initial version of the display (with central area darker than the surround), using an ordinary Kanizsa triangle display (black inducing elements on a white background) as a standard of 10. Subjects were also asked to rate the contour strength again after they had adjusted the brightness of the central area and surround to be an exact match. For the striped displays (Fig. 16.1a), ratings of the 11 (out of 12) subjects who could never eliminate the subjective figure ranged from 4 to 10, and averaged 6.8. This was slightly lower than their rating prior to brightness matching of 7.5. For the other display (Fig. 16.1b), ratings ranged from 5 to 10 and averaged 7.1 at the point of matched brightness, compared to initial ratings of 9.5.

The existence of robust subjective contours in the total absence of surface quality differences indicates configural causation. The results provide strong evidence for edge perception in both physically and perceptually homogeneous space. While abrupt surface quality differences across perceived boundaries may ordinarily be causes or consequences of those boundaries (Kanizsa, 1979; Koffka, 1935), they are not necessary for the visual perception of edges.

Our data suggest that brightness differences between center and surround can enhance the strength of perceived contours. This observation does not support a causal role for brightness factors in producing contours in physically homogeneous space. In these displays, local brightness differences due to contrast are confined by edges that also exist in their absence. These differences can serve to emphasize those edges.

Brightness Factors: A Direct Test

The existence of examples of configural causation of subjective contours does not logically exclude the possibility that induced brightness effects can also be causal. Can subjective edges ever be initiated by pools of induced brightness? Many have suggested that such brightness factors are in general the cause of subjective contours (Brigner & Gallagher, 1974; Day

[1] The single subject who reported contour disappearance seemed anomalous in several respects. His contour strength ratings were unusually low throughout the experiment. For the initial versions (before brightness matching) of the striped and solid inducer displays, he gave contour ratings of 3 and 3, respectively, while no other subjects gave ratings less than 7 and 6. Moreover, his reports of contour disappearance were not restricted to the range in which he reported that the brightness of the center matched the surround.

& Kasperczyk, 1983b; Frisby & Clatworthy, 1975; Jory & Day, 1979). Subjective contour-inducing elements, it is asserted, enhance or decrease the brightness of adjacent regions by retinal brightness contrast (perhaps along with some other local brightness effects). Regions of altered brightness are joined by the spreading of brightness, perhaps according to configural rules, across the areas between inducers. These claims have been difficult to evaluate directly. Many purported examples either contain configural confounds or appear to be somewhat different phenomena. For instance, displays may contain clearly enhanced brightness but rather diffuse edges (e.g., Kennedy, 1976a).

Recently, we completed a direct test of the causal status of brightness factors. Using a high-resolution CRT display, we created pools of enhanced brightness by "seeding" pixels of higher intensity in certain display areas, to see if these could then be assimilated into subjective figures. Individual pixels were undetectable and were spaced so as to create no abrupt boundaries. We used displays that do not normally give rise to subjective contours but contain the same amount of real contour, bounding a central square area, as do displays that ordinarily do give strong subjective contours. Specifically, we used four inducing elements in the shape of a cross or plus sign; such displays have previously been shown to be poor inducers of subjective contours (e.g., Kanizsa, 1979; Kellman & Cohen, 1982). We also used the same seeding patterns in the absence of any inducing elements. As control displays we obtained judgments about subjective figures and edges using the cross displays with no seeding. The seeding patterns were arranged to mimic the pools of brightness that would obtain in a subjective figure display with very effective inducers. We used seeding patterns that might arise from partial circle elements, as in Figure 16.1b, if all four inducers were black on a white background. To determine the induced brightness gradient, we developed a simple model of lateral inhibition in which inhibition of each retinal element was proportional to the luminance of the display point projected to that element, and inhibition declined quickly with inter-element distance. We also "cheated" in favor of brightness causation in two ways. First of all, we did not seed enhanced brightness anywhere outside of the central square area. Second, we seeded more than predicted by the model along the unspecified boundaries of the central square area. (This enhancement was added after no subjective contour induction was observed without it.)

The results showed very few reports of a subjective square or of subjective edges in the seeded displays, with or without the corners being specified by real contours. More frequent were uncertain reports from a few subjects of a "fuzzy diamond," "a circle," or a "plus sign." All of these reports occurred equally or more often with the unseeded cross-shaped inducing elements. There was no evidence in this experiment that subjective figures can be initiated by diffuse pools of enhanced brightness.

