(don't describe) what the form *looks like*. The key, however, is that images depict in a special way: Unlike pictures, images are perceptually organized. Thus the historical debate, asking whether mental images are more like "pictures" or more like "propositions," simply gave us the wrong alternatives. Images (like pictures) depict, but also (like propositions) are structured in a fashion that renders interpretation unambiguous.

References

Baddeley, A. D. and Andrade, J. (1995) The impact of concurrent articulation on auditory imagery. Unpublished manuscript; Applied Psychology Unit; Medical Research Council; Cambridge, England.

Baddeley, A. and Logie, R. (1992) Auditory imagery and working memory. In D. Reisberg (ed.), *Auditory Imagery*. Hillsdale, NJ: Erlbaum Associates.

Baddeley, A. and Wilson, B. (1985) Phonological coding and short-term memory in patients without speech. *Journal of Memory and Language* 24:490–502.

Besag, J. and Diggle, P. J. (1977) Simple Monte Carlo tests for spatial pattern. *Applied Statistics* 26:327–333.

Bishop, D. and Reisberg, D. (1996) The role of articulation in generating, maintaining, and manipulating speech-based representations. Unpublished manuscript; Applied Psychology Unit; Medical Research Council; Cambridge, England.

Crowder, R. and Pitt, M. (1992) Research in memory/imagery for musical timbre. In D. Reisberg (ed.), *Auditory Imagery*. Hillsdale, NJ: Erlbaum Associates.

Crutcher, R. (1994) Telling what we know: The use of verbal report methodologies in psychological research. *Psychological Science* 5:241–244.

Ericsson, K. and Simon, H. (1980) Verbal reports as data. *Psychological Review* 87:215–251.

Gupta, P. and MacWhinney, B. (1993) Is the phonological loop articulatory or auditory? *Journal of Memory and Language* 33:1–26.

Iverson, P. and Krumhansl, C. (1993) Isolating the dynamic aspects of musical timbre. *Journal of the Acoustical Society of America* 94:2595–2603.

Payne, J. (1994) Thinking aloud: Insights into information processing. *Psychological Science* 5:241–248.

Pylyshyn, Z. (1981) The imagery debate: Analogue media versus tacit knowledge. In N. Block (ed.), *Imagery* (pp. 151–206). Cambridge, MA: MIT Press.

Reisberg, D. (1996) The non-ambiguity of mental images. In Cornoldi, C., Logie, R., Brandimonte, M., Kaufmann, G., and Reisberg, D. (eds.), *Stretching the Imagination: Representation and Transformation in Mental Imagery* (pp. 119–172). New York: Oxford University Press.

Reisberg, D. and Logie, R. (1993) The ins and outs of working memory. In M. Intons-Peterson, B. Roskos-Ewoldsen, R. Blake, and K. Clayton (eds.), *Imagery, Creativity, and Discovery* (pp. 39–86). Hillsdale, NJ: Erlbaum Associates.

Reisberg, D., Smith, J. D., Baxter, D. A., and Sonenshine, M. (1989) "Enacted" auditory images are ambiguous; "Pure" auditory images are not. *Quarterly Journal of Experimental Psychology* 41A:619–641.

Schooler, J., Ohlsson, S., and Brooks, K. (1993) Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology* 22:166–183.

Smith, J. D., Reisberg, D., and Wilson, M. (1992) The role of inner speech in auditory imagery. In D. Reisberg (ed.), *Auditory Imagery* (pp. 95–119). Hillsdale, NJ: Erlbaum Associates.

Smith, J. D., Wilson, M., and Reisberg, D. (1996) The role of subvocalization in auditory imagery. *Neuropsychologia* 33:1433–1454.

Wilson, T. (1994) The proper protocol: Validity and completeness of verbal reports. *Psychological Science* 5:249–252.

Chapter 11

An Update on Gestalt Psychology

Philip J. Kellman

Long ago, in my first year of graduate school at Pennsylvania, I heard rumors about a mythical gathering of zealous researchers, whose weekly deliberations extended far into the night and disbanded only when exhaustion finally overcame insight. That year, the Gleitman research seminar was only myth, as Henry was ill. His recovery led to the seminar's revival, and both were sources of joy in the department.

When I joined the seminar, I was just beginning to learn about perception. Henry was and is a great interpreter of all things psychological, but in his heart, I believe, perception holds a special place. No doubt some of his interests in perception were traceable to his time on the faculty at Swarthmore College. He was there during the height of the Gestalt influence, interacting with Köhler, Wallach, and others. As Elizabeth Spelke and I began research on the developmental origins of principles from Gestalt psychology, Henry richly conveyed much of that tradition. He was always encouraging about our chances of answering some very old questions about perceptual organization; his support meant much to our fledgling project. Meanwhile, he kept testing my emerging ecological theoretical leanings with his own empiricist ones. In the seminar, Henry's insights and those of others improved many a research project. Lila, in particular, had me baffled. I wondered if there were another L. Gleitman who was famous in psycholinguistics, as her comments about perception research showed such depth and wisdom that that they could only have come from a specialist in perception.

Although we used them to guide our initial studies of the development of perceptual organization (e.g., Kellman and Spelke 1983), the Gestalt principles of object segregation, which had been applied to occlusion situations by Michotte, Thines, and Crabbe (1964), were vague and a bit numerous. In time, my own research has come back to these principles, trying to develop from them more precise ideas that could form part of computational and neural models of perception. In this gratifying and challenging enterprise, I have worked closely with Tim Shipley, whose dissertation Henry and I co-supervised in 1988. Much of what we have accomplished and much of what remains to be done can be characterized as an update of Gestalt psychology. Transforming the Gestalt insights into a detailed understanding of perceptual computations is important for diverse reasons: It advances our understanding of adult human visual perception, sheds light on perceptual development, and informs attempts to make artificial vision systems that could produce descriptions of physical scenes from information in reflected light. In this chapter, I will try to make clear what has become of various Gestalt principles in some current research.

The Gestalt principles have been applied in many domains, but their original and most familiar home is the domain of object perception, in particular the problems of visual segmentation and grouping. This is the domain I will consider in discussing the computational legacy of Gestalt ideas. The basic problems in segmentation and grouping are easy to describe and illustrate. In a sheaf of light rays arriving at the eye, no ray of light is physically connected to any others. Some image descriptions likewise preserve information separately for each physical location. A digitally encoded image might list for each location (pixel) intensity and chromatic values. There is no linkage between pixels 384 and 385, for example.

What we get *perceptually* from the light rays coming from a real scene or a digitized image is quite different. It is a description of objects and surfaces in a three-dimensional (3-D) space. The objects are seen as detached or separable from adjacent objects and surfaces, and each object unites many visual directions or pixels. The problems of segmentation and grouping involve a mapping from the optic array onto these representations of objects.

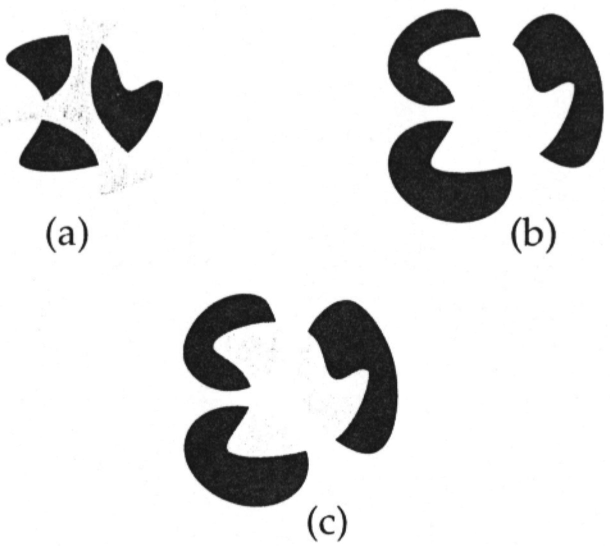

Figure 11.1.
Examples of boundary and surface interpolation. a) Partially occluded object. b) Illusory object. c) Apparent transparency.

How do we group together and separate regions to achieve these, usually accurate, representations of the objects in a scene? Perhaps the most vexing part of the problem is that parts of an object often reflect light to the eyes from several spatially-separated areas. In figure 11.1a, the black object appears as a single entity whose contours are partly occluded in several places. How can a human viewer or a computer vision system connect separate visible regions and represent their hidden contours, surfaces, and overall shapes? This set of problems will be our focus.

The Identity Hypothesis in Object Completion

There is one idea it will be helpful to introduce and place in the background: what we have called the *identity hypothesis* in object completion. There are a number of different-looking phenomena in which the visual system accomplishes segmentation and grouping by supplying hidden contours and connecting regions. Some of the phenomena are shown in figure 11.1. Along with the partial occlusion display in figure 11.1a, figure 11.1b shows what is usually called an illusory object or illusory figure, and figure 11.1c shows an apparently transparent (translucent, really) object. The identity hypothesis states that these different-looking perceptual completion phenomena are caused by the same underlying process. In these displays, the same parts of the central figure are defined by luminance edges, and the gaps across which edges are interpolated are the same. Those are formal similarities, but what I am suggesting is that the same gap-surmounting *process* is at work in all of these cases as well. The differences in our phenomenology for the various cases have to do not with differences in interpolation processes, but with how the interpolated edges and surfaces are situated relative to other surfaces in the array (especially whether they are in front or behind).

