

Original article

Visual cues and pictorial limitations for computer generated photo-realistic images

Christopher G. Barbour¹
and Gary W. Meyer

¹ Department of Psychology, University of Oregon,
Eugene, OR 97403, USA

² Department of Computer and Information
Science, University of Oregon, Eugene, OR 97403,
USA

The limitations of two-dimensional pictures as representations for reality are discussed. A review is made of the perceptual cues necessary to convey a sense of realism. These cues include, but are not limited to, binocular disparity, field of view, accommodation, vergence, and chromatic adaptation. Examples are given of how the physical characteristics of two-dimensional pictures limit the use of these cues in computer-graphic images. Techniques developed by artists and photographers to overcome some of these limitations are discussed.

Key words: Picture perception – Realism
– Image synthesis – Pictorial cues

Correspondence to: C.G. Barbour

1 Introduction

The objective of most work in realistic image synthesis is to create a picture that is optically correct. The focus, therefore, of many synthetic image-generation algorithms is to determine the array of electromagnetic energy that passes through a particular plane in space. The sophistication of this approach has continued to increase, drawing on work from areas as diverse as radar-scattering theory and radiation-heat transfer (Hall 1989). This work has now reached the point where the radiometric values computed by current simulation techniques match the physical quantities measured in an actual scene. A picture created from the results of the simulation has been successfully compared, under restricted viewing conditions, to an image of a real scene (Meyer et al. 1986).

Unfortunately, the picture produced from a physically based approach to image synthesis turns out not to be a perfect optical reproduction. Among the things that make the representation provided by a picture different than the view through a window are the facts that it is taken from a single point of view, has a limited dynamic range, lacks physical depth, and has a detectable surface finish. Some of these differences are intentional and are what make pictorial representation desirable. Others are the result of limitations in the reproduction media. For example, to overcome the limited dynamic range of the television monitor used in the above visual comparison experiment it was necessary to obscure the view of an exposed overhead-light source in the scene.

The physical limitations of pictures and the practical uses to which images are put makes it clear that optical identity is not the correct objective in realistic image synthesis (Mills 1985). However, if the reproduction is not optically indistinguishable from the original environment, the viewer sees a different two-dimensional array of light and their impression of the scene is altered. In addition, if the viewing conditions under which the reproduction is observed are different than the original viewing circumstances, additional distortions are introduced. For example, in the visual comparison experiment described above, differences caused by adaptation to the color of the viewing illuminant were avoided by providing identical observation conditions for both the real scene and the reproduction.

To overcome the limitations of pictures and the distortions that are introduced, an approach taking the perceptual experience of the viewer into account is required. First, it is important to know

which cues are used by our visual system to interpret the three-dimensional world around us. Next, it is essential to understand how the percept generated by the reproduction differs from that produced by the original scene. Finally, it is necessary to modify the image-synthesis algorithm to account for these distortions. Hints about how to solve these problems can be found by studying techniques that have been developed by artists and photographers. In most cases, however, research remains to be done in order to devise an algorithm that can correct for these problems.

This review article focuses on the first two of the above three steps leading to a perceptual approach to realistic image synthesis. The concentration is on full-color two-dimensional static images. [See Van de Grind (1986) for dynamic images and Hagen (1991) regarding perspective projection.] The limitations that prohibit a picture from presenting an observer with an optical field identical to that experienced in an actual scene are enumerated. Physical limitations, such as fixed point of view, finite size, and flatness, are covered first followed by lighting problems, such as limited dynamic range and the need for a viewing illuminant. Visual cues affected by each limitation are identified and the characteristics of these perceptual stimuli are discussed. Where it is available, the importance of each cue in the perception of pictures and three-dimensional scenes is described. Qualitative attempts by artists and photographers to overcome each limitation are presented.

2 Pictures are made from a fixed point of view

The creation of a computer-graphic image involves the perspective projection of a three-dimensional scene onto a two-dimensional image plane. This is accomplished by following projection lines from a fixed point of view back into the scene and then determining the point of intersection between these lines and the image plane. As such, each computer-generated image represents the world as seen from a single eyepoint. Because our visual system has two eyes, our brain receives information about the natural world from two slightly different viewpoints. A single computer-generated picture therefore lacks one of the cues that our visual system uses to help us decipher the three-dimensional world. This cue is known as binocular disparity.

2.1 Binocular disparity

When both eyes fixate on an object, some areas of the object project to corresponding points (same retinal positions) on the two retinas and some areas project to non-corresponding points (Fig. 1). This occurs because the eyes view an object from two different locations in space. Furthermore, as a result of the eyes being separated, an imaginary surface (called the horopter) is located in visual space at a distance from the observer that depends on where the eyes are fixated. Any object falling on this surface projects a retinal image that is located on corresponding points of both retinas (Goldstein 1984). Conversely, any object located off this surface projects to non-corresponding points. The difference in distance between these non-corresponding points on the two retinas is called *disparity* and the degree of disparity between the two points varies as a function of how far the object is from the horopter (Fig. 1). If the disparity is large enough, the percept of the object is doubled. Otherwise, a single fused image is perceived. Whether the image is doubled or fused, disparity is a depth cue our visual system uses to perceive a three-dimensional object from a two-dimensional retinal image (Goldstein 1984).

According to Graham (1965), the limiting range of binocular disparity as a cue can be considered as the greatest distance at which an object can be placed and still be considered nearer than an object at infinity. When the angular disparity in the separation between the two images is taken as 30 sec of arc and the interocular distance is taken as 65 mm, the limiting range has been calculated to be 495 yards. This underscores the large range of distances over which disparity can act as a cue for depth. Another way of measuring the limits of binocular disparity is to determine the threshold angular disparity between two images needed to produce stereopsis. This value has been found to be approximately 2.0 sec of arc, but can vary up to 40 sec, depending on the conditions of experimentation (Langlands 1926).

