CHAPTER 1

VOLUMETRIC SYSTEMS AND THE PROCESS OF VISUALIZATION

We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time*.

1.1 INTRODUCTION

Volumetric displays have been the subject of almost continual research activity for more than 50 years; during this time numerous systems have been proposed and some have been implemented in prototype form (Blundell and Schwarz 1995; Blundell et al 1993a).† Although simple in concept, volumetric displays have in general turned out to be complex when constructed, often taking the form of large systems of experimental apparatus surrounded by bundles of cables, optical components, and electromechanical devices operating in semidarkness. The construction of each prototype has generally represented a major undertaking in terms of the time devoted to the production and integration of mechanical, electronic and software subsystems.

The fascinating history of the research undertaken on volumetric display systems is outlined, perhaps for the first time, in this book. While reviewing historical aspects of volumetric display development we learn from the technical progress (and mistakes) of others and may also begin to understand why volumetric systems have failed to progress from the experimental prototype phase

* T.S. Eliot, "Little Gidding," in *The Four Quartets*, Harcourt Brace, NY, 1974.

† A very small number of systems have been developed as commercial products mainly for evaluation purposes.

to that of significant commercial realization. It is, after all, rather surprising that despite such a protracted period of research, the immense effort directed toward their development and the relative simplicity and availability of the technologies involved in their implementation, there has not yet been any major breakthrough in their commercialization.

Unfortunately, scientists and engineers working on the development of volumetric systems have tended to work in isolation. This has led to a major duplication of effort and a failure to adopt standardized terminology. This situation is reflected in scientific publications and patents where even the most basic components in display systems are often referred to by different names. Despite the extensive period during which work on this area has been carried out, we are forced to conclude that this area of research has not yet been adequately formalized and that past research has largely been conducted on an ad hoc basis. This situation may have occurred as a consequence of the perceived commercial potential of the technologies under study and a natural wish on the part of researchers to protect their commercial interests.

In this work we intend to review historical and current aspects of volumetric display system research, and perhaps more importantly, to establish standardized terminology and formalism which we hope will form a foundation for those studying this subject in the future. By means of these foundations, it is possible, in later chapters, to perceive with greater clarity the potential of volumetric displays as an aid to visualization and anticipate the next generation of high-definition volumetric systems.

In this chapter we introduce the most elementary principles of volumetric display system techniques and present a simple display model that will be used during the course of the book for the purpose of reducing some of the more abstract concepts to a simple form. Section 1.4 provides a brief description of the processes through which we perceive the three-dimensional (3-D) nature of our surroundings. Since many excellent texts exist concerning this fascinating subject, our discussions will, in general, be limited to matters that are especially relevant to volumetric visualization.

Each type of 3-D display technique is able to depict image scenes in such a way as to satisfy a number of the physiological and psychological depth cues. The inability of display techniques to satisfy all the various depth cues in a natural manner coupled with the varying emphasis that we place on these cues and the variety of situations in which displays may be operated precludes the possibility of developing a single universal 3-D display system. Volumetric displays will therefore complement a range of visualization systems. Although some display techniques have gained more widespread popularity and acceptance, there can be no question of direct competition between them, as each has a unique role to play within the spectrum of our visualization needs. In Section 1.5 we briefly outline some other types of 3-D display technologies. Rather than attempting a comprehensive discussion on each of these types of system, we provide references to some of the many books and journal articles that describe them in detail.

The use of specialized terminology in this chapter has been minimized so as to allow those visiting this area of research for the first time to concentrate on general concepts. Expressions have therefore been chosen for their descriptive merit and are not necessarily the same as those introduced in subsequent chapters which are more precisely defined.

1.2 VOLUMETRIC DISPLAY

Common to all volumetric display systems is a volume or region occupying three physical dimensions within which image components may be positioned, depicted, and perhaps manipulated. Some examples of simple volumetric images are illustrated in the set of color plates included in the book. Since these images are generated within a region of 3-D space, and may be observed directly therein, a number of depth cues (see Section 1.4) are automatically satisfied and their three-dimensionality is naturally perceived. The volume may therefore be considered to provide a 3-D tableau and in contrast to many other 3-D display techniques, the sensation of depth occurs automatically and is an intrinsic property of volumetric image depiction. Depth cue conflict may therefore be avoided and this denotes a considerable benefit of this type of display methodology. In principle, the computational overheads are small in comparison to other approaches, furthermore, images displayed within the volume may be observed naturally without the need for an operator to wear special glasses.

A number of the various volumetric display technologies that have been the subject of past research or which are currently under investigation impose very little restriction in the viewing angle. An operator may therefore move around the volume and, in the case of some volumetric architectures, view an image scene from practically any orientation. For those visiting this area of research for the first time, the concept and characteristics of the display volume may be appreciated by watching fish moving in an aquarium. The fish (which we may think of as image objects) can position themselves at any location within a clearly defined volume. As a consequence, the image scene is automatically 3-D and the spatial separation of the fish is readily apparent.

It is interesting to note (and will be a matter for further discussion later in this work) that the physical dimensions of the volume in which the fish move (as measured) differ from the apparent volume (as observed) due to the refraction of the light as it passes from the water and containing glass vessel to the surroundings. So although we consider that a volumetric display must, by definition, employ a physical volume within which images may be constructed, the apparent dimensions of this volume may not correspond to its physical measurements. Should we generate images within a volume whose refractive index differs from that of the surroundings, we may anticipate a bending of the light as it crosses the boundary of the volume. The shape (and curvature) of the volume will then become a matter for important consideration. In the case of a spherical volume with a high refractive index, the lensing effect may cause considerable image distortion, and it is suggested that the reader examine the consequences for volumetric imaging by observing fish swimming within a spherical bowl. As we will see, optical effects at the volume boundary may, for certain applications, be advantageous and for others quite the opposite.

In the case of all volumetric display systems known to the authors, the generation of images occurs within a containing vessel from which the observer(s) is excluded. Volumetric systems therefore provide a "God's-eye" view of any image scene. However, it is possible to envisage the potential of a display device able to project volumetric images into a region of space through which one or more observers are free to pass. The development of such an immersive environment may one day completely change the nature of volumetric display systems.

Finally we believe it to be appropriate at this point to provide a definition that characterizes the essential nature of any volumetric display system and which may therefore be used to differentiate between this type of display and other 3-D visualization techniques:

> *A volumetric display device permits the generation, absorption, or scattering of visible radiation from a set of localized and specified regions within a physical volume.*

1.2.1 Display System Classification

For many years, systems able to provide a 3-D tableau within which images may be constructed have been referred to as *volumetric displays*. The origins of this term are unclear, and it is perhaps unfortunate that the word *volumetric* was adopted in the first instance for the classification of these display devices. In the *Oxford English Dictionary** we find *volumetric* defined as follows:

volumetric: adj. Of, pertaining to, or noting measurement by volume.

Volumetric therefore clearly implies measurement by means of spatial information. Since all types of physical measurement lead to some variety of quantitative information, the term *volumetric display* may be considered to refer to display apparatus that permits quantitative measurement of visual information. This is misleading, as volumetric display systems are best suited to the qualitative visualization of spatial information and, as we will see in Section 8.2, are least suited to the depiction or provision of quantitative information.

Volumetric display systems have also been referred to as *multiplanar display systems* (Williams and Garcia 1988, Bains 1993) and *direct volume display devices* (DVDDs) (Clifton and Wefer 1993a, 1993b). The latter is perhaps a more appropriate title. However, it has not been adopted for use in this book, for two reasons: First, the use of this expression is not yet widespread, and second (and perhaps more importantly), the word *volumetric* is now firmly associated

* The *Oxford English Dictionary*, 2nd edition, Volume XIX, prepared by J.A. Simpson and E.S.C. Weiner (Oxford, Clarendon Press, 1989).

with certain types of image data sets. Although such data may be depicted on all types of display systems, in the case of volumetric displays they form an essential source of image information. Consequently, the common use of the word *volumetric* in the description of image data has led to our continued adoption of this word as a means of classifying the display systems that are the primary subject of this book.