The arguments and data suggesting a common interpolation process can be found elsewhere (Kellman and Shipley 1991; Kellman, Yin, and Shipley 1998; Ringach and Shapley 1996). Here I give one example to convey the general idea. In figure 11.2, we see yet another perceptual segmentation and completion phenomenon, called a self-splitting figure or SSO. The particular SSO shown is one constructed by Petter

(1956) and later discussed by Kanizsa (1979). The display has several interesting properties. As noted above, it resolves into two distinct objects—boundaries get constructed through homogeneously colored regions of the display. A second interesting property, and our immediate concern, is the depth relationship between the two objects. At the top of the display, the righthand ring appears to pass in front of the left, whereas at the bottom, the lefthand ring passes in front of the right. These effects of perceptual organization appear to be strong and consistent across observers.

At the top, where the righthand ring crosses in front, its contours are classic illusory contours and its surface is said to be *modally* completed (Michotte et al. 1964), meaning it has a sensory presence. In the same visual direction, the lefthand ring has a partly occluded surface and contours, sometimes called *amodal* completion. *Amodal* means that the hidden surfaces are perceived or represented, but they do not have local sensory presence. (You could not answer a question about the presence or absence of a smudge on the occluded surface because, after all, it is occluded.) The phenomenological difference between illusory and occluded contours and surfaces has led many to think that these are phenomena of very different character, the former explainable by sensory mechanisms and the latter involving cognitive processes. On the identity hypothesis, these involve, at least in part, the very same interpolation mechanisms, and the phenomenological difference concerns whether, in the final percept, the interpolated surface forms in front of or behind some other surface in the scene.

Here is where Petter's effect comes into the story. In displays such as the rings in figure 11.2, Petter observed that a simple rule governs which object will be seen as in front, having illusory contours, and which will be seen as going behind. The object that must be completed across the smaller gap always ends up in front, and the object that traverses the larger gap ends up behind. From this observation, which appears to be correct, we can make the following logical argument. If the final "illusory" or "occluded" status of a contour depends on some comparison with another interpolated contour, then some mechanism that interpolates contours must operate before the final status as illusory or occluded is determined. For an explicit or implicit comparison to take place, the visual system must recognize both contour completions crossing at that site. In other words, the mechanism that interpolates contours is not "modal" or "amodal." (For other phenomena and data that converge on the same point, see Kellman, Yin, and Shipley 1998.)

The idea of a common underlying mechanism producing phenomena whose subjective appearance differs so greatly is somewhat surprising, and there is residual controversy about what exactly is shared and what must differ in different-looking cases of visual completion. In what follows, the identity hypothesis will not be our focus, but it will allow us to move between experiments and data involving illusory and occluded objects and boundaries without distinguishing these cases.

Figure 11.2.
Self-splitting Object (SSO) after Petter (1956). Although the display is homogeneous in color, it is perceived as two bounded objects. The ring on the right tends to appear in front of the ring on the left at the top of the display but appears to pass behind the ring on the left at the bottom. (See text.)

Gestalt Principles and Unit Formation

Segmentation and grouping, illusory contours and transparency phenomena all involve issues of unit formation, determining what goes with what. The Gestalt psychologists first inquired into these problems, and Gestalt principles have been applied to all of these phenomena (Kanizsa 1982; Michotte et al. 1964). It is a nice consequence of the identity hypothesis that our "updates" of certain Gestalt principles will apply to all of them as well. We now look at particular principles and examine their legacies in more recent work.

Good Continuation

In his classic (1921) paper "Untersuchungen zur Lehre von der Gestalt" ("Laws of organization in perceptual forms"), Wertheimer gave a number of examples illustrating what he called the "Factor of Direction" or the "Factor of Good Curve." Despite offering these two formal names, another of Wertheimer's phrases used in passing—"good continuation"—has stuck as the name of this principle. Figure 11.3 shows some examples redrawn from Wertheimer (1921).

Despite the compelling and varied nature of the demonstrations, Wertheimer's definition of this principle is rather vague. In fact, the displays are meant to convey the following idea, without any formal definition:

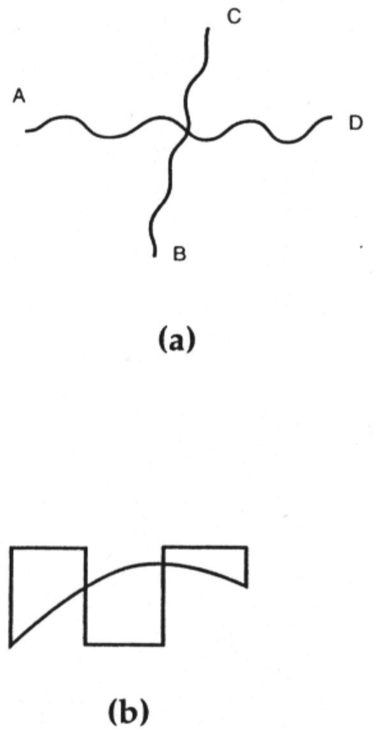

Figure 11.3.
Examples of the "Factor of Direction" (Good Continuation) from Wertheimer (1923). a) The segments labeled A and D appear to form a unitary object, as do those in C and B. b) Despite the possible appearance of three closed regions, the display is usually seen as containing a unitary curved edge and a square-wave.

> On the whole, the reader should find no difficulty in seeing what is meant here. In designing a pattern, for example, one has a feeling how successive parts should follow one another; one knows what a "good" continuation is, how "inner coherence" is to be achieved, etc.; one recognizes a "good Gestalt" simply by its own "inner necessity."

Despite its intuitive importance, it is hard to find in the seventy or so years since Wertheimer any explicit definition of the "good" in good continuation. One obvious candidate is to relate "good" to mathematical descriptions of contours. A function is often called "smooth" in mathematics if it has no discontinuities in the first derivative. A discontinuity would correspond to a sharp corner, that is, a point at which there is no unique slope of the function. But there are other notions of smoothness, involving higher derivatives. In the design of automobile bodies, for example, smooth might mean differentiable at least two or three times (Prenter 1989). Which notion captures the phenomena of human visual segmentation and grouping? Surprisingly, this issue has been the subject of little empirical investigation. Some of the issues, and some clues to the answers, are illustrated in figure 11.4. In (a), there is first-order continuity between parts A and B, but a second derivative discontinuity between them. Between A and C there is a first-order or tangent discontinuity (TD). Perceptually, A and B appear unitary whereas C appears separable. In (b), there is a smaller direction change between A and B than between A and C. Direction change might therefore predict that A and B will be linked more than A and C. On the other hand, both B and C have a TD with A. Perceptually, neither B nor C appears to have continuity with A, suggesting the importance of TDs in segmentation. In (c), parts A and B are not distinguishable as separate parts; as parts of a constant curvature arc, they agree in all derivatives. The case is different in (d). Here, a straight segment (B) meets a constant curvature segment (A). The two parts agree in the first derivative, but there is a second-order discontinuity. Nevertheless, the two parts appear smoothly joined. All of these examples suggest that TDs lead to segmentation and their absence—agreement in the first derivative—facilitates joining.

Apart from these considerations about continuous contours, we may ask what relationships between separated contours lead to their perceptual connections, as in partially occluded and illusory objects? Here again, the notion of good continuation has been invoked (e.g., by Michotte et al. 1964), but without any specific definition.

Parts (e) and (f) of the figure illustrate the same contour relations as (c) and (d) with gaps now caused by occlusion. Both displays produce the appearance of a unitary contour passing behind the occluder.

Formalizing Good Continuation: Ecological Constraints and Computational Theory

Seventy years after Wertheimer, the intuition behind the principle of good continuation is still important. Making this idea useful in models of human and computer vision requires first of all a precise mathematical specification. It also requires placing continuity in the context of the general problem of scene segmentation and object perception. Instead of starting with particular contours and patterns, we need to pose briefly the question of how objects reflecting light make available information that might be used to segment and group the world into discrete objects and surfaces. These are questions of *ecological optics* Gibson

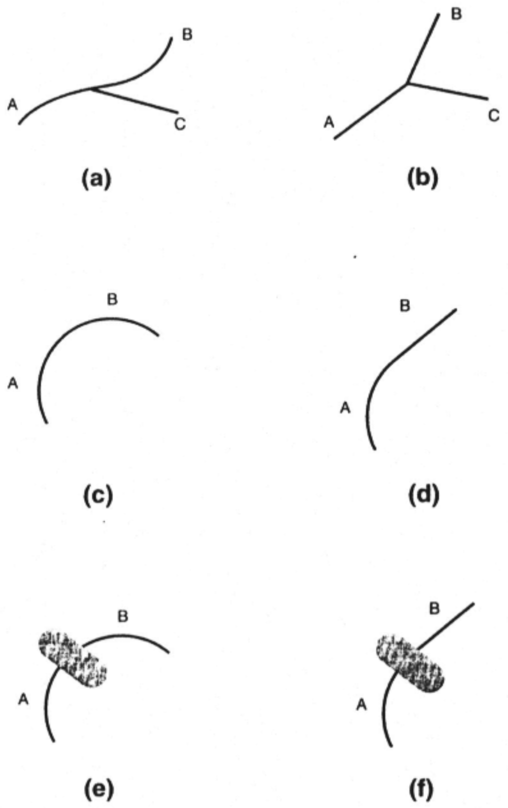

Figure 11.4.
Examples illustrating the importance of first-order continuity and discontinuity. a) Segments A and B, which are first-order continuous, appear connected moreso than A and C or B and C. b) A first-order or tangent discontinuity divides A, B and C. c) Apparently unbroken contour made from A and B segments, where all derivatives agree at the point of connection. d) Apparently unbroken contour made from A and B segments where the first derivatives of A and B agree at the point of connection, but there is a discontinuity in the second derivative. e) and f) Contours in c) and d) under partial occlusion. (See text.)

(1966, 1979) and computational analysis (Marr 1982). I consider this context first and then return to good continuation.