Pictures represent a three-dimensional object in a two-dimensional plane. If the viewer fixates on an area of the represented object, all other areas are in the same plane and fall on corresponding points, thereby eliminating any disparity that would have normally occurred if the object was actually three dimensional (Helmholtz 1881; Meinel 1973; Hochberg 1979, 1980). This lack of disparity not only

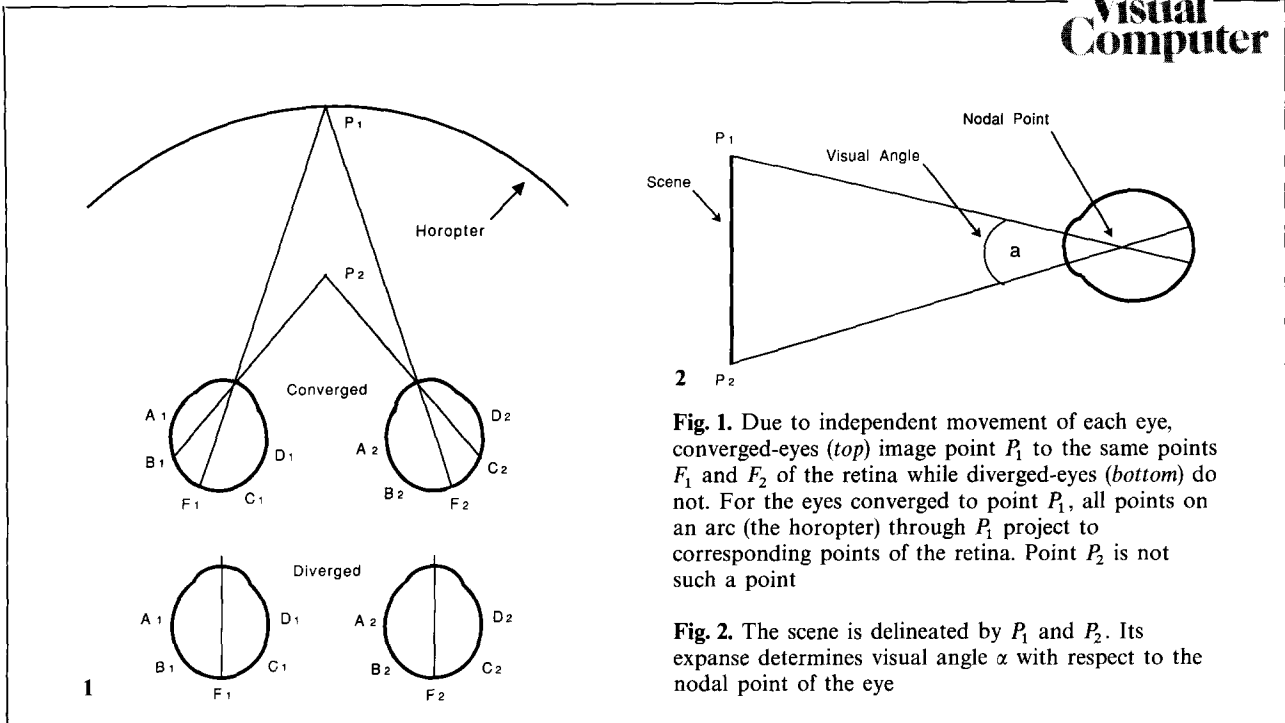


Fig. 1. Due to independent movement of each eye, converged-eyes (top) image point P_1 to the same points F_1 and F_2 of the retina while diverged-eyes (bottom) do not. For the eyes converged to point P_1 , all points on an arc (the horopter) through P_1 project to corresponding points of the retina. Point P_2 is not such a point

Fig. 2. The scene is delineated by P_1 and P_2 . Its expanse determines visual angle α with respect to the nodal point of the eye

reduces the number of depth cues available to an observer, but also cues the observer that the object is in fact two dimensional (Pirenne 1970; Hochberg 1979; Sedgwick 1980; Goldstein 1984). Despite the use of disparity as a depth cue by most individuals, it has been estimated that 5–10% of the general population cannot use it as such (Sekuler and Blake 1990). Therefore, it is possible that a flatness cue found in two-dimensional representations is not important to these individuals. However, this has not been studied empirically.

2.2 Overcoming limitations of monocular images

Meinel (1973), like Hochberg (1980), suggests that painters are not without possible solutions to offset the limitations caused by a single point of view. He refers to the work of Evan Walters, where a single picture is made that incorporates four perceptions usually only found in the double images resulting from large disparity. These four perceptions are: 1) various gradations of transparency of half of the double image in which one-half may vary between entirely transparent to entirely opaque, 2) creation of unique shapes due to the overlapping of dissimilar images, 3) halves of the double image seen at different vertical positions as a result of slight tilting of the head, and 4) color/

brightness effects resulting from the overlapping of two images that vary in color and/or luminance.

The creation of stereo-image pairs is the straightforward solution to the absence of the disparity depth cue in monocular images. Stereo image pairs have been created since the earliest days of computer-graphic image synthesis. Recent development of liquid-crystal shutters that can alternately present images to the left and right eyes has stimulated new interest in this area. The use of head-mounted displays as part of the effort to produce an entire *virtual world* for an observer (Chung et al. 1989) require that additional attention be paid to this subject.

Despite the evidence supporting the importance of disparity as a depth cue, Rock (1975) points out that this cue is often present with other monocular depth cues (see Sect. 4.4) that enhance the perception of depth. If these monocular cues are not present, the sensation of depth is more difficult to achieve for some observers, even though the disparity cue is still present. One method of overcoming the loss of binocular disparity in individual computer-graphic images is to make sure that these monocular cues are properly rendered and are emphasized wherever possible. The importance of monocular cues in the perception of depth is illustrated by experiments showing a pronounced sensation of depth when information that conflicts with the monocular depth information is not avail-

able to the viewer. This is accomplished by viewing the picture monocularly through a peep hole (see Sect. 3.2).

3 Pictures have a finite size

Most computer graphic images are displayed using color-television monitors or color-hardcopy devices. While there are projection systems such as IMAX that present very large images, these devices are impractical for the day-to-day working environment in which most computer-generated pictures are used. Therefore, a typical computer-graphic image is limited in its ability to convey a sense of realism, because it has a clearly identifiable boundary and does not occupy our entire field of view.

3.1 Field of view

An object at some finite distance from an observer projects an image onto the observer's retina. Because the size of the retinal image is directly related to both the size of the object and the distance of the object from the retina, it is useful to describe retinal-image size in terms of a single parameter: the visual angle that the object subtends with respect to the nodal point of the eye (De Valois and De Valois 1989, see Fig. 2). The entire scene viewed by an observer also projects an image to the retina, and the size of this retinal image is directly related to the expanse of the scene. The visual angle subtended by the scene is therefore a measurement of the expanse and is called field of view.

The maximum field of view for an individual using both eyes is approximately 200° visual angle (Darley et al. 1986; Van de Grind 1986) and in most cases the field of view of a real scene approximates this. However, a picture only occupies a small fraction of such a large field of view (Helmholtz 1881; Evans 1959; Pirenne 1970; and Gibson 1982). Gibson points out that this "narrowing" of the scene in a picture does not provide information to the periphery of the eye that is normally present during real-scene viewing and pictures are therefore limited in this regard.

