Behavior Research Methods, Instruments, & Computers
1997, 29 (1), 27-36

HIGH-PERFORMANCE COMPUTING AND HUMAN VISION I

Chaired by Albert Yonas, *University of Minnesota*

How the eye measures reality and virtual reality

JAMES E. CUTTING
Cornell University, Ithaca, New York

If virtual reality systems are to make good on their name, designers must know how people perceive space in natural environments, in photographs, and in cinema. Perceivers understand the layout of a cluttered natural environment through the use of nine or more sources of information, each based on different assumptions—occlusion, height in the visual field, relative size, relative density, aerial perspective, binocular disparities, accommodation, convergence, and motion perspective. The relative utility of these sources at different distances is compared, using their ordinal depth-threshold functions. From these, three classes of space around a moving observer are postulated: personal space, action space, and vista space. Within each, a smaller number of sources act in consort, with different relative strengths. Given the general ordinality of the sources, these spaces are likely to be affine in character, stretching and collapsing with viewing conditions. One of these conditions is controlled by lens length in photography and cinematography or by field-of-view commands in computer graphics. These have striking effects on many of these sources of information and, consequently, on how the layout of a scene is perceived.

We, as a species, seem to have been fascinated with pictures throughout our history. The paintings at Niaux, Altamira, and Lascaux (Clottes, 1995; Ruspoli, 1986), for example, are known to be about 14,000 years old, but with the recently discovered paintings in the Grotte Chauvet, the origin of representational art appears to have been pushed back even further (Chauvet, Brunel Deschamps, & Hillaire, 1995; Clottes, 1996), to 20,000 years ago if not longer.[1] Thus, these paintings date from about the time at which *homo sapiens sapiens* first appeared in Europe (Nougier, 1969). We should remember these paintings in the context of virtual reality; our fascination with pictures is by no means recent.

My intent is threefold: first, to discuss our perception of the cluttered layout, or space, that we normally find around us; second, to discuss the development of representational art up to our current appreciation of it; and third, to apply this knowledge to virtual reality systems. The first discussion focuses on the use of multiple sources of information specifying ordinal depth relations, within the theoretical framework that I have called *directed perception* (Cutting, 1986, 1991). The second discussion is embedded within the first, but it is steeped in neither history nor art history; instead, it is offered through the peculiar eyes of optics and psychology, particularly psychophysics. The third is addressed to one of the pressing problems of graphics and of virtual reality—how we perceive the layout of the environments that we simulate.

A FRAMEWORK FOR UNDERSTANDING THE PERCEPTION OF LAYOUT

I will focus on nine sources of information for the perception of layout (or depth) roughly in their order of discovery or use in various modes of depiction. They are the following: occlusion (often called interposition), relative size, height in the visual field (often called height in the picture plane, or angular elevation), relative density, aerial perspective (often called atmospheric perspective), binocular disparities, accommodation, convergence, and motion perspective. What follows draws and expands upon earlier work (Cutting & Vishton, 1995).

Methods and Assumptions

In using these different sources of information, the human eye and mind measure the world in different ways, as is suggested in Table 1. To compare these sources, I adopt the weakest common measurement scale—the ordinal scale—and consider the just noticeable difference (JND) in depth for two objects at different distances, given pre-

Supported by a John Simon Guggenheim Memorial Fellowship (1993-94) and by National Science Foundation Grants SBR-9212786 and ASC-9523483. Requests for information can be sent to the author at the Department of Psychology, Uris Hall, Cornell University, Ithaca, NY 14853-7601 (e-mail: jec7@cornell.edu).

Table 1
Assumptions and Scales for Each of Nine Sources of Information About Layout and Depth

Source of Information	Assumptions	Implied Measurement Scale
All	Linearity of light rays (see Burton, 1945; but also Minnaert, 1993, for exceptions). Luminance or textural contrast. In general, the rigidity of objects (rigidity implies object shape invariance).	—
1. Occlusion	Opacity of objects. Helmholtz's rule, or good continuation of the occluding object's contour (Hochberg, 1971; Ratoosh, 1949; but see Chapanis & McCleary, 1955).	Ordinal
2. Height in the visual field	Opacity of objects and of the ground plane. Gravity, or the bases of objects are on the ground plane. The eye is above the surface of support. The surface of support is roughly planar. (In hilly terrain, use may be restricted to the surface directly beneath the line of sight to the horizon.)	Ordinal, perhaps occasionally better
3. Relative size	Similarly shaped objects have similar physical size (Bingham, 1993) Objects are not too close. Plurality of objects in sight. (*Not* familiarity with the objects, which denotes "familiar size" [e.g., Epstein, 1963]).	Unanchored ratio possible, but probably ordinal
4. Relative density	Similarly shaped objects or textures have uniform spatial distribution. Plurality of objects or textures in the field of view.	Probably ordinal at best
5. Aerial perspective	The medium is not completely transparent. The density of the medium is roughly uniform.	Probably ordinal
6. Binocular disparities	The distance between eyes. The current state of vergence. Unambiguous correspondences.	Absolute (Landy et al., 1991), but perhaps only ordinal (van den Berg & Brenner, 1994)
7. Accommodation	Complex spatial frequency distribution (Fisher & Ciuffreda, 1988). The current state.	Ordinal at best
8. Convergence	The distance between eyes. The current state.	Ordinal
9. Motion perspective	A rigid environment. A spatial anchor of zero motion (horizon or a fixated object).	Absolute (Landy et al., 1991), unanchored ratio, but perhaps only ordinal

vious data and logical considerations. This procedure embraces scale convergence (Birnbaum, 1983), a powerful tool for perception and for science in general, and starts with weak assumptions (ordinality) in effort to converge on a near-metric (probably affine) representation of space.[2]

I will plot distance thresholds on graphical coordinates analogous to those of contrast sensitivity in the spatial-frequency domain. That is, in considering the distances of two objects, D1 and D2, I determine the JND of distance between them as scaled by their mean egocentric distance [2(D1 − D2)/(D1 + D2)], and then plot these values as a function of their mean distance from the observer ([D1 + D2]/2). Nagata (1991) was the first to present such diagrams, but my analysis differs from his in many respects. In relying on such plots, I will make several assumptions: (1) that a depth threshold of 10%, or 0.1 on the ordinate of Figure 1, is a useful limit in considering contributions to the perception of layout; (2) that the observer can look around, registering differences on the fovea and other regions of the retina, and retain that information; (3) that each source pertains to objects of appropriate size, so that the observer can easily resolve what he or she is looking at; and perhaps most importantly, (4) that threshold measurements are informative for everyday, suprathreshold considerations. In addition, each source itself is based on a different set of assumptions, which are also given in Table 1. This allows for sources to ramify each other, or for one source to falsify the assumptions of another.

