guage Development, N.A. Krasnegor, D.M. Rumbaugh, M. Studdert-Kennedy, and R.L. Schiefelbusch, Eds. (Erlbaum, Hillsdale, NJ, 1991).

5. E. Thelen, Dynamical systems and the generation of individual differences, in *Individual Differences in Infancy: Reliability, Stability, and Prediction*, J. Colombo and J.W. Fagen, Eds. (Erlbaum, Hillsdale, NJ, 1990).

6. E. Thelen and B.D. Ulrich, Hidden skills: A dynamic systems analysis of treadmill stepping during the first year, *Monographs of the Society for Research in Child Development, 56*(1, Serial No. 223, 1991).

7. E. Thelen, D. Corbetta, K. Kamm, J.P. Spencer, K. Schneider, and R.F. Zernicke, The transition to reaching: Mapping intention and intrinsic dynamics, *Child Development* (in press).

8. L.B. Smith and E. Thelen, Eds., *Dynamic Systems in Development: Applications* (MIT Press, Cambridge, MA, in press).

9. L.B. Smith and M.D. Sera, A developmental analysis of the polar structure of dimensions, *Cognitive Psychology, 24*, 99–142 (1992).

10. P. van Geert, A dynamic systems model of cognitive and language growth, *Psychological Review, 98*, 3–53 (1991).

11. S.S. Robertson, Behavioral chaos: Beyond the metaphor, in *Dynamic Systems in Development: Applications*, L.B. Smith and E. Thelen, Eds. (MIT Press, Cambridge, MA, in press).

Perceiving Objects Across Gaps in Space and Time

Philip J. Kellman and Thomas F. Shipley

Looking through a dense hedge, an observer may see only scattered spots of color from the scene behind. When the same observer looks while walking, however, the scene becomes clear. Somehow, the human visual system integrates spatially fragmentary information over time to achieve perception of objects and spatial layout. Such representations of bounded objects and surfaces in the three-dimensional (3-D) world are among the most remarkable products of visual perception. Obtaining them requires the solution of several difficult problems by the visual system.

One problem is detecting, from spatial variation in luminance, chromaticity, or texture, the projections of object boundaries. Work on edge-detection has produced a number of proposals about specific variables that might be useful and how they might be detected.[1] The outputs of such processing, however, are surprisingly remote from a description of the objects and boundaries in the world. The primary culprit, as in the hedge example, is occlusion. In the 3-D world, projection of an object to the eyes of an observer is more often than not interrupted by nearer objects. As a consequence, rarely do the projected shapes reaching the observer include all the surfaces of an object oriented toward the observer. Another consequence is that most projected edges mark locations where one surface ends, but the other passes behind the first. Locating edges, by itself, does not indicate which side is bounded.

Occlusion often causes a single object to project to the eyes in multiple, spatially separated regions. Conversely, sometimes homogeneous areas in the optical projection come from separate objects in the world. Figure 1 contains examples of both of these effects. The situation is complicated further when objects or observers move. Parts of an object that project to an observer's eyes at one moment become occluded, whereas previously hidden areas become visible. The momentary shapes of these projected areas bear little similarity to the shapes of the objects in the scene.

Despite the fragmentary nature of the input, human perceivers ordinarily have little difficulty detecting the unity and boundaries of objects. Explanations for this ability have been elusive. On some accounts, factors outside the stimulus, such as influences of familiarity, expectation, probability, or simplicity, are required. In terms of process, it is sometimes claimed that objects are determined by cognitive activity resembling general inference or problem-solving processes.[2] Our investigations of how spatial and temporal gaps are overcome in the perception of objects and boundaries offer support for a different perspective: Perception of objects and boundaries depends on particular stimulus relationships and a determinate process.[3]

The research points to a process common to a number of phenomena that on the surface seem diverse. We have identified the information used in the process and suggested a tentative account of the steps involved. The model applies to boundary interpolation in 2-D displays, but it also appears to extend straightforwardly to interpolation between surface boundaries in arbitrary 3-D positions. Some natural assumptions about visual processing of time-varying displays allow the model to account neatly for a number of results in spatiotemporal object perception, in which information is carried by motion. What follows is a

Philip J. Kellman is Associate Professor of Psychology at Swarthmore College and Adjunct Professor at the University of Pennsylvania. His research focuses on perception of objects, motion, and space in human adults and infants. **Thomas F. Shipley** is Assistant Professor of Psychology at the University of Georgia in Athens. Address correspondence to Philip J. Kellman, Department of Psychology, Swarthmore College, Swarthmore, PA 19081.