There is specific evidence as to how field of view affects perception of two-dimensional displays. For example, an experiment was conducted in which the "sensation of reality" was shown to vary as

a function of field of view. The measured psychological effects were continuity and unification between display space and observer space, feeling of expanse, naturalness, feeling of depth, and impressive powerfulness (Hatada et al. 1980). The results indicated that visual displays with horizontal visual angles between 30° and 100° and vertical visual angles between 20° and 80° produce psychological effects that increase the sensation of reality. Larger fields produced saturated responses. Another related study has shown that the subjective quality of a display, rated on a 1–100 scale, increases linearly with the logarithm of picture angle (Westerink and Roufs 1989).

In addition to a limited field of view, lack of a frame around a picture also minimizes the effectiveness of the picture as a high-fidelity surrogate, because the scenes represented do not have a margin or frame surrounding them (Helmholtz 1881; Evans 1959; Gombrich 1956; Pirenne 1970, 1975; Hochberg 1978a; Haber 1980a). Moreover, the presence of a frame has been shown to affect visual perception of an image. For example, the perceived orientation of a rod positioned inside a frame has been shown, among other things, to be a function of the degree of tilt of the frame (Witkin 1959). Furthermore, because perceived shape and form is affected by the simultaneous perception of other shapes and forms (see Evans 1959), a picture of a tree will look taller in a long skinny frame compared with a short squat frame.¹

3.2 Overcoming limitations of image size

Occluding the appearance of the frame has been shown to affect the perception of a picture. This style of viewing helps to eliminate those cues (e.g., frame of the picture, surface of the picture) informing the viewer that what is being looked at is a picture and not a real scene (Smith and Gruber 1958; Evans 1959; Pirenne 1970; Gibson 1982). Helmholtz (1881), Hochberg (1978b), and Pirenne (1970, 1975) report that viewing a picture through a peephole so that the frame is occluded from view often increases the sensation of depth in a picture. This increased perception of depth can be so com-

¹ It is also believed that the frame provides information to the observer concerning picture-surface attributes. This information is thought to affect perception of the picture, which further limits the picture as a high fidelity representation of real scenes (see Sect. 4.3)

elling that it has been likened to the 3-dimensional effect experienced when looking through a stereoscope (Pirenne 1970). In fact, this "peephole technique" has been used in art galleries and museums to enhance the sense of depth and thereby increase the visual impact of the picture (Pirenne 1970). In spite of the dramatic improvement in image realism obtained by viewing a picture through a peephole, it is not considered to be a practical technique to use when observing television or photographs.

A photographer will often use a view camera to overcome some of the problems caused by a picture having a border. A view camera is a device that allows the photographer to position the film plane so that it is not perpendicular to the axis of the lens. This provides additional freedom in composition, but introduces perspective distortion into the picture. One common use of the view camera is to adjust the perspective of tall objects, such as buildings, so that all verticals in the image are parallel. Without this modification, the frame of the picture tends to exaggerate the perspective convergence of vertical lines in a way that does not appear correct (Evans 1959; Stone 1987).

In Figs. 3 and 4 the computer-generated still-life picture of a book and a ball have been created using the ray tracing method. In Fig. 3, a simple perspective projection was produced by keeping the view plane perpendicular to the rays cast into the scene. The verticals of the book are not parallel to the frame of the picture, causing the book to appear as if it is tipping. In Fig. 4, the picture has been redone with the projection plane held parallel to the front of the book. This eliminates the problem with the book, but introduces some distortion of the sphere. While the effectiveness of these modifications are subject to individual interpretation, they are representative of the type of intentional distortions that photographers introduce into their work by using the view camera.

4 Pictures are flat

A number of consequences result from using a perspective-projection algorithm in realistic image synthesis to create a two-dimensional representation of a three-dimensional scene. Neural signals indicating the state of muscle contraction in the eye are missing when a picture lacks physical depth. This includes signals from the ciliary mus-

cles that focus the eye to objects at different depths (accommodation) and signals from the extraocular muscles that control independent movement of the two eyes while fixating on objects at unequal distances (vergence). One of the most important effects of image flatness is that the picture becomes an identifiable object in our environment rather than a window into a scene that lies beyond it. This is due to the surface perception capabilities of our visual system.

4.1 Accommodation

Accommodation is the process by which the eye changes the shape of the lens in order to focus on objects that are at different distances from the viewer. Experiments have shown that accommodation is an inaccurate source of depth information for distances more than a few feet (Pirenne 1970; Rock 1975). However, for shorter distances, this cue becomes more reliable (Rock 1975).² Information from accommodation should therefore be considered relevant for those pictures depicting short viewer-object distances.³

Because pictures are flat, information from accommodation is completely lost and a viewer is not able to accommodate to objects at different represented distances (Hochberg 1979). The missing information has been proposed to be in the form of proprioceptive⁴ feedback from sensory fibers in the ciliary muscles responsible for changing the shape of the lens (Rock 1975). Another form of lost information could be the absence of blurred non-accommodated objects. This blurring of objects in front of and behind the accommodated object could be used as a cue for depth (Evans 1959). Therefore, unless depth of field is provided during picture viewing (see Evans 1959), this potential source of depth information will not be present.

² Rock (1975) points out that for even these short distances, accommodation is not very reliable compared to other depth cues

³ Pirenne (1970, 1975) proposes that the information from accommodation is unimportant in pictorial representation. However, he does not address the issue of pictures representing objects that are short distances from the viewer

⁴ A proprioceptive cue refers to a possible neural signal arising from sensory fibers in the muscles of the eye. These muscles are responsible for converging or diverging the eyes and the sensory fibers in them would be able to provide a neural signal as to the state of degree of contraction of the eyes. This signal could then provide depth information concerning the fixated object (Rock 1975)

Computer-graphic algorithms have been developed to reproduce the depth of field effects created by photographic imaging devices (Potmesil and Chakravarty 1982; Cook et al. 1984).

4.2 Vergence

An observer can view an object with both eyes and maintain a perception of a single image. This is because both eyes can be converged to the angle that causes the image of the object to fall on corresponding retinal points (same retinal positions, see Sect. 2.1) of both retinas (Fig. 1). If an object is positioned at a greater or a lesser distance, a single image can still be maintained, because the eyes diverge or converge to a new angle that keeps the image positioned on corresponding retinal points (Fig. 1). Vergence also occurs as a reflex during accommodation and is called accommodative vergence. Controlling for accommodation, vergence is an effective depth cue up to distances of six feet and inconsequential for distances past 30 feet (Rock 1975). Furthermore, vergence is a more effective depth cue than accommodation (Swenson 1932).

