Vision-realistic Rendering: Simulation of the Scanned Foveal Image with Elimination of Artifacts due to Occlusion and Discretization

Brian A. Barsky

Computer Science Division and School of Optometry
University of California, Berkeley, California, 94720-1776, U.S.A.
`barsky@cs.berkeley.edu`
`http://www.cs.berkeley.edu/~barsky`

Abstract. We introduce the concept of **vision-realistic rendering** – the computer generation of synthetic images that incorporate the characteristics of a particular individual's entire optical system. Specifically, this paper develops a method for simulating the scanned foveal image from wavefront data of actual human subjects, and demonstrates those methods on sample images.

First, a subject's optical system is measured by a Shack-Hartmann wavefront aberrometry device. This device outputs a measured wavefront which is sampled to calculate an object space point spread function (OSPSF). The OSPSF is then used to blur input images. This blurring is accomplished by creating a set of depth images, convolving them with the OSPSF, and finally compositing to form a vision-realistic rendered image.

Although processing in image space allows an increase in speed, the images may have artifacts introduced due to *occlusion* or *discretization*. Two approaches for object identification to properly blur the scene are discussed.

Applications of vision-realistic rendering in computer graphics as well as in optometry and ophthalmology are discussed.

keywords vision-realistic rendering, optics, ray tracing, image synthesis, human visual system, blur, optometry, ophthalmology, LASIK, pupil, Point Spread Function (PSF),

1 Introduction

After the development of the fundamentals of raster graphics in the 1970's, advances in computer graphics in the 1980's were dominated by the quest for *photorealistic rendering*, and attention turned to *non-photorealistic rendering* in the 1990's. For the first decade of the the 21st century, we propose to extend this to *vision-realistic rendering* (VRR). VRR is the simulation of the actual human vision of a particular subject – not merely a model of vision, but the generation of images that incorporate characteristics of a particular individual's optical system.

Such an ambitious undertaking will require the integration of many fields of study, beyond the traditional ones for computer graphics, such as physiological optics, human visual perception, psychophysics of vision, visual neurophysiology, human color vision, binocular vision, visual sensory mechanisms, etc.

Fig. 1. Vision-realistic rendered image simulating vision based on actual wavefront data from a patient with keratoconus.

To embark upon this endeavor, we begin with the problem of simulating the retinal image formed by the optics of the eye. Since the goal is to simulate vision of an actual human subject, not just to use a model, we need data about the optics of the subject's eye.

Ideally, we would like to have data about the aberrations of these optics for each photoreceptor across the retina. That is, given a gaze direction, we would like to trace a ray from each photoreceptor, through the optical structures of the eye such as the internal crystalline lens and cornea, out into the environment, and measure the aberrations for that ray, including whatever would arise given the current accommodative state of the crystalline lens of the eye.

Unfortunately, such capability does not exist at the present time. That deficiency notwithstanding, it is exciting to note that we can achieve an approximation to such measurements using recently-developed technology motivated by the goal of improving laser corneal photoreactive vision correction surgeries such as LASIK (laser in-situ keratomileusis). This technology is *wavefront aberrometry*, that is, instruments that measure the wavefront emerging from the eye and quantify the amount of each different kind of optical aberration present. The limitation is that the instrument does so at only one point in the retina.

However, that limitation is not nearly as much of a problem as it may seem at "first glance". The reason is that the arrangement of photoreceptors on the retina is not at all the uniform structure that we are used to in raster graphics where pixels are arranged in neat, rectangular arrays. Rather, the cones are densely packed in a small area in the middle of the cornea, called the fovea, and are much more sparsely arranged towards the periphery. The fovea is approximately 600 microns wide and subtends an angle of view of about two degrees. When one looks at an object, the eye is oriented such that light comes to a focus in this foveal region of the retina. Consequently, if we use wavefront

aberrometry to measure the aberrations present for vision at a point in this foveal region, we will have a reasonable first approximation to the image perceived by the subject.