Nine Sources of Information and Their Relative Efficacy

1. *Occlusion* occurs when one object hides, or partly hides, another from view. As an artistic means of conveying depth information, partial occlusion has been used in paleolithic (see the images in Biederman, 1948;

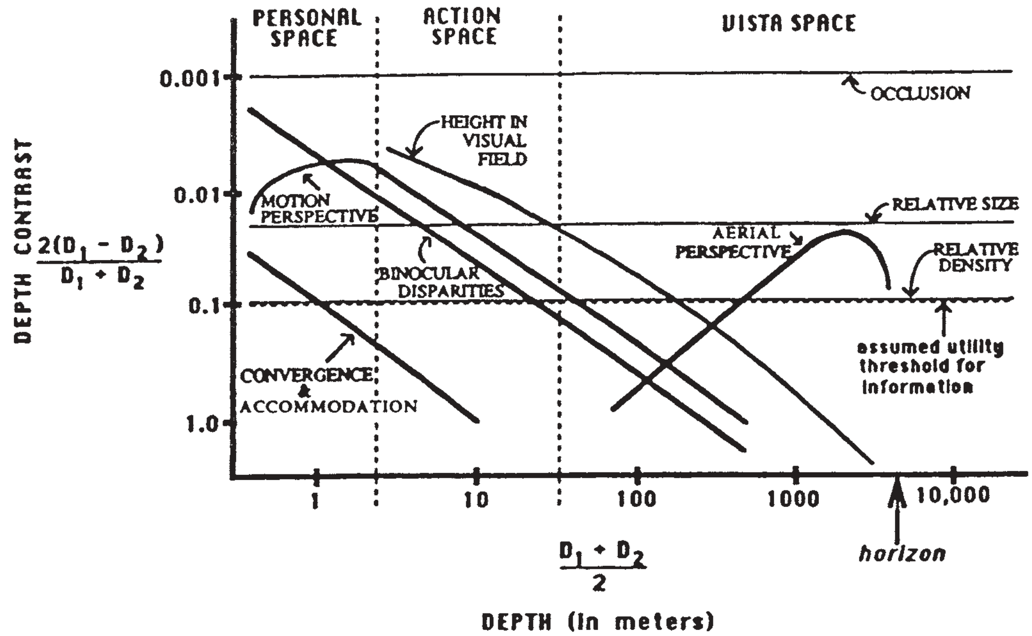

Figure 1. Just-discriminable ordinal depth thresholds as a function of the logarithm of distance from the observer, from 0.5 to 10,000 m, for nine sources of information about layout. I assume that more potent sources of information are associated with smaller depth-discrimination thresholds; and that these thresholds reflect suprathreshold utility. This array of functions is idealized for the assumptions given in Table 1. From "Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Contextual Use of Different Information About Depth," by J. E. Cutting and P. M. Vishton, 1995, in W. Epstein and S. Rogers (Eds.), *Perception of Space and Motion* (p. 80), San Diego: Academic Press, Copyright 1995 by Academic Press. Reprinted with permission.

Chauvet et al., 1995; Hobbs, 1991) and Egyptian art (see Hagen, 1986; Hobbs, 1991), where it is often used alone, with no other information to convey depth. Thus, one can make a reasonable claim that occlusion was the first source of information discovered and used to depict spatial relations in depth.

Because occlusion can never be more than ordinal information—one can only know that one object is in front of another, but not by how much—it may not seem impressive. Indeed, some researchers have rejected it as information about depth (e.g., Landy, Maloney, Johnston, & Young, 1995). But the range and power of occlusion is striking: As is suggested in Figure 1, it can be trusted at all distances without attenuation, and its depth threshold exceeds that of all other sources. Even stereopsis seems to depend on partial occlusion (Anderson & Nakayama, 1994). Normalizing size over distance, occlusion provides depth thresholds of 0.1% or better. This is the width of one sheet of paper against another at 30 cm, the width of a person against a wall at 500 m, or the width of a car against a building at 2 km. Cutting and Vishton (1995) have provided more background on occlusion along with justifications for this plotted function, as well as for those of the other sources of information discussed here.

2. *Height in the visual field* measures relations among the bases of objects in a 3-D environment as projected to the eye, moving from the bottom of the visual field (or image) to the top, and assuming the presence of a ground plane, of gravity, and the absence of a ceiling (see Dunn, Gray, & Thompson, 1965). Across the scope of many different traditions in art, a pattern is clear: If one source of information about layout is present in a picture beyond occlusion, that source is almost always height in the visual field. The conjunction of occlusion and height, with no other sources, can be seen in the paintings at Chauvet; in classical Greek art and in Roman wall paintings; in 10th-century Chinese landscapes; in 12th- to15th-century Japanese art; in Western works of Cimabue, Duccio di Buoninsegna, Simone Martini, and Giovanni di Paolo (13th–15th centuries); and in 15th-century Persian art (see Blatt, 1984; Chauvet et al., 1995; Cole, 1992; Hagen, 1986; Hobbs, 1991; Wright, 1983). Thus, height appears to have been the second source of information discovered, or at least mastered, for portraying depth and layout.

The potential utility of height in the visual field is suggested in Figure 1, dissipating with distance. This plot assumes an upright, adult observer standing on a flat plane. Since the observer's eye is at a height of about 1.6 m, no base closer than 1.6 m will be available; thus, the function is truncated in the near distance, which will have implications later. I also assume that a height difference of about 5′ of arc between two nearly adjacent objects is just detectable; but a different value would simply shift the function up or down. When one is not on a flat plane, the shape of the functions may change within small vertically sliced regions extending outward along any line of sight, but ordinality will be maintained.

3. *Relative size* is a measure of the angular extent of the retinal projection of two or more similar objects or textures. Well before the invention of globally coherent linear perspective, relative size was conjoined with occlusion and height in the visual field to help portray depth—for example, in Greek vases, or in the pre-Renaissance art of Giotto, Gaddi, and Lorenzini (14th century). And in many traditions within Chinese and Japanese art, all three sources can be found together, with no others (see Cole, 1992; Hagen, 1986). Thus, one can argue that relative size was the third source discovered and mastered to depict the layout of objects.