Because vergence provides information about depth for short viewer-object distances, the loss of this information during picture perception is an important factor to consider when creating a pictorial representation. Meinel (1973), Hochberg (1980), and Sedgwick (1980) point out that pictures do not require the observer to converge or diverge the eyes to objects of different represented distances. Therefore, any cue that vergence provides towards the perception of depth will be lost. This lost cue is most likely to be feedback from sensory fibers in the eye muscles that code for the degree of muscle contraction (Rock 1975). Other types of lost information have been proposed by Hochberg (1980) and Meinel (1973). They point out that not only is proprioceptive feedback lost but, because vergence is the same for all represented distances, the normal doubling of images (those objects projecting images to noncorresponding retinal points, see Sect. 2.1) is missing during picture viewing and should therefore provide information to the viewer that the image is flat (Hochberg 1980).⁵

⁵ This "doubling" occurs because objects that are at different relative distances than the object being fixated upon will create images that fall on non-corresponding points between the two retinas

4.3 Surface perception

A picture can be created using different kinds of media. Most of these media (e.g., paint and canvas, photographic emulsion, and computer-monitor screen) possess surface properties that can be distinguished from the image represented in the picture. For example, a painted picture has a distinguishable surface, in part, because it has the property of texture. Texture is present due to the uneven application of paint on the canvas or wood. Moreover, the canvas or wood itself also possesses a texture that contributes to the overall surface texture of the picture. To the extent that this surface is perceivable, it will be possible to distinguish it from the percept of the represented scene.

Numerous properties of a surface are important in determining its visual appearance (Gibson 1982). These include whether the picture radiates or reflects light, quality and quantity of illumination from an outside source, diffuse reflectivity, specular reflectivity, texture, and opacity (Gibson 1982; Foley et al. 1990). The visual response to these various properties has been studied (for example, Bartlett 1965; Brown and Mueller 1965; Graham and Brown 1965; Gibson 1982). However, the visual psychophysics of these properties are quite extensive and beyond the scope of this paper. What is important is whether the perception of the picture surface affects the perception of the represented scene and if so, in what way. In most cases, the surface of a picture is perceived along with the scene the picture is representing. To the extent this is true, the picture is limited. Therefore, the effects of surface perception (subsidiary surface awareness, see Pirenne 1970) on the perception of the represented image should be considered important in the context of realistic pictorial representation.

Detection of the picture's surface degrades the observer's percept of depth in the image (Evans 1959; Pirenne 1970; Haber 1980a, 1980b). For example, research has shown that when a real scene is given an artificial surface (sheet of cellophane placed between observer and scene) and other depth cues are eliminated (e.g., motion parallax), observers are unable to distinguish the scene from a 2-dimensional representation of the scene (Hochberg 1962). Other studies have shown that when surface information is absent in the represented scene, it appears more three dimensional. In fact, the perception of depth is likened to what is experienced under stereoscopic viewing (Helmholtz 1881; Pir-

enne 1970). Haber (1980a) points out that the surface of a picture “contributes massive amounts of information for flatness. The texture of the canvas or photographic paper can usually be seen, and it projects a zero gradient over the surface for the observer standing directly in front of the picture. If the picture is viewed from an angle, the gradient from the slant exactly matches that of the wall surface. Thus, the surface-perspective scale of space provides information about the spatial relations within the scene depicted by a picture and at the same time it specifies that the picture surface itself is flat” (Haber 1980a). Surface awareness also appears to affect the appearance of the represented scene’s shape.

It has been known for some time that a monocularly viewed picture seen from the wrong perspective does not appear distorted, although the theory of perspective drawing dictates that it should (Pirenne 1970; Cutting 1986). This paradox has been termed “La Gournerie’s paradox” (Cutting 1986). The viewer’s awareness of the picture surface has been suggested as the critical factor in maintaining the stable perception of the represented scene (Pirenne 1970). This hypothesis is supported by the fact that a photograph of a picture taken from the wrong viewing point results in a representation of that picture with perceptual deformations, presumably because the shape and position of the picture surface is no longer available to the viewer (Pirenne 1970).


4.4 Overcoming image flatness

The visual system uses static monocular depth cues to construct a perception of depth. These cues are interposition, relative size, relative height, atmospheric perspective, familiar size, texture gradient, shading, shadowing, and linear perspective (see Goldstein 1984). Artists and photographers intentionally introduce these cues into their pictures in order to enhance depth and overcome the inherent flatness of the image. For example, consider the collection of random blocks and spheres in the computer-generated image in Fig. 5. This picture was created using a simple illumination model and scan-line hidden-surface algorithm. A photographer might first arrange the objects in this picture so that they occlude one another (Fig. 6). This provides the cue of interposition, which gives the ob-

server knowledge about the relative ordering of the blocks and spheres in depth. Next the photographer could place the larger objects in front and smaller ones in the rear (Fig. 7). This amounts to a trick, because the actual size of the objects is unknown, but it is something that photographers do even when the objects can be identified (Evans 1959). This rearrangement exploits the cue of relative size, which is the effect perspective has on the size of distant objects. Also, the objects might be arranged so as to recede to a vanishing point, as in Fig. 8. This exaggerates the cue of linear perspective (for a thorough review of perspective see Hagen 1986).

Finally, it should be noted that in Figs. 5–8 shading and/or shadows have been introduced. These are important depth cues that, along with stereopsis, interposition, texture gradient, and linear perspective, are critically dependent on luminance contrast information. Livingstone and Hubel (1987, 1988) found that perceived depth in a picture was lost when any one of these cues was depicted by equiluminance chromatic contrast instead of luminance contrast. Moreover, loss of depth occurred within a range of luminance contrasts around equiluminance. This suggests that these depth cues are most effective when high luminance contrast is used. However, further research is needed to demonstrate quantitatively how perceived depth varies as a function of the cue’s luminance contrast. For example, it would be important to know the minimum luminance contrast necessary for the cue to remain an effective indicator of depth.


In addition to using static monocular depth cues, other techniques have been developed to help overcome image flatness (see Meinel 1973; Hochberg 1979). For example, Hochberg (1979, 1980) suggests that painters (e.g., Rembrandt) have developed a solution to offset the vergence problem by selecting a few areas in the real scene that have little depth as “focal regions.” These regions are then painted with high detail. Areas outside these regions contain a higher amount of depth information and are painted with large swatches of paint that provide little detail. Objects in these areas will only look normal and recognizable when viewed with the periphery of the eye due to the periphery’s low spatial acuity. If viewed with the fovea, they will appear as blurred, sketchy, blobs faintly resembling the objects they are meant to represent. As a consequence of this style of painting, the objects depicted only look normal when the viewer main-




3




4




5



6




7



8



9



10



11

Fig. 3. Boundary of picture causes book to appear as if it is tipping when the image plane is kept perpendicular to the line of sight [Figs. 3 and 4 after Upton and Upton (1981)]

Fig. 4. Making image plane parallel to the front of the book solves the tipping problem in Fig. 3, but causes the sphere to become distorted

Fig. 5. Original random arrangement of blocks minimizes sensation of depth in picture [Figs. 5–8 after Evans (1959)]

Fig. 6. Blocks rearranged to occlude one another and thereby make it easier to determine which block is in front of another

Fig. 7. Blocks organized by size to take advantage of the change in size with depth. This is in addition to being positioned so as to occlude one another

Fig. 8. Perspective convergence used in addition to size and occlusion to enhance depth

Fig. 9. Wall in back of table appears dark, because image lacks dynamic range [Figs. 9–11 after Evans (1948)]

Fig. 10. Increasing exposure lightens back wall, but over exposes other portions of the image

Fig. 11. Adding fill light to illuminate back wall distorts original lighting, but produces correct percept in final image

tains his fovea on the focal regions and periphery on the non-focal regions. The restriction of the viewer's gaze to only one or two areas in which the image appears normal and recognizable should therefore limit the flatness information to the observer (Hochberg 1980).