The reason that this approximation works so well is that when looking at a scene, a viewer naturally and involuntarily quickly scans around the scene at different objects. At any instant, the viewer is focused at only one object, using high resolution foveal vision. However, by scanning around the scene, the viewer gains the misleading impression that the entire scene has been viewed in this high resolution vision. But at any instant, in fact, it is only the object in the center of visual field that is seen in high resolution. The periphery of the scene is really being viewed in much lower resolution vision, even though that is not evident.

Ergo, our approach is to obtain the wavefront aberrometry from a point in the fovea, and then to simulate the vision as if the aberrations were constant across the visual field.

This paper describes a pipeline to simulate the scanned foveal image from wavefront data of actual human subjects, and shows some example images. These are the first images in computer graphics that are generated on the basis of the specific optical characteristics of actual individuals.

2   Optometry and Ophthalmology Motivation

In practice poor visual performance is often attributed to simple blur; however, our technique enables the generation of vision-realistic rendered images and animations that demonstrate specific defects in how a person sees. Such images of simulated vision could be shown to an individual's eye care clinician to convey the specific visual anomalies of the patient. Doctors and patients could be educated about particular vision disorders by viewing images that are generated using the optics of various ophthalmic conditions such as *keratoconus* (Figure 1) and *monocular diplopia*.

One of the most compelling applications is in the context of vision correction using laser corneal refractive eye surgeries such as PRK (photorefractive keratectomy) and LASIK (laser in-situ keratomileusis). Currently, in the United States alone, a million people per year choose to undergo this elective surgery. By measuring subjects pre-operatively and post-operatively, our technique could be used to convey to doctors what the vision of a patient is like before and after surgery (Figures 20 and 21). In addition, accurate and revealing medical visualizations of predicted visual acuity and of simulated vision could be provided by using modeled or adjusted wavefront measurements. Potential candidates for such surgery could view these images to enable them to make more educated decisions regarding the procedure. Still another application would be to show such candidates some of the possible visual anomalies that could arise from the surgery, such as glare at night. With the increasing popularity of these surgeries, perhaps the current procedure which has patients sign a consent form that can be difficult for a layperson to understand fully could be supplemented by the viewing of a computer-generated animation of simulated vision showing the possible visual problems that could be engendered by the surgery.

3   Previous and Related Work

For a discussion of camera models and optical systems used in computer graphics, the reader is referred to a pair of papers by Barsky et al. where the techniques have been separated into object space [1] and image space [2] techniques.

The first synthetic images with depth of field were computed by Potmesil and Chakravarty [3] who convolved images with depth-based blur filters. However, they ignored issues relating to occlusion, which Shinya [4] subsequently addressed using a ray distribution buffer. Rokita [5] achieved depth of field at rates suitable for virtual reality applications by repeated convolution with 3×3 filters and also provided a survey of depth of field techniques [6]. Although we are also convolving images with blur filters that vary with depth, our filters encode the effects of the entire optical system, not just depth of field. Furthermore, since our input consists of two-dimensional images, we do not have the luxury of a ray distribution buffer. Consequently, we handle the occlusion problem by the techniques described in Section 6 and in more detail by Barsky et al. [7] [8].

Stochastic sampling techniques were used to generate images with depth of field as well as motion blur by Cook et al. [9], Dippe and Wold [10], and Lee et al. [11]. More recently, Kolb et al. [12] described a more complete camera lens model that addresses both the geometry and radiometry of image formation. We also use stochastic sampling techniques for the construction of our OSPSF.

Loos et al. [13] used wavefront tracing to solve an optimization problem in the construction of progressive lenses. They also generated images of three dimensional scenes as viewed through a simple model eye both with and without progressive lenses. However, while we render our images with one point of focus, they chose to change the accommodation of the virtual eye for each pixel to "visualize the effect of the lens over the full field of view" [13]. Furthermore, our work does not rely on a model of the human optical system, but instead uses actual patient data in the rendering process.