Some caution is required, however. Negative results such as these do not rule out that under some other conditions, induced brightness could lead to subjective contour perception. Our primitive lateral inhibition model may not have provided appropriate patterns of seeding. Nevertheless, we believe that the brightness seeding method does isolate the central claim of brightness-based theories of subjective figures. Such theories must predict the existence of a class of brightness seeding patterns that do not contain abrupt luminance changes and are sufficient to produce subjective figures. On this point, it must be added that our informal manipulations of brightness seeding, apart from the specific model we tested, have been no more successful at producing subjective contours.

Taking these results at face value, we are very pessimistic about the possibility that subjective contours and figures are initiated by areas of altered brightness.

Overview of Configural and Brightness Factors

The results of our parallel investigations of configural and brightness causation of subjective contours and figures are clear and complementary. Certain configurations produce subjective figures in the total absence of physical and perceptual differences in surface qualities. Conversely, it has not been possible to obtain direct evidence that diffuse brightness alteration can initiate the phenomenon. Subjective contours seem to demand explanation in terms of the shapes and arrangements of visible areas. So far we have said little about what properties of shape or arrangement are relevant, a matter which itself has been controversial. After examining some other phenomena relevant to a theory of subjective figures, we consider this question later on in this chapter.

Subjective Figure Perception Across Space and Time: Kinetic Subjective Contours and Related Phenomena

Recently discovered phenomena indicate that unit formation processes operating across space in subjective figure perception have analogues in the temporal domain; subjective figures can be specified by figural information given over time by motion. Some of these phenomena have been reported previously, whereas others are described here for the first time. They provide additional evidence that subjective figures are only peripherally related to brightness perception. More importantly, they suggest connections to ordinary object perception, in which the important role of information given over time is well known (e.g., Gibson, Kaplan, Reynolds, & Wheeler, 1969; Kellman & Spelke, 1983).

Kinetic Subjective Contours

Kellman and Cohen (1984) reasoned that if information separated in space is used by the visual system to create subjective figures, then such interpolation across time might also be possible. Figure 16.2 shows how such "kinetic subjective contours" might be produced. When the inducing elements are partly occluded in sequence as by a rotating figure (of the same color as the surround), the figure is clearly seen. These percepts occur despite the design of the displays so that no subjective figure can be detected in any stationary view. In the most robust version of this phenomenon, the inducing elements rotate around the central area, and the subjective figure is seen in a constant position. No explanation in terms of brightness contrast, assimilation, etc. seems remotely plausible with these moving displays. It is more likely that kinetic subjective contours are related to the ordinary perception of moving objects or stationary objects by moving observers (see below).

Specification of Inducing Elements Outside of the Luminance Domain

The possibility that configural factors are causal but brightness factors incidental in subjective contours has led to the attempt to specify inducing elements outside of the luminance domain. Prazdny (1985) reported failure to obtain subjective contours when the inducers were specified only by stereopsis or by motion characteristics. In his displays, both the inducing elements and the surround were covered with black and white random dots. When the inducers were defined stereoscopically, they appeared in front or in back of the surrounding surface, and when defined by motion, the streaming of random dots in the inducing areas differed from the motion characteristics of the background. Although Prazdny concluded that his results were consistent with the need for subjective contour-inducing elements to be defined in the luminance domain, he noted that certain incidental properties of his displays may have prevented subjective contour formation (for example, streaming of random dots made inducing circles appear as holes in a surface).

We have recently produced robust subjective figures in which the inducing elements were specified outside of the luminance domain. We also used motion to indicate the boundaries of the inducers, but they were designed to accrete and delete background texture as they moved. They thus appeared in front of the background surface (Gibson, Kaplan, Reynolds, & Wheeler, 1959). Figure 16.2b illustrates the displays. The three square elements had the same random dot surface as the surround and were not visible when stationary. As they rotated, they became visible atop the surrounding surface, because the points of the square progressively occluded and disoccluded elements of the background. The inducers were themselves progressively occluded as they moved behind the points of a central triangle. A complete triangle with crisp boundaries was seen, although the displays were constructed so that none is physically present.