5 Pictures have a limited dynamic range

There is a limit to the maximum amount of light that the phosphors of a television monitor can emit or to the minimum density that the dyes of a photographic film can obtain. For each reproduction medium, the difference between the amount of light emitted (or reflected) by the blackest possible black and the whitest possible white is therefore bounded. This difference in light intensity is referred to as the dynamic range of the device. The human visual system is capable of operating over a much wider dynamic range than can be reproduced on any currently available display device. It accomplishes this by adapting to the average brightness level present in the scene. This means that no one reproduction medium can create the full range of light intensities over which the visual system operates, from the bright light of the mid-day sun to the dim light of the moonlit sky. Fortunately, brightness adaptation and brightness perception allow the visual system to adjust to the dynamic range available with a particular display device. There are differences that remain, however, between the percept created by the original range of intensities in the scene and the range that are available on the display device.

5.1 Brightness adaptation and perception

Before discussing luminance range, some terminology will first be established. A physical measure of light is radiance. Radiance is the amount of absolute electromagnetic energy emitted from or reflected off of an object. If the light provides a visual stimulus, this measure can be misleading, because the visual system is differentially sensitive to each of the wavelengths of light. Luminance is a perceptual measure of light that takes this differential sensitivity into consideration and thereby indicates the

effectiveness of light as a visual stimulus. Luminance is therefore used to measure the intensity of light independent of wavelength composition as experienced by a viewer under real and represented scene viewing (Riggs 1965).

Characteristics of the visual system change depending on the luminance value of the scene. For example, acuity, contrast sensitivity, and hue discrimination are dependent on the adapted state of the eye, which in turn is dependent on the level of luminance in the scene (other factors such as time of exposure and pre-adapted state of the eye are also critical factors in affecting these behaviors; see Bartlett 1965). Differences in brightness-magnitude estimation are also dependent on the level of stimulus luminance (Goldstein 1984). The potential effect that luminance can have on visual response becomes apparent when one considers the vast range of luminance values found for typical visual stimuli. According to Riggs (1965), visual stimuli in the real scene can vary from 10^9 millilamberts (sun's surface at noon) to 10^{-6} millilamberts (white paper in starlight).

Based on the above studies (see Brown and Mueller 1965; De Valois and De Valois 1988), the average level of luminance of the real scene and the individual level of luminances for objects in the real scene is an important factor in determining visual responses. Therefore, if pictures are to illicit the same visual response as the scene they represent, it is important to consider the extent to which pictures are limited in replicating identical luminance levels as the scene and to consider what changes occur in the visual response as a result of this limitation. Moreover, these changes are of particular importance with those pictures that attempt to represent scenes possessing luminance levels far outside the luminance range of the picture (e.g., bright daylight scenes and dark night scenes).

The range of brightnesses in a real scene can potentially vary by a factor of many hundreds to one (Helmholtz 1881; Hochberg 1979). However, the maximum luminance range of surface pigments is approximately 40-1 (Hochberg 1979). Other media, such as photographs, have higher ranges, but not to the extent of ranges found in the real scene. [Evans (1959) calculated the maximum range of a photograph to be approximately 300-1.] The inability of these media to meet the range of luminances in a given scene is further compounded by their inability to match the overall luminance of the scene (Helmholtz 1881; Gombrich 1956). For

example, this is true when night scenes are depicted by media dependent on light reflection (e.g., paintings). Luminous media, such as computer monitors, are not as limited in this regard. Indeed, if the brightness setting of the monitor is set sufficiently low, the overall luminance of the representation can be equated with the night scene. On the other hand, the monitor is limited if the mean luminance of the scene is too high. This can be corrected if light is added to the monitor (e.g., shining it on the screen). However, this decreases brightness contrast. As a result of these restrictions in luminance range and overall luminance level, pictures are limited in their capacity to elicit the same visual response as the scene they represent (Helmholtz 1881; Evans 1959; Gombrich 1956; Pirenne 1970, 1975; Hochberg 1978a, 1979; Haber 1980b; Mills 1985; Cutting 1986).

One consequence of this limitation is that most pictures reflect light at a luminance level that keeps the eye modestly light adapted (Helmholtz 1881). Therefore, most pictures that attempt to replicate outdoor scenes (especially bright, outdoor, daylight scenes or dark, moonlit, night scenes) will not be viewed with the same adapted state of the eye found during real-scene viewing (Helmholtz 1881; Evans 1959; Pirenne 1970; Hochberg 1979; Mills 1985). For example, Helmholtz (1881) points out that a painter attempting to paint a white object illuminated by the sun compared to a white object illuminated by the moon usually needs to use a pigment that has approximately the same reflectance for representing both objects. Furthermore, both represented objects are often viewed under the same light level. Therefore, the adapted state of the eye is approximately constant across both represented scenes. However, in the real scene, the white object in sunlight is approximately 100 million times brighter than in moonlight (Helmholtz 1881; Riggs 1965). During real-scene viewing then, the eye is extremely light adapted in the sunlit scene and extremely dark adapted in the moonlit scene.

This difference in behavior between the light-adapted and dark-adapted eye has been accommodated in the design of photographic films producing pictures to be viewed in either bright (reflection prints) or dark (transmission slides) environments (Bartelson and Breneman 1967). If the general level of illumination under which the reproduction is viewed matches the level of illumination surrounding the actual scene, it has been found that the

luminances of the reproduction should be linearly related to the luminance of the original scene by a simple scaling factor. However, if the level of illumination differs between the original and reproduction viewing conditions, optimum subjective appearance is produced when the relationship between the luminances is nonlinear (Bartelson and Breneman 1967; Mees and James 1966). For example, if the original scene is viewed under bright conditions and the reproduction is observed in dim surroundings, the line relating the logarithm of the scene luminance to the logarithm of the reproduction luminance should have a slope greater than one. The slope of this line is often referred to as the system gamma.