Relative size has the potential of yielding ratio information. That is, for example, if one sees two similar objects, one of which subtends one fourth the visual angle of the other, the former will likely be four times farther away. These, of course, could be at 10 and 40 cm or 10 and 40 km; no absolute anchor is implied. Nonetheless, as with the other sources, I will consider only its ordinality. Relative size has been studied in many contexts. Perhaps the clearest results are those of Teichner, Kobrick, and Wehrkamp (1955), who mounted large placards on two jeeps, drove them out into a desert, and measured observers' relative distance JNDs. Their data are replotted in Figure 1. Relative size can generally be trusted throughout the visible range of distances, providing a depth threshold of about 3%, or about 1.5 log units worse than occlusion.[3]

4. *Relative density* concerns the projected number of similar objects or textures per solid visual angle (see Barlow, 1978; Durgin, 1995). It appeared in Western art only with the near-full development of linear perspective. For example, the effects of relative density can be seen in the local (not fully coherent) perspective piazzas of Lorenzetti, but more strikingly in the global perspective tiled floors of Donatello, Massachio, and Uccello (15th century; see Cole, 1992). In fact, only with linear perspective are these first four sources of information—occlusion, height, size, and density—coupled in a rigorous fashion. In contemporary computer graphics, this coupling is accomplished through the hardware geometry engine and *z* buffering used to generate a display from a particular point of view.

The psychological effects of density were first discussed by Gibson (1950) and researched by Flock (1964), but their perceptual effects are not large (see Marr, 1981, p. 236). Indeed, Cutting and Millard (1984) showed that relative density was psychologically less than half as effective as the relative size in revealing exocentric depth. Thus, I have plotted it as weaker than relative size in Figure 1, two log units below occlusion, and at the 10% threshold.

5. *Aerial perspective* refers to the increasing indistinctness of objects with distance, determined by moisture and/or pollutants in the atmosphere between the observer and these objects. Its perceptual effect is a decrease in contrast with distance, converging to the color of the atmosphere. Although aerial perspective appears in art as early as Giotto, it was not systematically discussed and understood until Leonardo (15th–16th centuries; Richter, 1883/1970; see also Bell, 1993). In computer graphics, aerial perspective is understood in terms of participating media, typically generated by a "fog" command, associated with the geometry engine.

As is shown in Figure 1, the effectiveness of aerial perspective increases with the log of distance until luminance differences reach threshold (Nagata, 1991); but the effective range varies greatly, depending on air quality. Underwater visibility (Lythgoe, 1979) is analogous to aerial perspective, and the ecological roots of the perception of transparency (Gerbino, Stultiens, & Troost, 1990; Metelli, 1974) are to be found in aerial perspective.

6. *Binocular disparity* is the difference in relative position of an object as projected on the retinas of the two eyes. When disparities are sufficiently small, they yield stereopsis—or the impression of solid space. No source of information, other than perhaps motion (Rogers & Graham, 1979), can produce such a compelling impression of depth. When disparities are greater than stereopsis will allow, they yield diplopia—or double vision—which is also informative about relative depth (Duwaer & van den Brink, 1981; Ogle, 1952). The effect of disparities has been comprehensively studied (see Arditi, 1986; Gulick & Lawson, 1976, for reviews), and ordinal thresholds can be found throughout the literature (e.g., Nagata, 1991; Ogle, 1958). These are replotted in Figure 1. Binocular disparities have the potential of yielding absolute information about distances near the observer (Landy et al., 1995, Landy, Maloney, & Young, 1991), although they do not always appear to be used as such (van den Berg & Brenner, 1994). Stereo is also extremely malleable; it demonstrates large hysteresis effects (Julesz, 1971), and just one day of monocular vision can render one temporarily stereoblind (Wallach & Karsh, 1963).

Although it is clear that Leonardo and Dürer were aware of the problem of two eyes located in different positions, neither seemed to understand its implication. It took Wheatstone (1838) to exploit disparities to their fullest, and his and Brewster's stereoscope (Gulick & Lawson, 1976), to present them widely to the public. Typically produced with two cameras mounted as much as a meter apart at the same height, 19th century stereograms yielded vivid, if toy-like, impressions of the major cities of Europe. In America, stereograms of sequoias were not uncommon, but as something of a visual oxymoron. Because of the exaggerated disparities, the sequoias looked small; thus, people were added; but, of course, these looked small too. There is a century-old lesson here for computer graphics: Stereo can be produced easily through goggles that alternately present disparity images to the two eyes, and one has independent control over the spacing of the two graphics cameras, but beware!

7. *Accommodation* occurs with the change in the shape of the lens of the eye, allowing it to focus on objects near or far while keeping the retinal image sharp. Objects at other distances are blurred. Near and far points vary across individuals and, with age, within individuals. The effi-

cacy of accommodation alone is probably less than 2 m or so (Fisher & Ciuffreda, 1988), and it declines with age; but it can interact with other sources of information (Roscoe, 1984; Wallach & Norris, 1963). Although the effects of accommodation have been known at least since the time of Descartes, the artistic use of it may have first occurred with the Impressionists at the end of 19th century, only after the advent of photography (see Scharf, 1968). In computer graphics, problems with near accommodation are dealt with through infinity optics in head-mounted displays, but genuine image blur is computationally intensive. Raytracing techniques are typically done for the analogue of a pinhole camera.

8. *Convergence* is measured as the angle between foveal axes of the two eyes. When the angle is large, the two eyes are canted inward to focus near the nose; when it approaches 0º, the two eyes are aligned to focus near the horizon. Convergence is effective at close range, but not beyond about 2 m (Gogel, 1961; Hofsten, 1976; Lie, 1965). Although convergence has been known at least since the time of Berkeley, it also seems quite likely that it has no possible artistic use. In graphics, the effects of convergence result from their coupling with stereo. The limits of accommodation and convergence together are less than 3 m (Leibowitz, Shina, & Hennessy, 1972; see also Kersten & Legge, 1983; Morgan, 1968), as is suggested in Figure 1.

9. *Motion perspective* refers to the field of relative motions of objects rigidly attached to a ground plane around a moving observer (Helmholtz, 1867/1925; Gibson, 1950); it specifically does not refer to object motion. The first artistic uses of it were seen in films at the end of the 19th century (e.g., Toulet, 1988), in which cameras were mounted on cars, trolleys, and trains, and the effects were presented to appreciative audiences. In computer graphics, motion perspective is part of the cluster of information sources calculated by the geometry engine.