1.2.2 Volumetric Display System Family

When we examine various display system implementations, it is generally very easy to distinguish between those that are volumetric and those that fall outside this classification. However, for a small number of display system implementations, classification is more difficult. For example, varifocal mirror technologies employ a reflective surface, the curvature of which is rapidly varied. The surface therefore acts as a mirror of varying focal length. By means of an appropriate control system, portions of a 3-D image projected thereon appear to originate from different depths. The image thus occupies a virtual volume whose dimensions exceed the amplitude of the surface's motion. By the definition provided earlier in this section, such displays fall outside the scope of volumetric displays as considered in this book (since they employ a virtual rather than a physical display volume). Varifocal mirror techniques are discussed in a number of excellent publications [e.g., Traub 1967; Harris et al. (1986); U.S. patent 4,130,832; Sher (1993) and references therein]. The matter of their proper classification is left to future debate.

1.3 AN ELEMENTARY SYSTEM MODEL

Many issues relating to the implementation of a volumetric display may best be illustrated by reference to a simple example. In this section we therefore introduce an elementary model of a display device within which volumetric images may be generated. This model will be referred to during the course of the book and has been selected to reduce the underlying concepts and properties of an image depiction volume to their most basic and therefore the most comprehensible level. In the form described here, the model has a number of both obvious and subtle deficiencies and is therefore not intended to represent a feasible volumetric implementation. In subsequent sections we highlight some of the deficiencies associated with this model.

Graphical images are typically constructed from a set of image elements. In the case of conventional cathode ray tube (CRT)–based displays or those employing a liquid crystal display (LCD) panel, images are formed by the activation of picture elements (pixels) at appropriate locations within 2-D space. Each of these elements has a unique location (x, y) within a two-dimensional plane and a color and gray-scale value may be assigned to each. Volumetric display systems employ a similar method for image construction; elements [referred to as voxels (see Section 2.3)] are activated at the appropriate locations within the 3-D space contained within the volume. As is the case for conventional displays, each voxel has a unique position. However, unlike conventional display devices, this is defined within a 3-D volume (x, y, z). In common with conventional display systems, each voxel may be assigned a color and gray-scale value and perhaps for future display systems it will also be possible to assign to each voxel an opacity descriptor (see Sections 7.5.5 and 10.2.3).

Any volumetric display system must provide a mechanism for the generation of voxels within a volume, and in the case of our elementary model we employ for this purpose a simple 3-D array of light bulbs (filaments). The volume will be assumed to take the form of a cube with sides of length L containing N bulbs. Since each filament will be responsible for the generation of a single voxel, the display will be capable of generating up to a total of N spatially separated voxels. This example display model is illustrated in Figure 1.1. In view of the simple technique employed for the generation of each voxel, our elementary display will clearly support the depiction of only single color images. Controlling the current passing through each of the bulbs may accommodate gray scale.

Clearly, should we wish to generate small voxels, the display device will need to employ a large number of bulbs (and electrical connections) to achieve a useful display volume and so be able to depict images of an appropriate size. Furthermore, we can see that voxels may only be generated at unique locations within the volume, and as a consequence aliasing (Foley et al. 1990, Hill 1990) will occur as we attempt to generate lines that are neither horizontal nor vertical, nor which form a diagonal across the cubic image space.

Also worth consideration is the fraction of the total number of available bulbs that we may wish to activate simultaneously. Clearly, as more bulbs (and hence

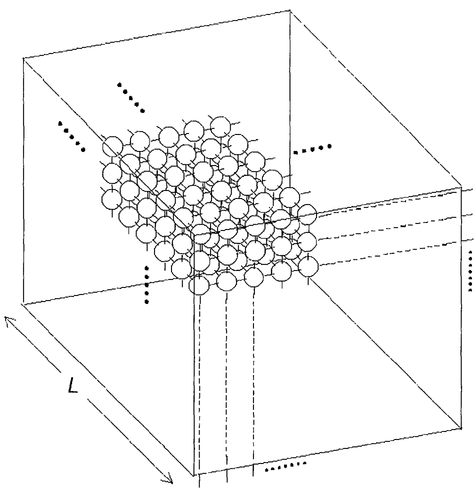

FIGURE 1.1 Elementary model volumetric display.The display volume has sides of length L and contains N regularly spaced lamps (filaments). Since this type of display does not incorporate any mechanical motion of components within the volume, it is classified as a static-volume display device (see Sections 3.1 and 6.2).

voxels) are illuminated, the display volume will become increasingly cluttered and it will become increasingly difficult to discern the spatial separation and form of the translucent image components contained therein. As will become clear in subsequent chapters, the number of bulbs illuminated at any one time should represent only a small fraction of the total number of available bulbs. This may be demonstrated by reference to our previous comparison between a display volume and tank containing fish. As we add increasing numbers of fish to the tank, their spatial separation becomes less apparent and some fish will become obscured by others that are closer to the observer and cross the observer's line of sight. If we continue adding fish, those closest to the observer will eventually obscure completely the motion of those located farther within the tank. In a similar way, if we increase the fraction of activated voxels in a volumetric display device, we will eventually reduce the benefits associated with the display volume's third dimension.*

1.4 PERCEPTION AND THREE-DIMENSIONAL SPACE

Human perception of the 3-D world is an extremely complex process that continues to be the subject of research. To depict images that may be viewed in a natural manner, electronic display devices take advantage of, and are perhaps limited by, the way in which our visual system perceives its surroundings. For example, the required brightness and speed at which images need to be updated for the realistic illusion of motion are both determined by parameters associated with the visual system. A study of 3-D display techniques must clearly encompass a knowledge of *depth cues*, the visual-based mechanisms through which we obtain depth information about the surrounding environment. An understanding of depth cues and the human visual system should form an essential starting point for all working on and developing display system technologies. Many excellent texts have been written on this fascinating subject. In this section we discuss some of the visual characteristics and depth cues that will be of importance in our subsequent discussions. For further information we refer the reader to some of the comprehensive works published in this area (Kaufman 1974, Dember and Warm 1979, Schiffman 1990, Wade and Swanston 1991, Yeh 1993).

1.4.1 Visual Acuity

Visual acuity is a measure of the level of detail the eye can distinguish. Most electronic displays present their images as a 2-D array of pixels, and volumetric systems do so as a 3-D distribution of voxels. In order that images appear sufficiently smooth, the required size and separation of these image elements, is dictated by the visual acuity of the eye in combination with the probable viewing

* This analogy is not quite accurate since the majority of volumetric display technologies give rise to translucent rather than opaque images. In this case considerable reliance is placed on the accommodation depth cue.

distance.* Moreover, the smallest size of image detail that can reliably be detected by an observer is also important when considering the likely or intended application(s) of a display system. This will govern the required spatial displacement of an image component relative to others in order that it may be observed reliably. It will also indicate the level of image detail that may be discerned.