These subjective figures demand explanation

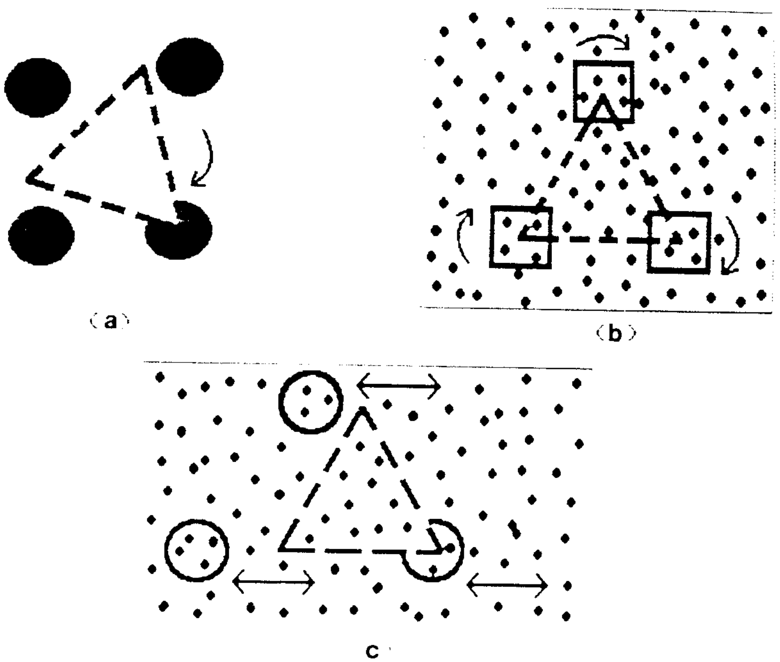

FIGURE 16.2. Schematic illustrations of subjective figures dependent on information given over time. a. A kinetic subjective figure. The black circles are partially occluded at different times as they would be by a triangle of the background color rotating in the plane. A clearcut subjective triangle is perceived. b. Kinetic specification of inducing elements. When stationary, a single surface with random-dot texture is seen. When the three squares rotate around their centers, they progressively occlude and reveal background texture, and are themselves occluded by vertices that could be part of a central triangle. A central triangle with clear edges is perceived. c. A kinetic subjective figure with kinetically specified inducing elements. Again the field is seen as a single surface when stationary, but when the circular areas translate back and forth, they become visible, and from their sequential occlusion, a central triangle is perceived.

in terms of configural factors, perhaps, as we argue below, in terms of relationships of figural discontinuities over space and time. Unless subjective figures given by kinetic specification of inducing elements are a wholly different phenomenon, it appears that the definition of inducers in the luminance domain as well as concomitant brightness effects in subjective figures are thoroughly incidental.

Although motion is involved in this new display, it is more of a static subjective figure, rather than a kinetic one, as we have used these terms. The reason is that all of the inflection points of the boundary of the subjective figure are specified simultaneously.

Kinetic Subjective Contours with Kinetically Specified Inducers

Can kinetic subjective figures also be produced using inducing elements specified outside of the luminance domain? We recently answered this question in the affirmative, using displays such as the one diagrammed in Figure 16.2c. Once again, the inducing forms had the same random dot surface appearance as the surround and were thus invisible when stationary. As the three circles translated laterally across the display, they accreted and deleted background texture and were themselves progressively occluded at different times by the corners of a central triangle. A clearly bounded, central triangle was seen, although once again, these boundaries did not exist in the actual display.

Implications of Motion Information in Subjective Figure Perception

These motion phenomena significantly extend the domain of discourse for attempts to explain subjective figure perception. The many similarities between static and kinetic subjective figures (Kellman & Cohen, 1984) suggest that a unified theoretical treatment might be possible. Moreover, taken together, they seem less likely to be incidental or isolated phenomena, and more likely to be indicators of a general unit formation process that functions in ordinary perception. In the next section, we suggest what that process might be.

Prospects for an Object Perception Theory: The Identity of Subjective Figure Perception and Ordinary Perception of Partly Occluded Objects

The wealth of previous research on subjective contours has produced few clear connections to ordinary perception. Subjective contours have remained primarily a laboratory curiosity; we almost never encounter them apart from our encounters with perceptionists.