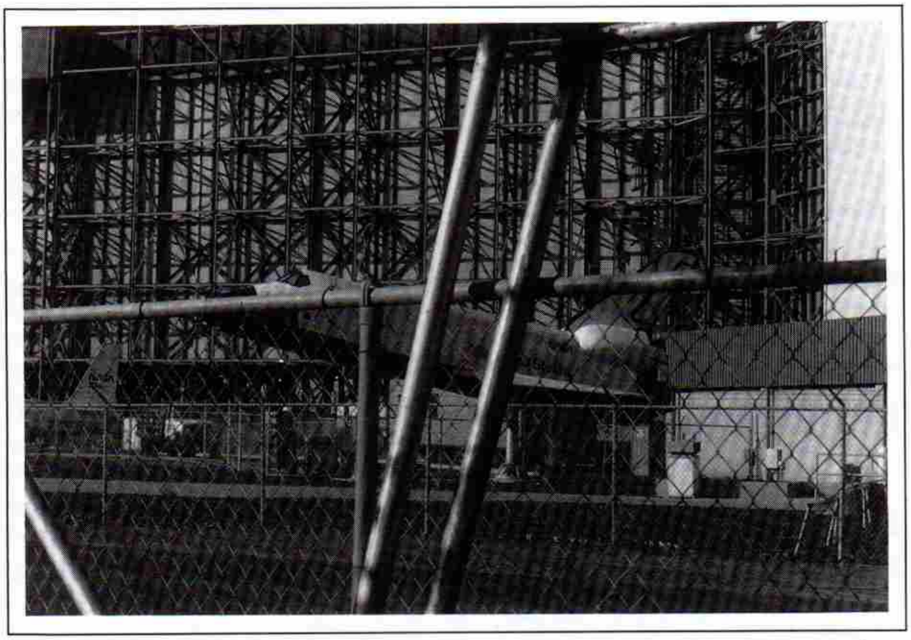

Fig. 1. The model of the space shuttle is visible despite projecting to the eyes from multiple locations interrupted by nearer objects. Also, parts of the fence in front and the structures behind have color and brightness similar to adjacent projected areas of the shuttle, yet are not seen as connected to it.

brief description of each of these aspects of the work.

DIVERSE PHENOMENA, SINGLE PROCESS

Perception of objects under partial occlusion is not the only case in which boundaries are perceived across gaps. Two others are shown in Figure 2, along with an example of partial occlusion (Fig. 2a). Figure 2b illustrates illusory contours. The perceived (interpolated) edges between the black inducing patterns are not defined by any physical gradient. The illusory figure seen in front of the inducing patterns is an example of what Michotte[4] termed *modal completion*. *Completion* refers to the perception of the whole object despite the gaps in the physical stimulus. The term *modal* signifies that the edges and surfaces supplied by the visual system come complete with sensory properties. One could, for example, judge the brightness of the interpolated white surface in the figure. In contrast, perception of occluded regions is *amodal*, lacking sensory attributes or modes. The black object in Figure 2a is perceived as a unit, but its occluded region produces no local sensations.

One key to our approach to understanding boundary interpolation is the idea that modal and amodal completion are not different processes. Instead, partly occluded boundaries and illusory contours are different-looking manifestations of an identical underlying process. When first proposed, this notion was surprising to some researchers because processes dealing with occlusion were often considered cognitive, whereas illusory figures, with their sensory accompaniments, were viewed as truly perceptual or sensory phenomena. In our view, the difference in appearance is superficial. What differs between illusory figures and partly occluded figures is whether the interpolated boundaries appear in front of or behind other surfaces in the scene. This difference in depth arrangement of perceived surfaces depends not on the process of unit formation,[5] but on depth information in the scene.