The acuity, or maximum resolving power, of the eye's foveal region,† is generally considered to be a visual angle ϕ_c approximately equal to 1 minute of arc. However, this acuity is broad and very approximate, encompassing a number of acuities involving different visual tasks (Poole 1966, Schiffman 1990). The resolving power of the eye is generally different for each acuity subtype and varies widely from person to person. Five types of visual acuity are relevant to the study of display systems: minimum visible, minimum perceptible, minimum separable, vernier, and stereoscopic acuity. These are discussed by Poole (1966) and are outlined briefly below in the context of volumetric displays:

- *Stereoscopic acuity*. This describes the minimum difference in distance (from the observer) of two objects that may be discerned. It depends not only on the absolute depth of the objects from the observer, but also on their spatial separation perpendicular to the line of sight. Consider two objects located at slightly different distances and offset from each other perpendicular to the line of sight, as illustrated in Figure 1.2. Each eye alone will perceive a slightly different angular separation of the two objects, denoted β_1 and β_2 in this diagram. The minimum difference $\eta = |\beta_1 - \beta_2|$ between the angular separation seen by each eye between two such objects provides a measure of the stereoscopic acuity. This is, in effect, a quantification of the binocular parallax depth cue (discussed in Section 1.4.4.). The stereoscopic acuity is typically about 10 arc seconds but may vary between observers from 2 to 100 arc seconds. This measure is of considerable importance to 3-D display systems, as it indicates the magnitude of relative depth displacement that may be distinguished reliably. In the case of volumetric (and autostereoscopic) systems, the viewing distance may be variable, but the availability of other depth cues will aid in appreciation of the 3-D structure of the displayed data. However, applications requiring a rapid response from the observer may rely more critically on stereoscopic acuity.
- *Minimum visible acuity*. This refers to the minimum size of an object that can be seen when it is brighter than its background‡ and is a function of intensity rather than size; it is detected by the triggering of one (or more) cones in the retina. Interestingly, this means that a bright object against a dark background may be visible even though it subtends an angle smaller

* This is because the acuity depends on the feature size as registered on the retina—the acuity is expressed in terms of a visual angle.

† The region containing a high concentration of cones in the center of the retina.

‡ In the context of volumetric displays, this could correspond to a luminous voxel representing, for example, aircraft position in air traffic monitoring applications.

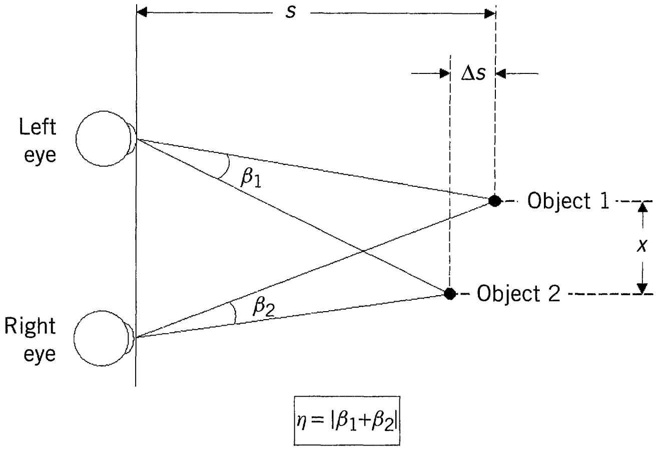

FIGURE 1.2 Stereoscopic acuity. Two objects located at distances s and $s - \Delta s$ from the observer are discerned as being at different depths if the difference between angles β_1 and β_2 subtended at each eye by the two objects is greater than the stereoscopic acuity. These angles depend on the absolute distance s, the relative distance Δs, and their separation x perpendicular to the line of sight. [After Poole (1966, p. 277) with permission.]

than the eye's minimum resolvable spot size (1 arc minute, determined by the size of the cones in the fovea) but it will *appear* to subtend 1 arc minute. This is because once a cone is triggered, the object appears as large as the cone.

- *Minimum perceptible acuity*. This is analogous to minimum visible acuity,* but for the case of a dark object against a bright background. This generally has a larger value than the minimum visible acuity, as a dark object will not be noticed if the surrounding bright region triggers all the cones in the corresponding area of the retina. The eye is thus better able to discern isolated bright points against dark backgrounds than vice versa.
- *Minimum separable acuity*. This corresponds to the ability of the eye to distinguish alphanumeric characters and is thus of considerable importance in the context of 2-D display systems. However, in the case of volumetric displays, the depiction of textual information may be considered at the present time to be of less importance, as discussed in Section 8.2. We refer the reader to Poole (1966, Chapter 15) for further discussion.
- *Vernier acuity*. This is the ability of the eye to distinguish an offset between two nearly colinear lines. The eye may distinguish such offsets at scales finer than 1 minute of arc; this is due to the random placement of the cones in the retina and the fact that a line will overlap to varying degrees a large number of cones.

---

* The minimum visible and minimum perceptible acuities may be considered together as minimum detectable acuity (Poole 1966).

An important consideration for volumetric displays is the role played by the background as well as the foreground in most acuity tasks. Recognition and interpretation of displayed information may thus be affected, in display units offering an all-round view, by objects and other people situated on the opposite side of the image space from the observer as well as by the ambient lighting conditions.*

Our ability to distinguish the physical dimensions of an object or detail depend on the distance from which it is viewed—the more distant an object, the smaller the angle it subtends at the eye. Consider an object of size x a distance s from the eye; as illustrated in Figure 1.3, the visual angle ϕ subtended by the object is given by 2 arctan $(x/2s)$. For small angles, standard approximations yield the dependence of the visual acuity on object size and distance as

$$\phi = \frac{x}{s} \tag{1.1}$$

Therefore, any detail of size x perpendicular to the line of sight can be resolved only while the viewing distance is less than $\approx x/\phi_c$.

Such considerations affect the design of electronic display systems, where the intended and possibly worst-case viewing positions may be considered. In the specific case of volumetric displays, the size, shape, and separation of individual voxels drawn within image space will determine to what extent, and for which viewing positions, the visual acuity limit is reached for a particular technique. This in turn may influence likely applications for the system. Although the discussion above has highlighted some of the issues governing the size and relative placement of image features, a sounder understanding of the precise implications of the various acuities as they apply to volumetric displays awaits empirical data from well-defined experimental trials.

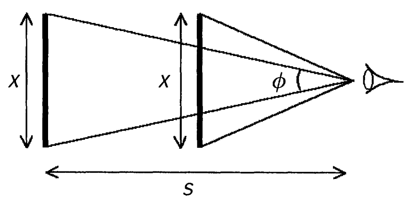

FIGURE 1.3 An illustration of Euclid's law (Burton 1945) in 2-D. An object of given size x subtends a smaller solid angle ϕ as its distance s from the observer increases.

* If the envelope surrounding the volume within which images are depicted is chosen so as to be semitransparent (rather than fully transparent), the image contrast relative to visual interference from background objects within the room may be improved. This is because external light crossing the volume must pass through this envelope twice and thus is attenuated by twice the amount compared to light originating from within the volume. However, the absolute image brightness is, of course, diminished.

1.4.2 Temporal Perception: Flicker

Most electronic display systems utilize transient luminescent phenomena for the creation of visible images. Images must therefore be refreshed at a frequency sufficient to ensure that the observer(s) perceives the visual information as continually present. For example, in the case of cathodoluminescence (von Ardenne 1939, Ozawa 1990), the light intensity emitted from an excited phosphorescent layer may decay substantially in a time period of perhaps 100 μs after the removal of the source of excitation. Any display employing this technique must therefore be refreshed at a rate sufficiently frequent that the continual variations in intensity are not apparent to the visual system. The minimum refresh frequency at which such an image appears to be continually present is known as the *critical flicker fusion frequency* (f_c). However, a precise value for this quantity is not easily determined, as it is dependent on the angular motion response of the eye (and hence the direction of gaze relative to the stimulus) and on the spatial and temporal nature of the flickering pattern. It also varies with the logarithm of the average luminance (L_{av}) of the stimulus according to the Ferry–Porter law (Sherr 1992):

$$f_c \sim a \log L_{av} + b \tag{1.2}$$

where a and b are constants. However, due to the other factors mentioned above, this equation does not in general provide an accurate value for f_c. The effects of flicker are more visible when the brightness levels of the stimulus are greater—this may be advantageous for displays operating in subdued lighting conditions. Since many volumetric display system technologies produce images of low intensity that are viewed in semidarkness and for only short periods, it has generally been possible to reduce the refresh frequency to levels that would otherwise be unacceptable.