Another consequence of a limited luminance range in pictures is the lack of color assimilation effects occurring under picture viewing (see Graham and Brown 1965). During real-scene viewing, if factors, such as contrast, saturation, and light level, are sufficient and the spatial frequency of surround and target is high, then the target will tend to appear the same color as the surroundings (see Graham and Brown 1965; Goldstein 1984). For example, a small dark shadow surrounded by bright yellow sunlight should appear slightly yellowish in the real scene even though the shadow is not projecting long wavelength light. However, a picture that represents bright yellow sunlight surrounding a dark shadow would not possess as high a luminance range and would therefore be unable to induce as great an assimilation effect (Graham and Brown 1965).

5.2 Overcoming limited dynamic range

Helmholtz (1881) points out that a painter needs to consider the different physiological conditions of the eye present during real-scene viewing (e.g., low visual acuity due to a dark adapted eye) and then “translate” these subjective phenomenon into the painting itself. Gombrich (1956) also agrees that the artist or photographer must attempt to suggest the presence of light (or the lack of it) in the picture by painting in the physiological reactions the observer naturally experiences under real-scene viewing. An example of this can be seen in Monet’s attempts to mimic the visual response of looking at a church with a light-insensitive eye (Monet: Rouen Cathedral, west facade, sunlight, see Mills 1985).

Hochberg (1979) suggests how simultaneous contrast can be used to construct pictures that simulate the perceptual response of the visual system. For example, he points out that early Impressionists, such as Corot, painted in color-contrast effects in order to simulate the effects of saturated colors on a light-adapted eye in a brightly lit scene. Furthermore, he suggests that artists, such as Rembrandt, Eakins, and Seurat, attempted to offset the limited luminance range in pictures by representing objects that have luminance levels outside this range (e.g., bright shiny highlights from light reflected off of a gold braid) using large swatches of light and dark in the outside regions of the picture. This was done in order to take advantage of simultaneous contrast effects (large dark regions surrounding bright regions cause the bright regions to appear brighter) that occur with low spatial frequency stimuli in the periphery of the eye.

Hochberg (1979) also proposes that the large swatches of light and dark employed by Rembrandt, Eakins, and Seurat take advantage of successive contrast effects that not only increase the perception of brightness, but the perception of saturation as well. He points out that an area of the retina stimulated by a dark region in the painting is somewhat dark adapted. Therefore, when a light region of the painting falls on this dark-adapted area as a result of minor eye movements, it appears brighter due to the somewhat increased sensitivity of that area of the eye (Hochberg 1979). In the case of increasing saturation, when a colored region stimulates a specific area of the retina, an afterimage is produced that is the complementary color of that region. If another colored region that is the complement of the first region were then to stimulate this same area of retina (as a result of minor eye movements), then the color of this second region should appear more saturated (Hochberg 1979; Goldstein 1984).

Photographers have also developed techniques to help them overcome the limited dynamic range available with their medium. Shadows are a particular problem, because our visual system is capable of seeing detail in real shadows, but a photograph doesn't have sufficient dynamic range to reproduce this detail. Photographers therefore try to flatten the lighting in a scene to eliminate deep shadows. The result is that the lighting on a television or a movie set does not appear correct when viewed on location, but looks correct when seen on a television monitor or in a movie theater. The same


problem exists in computer graphics as is illustrated by the synthetic images in Figs. 9 to 11. These pictures were produced using a radiometrically correct illumination model (Ward et al. 1988). In Fig. 9, the rear wall appears much darker than it would in the original scene (if we could be there). Trying to fix this by changing the exposure results in the overexposed picture in Fig. 10. The correct solution is to "distort" the original lighting by adding some additional illumination to the back wall so as to create a version that looks correct when observed in the final picture (Fig. 11).

6 Pictures are seen in a viewing illuminant


The effect of viewing illuminant on brightness adaptation has already been discussed in the preceding section about the limited dynamic range of pictures. In addition to affecting the level of illumination to which the visual system is adjusted, the viewing illuminant also has an impact on the color of the light to which the visual system is adapted. As in the case of brightness adaptation, it is possible for the color of the viewing illuminant to be different between the original scene and the reproduction. This process by which our visual system is able to discount the color of the illuminant and see the true color of objects in the environment is known as chromatic adaptation.

6.1 Chromatic adaptation

When the eye is moderately illuminated by a uniform colored light source, the true color of the light is initially perceived. In a short time, however, the eye adapts to the color and the light is accepted as white (as long as saturation is not too great, Evans 1959). This is called chromatic adaptation. Therefore, when an observer views an object that is illuminated by colored light, it first appears colored differently than its appearance under white light. Eventually, as the eye adapts to the color of the illuminant, the object color appears more like that found under white-light illumination. For example, Evans (1959) points out that for the adapted observer a white piece of paper viewed




12



13



14



15

Fig. 12. Torus with high specular reflectance appears shiny when illuminated by a single light source

Fig. 13. Same torus as in Fig. 12 looks dull when illuminated by four light sources

Fig. 14. Straightforward application of the laws of color science to produce an image of a scene illuminated by tungsten light

Fig. 15. Simulation of how tungsten-balanced film shifts the color balance to correctly reproduce a scene illuminated by tungsten light

under yellow incandescent light appears approximately the same color as an identical piece of paper viewed under white light. However, it should be noted that an observer, if asked to do so, can often perceive the color of the illuminant following adaptation (e.g., they can see the yellowishness of a paper illuminated by a yellow incandescent light or the bluishness when illuminated by skylight, Evans 1959).

Under normal photopic conditions, the laws of metameric color matching continue to hold even for a new state of chromatic adaptation (Jameson and Hurvich 1972). The traditional explanation for this phenomena is that the spectral quality of the illuminant produces a differential scaling of the three spectral sensitivity functions. This is called the von Kries coefficient law. There are several criticisms of the von Kries law, the principal one being that the proportionality rule of color matching is violated as the level of illumination varies (Jameson and Hurvich 1972). It is felt by some (Jameson and Hurvich 1972) that chromatic adaptation could be bet-

ter understood by considering changes in the opponent spectral sensitivity functions instead of the fundamental spectral sensitivity functions. This idea is supported by a recent computational approach to color-vision modelling (D'Zmura and Lennie 1986). The determination of the spectral power distribution of the illuminant and the spectral reflectances of the surfaces is the key in this and any other model of color adaptation. While a specific biological mechanism that could do this has not been found, several computational techniques (Buchsbaum 1980; Maloney and Wandell 1986) have shown that it is theoretically possible to determine these quantities from the spectral distribution of the light that reaches the eye.