One argument in favor of this idea comes from figures in which, because of weak information about depth ordering, edges switch from illusory contours to partly occluded boundaries, and vice versa. An example is the type of display we have called a spontaneously splitting object (SSO). Figure 2c shows an SSO.[6] Although the entire black region is homogeneous, the visual system carves it into two objects. The intersecting boundaries of these objects are illusory contours and partly occluded edges. If the vertically elongated object is seen in front, its edges are illusory contours, while the edges of the horizontally extended object are partly occluded edges. With prolonged viewing, however, something curious hap-

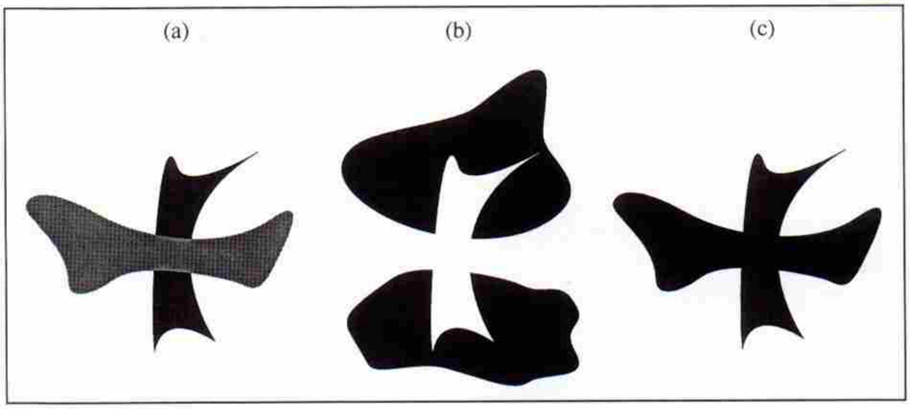

Fig. 2. Example of equivalent occlusion (a), illusory-figure (b), and spontaneously splitting object (c) displays.

pens. The object seen in front switches to behind, and vice versa. The shapes remain unchanged; what changes is their depth ordering. It would be possible, but cumbersome, to argue that there are two separate processes of modal and amodal completion involved. One would have to claim that at random times, each stops operating on one figure and switches to the other with perfect complementarity. A simpler interpretation is that the figural boundaries are given by a single interpolation process whose results do not change. What randomly alternates is the depth ordering of the figures, reminiscent of displays that undergo figure–ground reversals. Depth alternation is not surprising because there is no depth information specifying which object is in front. Such alternation does suggest that the visual system obeys the constraint that two different objects cannot occupy the same space at the same time.

Observations such as these led to the original proposal that these different examples of unit formation are manifestations of a single underlying process.[7] Empirical research has yielded strong support for the hypothesis. In brief, boundary interpolation seems to occur in virtually identical fashion in theoretically equivalent displays of illusory and occluded figures.[8] A number of variables, such as edge alignment, orientation, and spacing, affect edge interpolation equivalently in the two cases, across both simple and highly complex, randomly generated displays. Whether and how interpolation occurs depends on the relative positions of physically specified edges in the display, not on whether they are part of occlusion or illusory-figure displays. The phenomenological difference in these cases—what has been called amodal versus modal completion—has to do with the depth ordering of interpolated boundaries and other edges in the scene. Depth ordering, in turn, depends on available depth information. A striking consequence of the hypothesis is that it should be possible to construct displays in which an illusory contour connects to a partly occluded edge.[9] This is indeed the case, as Figure 3 illustrates.

CONDITIONS GOVERNING BOUNDARY INTERPOLATION

What relations among physically specified edges allow the visual system to interpolate new edges? A major issue among theories of visual completion is whether the relevant determinants are global or local. Global information might include the symmetry or simplicity of an object's overall shape. There is not, however, much evidence for global influences in unit formation, and we believe it to be primarily a local phenomenon.[10] Specifically, boundary interpolation depends on local edge orientations and their relations.

Some clues about the relevant information come from the ecology of perception. Of the several kinds of interpolated boundaries, partly occluded edges are most common in ordinary perception. Illusory contours, for example, are relatively rare outside perception laboratories. It is likely that the boundary interpolation process exists to deal with occlusion. Thus, the search for information governing the process might profitably consider the optics of occlusion.

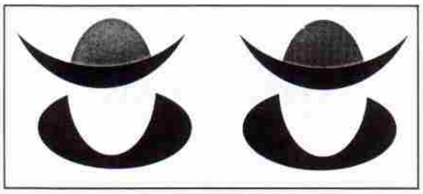

Fig. 3. Example of illusory contours joining partly occluded contours. The egg-shaped object has interpolated boundaries combining illusory-contour and partly occluded portions. Although the effect is visible in either the left or the right display, the depth relationships may be stabilized, and the effect made more striking, by viewing the two displays stereoscopically. The two views may be placed in a stereoscope or free-fused by crossing the eyes, so that each eye receives one of the two views.

A particularly helpful fact comes from projective geometry. Whenever one object partly occludes another, the optical projection contains sharp corners at the points of overlap. For example, in Figure 2a, there are four such sharp corners where the projected boundaries of black and gray regions meet. Formally, these corners are first-order, or tangent, discontinuities in the orientations of projected edges.