An image that is in direct gaze begins to appear free of flicker at refresh frequencies in excess of 25 Hz.* The higher the refresh frequency, the smaller is the perceived variation in brightness. The human eye is more sensitive to motion at the periphery of its vision (i.e., outside the fovea centralis); flicker unobservable when looking directly at an image is often noticed when the gaze is elsewhere. It is interesting to note that although an image may appear comfortably flicker-free to both the direct gaze and peripheral vision, it may still exhibit subliminal flicker that may lead to discomfort after prolonged viewing. To minimize the effects of flicker, computer monitors now commonly employ effective update frequencies of 70 Hz or greater, and time-sequential stereoscopic displays may have refresh frequencies as high as 120 Hz (MacDonald and Lowe 1997).

1.4.3 Illusion of Motion

The minimum refresh frequency required to ensure freedom from flicker is somewhat greater than that required to provide the illusion of motion. For a sequence

* This is an approximate value. Under favorable conditions it be may reduced.

of images to appear to change continually and smoothly, animation frames need only be changed at a frequency of approximately 10 Hz. Thus, while a higher image refresh rate is required to minimize flicker, it is not necessary to alter the displayed image frame at each refresh. For example, in the case of a display system in which the output device is updated at a frequency of 30 Hz, the image content need (in principle) change only after each third refresh cycle. As a consequence, during a continual animation sequence, the buffer that contains the image data for this display need only be updated each one-tenth of a second.*

1.4.4 Summary of Some Depth Cues

The 3-D space around us is, visually speaking, mostly void. To perceive and comprehend the spatial distribution of objects within these surroundings, we employ (at a subconscious level) various features of the 2-D projection of the objects in the visual field on our retinas. Furthermore, we also derive information from the physiological behavior of the eyes during the process of vision. The projection of a 3-D scene onto the retinas (or any 2-D surface) will at every point on the surface include information about the nearest opaque object in the line of the projection.

Elements of the visual perception of depth are known as *depth cues*. It is important to note that these cues do not act in isolation—visual input is augmented continually by tactile and auditory sensations. In addition to these external stimuli (which depend on the nature of the scene being observed), "internal" input from the brain, such as prior knowledge (experience), associations, or expectations, may also affect the process of visual perception. Past experience provides a calibration of the visual input and is an important sense of absolute depth information; most depth cues by themselves indicate only relative distance (Ittelson 1960, Kaufman 1974).

The depth cues outlined briefly below may be grouped according to several methodologies (Dember and Warm 1979). One grouping scheme divides them into two categories (Figure 1.4): the *physiological cues* comprise those that rely on feedback from muscle groups, and *retinal image cues* comprise the remainder, in which features of an image scene observed by the visual system provide the source of input. Retinal image cues may be further subdivided into static monocular (or "pictorial") cues, which function with a single retinal projection, and parallax cues, which make use of information from more than one projection.

- *Binocular parallax (or stereopsis).* Each human eye occupies a slightly different position in space (the interocular distance is, on average, about 6.5 cm), and hence each receives a slightly different perspective of any scene. The brain fuses the images from each eye and interprets the differences in the apparent relative positions of objects in each view as relative

* This is a generalization that makes assumptions regarding the dimensions of the display and the speed of the motion of any image components depicted therein.

Retinal image cues:		
	• Occlusion • Perspective • Aerial perspective • Shading • Height in the visual field	Static monocular cues
	• Binocular parallax • Motion parallax	Parallax cues
Physiological cues:	• Accommodation • Convergence	

FIGURE 1.4 One possible classification of depth cues, in terms of physiological cues (those based on feedback from muscle groups in the eye) and retinal image cues (those based on the visual projections received on the retinas).

depth (Wheatstone 1838; Julesz 1964, 1971.) In practice, for objects located at distances greater than approximately 10 m, the differences between the two perspectives are negligible, so beyond this distance the cue of binocular parallax provides no depth information. However, it is of interest to note that approximately 5% of the population may not possess this depth cue (Richards 1970, Yeh 1993).* Bardsley and Sexton briefly discuss benefits that may be derived from the use of stereoscopic video display techniques (MacDonald and Lowe 1997) and comment:

> An element of visual noise may be present in the video image due to environmental factors or poor quality image reception. The human visual system is particularly adept at filtering out uncorrelated visual noise from binocular scenes to give greater picture quality. This ability is invaluable in determining not only *where* objects are in unfamiliar or complex scenes, but frequently *what* they are. (Merritt 1983).

- *Motion parallax (or temporal parallax).* Just as each eye receives a different perspective of the 3-D environment, relative motion between the head and the 3-D scene provides a continually changing perspective. This provides depth information in a manner that is analogous to stereopsis. Objects at different distances will alter their position in the 2-D retinal projection(s) at different rates.
- *Linear perspective.* Objects at greater distances from the observer subtend a smaller angle in the visual field than do nearby objects, and hence appear smaller (Edgerton 1975) (cf. Section 1.4.1); this is known as *Euclid's law* (Figure 1.3). Thus the relative sizes of similar objects gives an indication of their spatial arrangement, as evidenced by the convergence to an

* Of the 150 subjects comprising the study sample of Richards (1970), 4% were unable to use disparity, and 10% did so only with great difficulty.

FIGURE 1.5 The linear perspective depth cue. The parallel straight railway lines appear to converge with distance from the camera, as they are not perpendicular to the line of sight. Also, similar-sized objects appear smaller as they become more distant.

apparent vanishing point of a railway track, fence, or road extending into the distance (Figure 1.5). When the objects in question are texture elements on the surfaces of other objects, the effect is sometimes referred to as *texture perspective* (Gibson 1950).

- *Occlusion (or interposition)*. An opaque object interposed between the observer and a second object will render at least part of the second object invisible. By means of this cue, the brain interprets partially occluded objects as lying farther away than interposing ones (Figure 1.6).
- *Shading*. The shading and shadowing patterns of an object provide an indication of its form, surface texture, and position relative to its surroundings. Since light propagates in straight lines, the shading gradients on the surface of an object enable the topography of its surface to be predicted by the brain; a gradual change in the darkness of a surface indicates a curved surface, whereas an abrupt change signals the presence of a sharp edge. In the absence of information to the contrary, the brain generally assumes a light source incident from above (Figure 1.7). In conventional computer graphics, this cue is heavily exploited in rendering techniques (Foley et al. 1994).
- *Aerial perspective*. Due to light scattering by the interceding atmosphere (including water and dust particles), distant objects tend to appear hazy.

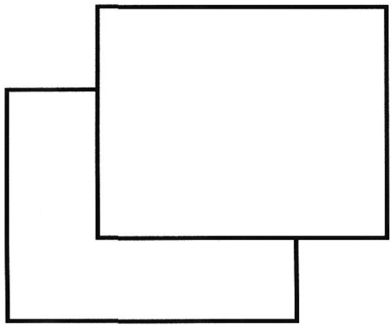

FIGURE 1.6 Occlusion depth cue. The square at the top right is perceived as being in front of the other because it partially occludes the latter.