Multiple Tasks in Object Perception

Object perception involves multiple computational tasks. The first is *edge detection*. Different physical objects, having different material composition, will tend to produce reflected light of differing luminance and spectral composition. Accordingly, abrupt changes in luminance and spectral characteristics are likely to indicate locations of object edges. Not all such changes are boundaries of objects, however. Some are shadows; others are textural markings on continuous surfaces, and so on. Some classification process must distinguish surface edges from these other cases.

Much of *edge classification* may be achieved by coordinating information from a luminance map of a scene with a depth map, gotten from stereoscopic information. For a moving observer, there will also be available a motion map, assigning to each location a velocity vector (see, e.g., Lee 1974). Discontinuities in the depth and motion maps will correspond to true surface edges with less ambiguity.

The most common type of edge emerging from these initial analyses is the *occluding edge*. It is a contour that bounds an object or surface on one side. Each occluding edge indicates where something ends in the scene, but each also marks a mystery. If a person is seen standing in front of a car, the image contour separating the visible surfaces of the car and the person is a boundary of the person but not the car. At this contour, the car's surface disappears behind. The mystery is where it goes.

Determining which side bounds the object is called *boundary assignment* (Koffka 1935). Nakayama, Shimojo, and Silverman (1989) suggested that an image contour be labeled *intrinsic* to a surface region if it bounds that region and *extrinsic* if it does not. Boundary assignment may not be implemented as a separate process. When depth or motion information is available, it is computationally simple to recover edges and the relative depth order of two surfaces at those locations. Because depth order determines boundary assignment (the nearer surface always owns the boundary), boundary assignment and edge classification may occur together.

To this point, we have a representation of occluding edges, partially bounding surface regions. Now we are in a position to consider how Gestalt notions of continuity can be implemented in perceptual processing. The story has two parts. The first involves particular locations in images in which edge continuity is disrupted—what I called above tangent discontinuities (TDs). A TD is nothing more than a sharp corner

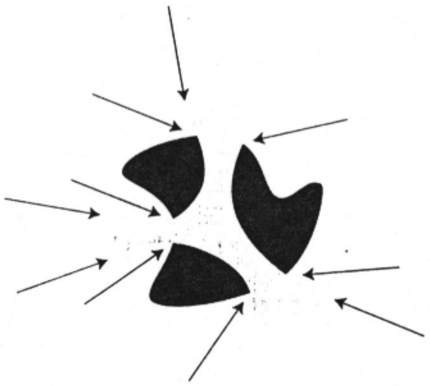

Figure 11.5.
Occlusion display from figure 11.1 with all tangent discontinuities indicated by arrows.

where contours meet.[1] At such a point there is no unique slope of the contour. TDs thus characterize all the standard types of contour junctions—"T," "X," "Y," arrow, and so on. To be a junction means to be a TD. In our updated notion of good continuation, a TD is the key concept. TDs in one of the displays from figure 11.1 are marked in figure 11.5. Referring to figure 11.1, it can be seen that all of the interpolated edges, in all displays, begin and end at TDs.

The ecological importance of TDs is straightforward. It can be proven (see Kellman and Shipley 1991, appendix B) that every instance in which an object boundary is partly occluded produces a TD at the place where the boundary goes out of sight. Thus TDs are potential loci of occlusion. They also mark the transition points from extrinsic to intrinsic contours of a surface region. For hidden parts of objects, TDs are where we pick up the trail of where their hidden edges might go.

Relatability

Some TDs are merely the visible corners of objects. Not all are loci of occlusion. Moreover, even when a TD is a locus of occlusion, there remains the question of where the occluded part of the boundary goes. More is needed to determine object boundaries. Here we come to the second part of the implementation of the Gestalt idea of good continuation, what we have called *relatability* (Kellman and Shipley 1991, 1992).

Relatability formalizes good continuation. It constrains unit formation based on an assumption that object boundaries tend to be smooth. Specifically, relatability expresses the conditions required to connect two edges by a smooth (at least once differentiable) and monotonic (singly

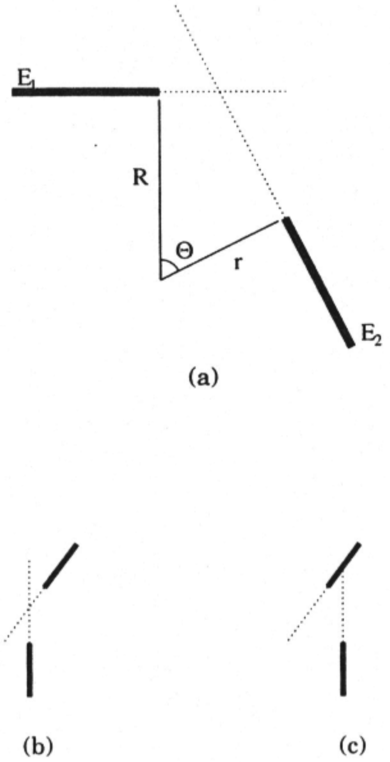

Figure 11.6.
Construction used to define relatability. a) E1 and E2 are surface edges; R and r are perpendiculars to the tips (points of tangent discontinuity) of E1 and E2, assigned so that $R > r$. Θ is the angle between R and r, E1 and E2 are relatable if $0 \leq R \cos \Theta \leq r$. b) Illustration of relatable edges. c) Illustration of nonrelatable edges. Either a doubly inflected curve or introduction of tangent discontinuities are required to connect two nonrelatable edges.

inflected) curve that agrees with the tangents of the two edges being connected at the point where each leads into a TD. We will define relatability with edges separated in the optical projection and show that the case of continuous edges (zero gap) is a limiting case. Relatability can be defined using the construction shown in figure 11.6a.

E1 and E2 are edges of surfaces. Let R and r be perpendiculars to these edges at the point where they lead into a TD. Let R be the longer of the two perpendiculars, and let the angle Θ be the angle of intersection of R and r. Intuitively, when relatability holds, there will always be a smooth, monotonic curve that can be constructed, starting from the endpoint of E1 (and matching the slope of E1 at that point) and proceeding through not more than a 90-degree bend to the endpoint of E2 (and matching the slope of E2 at that point). When $R \cos \Theta > r$, any connection between E1 and E2 would have to be doubly inflected (if it matched the slopes at E1 and E2) or would have to introduce sharp corners where the interpolated edge meets E1 and E2. (See figure 11.6c.) According to this model, visual boundary interpolation does not occur in such cases.

Formally, E1 and E2 are relatable iff:

$0 \leq R \cos \Theta \leq r$

This statement can be unpacked in two steps. The righthand side of the inequality simply states that the projection of R onto r ($R \cos \Theta$) falls within the extent of r. Whenever the length of r is less than projection of R onto r, the edges are not relatable. Second, the curve constructed to connect the two edges cannot bend more than 90 degrees. This limitation is expressed by the lefthand side of the inequality, because $\cos \Theta$ will be negative for $\Theta > 90$.

Below we will see that relatability should involve all three spatial dimensions, although we have defined it here in terms of two. A good deal of work, however, can be done with 2-D edge relations alone, because the smoothness of objects in the 3-D world has consequences for their 2-D projections. It can be shown using elementary projective geometry that collinear edges, smooth curves, and sharp corners in 3-space always project onto collinear edges, smooth cuves, and sharp corners in a 2-D projection (excluding degenerate cases, such as projection of a line to a single point). Thus, much of the information about object smoothness and edge relations is preserved in the optical projections reaching the eyes, even in a static, 2-D image.

Experimental Evidence about Relatability

A variety of experimental evidence supports relatability as a formal description of connections formed by the visual system under occlusion and in illusory contours (Kellman and Shipley 1991; Shipley and Kellman 1992a). Some of the best comes from an elegant paradigm introduced by Field, Hayes, and Hess (1993). Field et al. used arrays of oriented Gabor patches, small oriented elements consisting of a sinusoidal luminance pattern multiplied by a Gaussian window. A Gabor patch closely approximates the best stimulus for the oriented filters found in simple cells of V1, the first visual cortical area. Displays used by Field et al. contained randomly placed, spatially separated elements varying in orientation. Some displays contained a "path." A path was constructed by having the a sequence of several nearby elements having the same angular relationship, for example, successive elements were collinear, or successive elements differed by 15 degrees, etc. In the experiments, subjects on each trial judged which of two successively and briefly presented arrays contained a path. When the positional and angular relations satisfied the relatability criterion, subjects performed very well at this task. When the path consisted of a sequence of elements rotated 90 degrees, so that relatability was violated, performance was much poorer. It appears that certain edge relationships lead to edge connections which become salient, perhaps in parallel across large regions of the visual field. The study also supported the idea that edge connections decline as the angle varies from collinearity, with a cutoff around 90 deg.

Strength of interpolation also depends on the relative extents of the physically specified edges and gaps in a scene. Interpolation strength appears to be a linear function of the "support ratio": the ratio of physically specified edge lengths to total edge length (physically given edges plus gap length) over a wide range of display sizes (Shipley and Kellman 1992b; Lesher and Mingolla 1993). This relationship makes precise a version of the Gestalt law of *proximity*, that nearer elements are more likely to be grouped together.

Relatability in Cases of Minimal Gaps

We have defined and illustrated relatability in the context of occlusion and illusory contours—cases in which the visual system constructs connections across spatial gaps. In the classic Gestalt examples, good continuation was illustrated as determining the breakup of unoccluded displays, without appreciable gaps, into separate objects (as in figures 11.3 and 11.4). Unoccluded displays may be considered as a limiting case of relatability—the case where the gap is zero. (Actually, nearly zero. The contours of the perceived figures do overlap, producing minute occlusions and illusory contours.) In such cases, the "connection" of edges is the continuation of the edge that proceeds smoothly through a junction. We saw relevant examples in figure 11.4. These examples fit the definition of relatability in that smoothness resides in the first derivative. Connecting a straight segment (zero curvature) with a segment of positive curvature yields a well-defined first derivative at the point of connection but a discontinuity in the second derivative, yet figure 11.4d appeared to have perceptual continuity. In contrast, the sharp corner in figure 11.4b disrupts continuity of segment A with both B and C.