We have proposed that detection of this particular feature of optical projections is the starting point of the boundary interpolation process. The relevant tangent discontinuities (TDs) are those along the edges of extended regions.[11] The importance of TDs as starting points for the interpolation process has been supported by empirical findings.[12] The presence of TDs enhances the frequency of reports of illusory contours and the clarity of perceived contours, whereas smoothing TDs reduces or eliminates illusory contours.

TDs provide only the starting point, however. Although all occlusions give rise to TDs, not all TDs arise from occlusion. Some objects' shapes simply include sharp corners. What information could distinguish TDs arising from these two sources?

The visual system seems to solve this problem in combination with the problem of determining where occluded boundaries are located. Specifically, the solution depends on relative positions and orientations of the edges leading into TDs. When certain relations hold, edges are interpolated. We have labeled these relationships *conditions of relatability*. Two edges are relatable if they can be connected by a smooth (at least once differentiable), monotonic curve whose endpoints match the two edge tangents. We have proposed a formal criterion of relatability to express these relationships, which can be understood with reference to Figure 4a. If E_1 and E_2 are the edges of projected areas, leading into TDs, the tangent of each edge at the point where it meets the TD governs its relatability. If R and r are constructed, perpendicular to the tangents of E_1 and E_2, such that $R \geqslant r$ and θ is the angle between R and r, then E_1 and E_2 are relatable if and only if

$$0 \leqslant R \cos \theta \leqslant r$$

Figure 4b gives an example of two relatable edges; Figure 4c shows a pair of nonrelatable ones.

The relatability criterion embodies several principles. It implies that the visual system represents hidden edges according to a smoothness constraint[1] and in general does not construct corners or double inflections. Such a constraint exploits the ecological fact that objects tend to have (relatively) smooth boundaries and perhaps the fact that departures from smoothness in hidden edges would be difficult to anticipate. The utility of the criterion also depends on another principle from projective geometry: The optical projections of smooth curves and straight edges in the boundaries of objects will in general be smooth curves and straight edges. Thus, boundaries that are smoothly connected in the 3-D world will be relatable in their optical projections, even when partly occluded. Formally, the relatability criterion comprises necessary and sufficient conditions for a smooth (at least once differentiable) and monotonic curve to be constructed tangent to each of the two supporting edges.

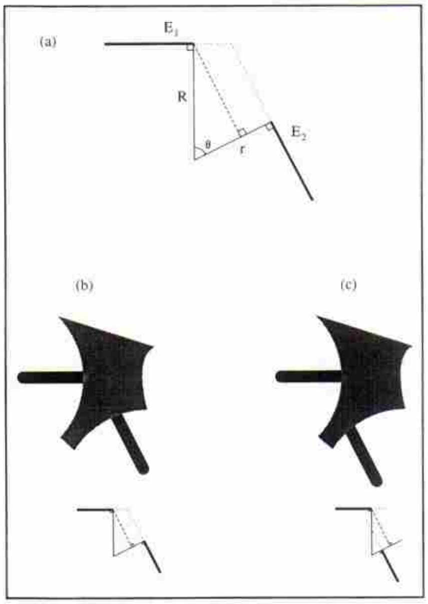

Fig. 4. Conditions of relatability. (a) Construction defining relatability (see text). (b) Example of relatable edges. (c) Example of nonrelatable edges.

The relatability criterion also contains the constraint that edges whose orientations differ by less than 90 deg ($R \cos \theta < 0$) are not relatable, a condition derived from empirical observation. There is little evidence that the boundary interpolation process creates edges bending through more than 90 deg. It should be noted that the limits of interpolation given formally are subject to perceptual thresholds. For example, misaligned parallel edges should not support interpolation, according to the model, but experimental evidence indicates that the process tolerates a small amount of misalignment, about 15 min of visual arc, before interpolation breaks down.