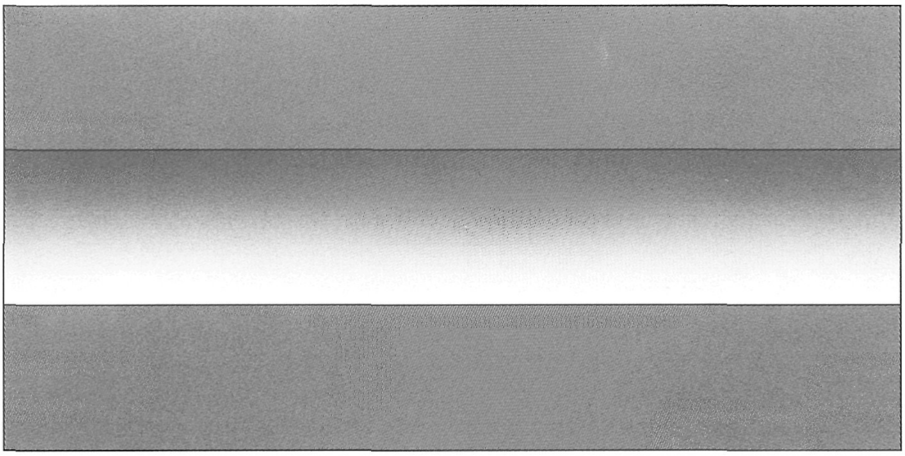

FIGURE 1.7 The central bar is generally seen as a concave trough, due to the brain assuming a light source from above. If the figure is viewed upside down, the same (subconscious) assumption leads to the pattern appearing convex.

As the distance of an object increases, it is seen less clearly, its coloration dulls, and it takes on a bluish tinge. Thus, assigning these characteristics to features of a rendered scene provides a sensation of distance. The aerial perspective depth cue may cause distant geographical features to appear to be in closer proximity on days when the air is clearer than usual. Under such circumstances, the effects of aerial perspective are minimized.

- *Height in the visual field.* Objects appearing higher in the retinal image are generally interpreted as being farther away, as the vertical positioning of objects in the visual field generally increases toward a horizon line (at eye level) as their distances approach that of the horizon. Similarly, objects above the eye level (e.g., clouds) appear to be farther away as they descend toward the horizon.
- *Accommodation.* This represents a physiological depth cue. To focus on the plane of attention, exertions of the ciliary muscle alter the shape of the lens

of the eye and so modify its focal length. The magnitude of the tension applied to these muscles is used by the brain to provide information relating to the depth in the field of view of the object under scrutiny. Beyond about 10 m, the lens shape is constant, so this cue is applicable only for objects that are closer than this distance. This cue is of particular importance in the case of volumetric displays which depict translucent images.

- *Convergence.* As with accommodation, this is also classified as a physiological cue in which the brain makes use of the muscular forces exerted on the two eyes so that they swivel inward by an amount necessary in order that they both center their gaze on an object of interest. As with binocular parallax and accommodation, this cue is effective only for distances up to approximately 10 m from the observer; as beyond this point the eyes are essentially parallel.
- *Familiarity, expectation, and context.* In the interpretation of a 3-D scene, the role of prior knowledge, experience, and the expectation of the attributes of an object in a particular context should not be underestimated. It has been demonstrated, for example, that assumptions about the size of a familiar object, based on the viewer's experience, can dictate the perceived distance of the object (Schiffman 1990).
- *Color.* The color of an object can also affect its perceived distance — colors toward the red end of the spectrum are perceived as being closer than those toward the green/blue. Although this effect has little consequence when regarding a real-world scene where a range of colors and shadings are present and other depth cues dominate, it is easily noticed when only a few discrete colors are present. The effect may be noticed on electronic text displays at some airports and other public buildings, where alternating lines of, say, red and green text are presented. From a distance, a clear depth offset between the text of the two colors may be noted. A similar situation applies to the current generation of volumetric displays, where only a few discrete colors of voxel may be generated. Volumetric displays are designed to be viewed sufficiently closely that cues such as motion and binocular parallax may be employed to their best advantage, but conflict with the color cue may be a potential source of eye fatigue. The color effect is, however, most noticeable in 2-D photographs or slides of volumetric images, where the additional depth cues do not apply.

1.4.5 Considerations Relating to Volumetric Displays

Many 3-D display techniques seek to add a sensation of depth by simulating either one or both of the parallax depth cues (see Section 1.5). In the case of volumetric display systems the situation is somewhat different. Rather than computing and presenting (via a 2-D medium) different perspectives of a 3-D data set, the entire 3-D data set is depicted simultaneously within a volume. Thus, an observer's eyes will each automatically see the image from a slightly different perspective. Furthermore, by moving relative to the volume, the observer changes the viewing position and hence the perspective. Both parallax cues are thus intrinsically satisfied in a natural manner. This is advantageous for viewers who do not possess the cue of stereopsis (see Section 1.4.4). The physiological cues are also satisfied, alleviating a possible source of eye fatigue. Furthermore, the depth cue information is presented without the need to undertake complex and often time-consuming computation.

Conventional high-resolution display devices are more able to satisfy some of the pictorial depth cues (such as shading) and better represent the surface detail of opaque objects than are current volumetric systems. Furthermore, the image space of a volumetric display represents directly the data displayed—apparent depth in a particular direction cannot be amplified by transforming the data using perspective and relative size, as is the case on a 2-D screen. The data displayed should be scaled uniformly in all spatial dimensions to ensure that it appears accurate from all viewing directions (see Section 10.7). The inability of most volumetric architectures to depict opaque objects results in a failure to properly support the cue of occlusion (see Section 10.2.3).

1.5 METHODS OF DISPLAYING 3-D INFORMATION

Any display system able to provide an observer with the sensation of depth can greatly enhance the impact and qualitative understanding of 3-D information. Rendering techniques used on conventional computer terminals utilize retinal image depth cues such as occlusion, perspective, shadowing, and possibly motion parallax (Fuchs et al. 1989, Coatrieux et al. 1990). The added realism that may be achieved by simulating the binocular parallax depth cue has spurred research into displays capable of replicating the slightly different views seen by each eye. [An excellent outline of some of the current research in this area may be acquired from McAllister (1993) and the *SPIE Proceedings*: for example, Fisher and Merritt (1990), Merritt and Fisher (1991, 1992, 1993), Fisher et al. (1994, 1995, 1996, 1997) and Bolas et al. (1998)]. Research is also being carried out (Okoshi 1976, 1980; Travis 1997) on display systems that permit a modified binocular view to be seen as the observer's head is moved. Examples of such systems include computer-generated holography (Tricoles 1987, Benton et al. 1993, Onural et al. 1994) and parallax barrier techniques (Akiyama et al. 1991). The majority of these autostereoscopic systems operate over a limited range of viewing angle.

1.5.1 Display Systems and Depth Cues

Any display system that attempts to depict spatial information creates images that satisfy a certain subset of the depth cues discussed in Section 1.4. Although cues such as binocular parallax and occlusion are often considered to be dominant in the perception of relief, realism is enhanced by satisfying as many cues as possible. In particular, it is important to avoid prolonged *conflict* between different depth cues. Situations in which two or more cues conflict and therefore provide the brain with differing depth estimates can lead (at the very least) to fatigue (Hart and Dalton 1990). For example, in the case of some stereoscopic techniques, the information provided by binocular disparity is inconsistent with that provided by accommodation. If the display system is not of sufficiently high resolution, the plane in which the image is generated may be perceived subliminally. The accommodation and convergence information may then be in conflict with that provided by the retinal image cues. Although volumetric displays automatically satisfy a number of depth cues in a natural manner and so avoid conflict, we will see in subsequent sections that they may exhibit other problematic characteristics.