Here we propose a new theory of why subjective contours exist and what more pervasive phenomena they can help us to understand. The central claim is this: Subjective contours are produced by the processes of unit formation that enable us to perceive objects despite occlusion.

The Problem of Occlusion

Perceiving objects and surfaces despite partial occlusion is a basic and ubiquitous aspect of normal perception. Consider the photograph of a simple scene in Figure 16.3. Although the scene is not especially cluttered, most of the objects are partly occluded by other objects. Light is reflected to an observer from only some of any object's surfaces, both because objects occlude parts of themselves and because occlusion of objects by others is pervasive. Many objects reflect light from spatially separated areas; in the photograph, this is true of each of the subject's feet, as well as the chaise lounge on which he is seated. We nevertheless perceive proximally separate light-reflecting areas to be part of a single distal object, with clear boundaries and shape. An additional complication is that adjacent light-reflecting areas that are parts of the same object may not be adjacent parts of the object. In the photograph we are able to detect that the subject's left leg does not merge with his right ankle but continues behind it.

The visual perception of objects and surfaces despite occlusion has been the subject of considerable study (Gibson, 1979; Kanizsa, 1979; Kellman & Spelke, 1983; Michotte, Thines, & Crabbe, 1964). Michotte et al. termed the perception of partly occluded areas *amodal completion*, meaning that the occluded areas are phenomenally present but do not possess the attributes of a sensory modality. For example, in the photograph, although one perceives the subject's left shoe as continuing behind the table leg, one would be unable to answer questions about whether there are scuff marks on the part of the shoe behind the table leg. Subjective figures, in contrast, are considered to be examples of modal completion, because the surface qualities of the completed surface are apparent to the perceiver.

Some have pointed out relationships between modal and amodal completion (Kanizsa, 1979; Michotte, Thines, & Crabbe, 1964). For example, it has often been observed that whereas subjective figures are examples of modal completion, the inducing elements are seen to continue behind them, i.e., the inducing elements are completed amodally (see Minguzzi, chap. 7; Sambin, chap. 14). Despite these connections,

FIGURE 16.3. *Portrait of Yasuo Kuniyoshi*. New York, Metropolitan Museum of Art. Copyright by Arnold Newman; reproduced by permission. (See text.)

perception of partly occluded objects and subjective figures have usually been considered as separable phenomena and studied independently. In the present analysis, we assert that subjective figure and partly occluded object perception are not simply connected in interesting ways, but that they are in fact identical phenomena. In particular, we suggest that the process of detecting a unit visually is the same in the two cases; whether a detected unit appears modally as a subjective figure or amodally as a partly hidden object depends on an additional and independent step having to do with the arrangement in depth of detected units.

Some Illustrative Phenomena

Two phenomena illustrate this claim. First, it has occasionally been pointed out (e.g., Bradley & Petry, 1977) that subjective figures can sometimes be seen as amodally, rather than modally, completed. Figure 16.4 shows a Kanizsa-type triangle, ordinarily seen as a modally completed triangle overlaying three black disks on a white background. It can also be seen as three points of a triangle *behind* the white surface showing through three windows or holes in the white surface. The black areas are then seen through the holes as a more distant surface. In Figure 16.4b some lines have been added to emphasize the "windows" interpretation. An especially vivid and intriguing example of this alternative appearance is the subjective Necker cube created by Bradley and Petry (1977). Once noticed, the two appearances tend to alternate in time. This bistability is also characteristic of kinetic subjective figures (Kellman & Cohen, 1984).

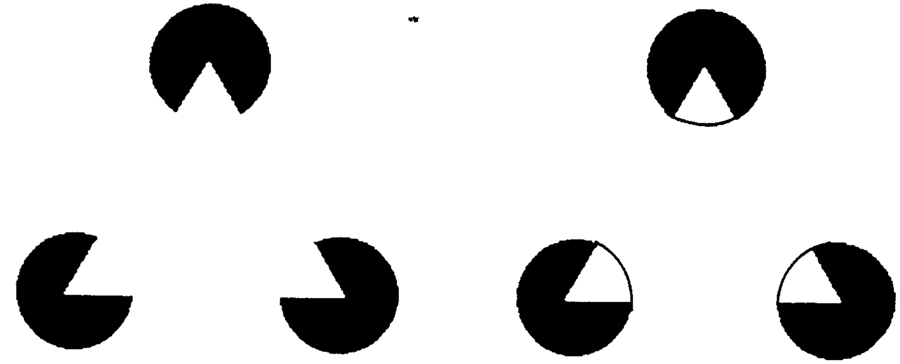

FIGURE 16.4. The bistability of ordinary subjective figures. a. A standard subjective figure display: after Kanizsa (1955a). b. The display in (a) with lines added to emphasize an alternative appearance: a unitary triangle is seen as partly hidden behind a nearer white surface, and the triangle's points are visible through "windows" or holes.