Ferris (1972) and Johansson (1973) demonstrated that, through motion perspective, individuals are quite good at judging absolute distances up to about 5 m, and ordinal accuracy should be high at greater distances as well (but see Gogel & Tietz, 1973). Based on flow rates generated at 2 m/sec and at the eye height of a pedestrian, the thresholds in Figure 1 are foveal for a roving eye, one undergoing pursuit fixation during locomotion (see Cutting, Springer, Braren, & Johnson, 1992). Graham, Baker, Hecht, and Lloyd (1948) and Zegers (1948) measured difference thresholds for motion detection; these values are used for a pedestrian at near distances (<10 m), and absolute motion thresholds (1′ of arc/sec) are used for distances beyond. However, this also assumes that one is looking at 90º to the direction of movement, somewhat unusual for a pedestrian; velocities nearer the heading vector are considerably slower. And of course, if one is riding in a car, or particularly in a fast train (such as the TGV in France), the function is moved to the right considerably.

Sources Not Included

In this list of nine sources of information about layout, I have left out at least six other candidates, some commonly accepted in the literature. Their omission has been purposeful, but needs some explanation. Let us consider each in turn.

Texture gradients were proposed by Gibson (1950) as important for the perception of layout. However, there are three such gradients—those of size, density, and compression. Two have already been considered, as relative size and relative density. On the basis of the slant results of Freeman (1966), Nagata (1991) suggested that compression should also be a strong depth source. However, Cutting and Millard (1984) demonstrated that compression is ineffective in revealing depth; instead, it serves as good information about object shape (see also Todd & Akerstrom, 1987). Compression is also a part of linear perspective (considered next) and, in computer graphics, is handled by the geometry engine.

Linear perspective is powerful in revealing depth (see, e.g., Kubovy, 1986), but I omit its discussion, because, as has been noted above, it is a system, not a unique source. It combines the three texture gradients with occlusion and height in the visual field, and then, following Leonardo (Richter, 1883/1970), is separated from natural perspective simply by the copious use of projections of parallel lines.

Brightness and shading have appeared on many lists about depth and layout (e.g., Boring, 1942; Landy et al., 1995; Nagata, 1991). In theatrical-like settings, increasing the luminance of an object can cause it to appear to move forward. However, in such cases, luminance seems to act as a surrogate for relative size, and it has been used so experimentally (Dosher, Sperling, & Wurst, 1986). More ecologically, the play of light on form yields shadows, which have also been construed as information about depth (see, e.g., Bülthoff & Mallot, 1988; Yonas & Granrud, 1985). However, with others (Cavanagh & Leclerc, 1990; Norman, Todd, & Phillips, 1995; Todd & Reichel, 1989), I think that attached shadows are more information about object shape than depth per se. If so, this has implications for computer graphics: Perhaps certain scenes can be usefully rendered with local shading models rather than global ones, each object shaded independently of the others. Cast shadows, on the other hand, have some capability of portraying depth (e.g. Gombrich, 1995; Wanger, Ferwerda, & Greenberg, 1992).

Kinetic depth is revealed through object motion (Wallach & O'Connell, 1953; see also Proffitt, Rock, Hecht, & Schubert, 1992) and has been thought to be information about depth (see, e.g., Landy et al., 1995). However, I suggest that, like compression and shading, this information concerns object shape, not layout. Moreover, in rotating wireframe objects it often fails ordinality—figures reverse in depth, even when perspective information is added (Braunstein, 1962, 1976).

Kinetic occlusion and disocclusion—also called accretion and deletion of texture—can yield compelling impressions of depth order (Kaplan, 1969; Yonas, Craton, & Thompson, 1987). However, kinetic occlusion is simply occlusion revealed through motion; it differs only in that it need assume no luminance or textural contrast at the occluding edge. In addition, it can also be construed as two values in a motion perspective function as combined with occlusion.

Gravity was proposed as a source of depth information by Watson, Banks, Hofsten, and Royden (1992), since retinal velocities of dropped or thrown objects will vary with distance. As clever an idea as this is, gravity as such could only be information for an isolated object, not layout.

PERCEPTUAL SPACES, GEOMETRIES, AND INFORMATION CONFLICT

The relations in Figure 1 suggest a framework for contextual use of these nine sources of information about layout. From their pattern of thresholds, the layout around a moving perceiver may be divided into three egocentric regions.

Personal space immediately surrounds the observer's head, generally within arm's reach and slightly beyond (see also Hall, 1966; Sommer, 1969). I delimit this space within about 1.5 m, and it is typically a space worked within by a stationary individual. Thus, motion perspective is not typically generated by the observer; instead, motion parallax and structure-from-motion information are generated by observer manipulation to reveal object shape. Within this region I claim that five sources of information are generally effective—occlusion, retinal disparity, relative size, and then convergence and accommodation—and when each is available, I claim that they should roughly dominate each other in that order.

In certain virtual settings, such as that of telemedicine, computer graphics applications portray personal space almost exclusively. This means that heads-up displays are compulsory, and there must be heavy use of the geometry engine (for occlusion and size, particularly), stereo goggles (for both disparities and convergence), and, at least potentially, the optical adjustments for accommodation. However, given that the ability to accommodate often declines to near zero by age 40, and that many surgeons will be older than that, it may not be necessary to plan for accommodatory processes other than to sharpen images.

Action space is a circular, general planar region beyond personal space. We move quickly within this space, talk within it, and if need be, we can throw something to a compatriot or at an animal. This space appears to be served by a different collection and ranking of sources: Occlusion, height in the visual field, binocular disparity, motion perspective, and relative size. Because the utility of binocular disparity and motion perspective decline to 10% roughly at about 30 m, this effectively delimits action space at the far boundary. The near boundary is delimited by the emergence of height in the visual field as a strong source, which also serves to limit this space to the ground plane. Viewing objects from above or below tends to make perception of their layout less certain.

In many other virtual settings, such as those presented in architectural walk-throughs, action space is the exclusive domain. This means that larger displays are possible and the observer can be placed at some distance from them. Heavily used are both the geometry engine (for occlusion, size, height, and motion perspective) and stereo goggles coupled to it (for disparities and convergence).

Vista space occurs beyond about 30 m, at least for a pedestrian, since the effectiveness of binocular disparity and motion perspective are negligible. In this space, the only effective sources are the traditional "pictorial cues"—occlusion, height in the visual field, relative size, relative density, and aerial perspective. Vista space can be strikingly portrayed in large *trompe l'oeil* paintings—such as on the Pozzo ceiling in the Church of St. Ignazio in Rome (Pirenne, 1970)—and in cinema, particularly in widescreen format. And of course, for several centuries, Renaissance and Baroque artists specialized in portraying action and vista spaces for stationary objects and observers.