Conventional 2-D display systems satisfy only the monocular (retinal image) depth cues. In the case of 3-D display techniques, some or all of the remaining depth cues (i.e., convergence, binocular parallax, accommodation, and motion parallax) are satisfied. Of these, binocular and motion parallax are most frequently employed to enhance the sensation of depth. Ideally, all the retinal image cues would also be satisfied; however, this is not always the case. If a display device cannot present certain depth cues to the viewer, its application tends to be restricted to areas in which the absent cues are not a major disadvantage. It is important to note that the relative emphasis that we assign to the various depth cues varies according to the nature of the image scene under observation. Current volumetric display systems satisfy most of the nonretinal cues but in general do not accommodate shading and full occlusion. However, for certain applications the absence of these cues is not a major disadvantage.

1.5.2 Monocular Displays

Following the development of mathematical prescriptions enabling objects and scenes to be accurately recreated according to the rules of linear perspective, it became possible to utilize 2-D renditions of 3-D objects within scientific works. The studies of Viator and Dürer in the fifteenth and sixteenth centuries were particularly influential in this regard (Ivins 1973). Coupled with the invention of printing (ca.1430), it became possible to include within scientific studies realistic sketches of objects ranging from abstract mathematical concepts and technical drawings to biological specimens. Even when a perspective reconstruction was not necessarily required, the philosophy of portraying objects "as seen by the eye" was extremely influential and marked a significant step forward for science in general (Edgerton 1991).

The invention of photography (ca. 1839) further facilitated the recording of items such as biological specimens. Perhaps the most significant development was the discovery of evidence for the existence of cathode rays (Plücker 1858, Hittorf 1869) using tubes in which an electron beam was generated by a high-voltage discharge in a gas (usually air) at a pressure of about 0.5 torr (Martin 1986). Crookes (1879) demonstrated that these "rays" comprised negatively charged particles that could be deflected by an electromagnetic field. Braun

(1897) invented the first CRT for use as an oscillograph (Blondel 1891). This also consisted of a gas discharge tube containing a cold cathode, requiring a high voltage ($\geq$10 kV) applied to an accelerating anode so as to maintain the discharge. Having passed by the anode the beam's spread was limited by an aperture plate before it impinged on a fluorescent screen. External coils were used initially to provide magnetic deflection, but later models incorporated additional electrodes within the tube so as to permit electrostatic deflection (Martin 1986).

Wehnelt (1904) invented the triode gun, which utilizes a control grid, so establishing the foundation for the modern CRT. Between 1904 and 1940 the fundamentals of electron optics were formulated by Busch (1926, 1927), Davisson and Calbick (1932), Knoll and Ruska (1932), Brueche and Scherzer (1934), and von Ardenne (1935, 1939, 1940). The cathode ray oscillograph enabled electrical signals to be depicted in real time. The development of computer systems from the 1940s onward provided further impetus for display systems* and it was soon realized that interaction with computers would be more efficient if printed output was augmented with devices able to display both text and graphics.

Initially, CRT-based computer graphics were generated primarily through the use of vectors; that is, the beam was directed from one visible point on the screen to the next and so moved in an irregular path; the path depending on the geometry of the image being depicted. The subsequent refinement of raster-scanned CRTs, the provision of full color, and the development of the associated hardware (including dedicated processing chips) have further increased graphics visualization capabilities. Ray tracing provides a very useful technique for increasing realism (in the photographic sense) of monocular computer images. This is a computationally expensive process in which the position and characteristics of one or more light sources are assumed and the luminance and color properties reflected in the direction of a chosen viewpoint are calculated (Foley et al. 1994). In essence, the more realistic the image, the greater the computation overhead, and this is reflected by the frequent inability of even today's high-performance workstations to render high-quality images in real time. Monocular displays are particularly suited to the depiction of surface detail. Furthermore, the ability to view data in cutaway form facilitates an appreciation of internal structure.

1.5.3 Stereoscopic Displays

The depth cue of binocular parallax, whereby the slightly different views of a scene obtained by each eye are combined in the brain to give an impression of relief (see Section 1.4.4), provides a powerful depth effect. There have been many successful attempts to incorporate binocular information into visualization hardware and a considerable amount of research continues in this area. In the case of stereoscopic displays, each eye is artificially presented with a slightly different view of an object or scene, so as to provide a strong sensation of depth.

---

* The earliest electronic (digital) computer systems also employed special-purpose CRTs for information storage and switching.

Although this effect was not noted until the nineteenth century; it had long been realized that each eye sees a different view of the visual field. Previously, philosophers had been more concerned with explaining why we do not see double and had assumed that the only depth cue arising from two eyes was that of convergence (Wade 1987). Wheatstone (1838) appears to have been the first to describe the phenomenon of stereopsis. Furthermore, he was able to demonstrate the powerful nature of this cue by using a device he termed the *stereoscope* which artificially presents each eye with its correct perspective of a scene or object. The object is then seen in relief and no longer appears contained within a 2-D plane. It is interesting to note that there was a certain amount of controversy during the nineteenth century regarding the true inventor of the stereoscope (Wade 1983, 1987). Brewster (1856), in particular, maintained not only that Wheatstone did not originate the device, but that the effect of relief was due to convergence alone. Brewster attempted to attribute the invention of the stereoscope to Elliot in 1834, whereas Wheatstone drew attention to earlier work by Mayo (1833). The relevant excerpts are quoted in Wade (1983, pp. 149, 173).

The stereoscope rapidly became extremely popular, and several different implementations were produced (Brewster 1856). With these devices, physically separate right and left eye views are presented to the eyes via a device comprising mirrors or lenses. A barrier is used to separate the visual fields of the two eyes. The Wheatstone and Brewster stereoscopes are illustrated in Figures 1.8 and 1.9. From the outset the educational and scientific applications of stereoscopic visualization equipment were emphasized (Brewster 1856, Helmholtz 1909).

Other stereoscopic techniques do not require the stereopair to be physically separated, so the viewer need only wear special glasses. This engenders significant freedom of head motion and enables the image to be viewed in a more natural manner. One such method, the *analglyph*, originating with the work of Rollman in 1853 (Southall 1962, p. 356), requires glasses comprising eyepieces of a different

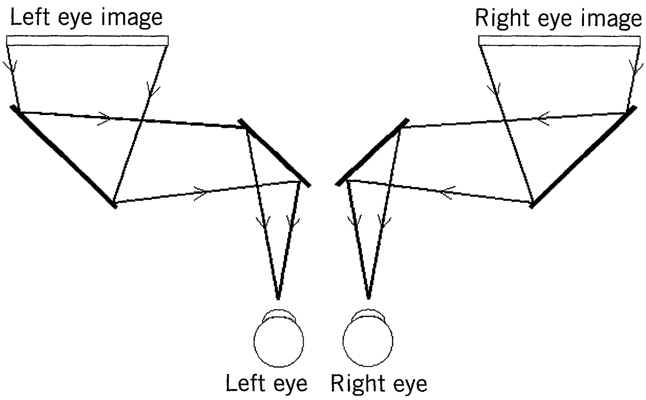

FIGURE 1.8 Arrangement of the Wheatstone (reflecting) stereoscope. Light from each view travels to the intended eye by the system of mirrors.