Before considering an interpretation of this bistability, consider another phenomenon, or rather, a set of three phenomena. Figure 16.5a depicts an occluded object; the unity of the triangle is easily seen. Figure 16.5b shows a subjective figure of the same shape. Figure 16.5c, after Arnheim (1974), shows an especially interesting intermediate case. The most fascinating aspect of Figure 16.5c is that, although it is a bounded, homogeneously colored figure, it perceptually resolves into two figures: a triangle and a rectangle. This phenomenon, the spontaneous resolving of homogeneous areas into separate entities, has received little attention. Arnheim (1974) mentions it as an indication of Gestalt tendencies toward simplicity. The second fascinating aspect of this display requires a longer period of fixation. You may notice that the triangle sometimes appears to be on top of the rectangle and sometimes behind it. In the former case, the triangle is modally completed, bounded by clear subjective contours, as in Figure 16.5b; in the latter case it becomes a partly occluded triangle, as in 16.5a, while the center part of the rectangle becomes modally present.

Separating Unit Formation and Depth Placement

If one considers modal and amodal completion to be separate phenomena, these examples of figures that alternate over time between subjective figures and partly occluded figures seem especially mysterious. We propose two complementary hypotheses that may serve to unify and clarify these phenomena:

1. The unit formation process for subjective figures and partly occluded objects is identical and is unvarying during the modal–amodal alternation. The units perceived to be in the array remain constant; in Figure 16.5c these units are the triangle and the rectangle.

2. The appearance of these units as in front of other surfaces (modal completion) or behind other surfaces (partial occlusion) depends on an independent depth placement process in which the depth relations of various units and surfaces is determined by available depth information.

When little or no depth information is available, the position of the unit relative to other surfaces will be bistable, as in Figure 16.5c, where no depth information specifies whether the triangle is in front of or behind the rectangle. When depth information is unequivocal, as in the case of most partly occluded objects in ordinary environments, an unambiguous position of the unit is perceived. The Kanizsa triangle in Figure 16.4a usually appears to be modally completed. This is probably due to the tendency of the circles to appear as figures on the white background, rather than as holes in it (Rubin, 1915). However, this fac-

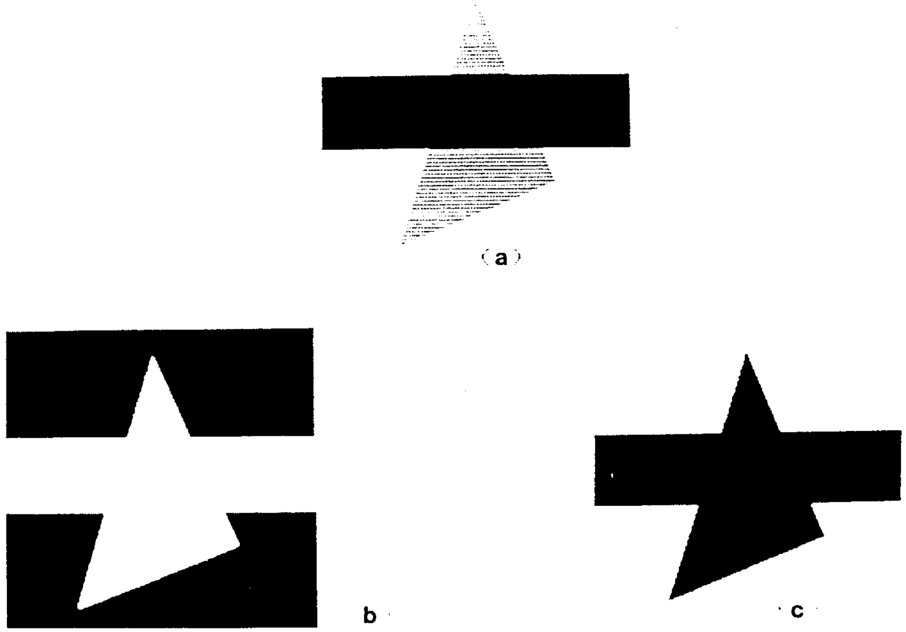

FIGURE 16.5. Equivalence of unit formation in subjective figures and partly occluded objects. a. A partly occluded object. b. An equivalent subjective figure. c. A spontaneously splitting figure. (See text.)

tor alone does not prohibit the alternative organization from being seen occasionally or with effort.