Geometries of Spaces

In general, two empirical findings occur when viewers estimate depth on the basis of multiple sources of information, often called "cues": Adding information generally increases the amount of depth seen and generally increases the consistency and accuracy with which depth judgments are made (Künnapas, 1968). Nonetheless, there is a tradition in what we now call cognitive science—stretching from Thomas Reid in the 18th century (Daniels, 1974) to the present, in both philosophy (Grünbaum, 1973; Suppes, 1977) and psychology (Battro, Netto, & Rozestraten, 1983; Blank, 1978; Foley, 1991; Indow, 1990; Luneburg, 1947; Tittle, Todd, Perotti, & Norman, 1995)—of reflection and data on the non-Euclidean nature of perceived space.

Not necessarily inconsistent with this view, however, are several other traditions that support the idea that layout is perceived reasonably accurately when measured in a Euclidean fashion. In one, individuals are placed in a cluttered environment within action space and are asked to make distance judgments among objects in that space (see, e.g., Kosslyn, Pick, & Fariello, 1974; Toye, 1986). The half-matrix of interobject judgments is then entered into a nonmetric multidimensional scaling program (NMDS), and the two-dimensional solution compares favorably to the real layout. Thus, whereas separate judgments about the distances among objects in a cluttered environment may vary, even widely, when taken together, they constrain each other in a manner well captured by NMDS procedures as they converge on a near metric solution (see, e.g., Baird & Wagner, 1983). A second, more direct tradition concerns psychophysical judgments of distance. The power function for such judgments by stationary observers is well described by an exponent with the value just under 1.0 (Baird & Biersdorf, 1967; Cook, 1978; DaSilva, 1985; Purdy & Gibson, 1955; Teghtsoonian & Teghtsoonian, 1969; Wagner, 1985), and for moving observers even this slight compression of depth often ceases to exist (Loomis, DaSilva, Fujita, & Fukushima, 1992; Rieser, Ashmead, Talor, & Youngquist, 1990).

In summary, then, it appears that personal space is the closest of the three spaces to being Euclidean (Durgin, Proffitt, Olson, & Reinke, 1995); action space and vista space have a more clearly affine character (Loomis, DaSilva, Philbeck, & Fukushima, 1996; Wagner, 1985). When the observer is allowed to move and act within the environment, however, action space also appears to approach a Euclidean nature.

Pictures as Situations of Information Conflict

In Figure 1, one finds rationale for, and food for thought about, certain well-known results in situations of "cue-conflict"—when two or more sources of information are presented to an observer, yielding discrepant information about depth. Within this framework, I suggest that, except for constraints of viewing pictures and cinema from the side (Busey, Brady, & Cutting, 1990; Cutting, 1987, 1988), there is nothing special about picture perception as compared with the perception of natural scenes, except that pictures are particular situations in which such conflicts are typically inherent.

Consider this conflict in three ways, with respect to pictures, cinema, and virtual reality. First, the cue conflict in viewing hand-held photographs is typically never resolved. Pictures have a dual aspect, and the existing information is used to support that duality; disparities, convergence, and accommodation tell observers they are holding a flat, small object, and the pictorial information tells them that a more extensive layout is represented. Each set of information sources thus serves as a channel for one of the two aspects of the picture, without resolution being necessary. Second, in the evolution and development of cinema, it seems that such conflicts were sought to be avoided. Today, as at the turn of the century, the viewer is often allowed to sit at distances greater than 15 and even 30 m. Such distances yield effects that are virtually no different perceptually than viewing a natural scene, except that it is delimited by the rectangular frame of the screen, and choices of lenses and shooting distances may make objects loom much larger that they would ordinarily. Third, it should be clear that the technical goal of virtual reality systems, if they are to mimic everyday perception, must be to try to remove *all* potential information conflicts. Any conflicts will compromise the generalization from perception within virtual reality to natural perception; but, as we know well from the history of art, photography, and cinema, the human visual system performs remarkably well, if not identically, despite such conflicts.

Thus concludes an outline of perceptual spaces and multiple sources of information. For further information, see Cutting and Vishton (1995). Let us consider, next, two caveats for virtual reality from the perspective of this approach, and, then, a summary.

SOME CAVEATS FOR VIRTUAL REALITY SYSTEMS

Expect Individual Differences

The first caveat stems from evolutionary considerations. The perception of layout is supported by at least nine sources of information—more than any other single property through any single modality in all of perception. Why? The answer must be, in part, that the perception of layout is extremely important to us and to other species—so important that it cannot be trusted to a single information source or to a small cluster of sources. In this manner, given the failure of any given source, others are likely to remain trustworthy, and the perception of layout will remain relatively stable (Todd, 1985). But this multiplicity has other potential consequences. Given that, in most environments, most information sources exist, individuals can, over the course of their own personal histories, come to rely on some sources more than others. The implication of this is that we should predict individual differences in information source use, and we are likely to discover considerable individual differences in the utility of virtual reality systems as a function of residual source conflict and of the particular environments simulated.

Expect Perceptual Effects of Lenses and Viewing Apertures

The second caveat stems from the generally affine character of the spaces around an observer. Most of us understand that, in photography, different lenses are used for various effects. Most of us, however, do not fully understand all of those effects. Consider three lenses for a 35-mm camera—a standard lens (typically 50 mm), a longer lens of 500 mm which magnifies the central field of view, and a shorter lens of 25 mm which minifies it. If pictures were taken of the same scene with these three lenses, developed, and printed at the same size, there would be other effects (see, e.g., Swedlund, 1974). Compared with the picture taken with the standard lens, that taken with the long lens will have depth (distance along the *z*-axis extended away from the picture plane) compressed by a factor of 10, or one log unit. That taken with the short lens will have depth expanded by a factor of two, or about one-third log unit (see Cutting, 1986, 1988; Farber & Rosinski, 1978; Lumsden, 1980). That we tend not to notice these particular spatial effects is a tribute to the generally affine character of perceived space; what we do notice are other effects such as the change in sizes of objects and the fact that portraits of people taken with standard and short lenses tend to give them large noses.

The compressions and dilations of space resulting from the use of long and short lenses, respectively, will create changes in the patterns of some, but not all, of the functions shown in Figure 1. On the one hand, the flat functions for occlusion, size, and density will remain constant, because the differences in objects perceived at different distances survive magnifications and minifications of distance very well and because the functions are flat when plotted on logarithmic coordinates. On the other hand, the nonflat functions will change. Consider the functions declining with distance—height in the visual field, motion perspective, and binocular disparity. With a 500-mm lens and the compression of depth by one log unit, these functions would effectively move rightward by one log unit.[4] With a 25-mm lens, they would move leftward about a third of a log unit. Such effects may not be overtly noticed by the perceiver, but they may play a role in his/her perception and judgments about the depicted space and may further be a function of the individual differences noted above.