This analysis of relatability at the limit sheds light on typologies of contour junctions in human and artificial vision (Clowes 1971; Waltz 1972). In a "T" junction, the contour that does not change direction indicates the boundary of a surface, whereas the other contour passes beneath. A "Y" junction is different in that no contour continues smoothly; all come to an end at that point in space. It has been suggested that the "Y" provides information for an object corner. Relatability subsumes these observations about contour junctions under a more general principle for connecting and segmenting visual arrays.

3-D Relatability: Depth Information in Object Completion

For convenience, we defined the notion of relatability in a plane. Perception of object unity and boundaries in the 3-D world requires taking into account 3-D relationships of contours, however. Over the years, several demonstrations of 3-D contour completion have been devised. One is shown below in figure 11.7. If this display is viewed stereoscopically (free-fuse by crossing or diverging the eyes), it gives rise to a 3-D illusory contour on one side and a 3-D occluded region on the other. Binocular disparity places the inducing edges at particular 3-D orientations, and contour interpolation processes build the connections, smoothly curving through three dimensions, across the gaps. The demonstration suggests that interpolation processes take 3-D positions and relations as their inputs and build connections across all three spatial dimensions.

Until recently, these phenomena have not been addressed experimentally. Recently, we carried out a series of experiments to test 3-D relations in object completion. A full report will appear elsewhere (Kellman, Yin, Shipley, Machado, and Li, in preparation); here I note some of the main results.

We used 3-D illusory object stimuli such as those shown in figure 11.8. Such displays appear to produce vivid 3-D illusory contours and surfaces. We hypothesized that these occur when the physically given contours satisfy a 3-D criterion of relatability. The extension from the 2-D case is this: Bounding contours are relatable in 3-D when they can be joined by a smooth, monotonic curve. This turns out to be equivalent to the requirement that, within some small tolerance, the edges lie in a common plane (not necessarily a frontoparallel plane), and within that plane, the 2-D relatability criterion applies. Another way of saying the same thing is that the linear extensions of the two edges meet in their extended regions in 3-D space (and form an angle greater than 90 degrees).

Three-dimensional relatability can be disrupted by shifting one piece in depth, as shown in figure 11.8b. Another relatable display and a corresponding shifted, nonrelatable display are shown in figures 11.8c and 11.8d.

The experimental paradigm used these displays as follows. Subjects were shown a stereoscopic display on each trial. Stereoscopic disparities were produced by outfitting the subject with liquid-crystal-diode

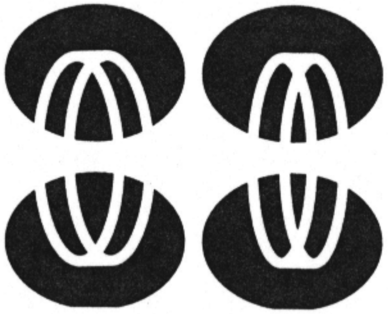

Figure 11.7.
Example of 3-D illusory and occluded contours. (Free-fuse by crossing or diverging the eyes.)

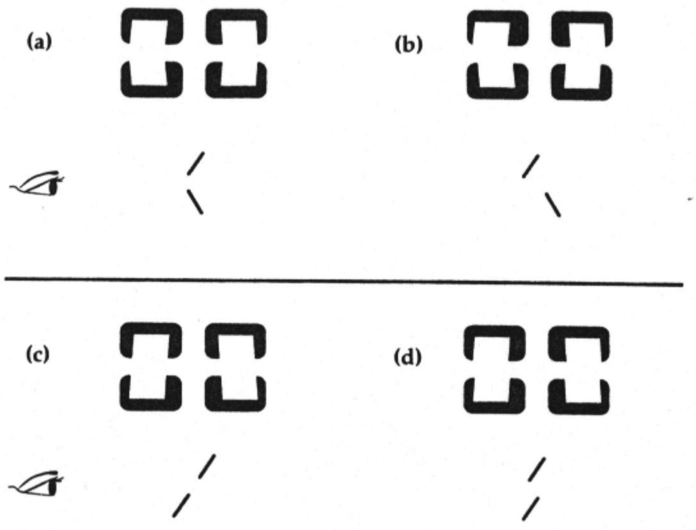

Figure 11.8.
Stimuli in depth relatability experiments. Each display is a stereo pair. (Free-fuse by crossing the eyes.) Below each stereo pair is a side view of the display with the relation to the observer's eye shown. a) 3-D relatable display. The top and bottom white areas lie in intersecting planes and appear connected by a 3-D illusory surface. b) Non-relatable display made by depth-shifting one inducing surface in (a) relative to the other. c) 3-D relatable display with top and bottom areas in a common plane. The top and bottom areas appear connected by a planar illusory surface, slanted in depth. d) Non-relatable display made by depth-shifting one inducing surface in (c) relative to the other. (From Kellman, Yin Shipley, Machado, and Li, in preparation.)

(LCD) shutter glasses, synchronized with alternating computer images. Subjects made a speeded judgment on each trial about the positions of the upper and lower parts of the display. Displays like those in figure 11.8a and 11.8b were said to be in intersecting or *converging* planes. Those in figure 11.8c and 11.8d were said to be in *parallel* planes (including coplanar). Note that the classification required from the subject on each trial was orthogonal to the display's status as relatable or nonrelatable. The key predictions were that (1) perception of a unified object would facilitate classification performance, and (2) perceived unity would depend on relatability. The former was expected based on results in 2-D displays showing that object completion produces an advantage in detecting boundary orientation (Shapley and Ringach 1996; Kellman, Yin, and Shipley 1998).

Results of the initial experiment (Kellman, Yin, Shipley, Machado, and Li, in preparation) are shown in figure 11.9, which shows discrimination sensitivity (d′) in a signal detection analysis by condition. Two values of depth displacement (used to disrupt relatability) were used. These corresponded to a 5 cm and a 10 cm shift in depth of one of the pieces from the observer's viewing distance (100 cm). Results indicate a clear superiority for the relatable displays. (Note that performance on parallel and converging displays are combined in the sensitivity analysis.) Response times reflected the same advantage: Both parallel and converging relatable displays produced faster responding.

On the surface, these results suggest that object completion produces a performance advantage in this task and that 3-D relatability, to a first approximation, predicts unit formation in these displays. Even the smaller value of depth shift disrupted performance markedly. As this is a new paradigm and new data, however, there are several alternative explanations to be considered. Some of these are still occupying us in the lab, but we can relate a couple of important results here.

First, it is possible that performance in our task might not really require object completion. Perhaps relatable displays were better processed because their pieces were more nearly at the same distance from the observer. Comparing two parts' orientations might be easier when the parts are equidistant. Our design allowed us to check this hypothesis using a subset of the data. As figure 11.8d illustrates, a subset of parallel displays used a shift away from the canonical (relatable) stimulus that actually made the two parts more nearly equidistant. We compared these displays (which had either 0 or 5 cm depth differences) with relatable parallel displays having parts that differed substantially in depth (10 cm for the largest slant condition). Results showed that relatability, not similarity in depth, produced superior accuracy and speed. More recently we have tested even more subtle alternatives to the idea that our effects are due to object completion. Results support the object completion hypothesis.

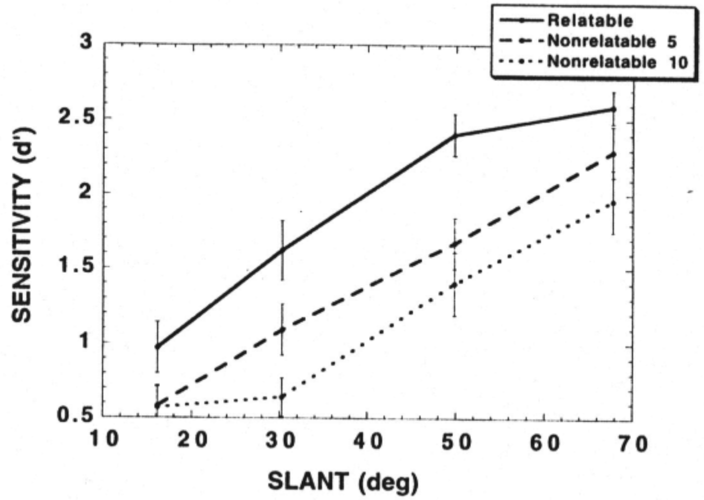

Figure 11.9.
Sensitivity as a function of slant in the depth completion experiment. Relatable displays were more accurately and rapidly classified, suggesting that the upper and lower inducing areas were processed as a connected unit. (From Kellman, Yin, Shipley, Machado, and Li, in preparation.)

But are these truly *three-dimensional* effects? Introducing binocular depth differences involves monocularly misaligning contours in each eye. Perhaps these monocular effects, not true depth effects, cause the performance decrement. It is known that misalignment of parallel or nearly parallel contours disrupts 2-D object completion (Shipley and Kellman 1992a; Kellman, Yin, and Shipley 1998).