If these phenomena emerge from the same mechanisms, it remains to be determined what these mechanisms are. Recognition of a basic identity between subjective figures, partly occluded objects, and spontaneously separating figures raises many questions but also opens a number of avenues for investigating the specifics of the unit formation process. From our preliminary investigations, we believe that it is possible to specify the relevant conditions for perception of a unit. It appears that the relevant information is relatively local (i.e., it does not directly involve considerations of overall symmetry), and it involves variables of boundary shape exclusively, rather than, e.g., surface similarity.

A Discontinuity Theory of Unit Formation in Subjective Figures and Partly Occluded Objects

Because of space limitations, we present only a sketch of the model here. For convenience, we first treat static optic arrays and afterward indicate how information over time may be incorporated. The theory presented here is intended to be strictly a computational one in Marr's (1982) sense. That is, we do not even hint at a specific algorithm to accomplish the pickup of the information we describe, nor do we have any commitment to a particular realization in neural hardware. With regard to the latter, however, we have noticed with interest possible relations between our model and the neurophysiological hypotheses of Shapley and

Gordon (chap. 11) concerning nonlinear operators.

It is fairly clear how regions whose edges are defined by abrupt luminance or spectral changes may be detected (e.g., Marr, 1982). The domain of our unit formation model includes the cases where boundaries are perceived in regions where no such surface quality changes are detected. The following may roughly comprise the steps in formation of such units:

1. *A necessary initial condition is a spatial discontinuity in a luminance boundary.*

By discontinuity we mean the mathematical notion: if a boundary is described by a function, it is not differentiable at the point of discontinuity. Intuitively, this includes any abrupt change in direction of a contour. However, the relevant class of perceptual discontinuities may be larger than this, including all segments exceeding a certain sharpness of curvature (see below). An example of a discontinuity on a subjective contour-inducing element, e.g., in Figure 16.4a, is a point at which the circular boundary ends. In the homogeneously colored figure in 16.5c, the points at which the triangular section begins to extend upward from the horizontal edge are discontinuities.

A discontinuity is a necessary condition for a new unit to be formed. That is, we hypothesize that only when an abrupt change in a bounding contour is present can a figure function as a subjective contour-inducing element or as an object part that is connectable to other parts in occlusion situations. Such a discontinuity is not a *sufficient* condition for unit formation or segregation. A single subjective contour-inducing element, in the absence of others, is seen as a self-contained, irregularly shaped figure. Likewise, if the lower protrusion in Figure 16.5c is removed, the upper triangular protrusion alone does not tend to split off from the rest of the figure; again an irregular, bounded figure is seen.

2. *A unit is formed when a discontinuity is continuously relatable to others, and when the connections generated among discontinuities yield a closed form.*

More specifically, to be continuously relatable means that an edge along one side of a discontinuity may be connected to another without any discontinuity occurring between them, i.e., they may be connected via a straight line or a smooth curve. A major, and as yet untested, empirical prediction of the theory is that a determination of what a perceptual discontinuity consists of for purposes of step 1 will also provide an accurate criterion for relatability in step 2. In other words, if for step 1 a curve must change direction more sharply than some threshold in order to function as a subjective contour-inducing element, then we would expect to find, regarding step 2, that for the edges along two discontinuities to be perceptually unified, something less than that threshold of curvature sharpness must be possible at all points between them. An additional constraint here is that the connection between edges must proceed monotonically from one to the other; otherwise any two edges would be connectable without any discontinuity occurring between them.[2]