These lens effects apply to computer graphics as well. After setting parameters for perspective viewing, the researcher typically has program control over the width of the viewing window onto the world that is simulated. Typically this might be 30º. If one sets this value to be 60º, one has effectively changed to a shorter lens length and all the objects in the scene become smaller; if one sets this value to 15º, one has effectively changed to a longer lens length and magnified the objects in the scene. Movements through or around this scene will similarly affect the motions and shapes that are generated. Particularly striking and disruptive are the projective changes in the simulated world's parallel lines during pans with a wide-angle viewing aperture. These lines can often be seen to hinge against one another, creating an effect of perceived nonrigidity. In this and other manners, the adjustment of a few parameters may have marked effects on how people view simulated scenes.

CONCLUSION

Despite appearances, our perception of the cluttered space around us is not homogenous. Instead, through the differential use of at least nine different sources of information, I suggest that, for a pedestrian, at least three different egocentric spaces grade into one another—personal space (out to about 1.5 m), action space (from 1.5 to about 30 m), and vista space (beyond about 30 m). Personal space seems nearly Euclidean in nature, but the latter two seem largely affine, and subject to compressions and dilations along the depth axis. Travel in modern conveyances, variations in environments, and differences among people will alter the functions of the different sources and consequently alter these spaces, as will the uses of different lenses in photography and different viewing ports in computer graphics.

REFERENCES

Anderson, B. L., & Nakayama, K. (1994). Towards a general theory of stereopsis: Binocular matching, occluding contours and fusion. *Psychological Review*, **101**, 414-445.

Arditi, A. (1986). Binocular vision. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.) *Handbook of perception and human performance* (Vol. 1, pp. 23:1 to 23:41). New York: Wiley.

Baird, J. C., & Biersdorf, W. R. (1967). Quantitative functions for size and distance. *Perception & Psychophysics*, **2**, 161-166.

Baird, J. C., & Wagner, M. (1983). Modeling the creation of cognitive maps. In H. L. Pick, Jr., & L. P. Acredolo (Eds.), *Spatial orientation: Theory, research, and application* (pp. 321-366). New York: Plenum.

Barlow, H. B. (1978). The efficiency of detecting changes of density in random dot patterns. *Vision Research*, **18**, 637-650.

Battro, A. M., Netto, S. P., & Rozestraten, R. J. A. (1983). Depth perception of surfaces of variable curvature in visual space: Looking for conventions in Pandora's box. *Perception*, **5**, 9-23.

Bell, J. C. (1993). Zaccolini's theory of color perspective. *Art Bulletin*, **75**, 91-112.

Biederman, C. (1948). *Art as the evolution of visual knowledge*. Red Wing, MN: Charles Biederman.

Bingham, G. (1993). Form as information about scale: Perceiving the size of trees. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 1139-1161.

Birnbaum, M. H. (1983). Scale convergence as a principle for the study of perception. In H. Geissler (Ed.), *Modern issues in perception* (pp. 319-335). Amsterdam: North-Holland.

Blank, A. A. (1978). Metric geometry in human binocular vision: Theory and fact. In E. E. J. Leeuwenberg & H. F. Buffart (Eds.), *Formal theories of visual perception* (pp. 82-102). Chichester, U.K.: Wiley.

Blatt, S. J. (1984). *Continuity and change in art*. Hillsdale, NJ: Erlbaum.

Boring, E. G. (1942). *Sensation and perception in the history of experimental psychology*. New York: Appleton-Century-Crofts.

Braunstein, M. L. (1962). Depth perception in rotating dot patterns: Effects of numerosity and perspective. *Journal of Experimental Psychology*, **64**, 415-420.

Braunstein, M. L. (1976). *Depth perception through motion*. New York: Academic Press.

Bülthoff, H. H., & Mallot, H. A. (1988). Interpreting depth modules: Stereo and shading. *Journal of the Optical Society of America A*, **5**, 1749-1758.

Burton, H. (1945). The optics of Euclid. *Journal of the Optical Society of America*, **35**, 357-372.

Busey, T. A., Brady, N. P., & Cutting, J. E. (1990). Compensation is unnecessary for the perception of faces in slanted pictures. *Perception & Psychophysics*, **48**, 1-11.

Cavanagh, P., & Leclerc, Y. G. (1990). Shape from shadows. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 3-27.

Chapanis, A., & McCleary, R. A. (1955). Interposition as a cue for the perception of relative distance. *American Journal of Psychology*, **48**, 113-132.

Chauvet, J.-M., Brunel Deschamps, E., & Hillaire, C. (1995). *La grotte Chauvet à Vallon-Pont-d'Arc*. Paris: Seuil.

Clottes, J. (1995). *Les cavernes di Niaux*. Paris: Seuil.

Clottes, J. (1996, January). Le grotte ornate del paleolitico. *Le Scienze: edizione italiana di Scientific American*, **329**(1), 62-68.

Cole, A. (1992). *Perspective*. London: Dorling Kindersley.

Cook, M. (1978). Judgment of distance on a plane surface. *Perception & Psychophysics*, **23**, 85-90.

Cutting, J. E. (1986). *Perception with an eye for motion*. Cambridge, MA: MIT Press, Bradford Books.

Cutting, J. E. (1987). Rigidity in cinema seen from the front row, side aisle. *Journal of Experimental Psychology: Human Perception & Performance*, **13**, 323-334.

Cutting, J. E. (1988). Affine distortions in pictorial space: Some predictions for Goldstein (1987) that La Gournerie (1959) might have made. *Journal of Experimental Psychology: Human Perception & Performance*, **14**, 305-311.

Cutting, J. E. (1991). Four ways to reject directed perception. *Ecological Psychology*, **3**, 25-34.

Cutting, J. E., & Millard, R. M. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General*, **113**, 198-216.

Cutting, J. E., Springer, K., Braren, P., & Johnson, S. (1992). Wayfinding on foot from information in retinal, not optical, flow. *Journal of Experimental Psychology: General*, **121**, 41-72.

Cutting, J. E., & Vishton, P. M. (1995). Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth. In W. Epstein & S. Rogers (Eds.), *Perception of space and motion* (pp. 69-117). San Diego: Academic Press.

Daniels, N. (1974). *Thomas Reid's inquiry*. New York: Burt Franklin.

DaSilva, J. A. (1985). Scales of perceived egocentric distance in a large open field: Comparison of three psychophysical methods. *American Journal of Psychology*, **98**, 119-144.