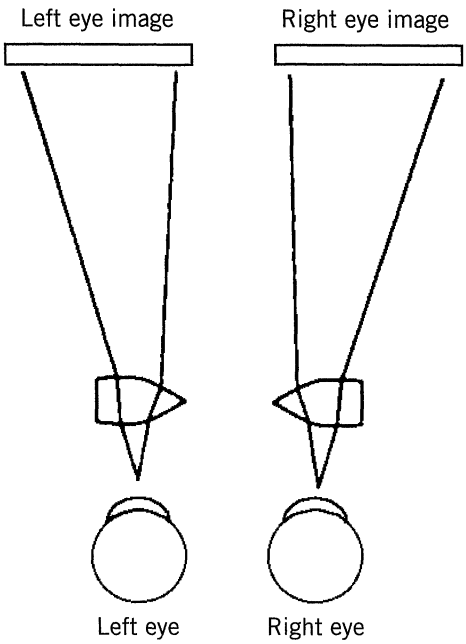

FIGURE 1.9 Arrangement of the Brewster (refracting or ocular) stereoscope. The two image projections, corresponding to the views seen by the left and right eyes, are directed to each eye by means of prisms.

color (e.g., red and blue). If a stereopair of images are overlaid in red and blue, each eye will then only see the image of the opposite color. This process has been applied to cinematography at various times since the beginning of the twentieth century (Dudley 1951) and is still used today for an economical 3-D effect in scientific visualization and entertainment. One problem with this technique is the provision of natural color to the stereoscopic images. Color may be retained to a certain degree by ensuring the colored filters on the eyepieces are wideband and admit a range of colors rather than a narrow range of wavelengths. Thus, one of the images would contain the low-frequency (yellow, orange, red) color components, and the other the higher-frequency (blue, green) colors. When viewed through corresponding filters, the colors are blended with the fusion of the two images, and a reasonable color reproduction may be obtained (Dudley 1951).

In an alternative embodiment the eyepieces of the viewer's glasses take the form of perpendicularly polarized glass or plastic. An image stereopair is displayed in which the light of one of the pair is perpendicularly polarized with respect to the other. According to Dudley (1951), Anderton first proposed this process in 1890, although difficulties in obtaining the necessary materials delayed practical implementation for nearly half a century. In the context of modern scientific visualization, consecutive frames on a computer monitor may form stereopairs and be polarized alternately. In this instance the members of the stereopair are temporally separated (frame sequential), but since visual information received within about 100 ms is processed simultaneously, the consecutive left/right eye images are, in effect, concurrent.

1.5.4 Autostereoscopic Displays

A limitation of stereoscopic displays is that motion parallax is not provided automatically; moving one's head from side to side does not necessarily produce a different viewpoint of the scene depicted. Display devices that are able to provide the correct stereoscopic perspective to each of the viewer's eyes over a range of viewing positions may be defined as *autostereoscopic*. Volumetric displays fall into this category. Below, two other broad classes of autostereoscopic display techniques, which do not require special glasses or headgear, are outlined.

The stereoscopic displays described above, in combination with powerful computer systems, can also provide an updated stereoscopic view in response to motion of the observer's head if an appropriate position or motion sensor is provided [see Travis (1997) and references therein]. However, the processing of the sensing device feedback and recomputation of the image so as to update the stereoscopic views of the virtual scene continuously and smoothly greatly increase the burden on the associated graphics engine. Nevertheless, increasing computational speed coupled with more efficient image data-processing algorithms and feedback from applications is steadily improving this technology. When the viewing apparatus takes the form of a head-mounted display, providing an immersive environment that the operator "enters", such systems are often known as *virtual reality* or *virtual environment* systems. Despite the computational overheads, the realism associated with this approach (which may be enhanced by the incorporation of tactile feedback in sensor equipment (such as gloves) ensures the continued development of this type of display methodology.

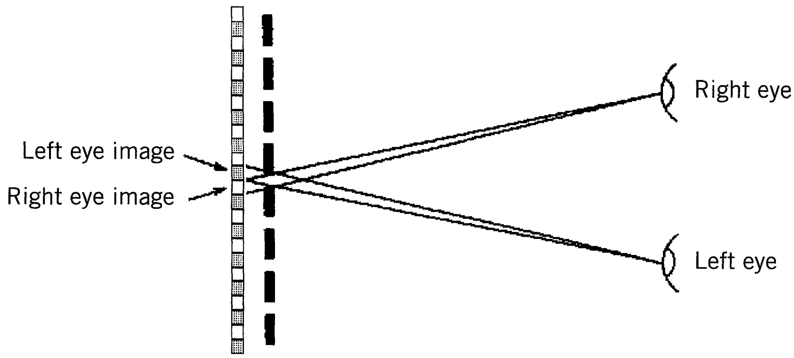

FIGURE 1.10 Parallax barrier technique for presenting binocular images. The viewer's head position needs to be quite constant relative to the display, or the left and right eye images swap and the relief of the figure is inverted.

Parallax Barrier and Lenticular Screen Techniques

The parallax stereogram, illustrated in Figure 1.10, was originally devised by F. E. Ives in 1903 (Okoshi 1976). A sheet containing fine vertical slits is placed in front of a specially prepared picture containing alternate left and right eye views of an image. The slits control the direction of the light emitted from the picture so that each eye sees a different perspective of the image. However, this original method allowed little freedom in viewing position. In 1918, Kanolt improved the system in this regard by decreasing the ratio of the image strip width s to the slit width p from about $\frac{1}{2}$ to $\frac{1}{10}$ (Figure 1.11).

A difficulty associated with this technique was the creation of the picture sheets with the alternate-view strips. Various ingenious schemes to achieve this, utilizing one or more cameras, were devised by H. E. Ives (the son of F. E. Ives) in the late 1920s (Ives 1928, 1929, 1930a, 1930b; Okoshi 1976). However, several major drawbacks to the parallax barrier method resulted in the technique losing favor by the end of the 1950s. The darkening of the images due to the presence of the barrier increases as the slit-to-strip ratio decreases, and hence as motion parallax range increases. Moreover, the slits must be sufficiently narrow that they are not observable. But this then can lead to unacceptable spreading of the image light due to diffraction.

The technique of integral photography, invented by Lippmann (1908), provides a method of recording a spatial image on a single photographic plate. The process

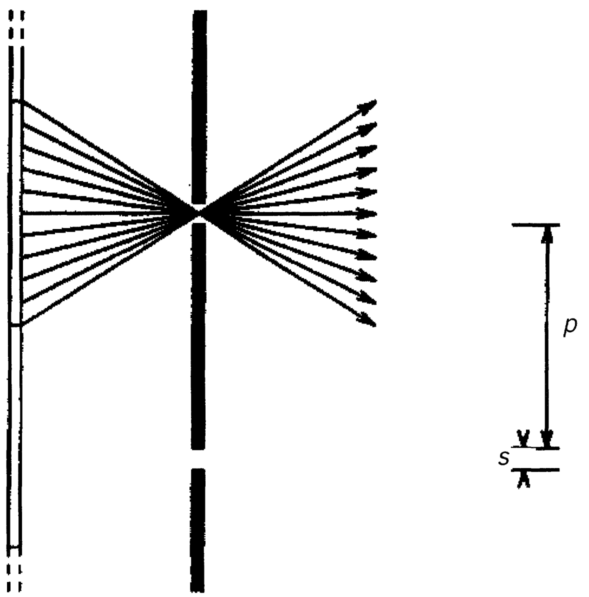

FIGURE 1.11 Parallax panoramagram, in which the slit/strip width ratio s/p is decreased from about $\frac{1}{2}$ to $\frac{1}{10}$. This gives more freedom of lateral movement to the viewer.