As a result of chromatic adaptation, objects illuminated with colored light appear similar to that found under white-light illumination. However, other factors, such as the test reflectance, background reflectance, and size of the background (Graham and Brown 1965) influence chromatic adaptation affects. These factors vary depending on

the type of scene in which the object is viewed. As a result, the color of the object often appears different under real-scene viewing compared with its appearance under picture viewing. Therefore, to the extent that chromatic adaptation differs between these two types of viewing, pictures are limited as realistic representations of real scenes (Helmholtz 1881; Gombrich 1956).

Evans (1959) agrees that different chromatic adaptation effects are a limitation for pictures and discusses these limitations in terms of photographs. For example, he discusses a situation in which a photograph is taken of two adjacent pieces of paper illuminated by a yellow incandescent light. One piece is white with high reflectance and the other is grey with low reflectance. Chromatic adaptation to the yellow illuminant would cause the two pieces of paper to lose most of their yellowish appearance. However, the differential reflectance between the two pieces of paper would dictate the grey paper to appear less yellowish than the white paper. A photograph of these two pieces of paper could be made so that the white paper is reproduced as yellowish, but the grey would need to be reproduced as more neutral. Evans (1959) points out that this is not ordinarily possible.

According to Evans (1959), "The requirement for satisfactory reproduction of a scene is that the photograph under the condition of viewing shows the proper hue and saturation differences from the *adaptation state of the observer*." Therefore, the hue and saturation of colors in a picture should be dependent on, among other things, the chromatic adapted state of the observer. According to this then, alterations in colors of objects represented in the picture are needed for each type of viewing condition (at least those that alter the chromatic adapted state of the observer) if a person wants to maintain color response fidelity.

6.2 Overcoming effects of viewing illuminant

Figures 12 and 13 illustrate the difficulties that can be encountered due to differences between the illumination simulated in the computer-generated picture and the illumination under which the synthetic image is viewed. The torus in Fig. 12 appears to have a shiny surface, while the torus in Fig. 13 looks as if it has a matte finish. In actuality, the same coefficients are used in the illumination model

governing the amount of diffuse and specular reflection for each of the tori. It is the difference in illumination between the two pictures that causes the change in surface appearance. We are unable to recognize there is only one light source in Fig. 12 and four light sources in Fig. 13, because our viewing illumination is constant and different from that in either picture. We are also at a disadvantage, because none of the illumination from the lights in the picture spills out into the room as it would if the picture plane were a real window.

A solution to the chromatic adaptation problem has been developed in photography. Two different types of film are employed: one for use outdoors in daylight and the other for use indoors with tungsten light. Figures 14 and 15 are a computer-graphic simulation of the chromatic-adaptation problem and how tungsten-balanced film solves the problem for indoor scenes. In Fig. 14, a picture has been made by a straightforward application of the laws of color science (Meyer 1989). Note the very yellow appearance this image has. In Fig. 15, the effect of tungsten-balanced film is simulated. The overall yellow color of the image has disappeared. Thus, the film concentrates on recording the correct visual percept instead of trying to capture an optical identity.

7 Summary and conclusions

While it is possible to synthesize a computer-graphic image that is optically identical to a real scene, it is impossible to display this picture given the color-reproduction technology available today. As has been shown in this article, pictures have limitations in terms of their physical size, fixed-point perspective, dynamic range, and circumstances under which they are observed. These limitations restrict the extent to which pictures can reproduce the cues that our visual system uses to interpret three-dimensional scenes. These cues include the differences between the images projected onto each retina, amount by which each eye must be adjusted to focus on the center of attention, and level of adaptation necessary to adjust to the level and color of illumination present in the scene.

Artists and photographers have developed many techniques for overcoming the limitations of pictorial media. In the absence of binocular cues to depth, they have made sure that monocular cues

are properly rendered. They have employed view cameras to adjust perspective in a way that limits the effect of the frame surrounding the picture. Composition has been carefully considered to emphasize such things as object occlusion and the effect of perspective on object size. They have intentionally painted color contrasts into pictures in order to overcome the limited dynamic range provided by pigments. Individuals involved in the synthesis of realistic images can learn from the techniques developed by artists and photographers. An optically identical representation for a scene cannot be displayed using existing color-reproduction technologies. The reproduction must be manipulated and, in many cases, the original scene must be changed so that the correct perception is produced for the observer.

While employing artistic and photographic techniques can improve the quality of today's photorealistic computer-graphic images, future image synthesis algorithms should incorporate the characteristics of the human visual system directly into the image creation process. With the flexibility that computer graphics provides, it makes no sense to slavishly simulate a camera and continue to reproduce imaging problems that the camera creates or is unable by itself to eliminate. By determining the relative priority of each visual cue and by developing an initial image representation that is perceptual in nature, efficiency can be improved, because computational effort is directed toward those things that are perceptually important, and device independence can be achieved, because the image has not been created for a particular reproduction device. Perhaps most importantly, this approach concentrates attention on the human being for whom the picture is being made and away from the computational physics that has recently dominated computer graphics. This change of focus is critical if we are to eventually synthesize pictures that are not just realistic, but also communicate information to people.

Acknowledgements. Chet Haase wrote the software and created the images shown in Figs. 3 and 4. Darren Anderson generated Figs. 5–15. He used the Radiance software package to produce some of them (Ward et al. 1988). This work was funded by the National Science Foundation under grant number CCR 90-08445 and by the National Institute of Health Systems Physiology Training Program grant number GM070257.