Dosher, B. A., Sperling, G., & Wurst, S. A. (1986). Tradeoffs between stereopsis and proximity luminance covariation as determinants of perceived 3D structure. *Vision Research*, **26**, 973-990.

Dunn, B. E., Gray, G. C., & Thompson, D. (1965). Relative height on the picture-plane and depth perception. *Perceptual & Motor Skills*, **21**, 227-236.

Durgin, F. H. (1995). Texture density adaptation and the perceived numerosity and distribution of texture. *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 149-169.

Durgin, F. H., Proffitt, D. R., Olson, T. J., & Reinke, K. S. (1995). Comparing depth from motion with depth from binocular disparity. *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 679-699.

Duwaer, A. L., & van den Brink, G. (1981). What is the diplopia threshold? *Perception & Psychophysics*, **29**, 295-309.

Epstein, W. (1963). The influence of assumed size on apparent distance. *American Journal of Psychology*, **76**, 257-265.

Farber, J., & Rosinski, R. R. (1978). Geometric transformations of pictured space. *Perception*, **7**, 269-282.

Ferris, S. H. (1972). Motion parallax and absolute distance. *Journal of Experimental Psychology*, **95**, 258-263.

Fisher, S. K., & Ciuffreda, K. J. (1988). Accommodation and apparent distance. *Perception*, **17**, 609-612.

Flock, H. R. (1964). A possible optical basis for monocular slant perception. *Psychological Review*, **71**, 380-391.

Foley, J. M. (1991). Stereoscopic distance perception. In S. R. Ellis, M. K. Kaiser, & A. C. Grunwald (Eds.), *Pictorial communication in virtual and real environments* (pp. 559-566). London: Taylor & Francis.

Freeman, R. B. (1966). Absolute threshold for visual slant: The effect of size and retinal perspective. *Journal of Experimental Psychology*, **71**, 170-176.

Gerbino, W., Stultiens, C. J., & Troost, J. M. (1990). Transparent layer constancy. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 3-20.

Gibson, J. J. (1950). *Perception of the visual world*. Boston: Houghton Mifflin.

Gogel, W. C. (1961). Convergence as a cue to absolute distance. *Journal of Psychology*, **52**, 287-301.

Gogel, W. C., & Tietz, J. D. (1973). Absolute motion parallax and the specific distance tendency. *Perception & Psychophysics*, **13**, 284-292.

Gombrich, E. H. (1995). *Shadows: The depiction of cast shadows in Western art*. New Haven, CT: Yale University Press.

Graham, C. H., Baker, K. E., Hecht, M., & Lloyd, V. V. (1948). Factors influencing thresholds for monocular movement parallax. *Journal of Experimental Psychology*, **38**, 205-223.

Grünbaum, A. (1973). *Philosophical problems of space and time* (2nd ed.). Boston: Reidel.

Gulick, W. L., & Lawson, R. B. (1976). *Human stereopsis*. New York: Oxford University Press.

Hagen, M. A. (1986). *Varieties of realism: Geometries of representational art*. Cambridge: Cambridge University Press.

Hall, E. T. (1966). *The hidden dimension*. New York: Doubleday.

Helmholtz, H. von (1925). *Physiological optics* (3rd ed., Vol. 3; J. P. C. Southall, Trans.). Menasha, WI: The Optical Society of America. (Original work published 1867)

Hobbs, J. A. (1991). *Art in context* (4th ed.). Fort Worth, TX: Harcourt Brace & Jovanovich.

Hochberg, J. (1971). *Perception*. In J. W. Kling & L. A. Riggs (Eds.), *Handbook of experimental psychology* (3rd ed., pp. 396-550). New York: Holt, Rinehart & Winston.

Hofsten, C. von (1976). The role of convergence in space perception. *Vision Research*, **16**, 193-198.

Indow, T. (1990). A critical review of Luneburg's model with regard to global structure of visual space. *Psychological Review*, **98**, 430-453.

Johansson, G. (1973). Monocular movement parallax and near-space perception. *Perception*, **2**, 136-145.

Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.

Kaplan, G. A. (1969). Kinetic disruption of optical texture: The perception of depth at an edge. *Perception & Psychophysics*, **6**, 193-198.

Kersten, D., & Legge, G. (1983). Convergence accommodation. *Journal of the Optical Society of America*, **73**, 322-388.

Kosslyn, S. M., Pick, H. L., & Fariello, G. R. (1974). Cognitive maps in children and men. *Child Development*, **45**, 707-716.

Kubovy, M. (1986). *The psychology of perspective and renaissance art*. Cambridge: Cambridge University Press.

Künnapas, T. (1968). Distance perception as a function of available visual cues. *Journal of Experimental Psychology*, **77**, 523-529.

Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination. *Vision Research*, **35**, 389-412.

Landy, M. S., Maloney, L. T., & Young, M. J. (1991). Psychophysical estimation of the human depth combination rule. In P. S. Shenker (Ed.), *Sensor fusion III: 3-D perception and recognition* (Proceedings of the SPIE, Vol. 1383, pp. 247-254).

Leibowitz, H. W., Shina, K., & Hennessy, R. T. (1972). Oculomotor adjustments and size constancy. *Perception & Psychophysics*, **12**, 497-500.

Lie, I. (1965). Convergence as a cue to perceived size and distance. *Scandinavian Journal of Psychology*, **6**, 109-116.

Loomis, J. M., DaSilva, J. A., Fujita, N., & Fukushima, S. S. (1992). Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 906-921.

Loomis, J. M., DaSilva, J. A., Philbeck, J. W., & Fukushima, S. S. (1996). Visual perception of location and distance. *Current Directions in Psychological Science*, **5**, 72-77.

Lumsden, E. A. (1980). Problems of magnification and minification: An explanation of the distortions of distance, slant, and velocity. In M. Hagen (Ed.), *The perception of pictures* (Vol. 1, pp. 91-135). New York: Academic Press.

Luneburg, R. K. (1947). *Mathematical analysis of binocular vision*. Princeton, NJ: Princeton University Press.

Lythgoe, J. N. (1979). *The ecology of vision*. Oxford: Oxford University Press.

Marr, D. (1981). *Vision*. San Francisco: Freeman.

Metelli, F. (1974). The perception of transparency. *Scientific American*, **230**(4), 90-98.

Minnaert, M. (1993). *Light and color outdoors*. New York: Springer-Verlag.

Morgan, M. W. (1968). Accommodation and convergence. *American Journal of Optometry & Archives of American Academy of Optometry*, **45**, 417-454.