In designing the original study, we aimed to produce significant depth shifts using misalignments that remained within the tolerances for 2-D completion. It has been estimated that contour completion breaks down at about 15 minutes of misalignment of parallel edges (Shipley and Kellman 1992a). Our misalignments were on the order of about 10 minutes in the maximum depth shift condition. To check the effect of monocular misalignment, we carried out a separate experiment. In our binocular, depth-shifted displays, each eye had the same misalignment with opposite sign. In this experiment, we used the same displays, but gave misalignment of the same sign in both eyes. Thus the amount of monocular misalignment was exactly identical in every display as in the original experiment. Because both members of each stereo pair had misalignments of the same sign, shifted displays appeared to be at the same depths as relatable displays, but with some lateral misalignment. Results showed no reliable accuracy or speed differences between shifted and relatable displays in this experiment. This outcome is consistent with the idea that perceived depth relationships affected object completion in the first study. The effects are not explainable by monocular misalignment.

This line of research is just beginning, but it suggests that our updated notion of good continuation—contour relatability—applies in three spatial dimensions.

Good Form

The principle of good form (or more generally, *Prägnanz*) describes the tendency of perceptual processing to maximize simplicity and or regularity. Whether perceptual systems act in accordance with such a principle remains controversial. The principle has been difficult to define precisely, in part because it seems to refer to perceptual outcomes rather than stimulus relationships. Some attempts have been made to formalize the notion of overall figural simplicity (e.g., Buffart, Leeuwenberg, and Restle 1981).

It is difficult to separate good form from other factors. Common illustrations almost invariably involve edge continuity besides good form. Figure 11.10 shows two illustrations of good form redrawn from a textbook on perception. Both can be explained in terms of edge relatability. In the display in (a), the edges leading into the TDs are relatable so that the physically specified plus interpolated edges produce two closed forms—the triangle and the rectangle. The second example involves a case of relatability across minimal gaps. At each contour intersection, edges entering and leaving with no TD in between are classified visually as connected. In contrast, a TD between entering and leaving contours indicates a possible breakpoint. In the figure, the continuity of edges gives the two closed forms shown. Kanizsa (1979) argued that that global symmetry is a questionable or weak determinant of object completion, using demonstrations that pitted global factors against local edge continuity. Two of these are redrawn in figure 11.11.

The debate about local vs. global determinants of segmentation and completion has persisted, however. Sekuler, Palmer, and Flynn (1994), for example, reported evidence from a priming paradigm suggesting that global completion occurs in displays like the one shown in figure 11.12a. (Global completion entails seeing a fourth articulated part behind the occluder, making the display radially symmetric.) Others have reported evidence for both global and local completions using priming (Sekuler 1994; van Lier, van der Helm, and Leeuwenberg 1995). Van

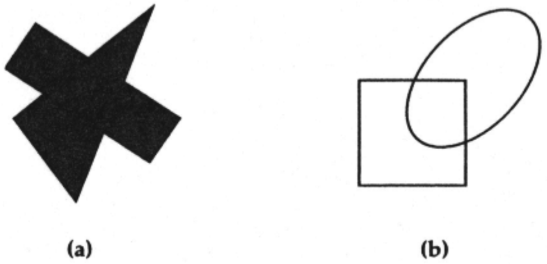

Figure 11.10.
Putative examples of good form or Pragnanz. a) A triangle and a rectangle are seen. b) an ellipsoid and a square are seen. Both outcomes are explainable by relatability with no additional principle of good form or Pragnanz. (Redrawn from Goldstein 1995).

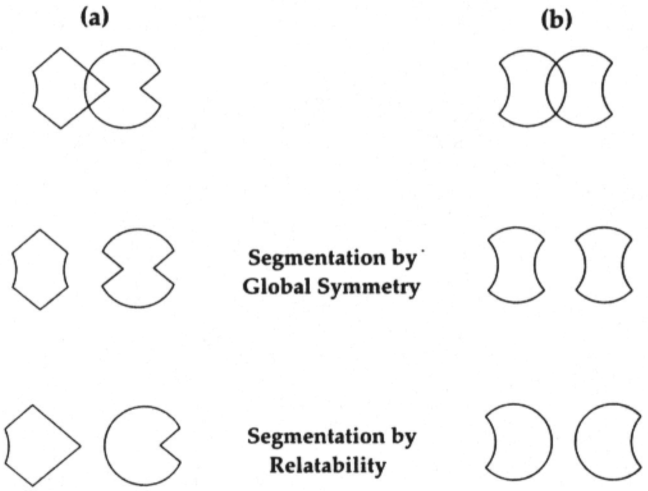

Figure 11.11.
Kanizsa's Demonstrations pitting local continuity against global symmetry. a) (Redrawn from Kanizsa 1979.)

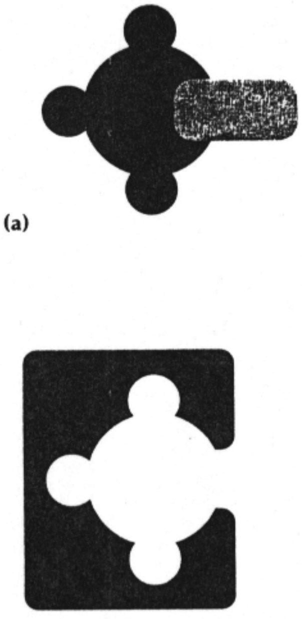

Figure 11.12.
Displays pitting local continuity and global symmetry. a) Occluded object for which local and global completion hypotheses make differing predictions. b) Illusory object version of a. Although subjects are willing to report a global (symmetric) completion in the occluded version, the symmetric completion is not seen in the illusory object display.

Lier et al. interepreted their results in terms of dual or multiple representations activated by partly occluded displays.

This suggestion is close to our own hypothesis: Various experimental effects reflect two distinct categories of processing. One is a bottom-up, relatively local process that produces representations of boundaries according to the relatability criterion. This process is *perceptual* in that it involves a modular process that takes stimulus relationships as inputs and produces boundaries and forms as outputs. The other process is more top-down, global, and cognitive, coming into play when familiar or symmetric forms can be recognized. For lack of a more concise label, we call it *recognition from partial information* (RPI).

One factor pointing toward such a distinction involves the identity between partly occluded and illusory objects, which we have already described. The identity hypothesis has received considerable support (Kellman, Yin, and Shipley 1998; Ringach and Shapley 1996; Shipley and Kellman 1992a), and certain types of displays, such as the Petter effect which we considered earlier, suggest that an identity at some point in processing is logically required (Kellman, Yin, and Shipley 1998).

If true, the identity hypothesis sheds light on the global-local controversy, for this reason. Global completion phenomena are not observed in illusory object displays. Figure 11.12b shows the illusory object display with physically defined edges equivalent to those in figure 11.12a. The reader may observe that there is no appearance of a fourth articulated part in the illusory figure display.

If the identity hypothesis is true, why should global completion occur in occluded but not illusory object displays? The answer may be that the displays are the same in terms of the perceptual processes of contour and surface interpolation but different in terms of RPI. An occluded surface is an interpolated surface that is not the nearest to the observer in some visual direction (i.e., there is something in front of it). An illusory surface is nearest to the observer among all surfaces in a certain visual direction. The crucial consequence of this difference is this: An observer viewing an occluded display is aware that part of the object is hidden from view. This allows certain kinds of reasoning and responses that are not sensible when no part of an object is occluded. In particular, despite any local completion process, *the observer can notice what parts are visible (unoccluded) and whether they are consistent with some familiar or symmetric object.*

Consider a concrete example. If the tail rotor of a helicopter is seen protruding from behind a building, an observer may easily recognize and report that such a helicopter is present, even though the particular contours and surfaces of the hidden parts are not given perceptually. A stored representation of the helicopter may be activated and a belief about the presence of the helicopter may be formed. But RPI differs from perceptual processes that actually specify the positions of boundaries and surfaces behind an occluder.

This separation of processes might explain conflicting reports about global and local processing. First, the only objective data supporting global outcomes come from priming studies. It is well known that priming occurs at many levels, from the most basic representation of the stimulus to higher conceptual classifications involving the stimulus (e.g., Kawaguchi 1988). Unfortunately, there have been no attempts to distinguish these influences in the priming literature on occlusion. Studies reporting global completion have typically used large numbers of trials with a small set of familiar and/or symmetric figures, such as circles and squares. Even if the subjects start out with little familiarity or do not notice the possibility of symmetry under occlusion, repeated exposure may produce familiarity or symmetry responses.

The Dot Localization Paradigm

Priming may not be suitable for separating perceptual processes of boundary and surface completion from more cognitive influences. To test the possibility of different processes, we developed a new experimental paradigm. We focused on the idea that perceptual boundary completion processes lead to specific perceived boundary locations whereas RPI will not in general do so, as in our occluded helicopter example. We measured the precision of boundary location by showing an occluded display and briefly flashing a probe dot in front of the occluder. Subjects were instructed to respond on each trial whether the probe dot fell inside or outside the occluded object's boundaries (i.e., whether the projection of the occluded object to the eye would or would not encompass the dot).

We used an adaptive staircase procedure. In this procedure, the stimulus value for each trial changes depending on the subject's responses. Systematic changes allow a single point on the subject's psychometric function to be estimated. For each display, we used both a "two-up, one down" and a "one up, two down" staircase to estimate two points: the 0.707 probability of seeing the dot as outside the boundary and 0.707 probability of seeing the dot inside the boundary (= 0.293 probability of outside). We took the *difference* between these estimates as a measure of the precision of boundary perception, and the *mean* of these estimates as an estimate of the perceived location of the boundary. Staircases for several stimulus patterns were interleaved, that is, patterns appeared in a random order, and screen position was varied randomly.

We realized that competing perceptual and recognition processes might lead to different strategies across subjects. Therefore, we gave subjects explicit strategy instructions. In the *global instruction condition*, we told subjects that they should see the display as symmetric; for the display in figure 11.12a, for example, they were told that there was a fourth protrusion behind the occluder identical to the three visible protrusions around the circle. In the *local instruction condition*, we told them that we wanted them to see the display as containing a simple curve connecting the two visible edges. In this manner, we sought to find subjects' best abilities to localize boundaries under a global or local set.