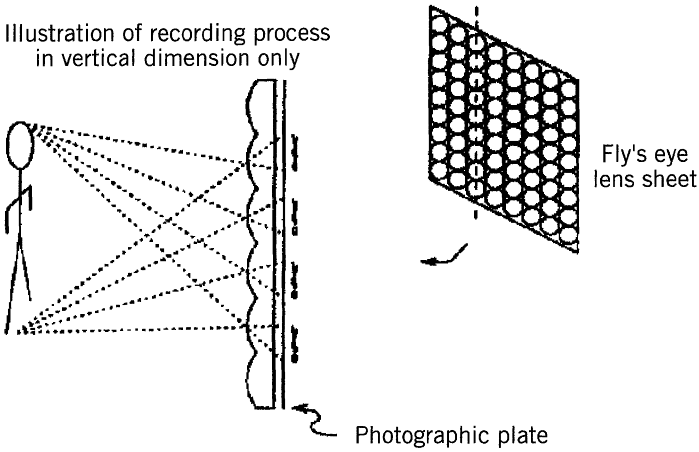

FIGURE 1.12 Technique of integral photography. The array of lenses provides an array of photographic replicas of the target object. When developed appropriately and viewed through an equivalent lens array, a 3-D image ensues as each lens projects a slightly different view of the object to the observer. [after Okoshi (1976, p. 22) with permission.]

utilizes a large array of tiny convex lenses, known as a *fly's-eye lens sheet.* An object in the focal plane of this sheet will be imaged by each lens from a slightly different viewpoint. A photographic plate behind the sheet records a large number of small images of the object, each viewed from a different direction. By developing and reilluminating the photograph, and viewing through the fly's-eye sheet, a 3-D image of the object, satisfying binocular and motion parallax (both horizontally and vertically) is observed (Figure 1.12). However, the resulting image will be *pseudoscopic*, that is, reversed in depth. This difficulty may be overcome by taking a second integral photograph of the first image (Ives 1931), although this may result in resolution being degraded (Okoshi 1976). The first experiments with this technique were performed by Sokolov in 1911, using an array of pinholes to approximate the small lenses (Valyus 1966, Okoshi 1976); lens sheets of sufficient quality were not available until after World War II.

Interest in a simplified version of the technique, dating back to the work of F. E. Ives and H. E. Ives, underwent a resurgence in the United Sates in the 1960s (Burckhardt 1968a). The 2-D array of lenslets was replaced by a linear array of cylindrical lenses, known as a *lenticular sheet* (Figure 1.13). In this case, the parallax is obtained in the horizontal direction only. Arrays of diffraction gratings may be employed in place of lenslet arrays (Nordin et al. 1994, Sakamoto et al. 1996, Toda et al. 1996). These provide a field of view of greater angular extent than the approximately 15° available with a multilens array. However, both of these methods require a pixel resolution corresponding to the product of the resolution of each view and the number of views (Travis 1997). This places strong demands on both the manufacture of the display screen and on the computation and data throughput capabilities required of the graphics

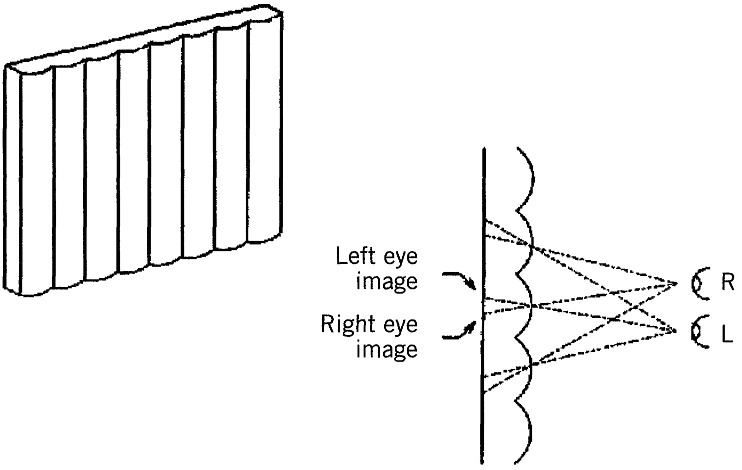

FIGURE 1.13 Principle of lenticular sheet imaging. The principle is the same as that of integral photography, except that parallax is retained in the horizontal direction only. The array of strip lenses is more easily fabricated than an array of fly's-eye lenslets (Figure 1.12).

engine. A further possibility for autostereoscopic display is that of spatially or temporally multiplexed video images with the aid of a lens array (Travis 1997).

Lenticular sheet–based electronic displays still form an active area of research interest. An LCD panel illuminated from the behind and equipped with a lenticular sheet permits the presentation of autostereoscopic images across a range of viewing positions (Isono et al. 1993, Sheat et al. 1993).

Computer-Generated Holography

The method of holography was originally proposed by Gabor (1948) to overcome resolution limitations in electron microscopy. The potential application of Gabor's holographic technique to 3-D imaging was heralded by the work of Leith et al. (Leith and Upatnieks 1964, 1965; Leith et al. 1965), which pioneered the two-beam method of hologram formation. This required a stable optical bench, on which the coherent light of a laser source was split into two beams, one of which was scanned across an object and subsequently interfered with the other beam to produce an interference pattern, characteristic of the object, and which could be captured on film. When reilluminated by laser light of the same wavelength, the interference pattern diffracts the light in an identical fashion to the object, so that a 3-D image of the object is made visible and appears suspended in the air between the observer and the holographic plate. Images generated in this manner satisfy both the binocular and motion parallax depth cues. Moreover, the cue of accommodation is present — the eye must change focus depending on the depth of the area of attention within the hologram. An important advance of this technique was the development of white-light holograms (Leith 1976, Caulfield 1979).

In employing holograms as an electronic display, it is necessary to compute the holographic interference pattern for each frame, which requires a large amount of information. Each point in the hologram may be calculated using the sum of the partial waves from all points in the object. For the computation of holographic images to be practical on a near real-time basis, the information requirements must be reduced substantially (Burckhardt 1968b). To this end, a number of simplifications and efficiencies have been adopted (Benton et al. 1993, Lucente 1993) and advanced prototypes are under development.

1.6 DISCUSSION

Volumetric displays provide 3-D images that are in many ways unique among visualization techniques. Being cast within a physical volume, the images naturally satisfy depth cues such as binocular and motion parallax without the need for additional viewing apparatus, and generally, there is very considerable freedom in viewing orientation. However, as will become apparent during the course of this book, a bewildering variety of technologies and physical mechanisms have been proposed and/or developed (in prototype form) with the aim of achieving these goals. In this book, we attempt to formulate concepts that are common across the wide range of the technologies and embodiments to be found in the volumetric field.

Although the ideal volumetric image space may be imagined as a uniform, isotropic, high-resolution tableau, this represents a formidable and as yet unrealized technical challenge. A number of aspects of the image space, including its dimensions, voxel size, and voxel spacing, affect the degree of detail and spatial structure that may be depicted therein. Conversely, at the design stage, the intended application will place requirements on these and other parameters of the image space and display system.* The manner in which the observer perceives the displayed data plays a crucial role in this interrelationship.

Almost all of the research that has been carried out to date regarding human depth perception has concerned our perception of the world around us (being the environment in which our visual system has evolved) and more recently the perception of images on 2-D or stereoscopic displays. Volumetric display units provide quite a different medium for the presentation of computer-generated data, and a number of perception issues relevant to volumetric display systems remain open. For example, the spatial acuity (and viewing distance) of the eye affect the required voxel size and separation. Furthermore, the relationship of the spatial acuity to the contrast between a feature of interest and its surroundings and to its degree of motion is complex. Sound knowledge of the spatial acuity in

* The application dependence of the image space characteristics is an important issue (see Chapter 10).

relation to volumetric images would therefore be invaluable in the development of volumetric systems, but to the authors' knowledge, has not yet been established.

Another important issue is that of conflict between depth cues. The physically 3-D nature of volumetric images provides consistency between the information obtained from a number of the depth cues (in particular, the physiological and parallax cues). However, the perception of volumetric image components comprising different discrete colors in systems providing a very limited color palette may provide some confusion. As the color of the component itself can provide a cue to depth, this may conflict with the information provided by other cues. This effect is stronger in 2-D representations of volumetric images (such as photographs), where the important parallax and physiological cues have been removed. The perception of images depicted on volumetric systems remains a largely unexplored area of research. It is to be hoped that some of the questions in this area will be addressed, and resolved, in the not too distant future.