THE PROCESS OF UNIT FORMATION

Figure 5 summarizes our model of the process by which boundaries are interpolated and combined with locally specified boundaries to produce new perceived objects. The first two steps are not part of the model per se; they involve two earlier stages of visual processing that provide the input representation. The first, yielding a representation like Marr's primal sketch,[1] locates edges based on certain luminance changes across the optical projection. The second, as in Marr's 2.5-D sketch, assigns depth information to each point. Step 3 is the search of this representation for TDs, and orientations of edges leading into TDs are determined in Step 4. The test for relatability occurs in Step 5, and, in Step 6, smooth, monotonic connections are constructed between edge pairs passing the test.[13] Construction of an edge has the consequence that some TDs in the optical projection are perceived as loci of occlusion rather than as sharp corners of objects. Step 7 tests whether locally given plus interpolated edges form an enclosed area, in which case a unit (object, figure, or aperture) is perceived. Formation of such an area can lead to altered depth relations among surfaces at similar depths (Step 8), as a result of the operation of certain depth cues. In particular, the depth cue of interposition can be formalized as a set of rules operating on physically given and interpolated edges.[3]

BOUNDARY INTERPOLATION IN THREE-DIMENSIONAL SPACE

Interpolation processes have most often been illustrated and studied in 2-D, pictorial contexts. Recent observations and experiments suggest that depth information in the input representations plays an important role, however.[3,14] First, a pair of edges might line up in a 2-D projection, but, because of differences in their depth, they might not be smoothly connectable in 3-D. Interpolation does not seem to occur in these cases. Second, the interpolation process can produce edges that extend or curve in depth. These have been produced for both occluded objects and illusory contours, indicating that the process is sensitive to the 3-D orientations of the input edges. An example is

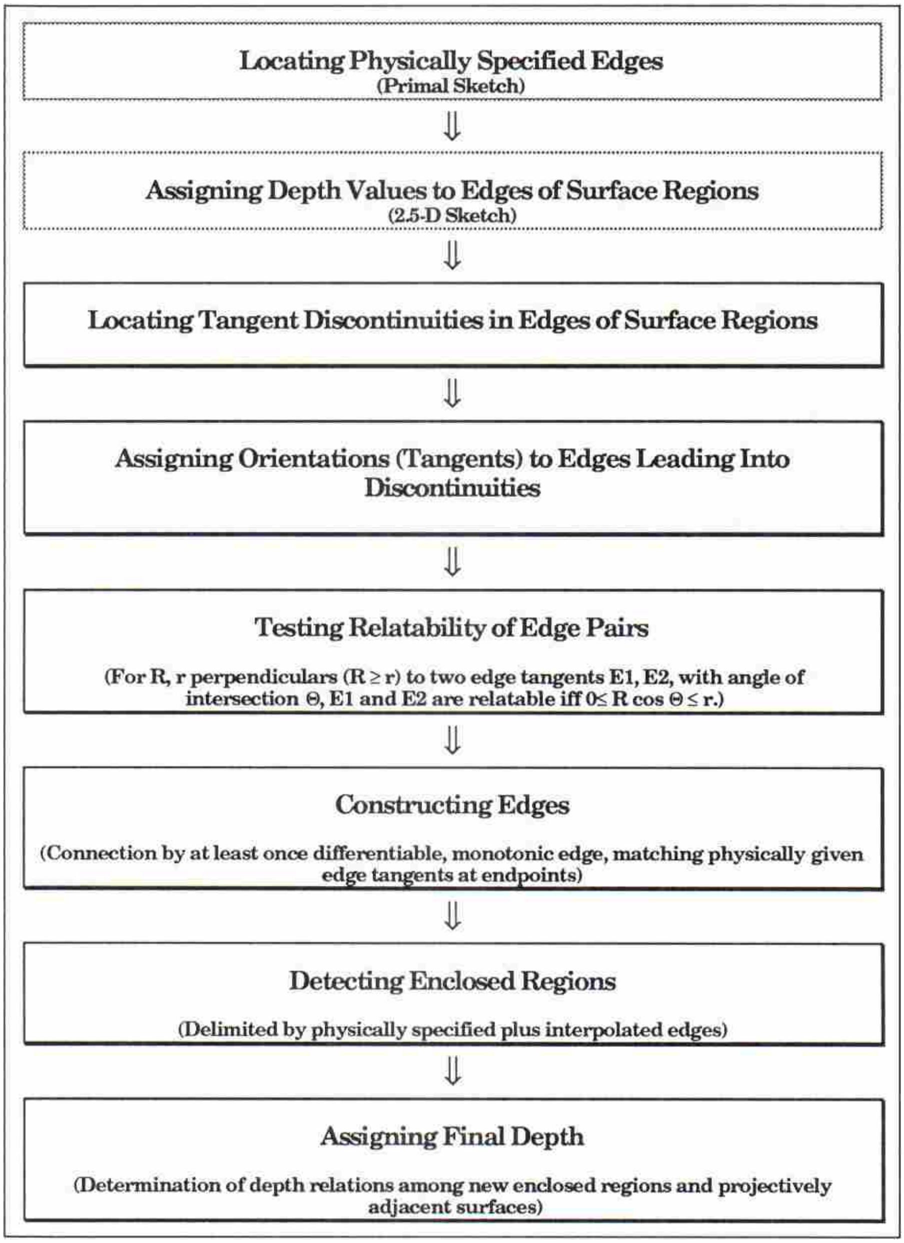

Fig. 5. Schematic of the unit formation model. (See text.)