A number of interesting findings have emerged (Kellman, Shipley, and Kim 1996). Localization of boundaries in displays where completion is predicted by relatability is extremely precise. This is true for straight (collinear) and curved completions. A very different outcome occurs in cases where completion is predicted to follow global symmetry. Here, the precision (difference between "out" and "in" thresholds) is an order of magnitude worse. It is about 15 mm in a display of about 70 cm diameter (in visual angle, about 20 arcmin in a display 87 arcmin in diameter). Moreover, the midpoint of the range is close to 1 cm away from the theoretically predicted location of the boundary. This result has shown up consistently in a range of displays testing symmetry and related global notions of object completion. There are a number of issues still under investigation in this new paradigm. What is already clear is that global influences do not lead to specification of precise boundary position in the way local perceptual completion does. These outcomes are consistent with the idea of separate perceptual completion and more cognitive RPI processes.

Similarity

An interesting feature of edge relatability is that it does not seem to be sensitive to similarity of surface quality (e.g., lightness, color, or texture). Figure 11.13 gives two examples. In (a) the visible parts are seen as a unified object despite differences in their surface lightness and contrast polarity from the occluding object. In (b) an illusory figure is formed from connections between pieces of very different luminances. Shipley and Kellman (1992a) found that magnitude estimations of object completion under occlusion in a large sample of randomly generated figures showed no reliable differences whether the relatable pieces were the same or different in luminance and chromatic color. The Gestalt principle of similarity thus seems to have little effect on relatability or the boundary interpolation process in general.

Does this mean that there is no role for similarity in object completion? Kellman and Shipley (1991) proposed a surface-spreading process

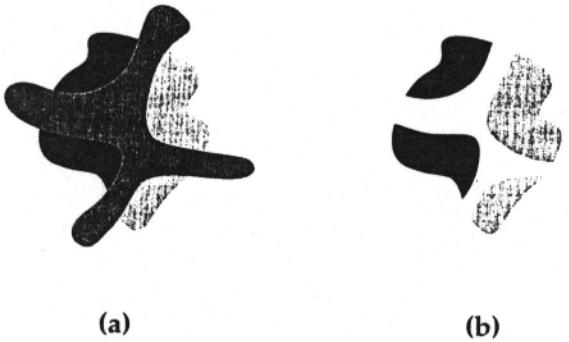

Figure 11.13.
Surface color insensitivity of boundary interpolation. a) A unitary partly occluded object is seen despite differences in lightness of its visible regions. b) Illusory contours form between surfaces of different lightnesses.

that complements boundary interpolation (cf. Yarbus 1967; Grossberg and Mingolla 1985). Surface quality spreads within physically specified and interpolated boundaries. In figure 11.14a the circle appears as a spot on a background. In figure 11.14b, the righthand circle still looks the same way but the lefthand circle may appear as a hole in the occluding surface. This effect appears to be dependent on similarity between the surface lightness and texture of the circle and the partly occluded ellipse. Because the circle has no TDs, it does not participate in the boundary interpolation process. What connects the circle with the surface behind the occluder appears to be a separate connecting process related to surface similarity. This surface process appears to be confined within the boundaries of the completed partly occluded figure in figure 11.14b. Figure 11.14c suggests, however, that surface spreading also occurs within the extended tangents of the boundaries of a partly occluded area (the half of the ellipse above the occluder), even when they are not relatable to others.

In her dissertation, Carol Yin tested these two hypotheses—that surface quality spreads within relatable edges and also within extended tangents of nonrelatable edges continuing behind occluding surfaces (Yin, Kellman, and Shipley 1997). In a series of experiments, subjects made a forced choice of whether a circular area appeared to be a hole in a surface or a spot on top of the surface in a number of displays varying edge and surface similarity relations. In a variant of the method, subjects made forced-choice responses of which of two displays looked more like it contained a hole for all possible pairs of displays in a particular experiment. These studies confirmed the hypotheses of surface spreading within relatable edges and tangent extensions. Yin also studied the surface completion process from an objective performance paradigm, pitting the effects of surface completion in making a circle look like a hole or a spot against small amounts of stereoscopic disparity. She found that surface completion interactions reduced sensitivity to stereoscopic depth (Yin, Kellman, and Shipley in press).

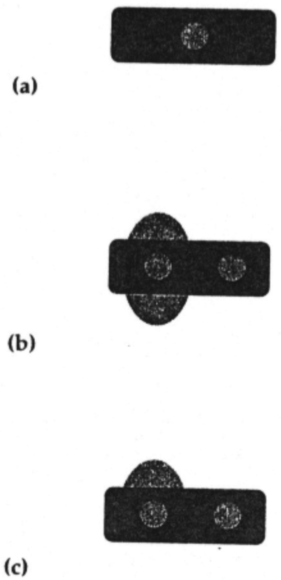

Figure 11.14.
Examples illustrating the surface completion process. a) The circle appears as a spot in front of a background. b) The lefthand circle now appears as a hole, due to surface completion, based on similarity of lightness and texture. c) Surface completion can occur even without edge relatability. (See text.)

Surface similarity and edge relatability seem to play complementary roles in object perception. Interpolated edges establish connections under occlusion, and surface qualities (lightness, color, and texture) spread within physically given and interpolated boundaries.

Common Fate

Wertheimer (1921) defined the "Factor of Common Fate" in this way. Suppose one sees a row of dots in which some are closer to others, leading to grouping by proximity. Now suppose some dots are shifted upward while others remain at rest. The shift will seem more disruptive if only dots that were initially grouped together are moved. If the shift involves some dots from different groups, it appears to change the grouping.

The principle of common fate received little emphasis in later Gestalt discussions of perceptual organization. In Koffka's (1935) treatise, for example, the principle is not even mentioned. In some ways, however, the nugget of insight in the principle of common fate connects to the most important modern developments in understanding perception. Owing in part to the development of ecological analyses of perception (Gibson 1966; Johansson 1968), we know that motion relationships provide a wealth of information about object structure and spatial layout.

For perceiving unity under occlusion, there are two distinct types of information (Kellman and Shipley 1991). One, a direct descendant of Wertheimer's common fate, we have called the *edge-insensitive process*. Certain motion relationships lead two visible parts to be seen as connected. This connecting principle does not require any particular relationships among the visible edges of the parts for unity to be seen. Computational and psychophysical research has revealed processes that can determine whether particular 2-D motion patterns are consistent with a rigid 3-D structure, and if so, what structure it is. Wertheimer's notion of common fate includes at least the stimulus relationships that allow recovery of rigid structure (Ullman 1979; Todd 1981). They may also include many nonrigid motions, such as the jointed motions characteristic of a moving human body, and elastic motions, characteristic of organisms or inanimate objects that stretch and contract during movement (Johansson 1975).

Spatiotemporal Relatability of Edges

A complementary process—the *edge-sensitive process*—does involve edge relationships in information given over time by motion. If a stationary observer looks through dense foliage, she may see meaningless fragments of color from the scene behind. If the observer moves while looking, however, the objects and spatial layout behind the foliage may be revealed. Sequential projection of parts seems to allow visual perception of complete objects, although this ability has not been much studied. There is evidence that sequential perception of inducing elements can produce illusory contours and figures (Kellman and Cohen 1984; Bruno and Bertamini 1988). Perception under these circumstances requires not only integration of information over time, but interpolation, because some parts of the object never project to the eyes. The situation is one encountered often in ordinary perception.

What stimulus relationships in both space and time lead to perception of complete objects? With the extra degree of freedom given by motion, attempting to answer this question might seem daunting. It might be possible, however, to extend the criterion of spatial relatability to account for completion in dynamic scenes. A simple hypothesis about how this might be done is illustrated in figure 11.15. In (a), a moving opaque panel containing two apertures moves in front of an object. Suppose one part of the figure becomes visible through an aperture at time t1 and another part becomes visible at time t2. If the position and edge orientation of the part seen at t1 is encoded in a buffer and persists until the part at t2 appears, the standard relatability computation can be performed on the currently visible part and the earlier encoded part. The situation in (b) adds a step. Here the object moves, revealing one part through the bottom aperture at t1 and another through the top aperture at t2. Here the hypothesis is that when the part appears at t1, the visual system encodes not only its position and edge orientation but a velocity signal. This velocity signal could be used to update the spatial position of the earlier visible part over time, either in a special-purpose buffer or by triggering a pursuit eye movement. When the second part becomes visible, it is combined with the updated position of the first part in the standard spatial relatability computation.

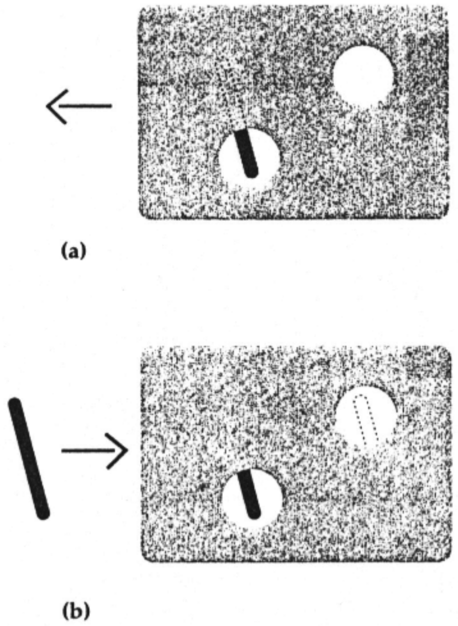

Figure 11.15.
Spatiotemporal relatability. a) A moving occluding panel with two windows passes in front of an object, projecting parts of the object to the eyes at different times. If a trace of the first visible part can be preserved until the second appears, spatial relatability can operate. b) A moving object's parts are projected at two different times in two different places. If velocity information is available, the position of the initially viewed part can be updated (by an eye movement or in a visual buffer) so that it's position relative to the second visible part can be extrapolated. Spatiotemporal relatability applies the spatial relatability computation to the currently visible and previously visible, positionally extrapolated parts. (From Palmer, Kellman, and Shipley, in preparation.)