References

- Bartelson CJ, Breneman EJ (1967) Brightness reproduction in the photographic process. *Photogr Sci Eng* 11:254–262
- Bartlett NR (1965) Dark adaptation and light adaptation. In: Graham CH, Bartlett NR, Brown JL, Hsia Y, Mueller CG, Riggs LA (eds) *Vision and visual perception*. John Wiley, New York
- Brown JL, Mueller CG (1965) Brightness discrimination and brightness contrast. In: Graham CH, Bartlett NR, Brown JL, Hsia Y, Mueller CG, Riggs LA (eds) *Vision and visual perception*. John Wiley, New York
- Buchsbaum G (1980) A spatial processor model for object colour perception. *J Franklin Inst* 310:1
- Chung JC, Harris MR, Brooks FP, Fuchs H, Kelley MT, Hughes J, Ouh-young M, Cheung C, Holloway RL, Pique M (1989) Exploring virtual worlds with head-mounted displays, three-dimensional visualization and display technologies. In: Robbins WE, Fisher SS (eds) *Proc SPIE 1083:42–52*
- Cook RL, Porter T, Carpenter L (1984) Distributed ray tracing. *Comput Graph* 18:137–145
- Cutting JE (1986) *Perception with an eye for motion*. MIT Press, pp 1–40
- Darley JM, Glucksberg S, Kinchla RA (1986) *Psychology*. Prentice-Hall, Englewood Cliffs, p 86
- De Valois RL, De Valois KK (1988) *Spatial vision*. Oxford University Press, New York, pp 27–31
- D'Zmura M, Lennie P (1986) Mechanisms of color constancy. *J Optic Soc Am [A]* 3:1662–1672
- Evans RM (1948) *An introduction to color*. John Wiley, New York
- Evans RM (1959) *Eye, film and camera in color photography*. John Wiley, New York
- Foley JD, van Dam A, Feiner SK, Hughes JF (1990) *Computer graphics: principles and practice*. Addison-Wesley, Massachusetts
- Gibson JJ (1982) *Reasons for realism: selected essays of James J. Gibson*. Reed E, Jones R (eds) Lawrence Erlbaum Associates, London
- Goldstein B (1984) *Sensation and perception*. Wadsworth, California, pp 203–287
- Gombrich EH (1956) *Art and illusion*. Princeton Press, Princeton
- Graham CH (1965) Visual space perception. In: Graham CH, Bartlett NR, Brown JL, Hsia Y, Mueller CG, Riggs LA (eds) *Vision and visual perception*. John Wiley, New York
- Graham CH, Brown L (1965) Color contrast and color appearances: brightness constancy and color constancy. In: Graham CH, Bartlett NR, Brown JL, Hsia Y, Mueller CG, Riggs LA (eds) *Vision and visual perception*. John Wiley, New York
- Haber RN (1980a) How we perceive depth from flat pictures. *Am Sci* 68:370–380
- Haber RN (1980b) Perceiving space from pictures: a theoretical analysis. Hagen M (ed) In: *The perception of pictures*. Springer, New York
- Hagen MA (1986) *Varieties of realism: geometries of representational art*. Cambridge University Press, New York
- Hagen MA (1991) How to make a visually realistic 3D display. *Computer Graphics* 25:76–81
- Hall R (1989) *Illumination and color in computer generated imagery*. Springer, New York
- Hatada T, Sakata H, Kusak H (1980) *Psychophysical analysis*

- of the "sensation of reality" induced by a visual wide-field display. *SMPTE J* 89:560-569
- Helmholtz H von (1881) On the relation of optics to painting. In: Popular scientific lectures (translated by E Atkinson). Appleton, New York
- Hochberg JE (1962) The psychophysics of pictorial perception. *Audio-Visual Comm Rev* 10:22-54
- Hochberg J (1978a) Art and perception. In: Carterette EC, Friedman MP (eds) *Handbook of perception*, vol 10. Academic Press, New York, 10:225-259
- Hochberg JE (1978b) *Perception* (2nd edn). Prentice Hall, New York
- Hochberg J (1979) Some of the things that paintings are. In: Nodine CF, Fisher DF (eds) *Perception and pictorial representation*. Praeger, New York
- Hochberg J (1980) Pictorial functions and perceptual structures. In: Hagen M (ed) *The perception of pictures*. Springer, New York
- Jameson D, Hurvich LM (1972) Color adaptation: sensitivity, contrast, and afterimages. In: Jameson D, Hurvich LM (eds) *Handbook of sensory physiology*, (vol. 7, part 4). Springer, Berlin Heidelberg New York
- Langlands HMS (1926) Experiments in binocular vision. *Trans Opt Soc (London)* 28:45-82
- Livingstone M, Hubel H (1987) Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *J Neurosci* 7(11):3416-3468
- Livingstone M, Hubel H (1988) Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science* 240:740-749
- Maloney LT, Wandell A (1986) Color constancy: a method for recovering surface spectral reflectance. *J Opt Soc Am [A]* 3:29-33
- Mees GEK, James TH (1966) *Theory of the photographic process* (3rd edn). Macmillan, New York
- Meinel E (1973) Peripheral vision and painting. *Br J Aesthetics* 13(3):287-297
- Meyer GW (1989) Reproducing and synthesizing colour in computer graphics. *Displays: technology and applications* 10:161-170
- Meyer GW, Rushmeier HE, Cohen MF, Greenberg DP, Torrance KE (1986) An experimental evaluation of computer graphics imagery. *ACM Trans Graph* 5:30-50
- Mills MI (1985) Image synthesis, optical identity or pictorial communication. In: Magnenat-Thalmann N, Thalmann D (eds) *Computer generated images, the state of the art*. Springer, New York
- Pirenne MH (1970) *Optics, painting and photography*. Cambridge University press, London
- Pirenne MH (1975) Vision and art. In: Carterette EC, Friedman MP (eds) *Handbook of perception* (vol 5). Academic Press, New York
- Potmesil M, Chakravarty I (1982) Synthetic image generation with a lens and aperture camera model. *ACM Trans* 1:85-108
- Riggs LA (1965) Light as a stimulus for vision. In: Graham CH, Bartlett NR, Brown JL, Hsia Y, Mueller CG, Riggs LA (eds) *Vision and visual perception*. John Wiley, New York
- Rock I (1975) An introduction to perception. Macmillan, New York, pp 79-153
- Sedgwick H (1980) The geometry of spatial layout in pictorial representation. In: Hagen M (ed) *The perception of pictures*. Springer, New York
- Sekuler R, Blake R (1990) *Perception*. McGraw-Hill, New York
- Smith OW, Gruber H (1958) Perception of depth in photographs. *Am J Psych* 8:307-313
- Stone J (1987) *A user's guide to the view camera*. Little Brown, Boston
- Swenson HA (1932) The relative influence of accommodation and convergence in the judgement of distance. *J G Psych* 7:360-380
- Upton BL, Upton J (1981) *Photography*. Little Brown, Boston
- Van de Grind WA (1986) Vision and the graphical simulation of spatial structure. *Proc ACM Interactive 3D Graphics*, pp 197-235
- Westerink JHDM, Roufs JAJ (1989) Subjective image quality as a function of viewing distance, resolution, and picture size. *SMPTE J* 98(2):113-119
- Witkin HA (1959) The perception of the upright. *Sci Am* 200(2):50-56



CHRISTOPHER G. BARBOUR is a doctoral candidate in the psychology department at the university of Oregon. He received his BS from the University of California at Davis in 1986 and his MS in cognitive psychology at the University of Oregon in 1988. He is interested in visual psychophysics with special interest in visual information channels, neural coding, flicker, and pattern detection. He is a student affiliate of the Human Factors Society.



GARY W. MEYER is an associate professor in the Department of Computer and Information Science at the University of Oregon. His research interests include color reproduction and color selection for the human-computer interface, perceptual issues related to synthetic image generation, and the application of computer graphics to scientific computing. Meyer has been a member of the technical staff at Bell Telephone Laboratories. He received a BS from the University

of Michigan, an MS from Stanford University, and a PhD from Cornell University. He is a member of ACM SIGGRAPH, IEEE Computer Society, SMPTE, and OSA.