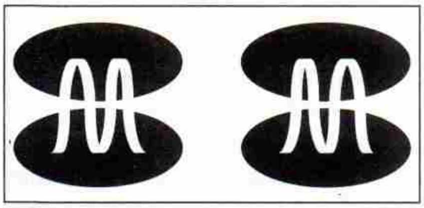

Fig. 6. Example of three-dimensional interpolation. The two views may be placed in a stereoscope or free-fused by crossing (or diverging) the eyes, so that each eye receives one of the two views.

shown in Figure 6. Third, work in progress suggests that subjects perceive interpolated edges in 3-D only when relatively little torsion (twisting) is required to connect surface edges.

A simple generalization of the model allows it to apply to surface edges in 3-space. The same relatability criterion determines whether two surface edges can connect, but the criterion can apply to edges in any plane, not just a frontoparallel plane. It follows that 3-D interpolation should not occur between two edges that lie beyond some relatively small deviation (i.e., some threshold) from co-planarity. Although further study is needed, this generalization fits available data.

SPATIOTEMPORAL UNIT FORMATION

Perception of hidden edges and surfaces in ordinary circumstances often involves the relative motions of objects and observers. When motion and occlusion combine, perceiving objects requires the integration of fragmentary information over time. Consider the case of a stationary observer viewing a stationary target while objects move in between. (This case also approximates that of a moving observer viewing a distant scene through foliage.) In this case, there is no relative motion between the occluded scene and the observer, but various parts of the scene become visible at different times, and perhaps some parts never become visible. How might the visual system produce a representation of objects and surfaces by integrating available surface and edge fragments? A straightforward extension of our model might explain perception across time: Edge parts accumulated over time support interpolation when they meet the relatability criterion.

A second integration problem arises when objects move relative to the observer. Consider a running dog seen through a hedge. Partial information is given not only at different times, but in constantly changing locations. How can these fragments be pieced together? Perhaps the simplest possibility is that velocity information is used to infer where moving edge pieces will be later in time. If so, our model could be applied to the updated (extrapolated) positions of the fragments tracked over time. The latter process would require stimulus information about velocity as a basis for the temporal extrapolation.

Both of these extensions of the model have been confirmed empirically.[3] One line of studies involved the experimental apparatus shown schematically in Figure 7a. A moving occluding surface contained two slits, and displays were placed behind the occluding panel. Any momentary view provided minimal information about the figures behind the occluder, and a section between the two slits was always occluded. Using displays like those in Figures 7b and 7c, we varied edge relatability. Back-and-forth movement of the occluding panel, above a threshold speed, readily allowed subjects to detect the figural areas behind the occluder. The crucial question was how these edges would be used to interpolate form and boundaries across the vertical gap between the two slits. Subjects' responses about unity (forced choice—one vs. two display objects) and perceived form were predicted accurately by the relatability criterion applied to the physically specified edges accumulated over time.

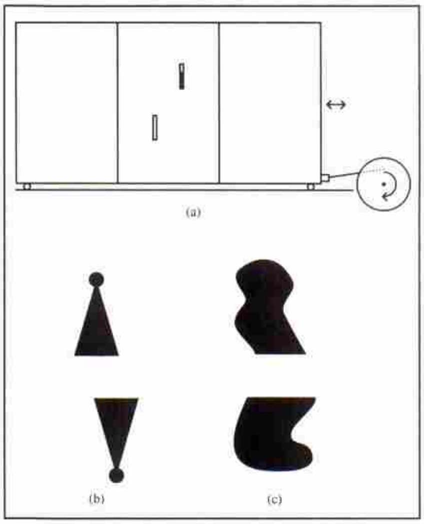

Fig. 7. Experimental investigation of spatiotemporal unit formation. (a) Schematic of apparatus used. Edge parts are revealed over time through two slits in a moving occluder. No object information is available at any time in the area between slits. (b) Nonrelatable display. (c) Relatable display.