The discontinuities we have been discussing are retinally specifiable. Unit formation in our theory has the consequence that some retinal discontinuities do not correspond to discontinuities in the boundaries of perceived objects. For example, when a subjective triangle is seen in the standard Kanizsa display, two of the three retinal discontinuities in the boundaries of each inducing circle are seen as resulting from occlusion rather than from abrupt changes in the boundaries of an object. A potentially fruitful way of expressing step 2—that units are formed if discontinuities are continuously relatable to others—is that the visual system uses these rules to minimize the number of discontinuities assigned to perceived objects' boundaries.[3]

3. *A unit formed by steps 1 and 2 is positioned in depth relative to other surfaces in the array as determined by available depth information.*

In subjective figures, the unit appears in front of retinally adjacent surfaces. In the case of partly occluded objects, the unitary figure appears behind an intervening object or surface. In ordinary perception of partly hidden objects, in contrast to pictorial displays, the relative

[2] We are indebted to Martin S. Banks for pointing out the need for this constraint.

[3] We are indebted to Henry Gleitman for this insight.

depths of occluding and occluded objects are often very well specified.

Events as Discontinuities

When information is given over time, the unit formation process may follow the same framework, although both the initiating and relatability conditions must be broadened to include certain *events*. In the case of kinetically specified inducing elements, the application of the theory is straightforward. Initiating spatial discontinuities are specified by accretion and deletion of texture elements, rather than by changes in a luminance boundary. These spatial discontinuities function in the same way as those arising from luminance boundaries. The luminance domain per se has no special role, but serves as one way (the most common one) of specifying discontinuities.

Kinetic subjective figures and partly occluded objects detected over time raise some interesting issues, since the discontinuities are specified over time and are also separated over time as well as space. Regarding the initiating events, we propose that certain figural changes given over time function as discontinuities, analogous to abrupt changes in contour. We have not yet fully developed these ideas, but a basic parallel may hold between space and time. Continuity over time may be definable in terms of the class of projective changes on the retina that can occur from an object translating or rotating in space, or a stationary object viewed by a moving observer (Gibson, Kaplan, Reynolds, & Wheeler, 1969; Johansson, 1970). Other retinal changes that are not projectively invariant may constitute temporal discontinuities. Such discontinuities in isolation may be seen as figural deformation, although they are more likely to be seen as occlusion (Michotte, Thines, & Crabbe, 1964).[4]

Relatability of the edges forming a discontinuity to others may depend on the same requirements of smooth curvature as in the static case. When the objects' parts are changing positions in space, however, the relatability criteria must apply to the real spatial positions of these parts at a given time, not to the retinal locations at which a part was last specified. Constraints probably exist regarding the continuity of figural motion necessary for temporally separated discontinuities to form kinetic subjective figures (Kellman & Cohen, 1984). This problem is minimized when kinetic subjective figures are specified in a stationary position by moving inducing elements, which may account for the slightly more robust subjective figures in this case (Kellman & Cohen, 1984).

As in static arrays, the formation of units from relatable discontinuity events minimizes the number of discontinuities—spatial and temporal—that must be assigned to perceived objects.

Prospects and Problems for a Discontinuity Theory

The present theory accrues some large burdens but also points toward a number of interesting predictions and some new approaches to problems in object perception. Space permits only brief mention of some of the main empirical and conceptual lines being pursued.

Much of our recent work has examined the hypothesis of identical unit formation mechanisms in subjective figures, partly occluded objects, and "spontaneously splitting" homogeneous figures. In our view, any subjective figure display should be transformable by certain simple rules into an equivalent partly hidden object display, and vice versa. Likewise, either of these displays should be transformable into a homogeneously colored display that spontaneously resolves into more than one figure. Our investigation of these equivalences has been highly confirming of these predictions. The transformation rules are readily discoverable by the reader; computer-generated displays are helpful for making rapid changes and comparisons.

These equivalences furnish convergent measures for empirical study of the initiating configurations and events in unit formation. A major challenge is to develop a measure of abruptness of curvature and find some range of values of such a variable that defines a spatial discontinuity with the explanatory power we hypothesize. Another source of convergence will be comparison of data on initiating contour formation with data about relatability. There are also questions about whether the abruptness of curvature can be defined in purely retinal terms. While a mathematical discontinuity in an object boundary consistently projects a retinal discontinuity, the sharpness of certain curved contours may depend on their distance, projected size, and orientation.