The Dynamic Occlusion Paradigm

Evan Palmer, Tim Shipley, and I recently developed an experimental paradigm to test these ideas (Palmer, Kellman, and Shipley 1997). The paradigm works as follows. On each trial, an object passes behind an occluder with several narrow slits, vertically separated so that some parts of the object never project to the eyes. This feature makes the task a completion or interpolation task as opposed to only an integration task (where visible parts are integrated over time). On each trial an object passes once back and forth behind the occluder. Subjects then make

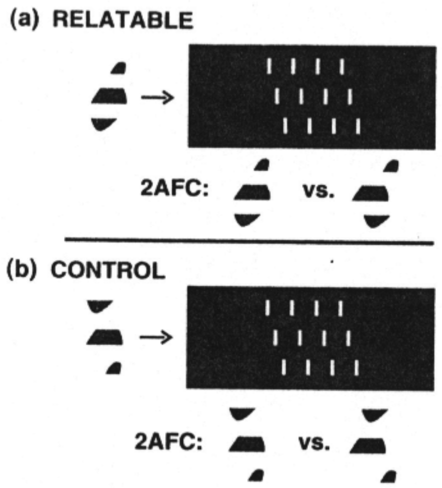

Figure 11.16.
Design for studying dynamic object completion. A target array consisting of three visible parts moves behind the occluder, visible only through narrow apertures. After each presentation, the subject makes a forced choice between two displays. a) Relatable display. b) Nonrelatable display. (See text.) (From Palmer, Kellman, and Shipley, in preparation.)

a forced choice between two test displays, choosing which matched the moving target display. The design is illustrated in figure 11.16.

Two display conditions were used. Relatable displays (apart from the shift manipulation; see below) met the criterion of spatiotemporal relatability. The upper test display in figure 11.16 is an example. The other test display differs from the first by having one of the three fragments shifted by some amount. Five different amounts of shift (ranging from 1.67 arcmin to 8.33 arcmin of visual angle) were used. The target matched the unshifted test display on half of the trials and the shifted display on the other half.

We predicted that relatability would facilitate encoding of the visible parts in the target display. If three parts moving behind slits were grouped into a single, coherent object, this might lead to more economical encoding and memory than for control displays (see below) in which three detached pieces were encoded. For simplicity, I will consider here only the cases in which either a test display or both the target and a test display were relatable. In these cases, it was predicted that the greater ease of encoding a relatable display would lead to better performance.

Displays in a second condition were compared to the first. These nonrelatable displays consisted of the identical three pieces as in the relatable condition, but the top and bottom pieces were permuted. (See figure 11.16b.) With these nonrelatable displays, it was hypothesized

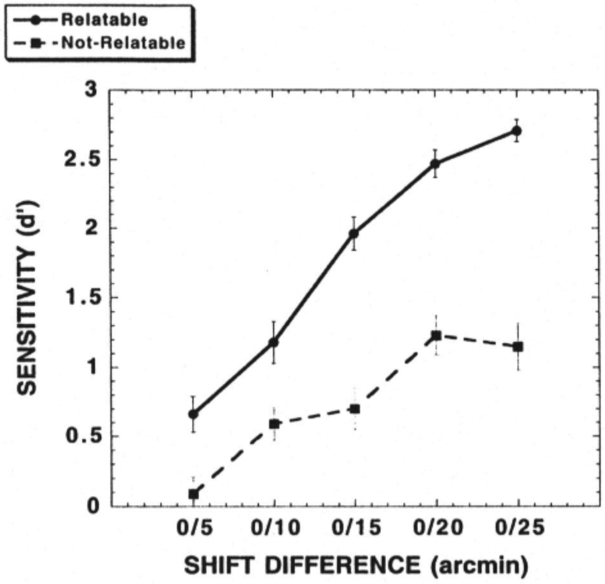

Figure 11.17.
Results of dynamic object completion experiment. Sensitivity is shown as a function of the misalignment difference between the canonical display and the other test choice. Separate plots are given for relatable and nonrelatable displays. (From Palmer, Kellman, and Shipley, in preparation.)

that visual completion would not occur; each nonrelatable target might have to be encoded as three distinct pieces, which would lead to greater encoding demands and lower sensitivity to the relative spatial positions of the three parts.

These experiments are just beginning, but we can present some early results. Figure 11.17 shows accuracy data (discrimination d') from 16 subjects for relatable and nonrelatable displays as a function of shift. Relatable displays were far more accurately discriminated than displays made of the identical physical parts but placed in nonrelatable positions. The results provide tentative support for generalizing the notion of relatability from the spatial to the spatiotemporal domain. There are a whole range of issues raised but not yet addressed by the results. For example, we did not control fixation, and it is unclear whether eye movements based on velocity signals from the moving fragments facilitate spatiotemporal object completion. Likewise, we have not yet investigated effects of a number of other parameters. One of special importance is velocity. We suspect from other research (Shipley and Kellman 1994) that spatiotemporal completion will occur within a restricted temporal window of integration, around 165 msec. So the results of our initial studies of dynamic occlusion raise more questions than they answer. They do provide some basis for connecting dynamic object perception to previous work with static displays, by means of the extended notion of relatability.

Neural Models

The theoretical ideas about boundary interpolation and surface filling that I have sketched are largely formal or computational in nature. That is, they characterize stimulus relationships that underlie object completion. They provide only hints about a precise process model or neural realization. I think it is worth concluding by mentioning some clues in these areas that are central to some of our current thinking and work in progress, as well as some work by others.

We defined relatability in edge interpolation as a simple mathematical relationship between edge pairs. A number of considerations are leading us to consider interpolation effects as resultants of excitation fields that arise from individual edges. For example, there is some evidence that edges and the surface of a single region continue behind an occluder even when they do not connect to any other region (Kanizsa 1979; Nakayama and Shimojo 1992). We call this edge *continuation* to distinguish it from edge completion or interpolation. In this case, edges seem to continue along linear extensions of edge tangents at the point of occlusion. Surface spreading along such tangent extensions was found in Yin's research, described above.

One way to account for edge continuation and interpolation is to assume that each physically specified edge at its endpoint gives rise to a field of excitations at nearby locations. A vector field would identify with each spatial location and at each orientation (perhaps in a 3-D network) a certain excitation. Excitation would decrease with distance and would also depend on the orientation and positional relations as specified in the geometry of relatability. An interpolated boundary in this scheme arises when excitation fields from two separate physically specified edges meet, with a winner-take-all inhibition scheme preventing multiple completions. The temporal component of spatiotemporal relatability could be realized by adding the dimension of time to the vector field.

Our research group and others are working on the specifics of this kind of model. For now it may be sufficient to note that this approach is consistent with some other psychophysical work, including that of Field and colleagues, Polat and Sagi (1994), Das and Gilbert (1995), and others. Both neurophysiological and psychophysical experiments suggest that cortical cells sensitive to orientation trigger the kinds of spatial interactions that could implement relatability. There is, of course, more work to do in pursuing these general ideas. A meaningful theory will build on previously proposed frameworks (Grossberg and Mingolla 1985; Grossberg 1994; Heitger and von der Heydt 1993) but specific quantitative relationships faithful to psychophysical data must be added. New dimensions must also be added. Our research suggests that successful models must incorporate relationships across all three spatial dimensions and relationships in information given over time. As daunting as the theoretical task appears, it may be made tractable by precisely characterizing the grammar of object completion. In particular, we are encouraged by the idea that a simple piece of geometry—the notion of relatability—may provide a common thread knitting together pictorial, 3-D, and spatiotemporal object completion. This unifying idea may provide a platform for precise process modeling and investigations into the underlying neural mechanics.

Conclusion

Understanding perceptual organization—and segmentation and grouping in particular—still poses deep mysteries to researchers in biological and artificial vision. Yet often, when progress is made, we can trace its roots to insights made more than a generation ago by the Gestalt psychologists. It is amazing to realize that not only did the Gestaltists provide some of the clues about how to solve these problems, but they were the first to articulate clearly that these problems existed at all. At the same time, it must be admitted that their principles lacked precision and coherence. That these principles can still be recognized in more recent computational models, however, attests to the robustness of the original insights. In this chapter, I have attempted to make explicit some of these connections between the old and the new.

A simple piece of geometry—the relatability criterion—appears to capture much of the grammar of edge interactions that lead to object completion. With rather simple extensions, relatability can be applied to contour interactions in depth and to dynamic object completion. Underlying this principle—and the Gestalt idea of good continuation—is the idea that object boundaries tend to be smooth. An alternative ecological interpretation might be that objects are not all that smooth, but for making inferences about where objects go under occlusion, smoothness is the best general assumption for a visual processor to use. Relatability might be implemented by simple interactions of units responding to oriented edges. Evidence is beginning to suggest that such interactions occur surprisingly early in cortical visual processing.

Complementary to the boundary completion process is the spreading of surface quality within boundaries. Here, the Gestalt principle of similarity lives on. Some other principles, such as an idea of Prägnanz or global symmetry, may turn out not to be determinants of perceptual representations per se, but may exert their effects more in memory and recognition.