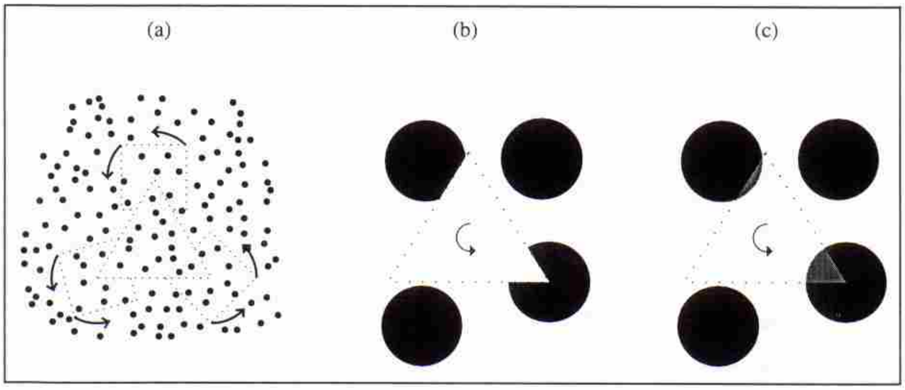

Fig. 8. Spatiotemporal unit formation phenomena. (a) Diagram of an illusory figure with motion-specified inducing elements. When the display is stationary, no figures are visible. When the square areas indicated by dotted lines rotate around their centers, they become visible as a result of accretion and deletion of background texture. The corners of the central triangle indicated by dotted lines then become visible because they lie in front of the squares and cause their own accretion and deletion of texture. Finally, these corner edges support interpolation of unspecified edges between them to make the complete triangle visible. (b) Schematic of kinetic illusory figure. The dotted lines indicate boundaries of a white triangle that is not visible against the white background. When the triangle rotates, sequential interruptions in the black circles allow the triangle to be perceived. (c) Schematic of kinetic occlusion. The same sequential interruptions in the black circles as in (b) appear as gray areas, rather than white, causing a gray rotating triangular figure to appear to be behind the white surface, and to be visible through holes or windows in that surface.

A related finding is that the physically specified edges participating in boundary interpolation can be defined exclusively by information given over time. Figure 8a shows a schematic of an illusory-figure display in which the input edges are given by accretion and deletion of texture elements, and no edge orientation information is available in any momentary projection.

Other research has involved moving objects revealed in fragments over time. In this case, the visual system appears to extrapolate the position of briefly given edge pieces over time, and relationships between these support interpolation. Examples of equivalent kinetic illusory-figure and kinetic occlusion displays[15] are shown in Figures 8b and 8c. In the case of the illusory figure (Fig. 8b), a rotating triangle of the background brightness and color is defined only by sequential partial occlusion of the separated circles. The result of these changes is perception of a unitary, rotating triangle with crisp edges. In the occlusion case (Fig. 8c), the gray pieces are seen as parts of a triangle rotating behind a white surface with holes in it. The spatial relations given in the model together with the assumption of spatiotemporal tracking of the briefly given edge parts give a good account of these phenomena.

Although many motion phenomena remain to be investigated, available evidence supports the idea that spatiotemporal unit formation relies on the same basic process that accomplishes spatial interpolation. The additional requirements of spatiotemporal unit formation involve the accumulation of inputs to the interpolation process over time. Our investigations to date suggest three ways in which this accumulation occurs: (a) The edges on which interpolation is based can be created by motion-carried information, such as accretion and deletion of background texture. (b) Interpolation occurs between edges that are not specified simultaneously; that is, input edges can be registered sequentially. (c) A process of edge extrapolation in space allows the relatability criterion to be applied to the current positions of moving edges rather than their locations at the time they were registered.

CONCLUSION

Perception of objects from information that is fragmentary in space and time tells much about the character of perception generally. It illustrates, as the gestalt psychologists emphasized, that the outputs of perceptual processes do not mirror the local sensory inputs. Perception of unitary objects and continuous boundaries despite occlusion depends on spatial and temporal relationships in the input, and leads to abstract representations (completed objects and boundaries) in the output. Until recently, however, ideas about the process have been vague. Claims that perception involves simplicity, inference, hypothesis testing, or prior knowledge, or a combination of these, have placed few bounds on the processes and knowledge potentially involved.