Initial attempts to refine the discontinuity notion are focusing on stationary arrays. If precise characterization of the spatial discontinuity notion proves possible, a further challenge will be incorporating it with information about the spatiotemporal constraints on event-specified discontinuities. One intriguing hypothesis is that the spatial discontinuities in the static case may simply be limiting (or mildly degenerate) cases of more fundamental event-specified discontinuities. Such a view would fit with theories emphasizing the primacy of optical transformations in human visual perception (Gibson, 1979; Johansson, forthcoming).

Relations to Other Perspectives

The present theory bears important resemblances to previous views of object and edge perception. Because we are treating both subjective contour phenomena and perception of partly occluded objects, these ideas intersect with many older accounts. As noted above, however, few discussions have explicitly connected these domains. Kanizsa (1979) treats these phenomena separately for the most part, although he repeatedly notes similarities in the perceptual tendencies operating in the two domains. Michotte, Thines, & Crabbe (1964), while not explicitly discussing subjective contours, consider examples of modal along with amodal completion, suggesting that both follow Gestalt principles. The claim that subjective contours and partly hidden objects arise from the same unit formation process and differ because of an independent depth-placement step is, as far as we can determine, novel.

The specifics tentatively proposed for this common unit formation process have much in common with certain traditional Gestalt principles, especially good continuation (Michotte, Thines, & Crabbe, 1964; Wertheimer, 1923a). This is reassuring, since we believe that the older accounts and demonstrations concerning good continuation and simplicity tendencies identified important aspects of the problem. However, these notions have remained vague and have shown little predictive utility. Moreover, determining how such principles may be realized in a visual system, human or artificial, has been problematic. How might a rapidly functioning system consider all possible arrangements of an array and choose the simplest or best? Our tentative computational ideas have the virtues of suggesting precisely definable and relatively local stimulus attributes as the basis for unit formation. By developing a mathematically specifiable and local notion of contour discontinuity it may be possible to give precision to more elusive notions such as "good" continuation. Such a theory differs from a list of contributing principles, (e.g., Michotte, Thines, & Crabbe, 1964); for example, similarity of surface color and texture have no role in our model. Because our approach emphasizes local boundary properties, it contrasts with any approach emphasizing global symmetry, e.g., in the overall shape of subjective contour-inducing elements (Kanizsa, 1979). Along with others, this would seem to be a clear difference from coding theories (Leeuwenberg, 1981; Restle, 1981) which have been applied to occlusion cases. Some of the most heartening results of our investigations so far have been cases in which more global attributes such as symmetry and overall simplicity turn out to be the consequences of the application of local rules (cf. Marr, 1982).

Among subjective contour theories, our approach is perhaps closest to Coren's (1972) view that certain boundary configurations function as implicit interposition cues. We see our work as extending Coren's insights in several ways. We have suggested specifically what the initiating stimulus variable (a contour discontinuity) might be and we are developing some ways of determining its scope. Moreover, we hypothesize a relationship between the characteristics of the initiating discontinuities and the conditions that govern their connections to one another. Furthermore, we are not sure that the discontinuity notion must be characterized as initially involving depth, although its outcome certainly has consequences for the units that must be segregated in depth. While clarifying

[4] Another possibility is that the important effects of temporal changes may be to introduce momentary spatial discontinuities in the inducing figures.

these basic notions, we are also generalizing the theory to the problem of partly occluded objects and to perception based on discontinuities given over time as well as in space.

Much remains to be learned, but we see great promise in assimilating the subjective contour phenomenon to visual processes whose main function is the ordinary perception of unitary objects in a visual world pervaded by occlusion. By considering boundary discontinuities across space and time, a comprehensive understanding of occluded object perception, encompassing objects and observers both moving and stationary, may be possible.

Acknowledgments. This research was supported by NSF Grant BNS-10110 and a Swarthmore College Faculty Research Grant to Philip J. Kellman. The authors thank Tim Shipley, Kirk Swenson, Henry Gleitman, Elizabeth Spelke and Martin S. Banks for helpful discussions. The authors also thank Steve Platt for generously lending computer facilities, George Flickinger for expert photographic assistance, and James Crowell and Marian Staats for general assistance. Reprint requests should be addressed to Philip J. Kellman, Department of Psychology, Swarthmore College, Swarthmore, Pennsylvania 19081.