Of the original Gestalt principles, it is the notion of good continuation that emerges as having the most important legacy in models of object perception. This is the principle that also stands out when I reflect on the impact of the Gleitmans and the Gleitman Research Seminar. These many years later, Henry's and Lila's insight, dedication, and high standards continue to help all of us in our academic endeavors. That we seek to emulate them in our own research and teaching is perhaps the best principle of good continuation.

Acknowledgments

Portions of this research were supported by National Science Foundation grant SBR-9496112. I thank Thomas Shipley, Carol Yin, Sharon Guttman, and Evan Palmer for useful discussions, and John Jonides and Dan Reisberg for helpful comments on an earlier draft of this chapter. Address reprint requests to Philip J. Kellman, Department of Psychology, UCLA, 405 Hilgard Avenue, Los Angeles, CA 90095–1563 or by email to <Kellman@cognet.ucla.edu>.

Note

1. Even the language we use to describe the idea contains the idea implicitly. We say a TD is a point where "contours meet," but the presence of the TD is what makes it sensible to say "contours" (plural). Without the TD there is only a single contour.

References

Bruno, N. and Bertamini, M. (1990) Identifying contours from occlusion events. *Perception and Psychophysics* 48 (4):331–342.

Buffart, H., Leeuwenberg, E., and Restle, F. (1981) Coding theory of visual pattern completion. *Journal of Experimental Psychology: Human Perception and Performance* 7(2):241–274.

Clowes, M. B. (1971) On seeing things. *Artificial Intelligence* 2:79–112.

Das, A. and Gilbert, C. D. (1995) Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature* 375(6534):780–784.

Field, D., Hayes, A., and Hess, R. F. (1993) Contour integration by the human visual system: Evidence for a local "association field." *Vision Research* 33 (2):173–193.

Gibson, J. J. (1966) *The Senses Considered as Perceptual Systems.* Boston: Houghton-Mifflin.

Gibson, J. J. (1979) *The Ecological Approach to Visual Perception.* Boston: Houghton-Mifflin.

Grossberg, S. (1994) 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics* 55 (1):48–120.

Grossberg, S. and Mingolla, E. (1985) Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review* 92:173–211.

Heitger, F. and von der Heydt, R. (1993) A computational model of neural contour processing: Figure-ground segregation and illusory contours. *Proceedings of the Fourth International Conference on Computer Vision.* Los Alamitos, CA: IEEE Computer Society Press, 32–40.

Johansson, G. (1970) On theories for visual space perception: A letter to Gibson. *Scandinavian Journal of Psychology* 11(2):67–74.

Johansson, G. (1975) Visual motion perception. *Scientific American* 232(6):76–88.

Kanizsa, G. (1979). *Organization in Vision.* New York: Praeger.

Kawaguchi, J. (1988) Priming effect as expectation. *Japanese Psychological Review, Special Issue: Problems of repetition in memory* 31(3):290–304.

Kellman, P. J. and Cohen, M. H. (1984) Kinetic subjective contours. *Perception and Psychophysics* 35(3):237–244.

Kellman, P. J., Machado, L., Shipley, T. F. and Li, C. C. (1996) 3-D determinants of object completion. *Investigative Ophthalmology and Visual Science* 37(3):685.

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

Kellman, P. J. and Shipley, T. F. (1992) Visual interpolation in object perception. *Current Directions in Psychological Science* 1(6):193–199.

Kellman, P. J., Shipley, T. F., and Kim, J. (1996) Global and local effects in object completion: Evidence from a boundary localization paradigm. Paper presented at the 32nd Annual Meeting of the Psychonomic Society, St. Louis, Mo., November 1996.

Kellman, P. J. and Spelke, E. S. (1983) Perception of partly occluded objects in infancy. *Cognitive Psychology* 15:483–524.

Kellman, P. J., Yin, C., and Shipley, T. F. (1998) A common mechanism for illusory and occluded object completion. *Journal of Experimental Psychology: Human Perception and Performance* 24(3):859–869.

Kellman, P.J., Yin, C., Shipley, T.F., Machado, L. and Li, C.C. The 3-D geometry of object completion. Manuscript in preparation.

Koffka, K. (1935). *Principles of Gestalt Psychology.* New York: Harcourt.

Lee, D. N. (1974) Visual information during locomotion. In R. B. MacLeod and H. L. Pick (eds.), *Perception: Essays in Honor of James J. Gibson.* Ithaca, NY: Cornell University Press.

Lesher, G. W. and Mingolla, E. (1993) The role of edges and line-ends in illusory contour formation. *Vision Research* 33(16):2253–2270.

Marr, D. (1982) *Vision.* San Francisco: Freeman.

Michotte, A., Thines, G., and Crabbe, G. (1964) Les complements amodaux des structures perceptives. *Studia Psycologica.* Louvain: Publications Universitaires de Louvain.

Nakayama, K., and Shimojo, S. (1992) Experiencing and perceiving visual surfaces. *Science* 257(5075):1357–1363.

Nakayama, K., Shimojo, S., and Silverman, G. (1989) Stereoscopic depth: Its relation to image segmentation, grouping and the recognition of occluded objects. *Perception* 18(1):55–68.

Palmer, E., Kellman, P. J., and Shipley, T. F. (1997) Spatiotemporal relatability in dynamic object completion. *Investigative Ophthalmology and Visual Science* 38(4):256.

Palmer, E., Kellman, P. J. and Shipley, T. F. Spatiotemporal relatability in dynamic object completion. Manuscript in preparation.

Petter, G. (1956) Nuove ricerche sperimentali sulla totalizzazione percettiva. *Rivista di psicologia* 50:213–227.

Polat, U. and Sagi, D. (1994) The architecture of perceptual spatial interactions. *Vision Research* 34(1):73–78.

Prenter, P. M. (1989) *Splines and Variational Methods.* New York: Wiley.

Ringach, D. L. and Shapley, R. (1996) Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research* 36:3037–3050.

Sekuler, A. B. (1994) Local and global minima in visual completion: Effects of symmetry and orientation. *Perception* 23(5):529–545.

Sekuler, A. B., Palmer, S. E., and Flynn, C. (1994) Local and global processes in visual completion. *Psychological Science* 5(5):260–267.

Shipley, T. F. and Kellman, P. J. (1992a) Perception of partly occluded objects and illusory figures: Evidence for an identity hypothesis. *Journal of Experimental Psychology: Human Perception and Performance* 18(1):106–120.

Shipley, T. F. and Kellman, P. J. (1992b) Strength of visual interpolation depends on the ratio of physically-specified to total edge length. *Perception and Psychophysics* 52(1):97–106.

Shipley, T. F. and Kellman, P. J. (1994) Spatiotemporal boundary formation. *Journal of Experimental Psychology: General* 123(1):3–20.

Todd, J. T. (1982) Visual information about rigid and nonrigid motion: A geometric analysis. *Journal of Experimental Psychology: Human Perception and Performance* 8(2):238–252.

Ullman, S. (1979) *The Interpretation of Visual Motion*. Cambridge, MA: MIT Press.

van Lier, R. J., van der Helm, P. A., and Leeuwenberg, E. L. J. (1995) Competing global and local completions in visual occlusion. *Journal of Experimental Psychology: Human Perception and Performance* 21(3):571–583.

Waltz, D. L. (1972) Generating semantic descriptions from drawings of scenes with shadows (Tech. Rep. AI-TR-271). Cambridge, MA: MIT.

Wertheimer, M. (1921) Laws of organization in perceptual forms. In W. D. Ellis (ed.), *Readings from Gestalt Psychology*. New York: Harcourt Brace, 1938.

Yarbus, A. L. (1967) *Eye Movements and Vision*. New York: Plenum Press.

Yin, C., Kellman, P. J., and Shipley, T. F. (1997) Surface completion complements boundary interpolation. *Perception*, special issue on surface appearance, 26:1459–1479.

Yin, C., Kellman, P.J., and Shipley, T.F. (in press). Surface integration influences depth discrimination. *Vision Research*.

Chapter 12

Beyond Shipley, Smith, and Gleitman: Young Children's Comprehension of Bound Morphemes

Katherine Hirsh-Pasek

In the fall of each year, as leaves turn bright against the New England landscape, psycholinguists make their annual pilgrimage to the Boston Language Conference. One of the highlights of the trip to Boston is the Gleitman dinner, a gathering of all those fortunate enough to be Lila and Henry's intellectual children, grandchildren, and great grandchildren. As you look around the dining room, you can't help but be impressed by the large number of scientists who have been touched by the Gleitman tradition, a tradition characterized by outstanding scholarship, first-rate teaching, and personal friendship.

There is no match for the scholarship that we witnessed during our graduate years. Lila always understood the big picture of language development, constantly reframing our narrow questions into ones that addressed major issues in the field. I remember marveling at the way in which she made our first-year research projects seem so much more important than we had imagined. (She magically molded my research on young children's understanding of jokes into a key project on the relationship between metalinguistic processing and reading.) Lila also had (and still has) the insight and common sense to know just where to look to test her account of a developmental story. She has that rare ability to integrate data from linguistic and psychological journals with examples from the TV guide, Star Trek, and a neighborhood two-year-old. While Lila helped us ask the questions, however, it was Henry who would sculpt those questions into psychologically interesting research. The result was a constant stream of papers in child language, each of which fit into a larger program of research, many of which became classics in the field.

Their scholarship is unquestioned, yet their style of teaching and advising stand out as the shining light of my graduate years. When my thirteen-year-old son recently asked Henry what he would describe as his greatest accomplishment in psychology, he answered without hesitation, "My students." No one who worked with Henry or Lila would be surprised by that answer. The Thursday night cheese seminars at the