Nagata, S. (1991). How to reinforce perception of depth in single two-dimensional pictures. In S. R. Ellis, M. K. Kaiser, & A. C. Grunwald (Eds.), *Pictorial communication in virtual and real environments* (pp. 527-545). London: Taylor & Francis.

Norman, J. F., Todd, J. T., & Phillips, F. (1995). The perception of surface orientation from multiple sources of optical information. *Perception & Psychophysics*, **57**, 629-636.

Nougier, L. R. (1969). *Art préhistorique*. Paris: Librarie Générale Française.

Ogle, K. O. (1952). On the limits of stereoscopic vision. *Journal of Experimental Psychology*, **48**, 50-60.

Ogle, K. O. (1958). Note on stereoscopic acuity and observation distance. *Journal of the Optical Society of America*, **48**, 794-798.

Pirenne, M. H. (1970). *Optics, painting, & photography*. Cambridge: Cambridge University Press.

Proffitt, D. R., Rock, I., Hecht, H., & Schubert, J. (1992). Stereokinetic effect and its relation to the kinetic depth effect. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 3-21.

Purdy, J., & Gibson, E. J. (1955). Distance judgments by the method of fractionation. *Journal of Experimental Psychology*, **50**, 374-380.

Ratoosh, P. (1949). On interposition as cue for the perception of distance. *Proceedings of the National Academy of Sciences*, **35**, 257-259.

Richter, J. P. (1970). *The notebooks of Leonardo da Vinci*. New York: Dover Press. (Original translation published 1883)

Rieser, J. J., Ashmead, D. H., Talor, C. R., & Youngquist, G. A. (1990). Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception*, **19**, 675-689.

Rogers, B. J., & Graham, M. (1979). Motion parallax as an independent cue for depth. *Perception*, **8**, 125-134.
Roscoe, S. N. (1984). Judgments of size and distance with imaging displays. *Human Factors*, **27**, 615-636.
Ruspoli, M. (1986). *Lascaux: Un nouveau regard*. Paris: Bordas.
Scharf, A. (1968). *Art and photography*. London: Penguin.
Sommer, R. (1969). *Personal space*. Englewood Cliffs, NJ: Prentice Hall.
Suppes, P. (1977). Is visual space Euclidean? *Synthèse*, **35**, 397-421.
Swedlund, C. (1974). *Photography: A handbook of history, materials, and processes*. New York: Holt, Rinehart & Winston.
Teghtsoonian, R., & Teghtsoonian, M. (1970). Scaling apparent distance in a natural outdoor setting. *Psychonomic Science*, **21**, 215-216.
Teichner, W. H., Kobrick, J. L., & Wehrkamp, R. F. (1955). The effects of terrain and observation distance on relative depth perception. *American Journal of Psychology*, **68**, 193-208.
Tittle, J. S., Todd, J. T., Perotti, V. J., & Norman, J. F. (1995). The systematic distortion of perceived 3-D structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 663-677.
Todd, J. T. (1985). Perception of structure from motion: Is projective correspondence of moving elements a necessary condition? *Journal of Experimental Psychology: Human Perception & Performance*, **11**, 689-710.
Todd, J. T., & Akerstrom, R. A. (1987). Perception of three-dimensional form from patterns of optical texture. *Journal of Experimental Psychology: Human Perception & Performance*, **13**, 242-255.
Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, **96**, 643-657.
Toulet, E. (1988). *Cinématographie: Invention du siècle*. Paris: Découvertes Gallimard.
Toye, R. C. (1986). The effect of viewing position on the perceived layout of space. *Perception & Psychophysics*, **40**, 85-92.
van den Berg, A.V., & Brenner, E. (1994). Why two eyes are better than one for judgments of heading. *Nature*, **371**, 700-702.
Wagner, M. (1985). The metric of visual space. *Perception & Psychophysics*, **38**, 483-495.
Wallach, H., & Karsh, E. B. (1963). Why the modification of stereoscopic depth-perception is so rapid. *American Journal of Psychology*, **76**, 413-420.
Wallach, H., & Norris, C. M. (1963). Accommodation as a distance cue. *American Journal of Psychology*, **76**, 659-664.
Wallach, H., & O'Connell, D. N. (1953). The kinetic depth effect. *Journal of Experimental Psychology,* **45**, 205-217.
Wanger, L. R., Ferwerda, J. A., & Greenberg, D. P. (1992, May). Perceiving the spatial relationships in computer-generated images. *IEEE Computer Graphics*, **12**, 44-59.
Watson, J. S., Banks, M. S., Hofsten, C. von, & Royden, C. S. (1992). Gravity as a monocular cue for perception of absolute distance and/or absolute size. *Perception*, **12**, 259-266.
Wheatstone, C. (1838). Contributions to the physiology of vision: I. On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London*, **128**, 371-394.
Wright, L. (1983). *Perspective in perspective*. London: Routledge & Kegan Paul.
Yonas, A., Craton, L. G., & Thompson, W. B. (1987). Relative motion: Kinetic information for the order of depth at an edge. *Perception & Psychophysics*, **41**, 53-59.
Yonas, A., & Granrud, C. E. (1985). The development of sensitivity to kinetic, binocular, and pictorial depth information in human infants. In D. Ingle & D. N. Lee (Eds.), *Brain mechanisms and spatial vision* (pp. 113-145). Dordrecht: Martinus Nijhoff.
Zegers, R. T. (1948). Monocular movement parallax thresholds as functions of field size, field position, and speed of stimulus movement. *Journal of Psychology*, **26**, 477-498.

NOTES

1. There is some controversy over the dating of the Chauvet paintings. An early report of carbon dating in the press stated that they were 30,000 years old ("Les peintures de la grotte Chauvet datent de 30 000 ans avant nôtre ère," *Le Monde*, Juin 4/5, 1995). Later, however, Clottes (1996) reported that they were 20,000 years old. Regardless of which dating is correct, the Chauvet paintings are the oldest large collection of paleolithic art yet found.

2. This approach contrasts with others in the literature. For example, Landy, Maloney, Johnston, and Young (1995) start by considering the sources that imply the strongest measurement scales, and use them to modify those with weaker scales.

3. Relative size should not be confused with familiar size (see, e.g., Epstein, 1963), which relies on knowledge of the observer; relative size assumes only the presence of many similarly shaped and sized objects.

4. The disparity function moves because depth is compressed while the distance between the eyes remains the same. This effectively makes disparities useful at greater depth in the scene. It could also serve to make things look a bit smaller, as discussed in connection with the stereograms of sequoias.