Our work suggests that a number of object perception phenomena can be explained in terms of certain spatial and temporal relationships and a relatively local, autonomous process that progresses from these inputs to bounded objects in determinate ways. Although simple shapes and smooth boundaries can be outcomes of this process, simplicity and regularity principles are not its causes, except, perhaps, in a deeper sense. Perceived objects correspond to physical objects because visual processing exploits regularities about the physical world and laws determining its projection to the observer.

Acknowledgments—This research was supported by National Science Foundation Research Grant BNS 89-13707 to the authors. We thank Ken Nakayama and an anonymous reviewer for helpful comments.

Notes

1. D. Marr, *Vision* (Freeman, San Francisco, 1982); R.J. Watt, *Visual Processing: Computational, Psychophysical and Cognitive Research* (Erlbaum, London, 1988); J. Beck, Textural segmentation, in *Organization and Representation in Perception*, J. Beck, Ed. (Erlbaum, Hillsdale, NJ, 1982).

2. See, e.g., R.L. Gregory, *The Intelligent Eye* (McGraw-Hill, New York, 1970); I. Rock, *The Logic of Perception* (Bradford Books/MIT Press, Cambridge, MA, 1983).

3. P.J. Kellman and T.F. Shipley, A theory of visual interpolation in object perception, *Cognitive Psychology, 23,* 141–221 (1991).

4. A. Michotte, G. Thines, and G. Crabbe, Les complements amodaux des structures perceptives, *Studia Psycologica* (Publications Universitaires de Louvain, Louvain, Belgium, 1964); see also G. Kanizsa, *Organization in Perception* (Praeger, New York, 1979).

5. We use the phrase *unit formation* after Koffka to encompass the connecting of visible parts into unitary objects, but also more generally, because the results of boundary interpolation can also be completion of holes in surfaces or single contours; K. Koffka, *Principles of Gestalt Psychology* (Harcourt Brace, New York, 1935).

6. A number of interesting examples of this type of display can be found in Kanizsa, note 4.

7. P. Kellman and M. Loukides, An object perception approach to static and kinetic subjective contours, in *The Perception of Illusory Contours*, S. Petry and G. Meyer, Eds. (Springer-Verlag, New York, 1987).

8. T.F. Shipley and P.J. Kellman, Perception of partially occluded objects and illusory figures: Evidence for an identity hypothesis, *Journal of Experimental Psychology: Human Perception and Performance, 18,* 106–120 (1992).

9. We thank Nancy Kanwisher for suggesting this possibility.

10. Global (or top-down) influences may have a much larger role in object recognition. Recognition from partial information, as when one recognizes a moving blur underneath the sofa as one's pet cat, is separable, conceptually and psychologically, from perception of boundaries and shape (Marr, note 1), although studies have often confounded the two processes.

11. Corners formed by thin lines, although useful for representing some aspects of shape, do not trigger the interpolation process.

12. T.F. Shipley and P.J. Kellman, The role of discontinuities in the perception of subjective contours, *Perception & Psychophysics, 48,* 259–270 (1990).

13. A specific form of edge arising naturally from the model extends one physically specified edge as a straight line and the other as an arc of constant curvature to a point where their tangents match. Thus, the predicted curve is once differentiable but not twice differentiable (the second derivative is undefined at the point of intersection between curved and straight segments). An alternative form of the connection would be a cubic spline, which by definition is everywhere twice differentiable. There are not currently precise enough data allowing a choice between these possibilities.

14. K. Nakayama and S. Shimojo, Toward a neural understanding of visual surface representation, *Cold Spring Harbor Symposia on Quantitative Biology, 40,* 911–924 (1990).

15. P.J. Kellman and M.H. Cohen, Kinetic subjective contours, *Perception & Psychophysics, 35,* 237–244 (1984).

Filling in Gaps in Perception: Part I[1]

V.S. Ramachandran

Neurologists have long known that a systematic two-dimensional map of the retina exists in the visual cortex, and that any sharply localized damage to the visual cortex always results in an island of blindness in the visual field called a *scotoma*.[2] One enigmatic aspect of scotomas, however, is that the patients themselves are often blissfully unaware of them. A patient who gazes at, say, a pink wall does not see a dark hole corresponding to the scotoma even though visual information from this region does not reach the brain. The wall looks homogeneously pink. Indeed, even if the patient gazes at a pattern of wallpaper, no gap or hole is seen—the wallpaper seems to

V.S. Ramachandran, M.D., Ph.D., is Professor with the Neuroscience Program and the Department of Psychology, University of California, San Diego. Address correspondence to V.S. Ramachandran, Department of Psychology, University of California, San Diego, La Jolla, CA 92093.