Light field rendering [14] and lumigraph systems [15] were introduced in 1996. These techniques represent light rays as a pair of interactions of two parallel planes. This representation is a reduction of the plenoptic function, introduced by Adelson and Bergen [16]. The algorithms take a series of input images and construct the scene as a 4D light field. New images are generated by projecting the light field to the image plane. Although realistic object space techniques consume a large amount of time, Heidrich et al. [17] used light fields to describe an image-based model for realistic lens systems that could attain interactive rates by performing a series of hardware accelerated perspective projections. Isaksen et al. [18] modeled depth of field effects using dynamically reparameterized light fields. We also use an image-based technique, but do not use light fields in our formulation.

There is a significant and somewhat untapped potential for research that addresses the role of the human visual system in computer graphics. One of the earliest contributions, Upstill's Ph.D. dissertation [19], considered the problem of viewing synthetic images on a CRT and derived post-processing techniques for improved display. Spencer et al. [20] investigated image-based techniques of adding simple ocular and camera effects such as glare, bloom, and lenticular halo. Bolin and Meyer [21] used a perceptually-based sampling algorithm to monitor images as they are being rendered

for artifacts that require a change in rendering technique. [22–26] and others have studied the problem of mapping radiance values to the tiny fixed range supported by display devices. They have described a variety of tone reproduction operators, from entirely ad hoc to perceptually based. For a further comparison of tone mapping techniques, the reader is referred to [27]. Meyer and Greenberg [28] presented a color space defined by the fundamental spectral sensitivity functions of the human visual system. They used this color space to modify a full color image to represent a color-deficient view of the scene. Meyer [29] discusses the first two stages (fundamental spectral sensitivities and opponent processing) of the human color vision system from a signal processing point of view and shows how to improve the synthesis of realistic images by exploiting these portions of the visual pathway. Pellacini et al. [30] developed a psychophysically-based light reflection model through experimental studies of surface gloss perception. Much of this work has focused on human visual perception and perceived phenomena; however, our work focuses exclusively on the human optical system and attempts to create images like those produced on the retina. Perceptual considerations are beyond the scope of this paper.

In human vision research, most simulations of vision [31, 32] have been done by artist renditions and physical means, not by computer graphics. For example, Fine and Rubin [33, 34] simulated a cataract using frosted acetate to reduce image contrast. With the advent of instruments to measure corneal topography and compute accurate corneal reconstruction, several vision science researchers have produced computer-generated images simulating what a patient would see. Principally, they modify 2D test images using retinal light distributions generated with ray tracing techniques. Camp et al. [35, 36] created a ray tracing algorithm and computer model for evaluation of optical performance. Maguire et al. [37, 38] employed these techniques to analyze post-surgical corneas using their optical bench software. Greivenkamp [39] created a sophisticated model which included the Stiles-Crawford effect [40], diffraction, and contrast sensitivity. A shortcoming of all these approaches is that they overlook the contribution of internal optical elements, such as the crystalline lens of the eye.

Garcia, Barsky, and Klein [41–43] developed the *CWhatUC* system, which blurs 2D images to produce an approximation of how the image would appear to a particular individual. The system uses a reconstructed corneal shape based on corneal topography measurements of the individual. Since the blur filter is computed in 2D image space, depth effects are not modeled.

The latter technique, like all those that rely on ray casting, also suffers from aliasing problems and from a computation time that increases with scene complexity. These problems are exacerbated by the need to integrate over a finite aperture as well as over the image plane, driving computation times higher to avoid substantial image noise. Since our algorithms are based in image space, they obviate these issues. That notwithstanding, the source of our input images would still need to address these issues. However, since our input images are in sharp focus, the renderer could save some computation by assuming a pinhole camera and avoiding integration over the aperture.

Vision-Realistic Rendering was introduced to the computer graphics community by the author in [44] and [45], and is presented in more detail here.
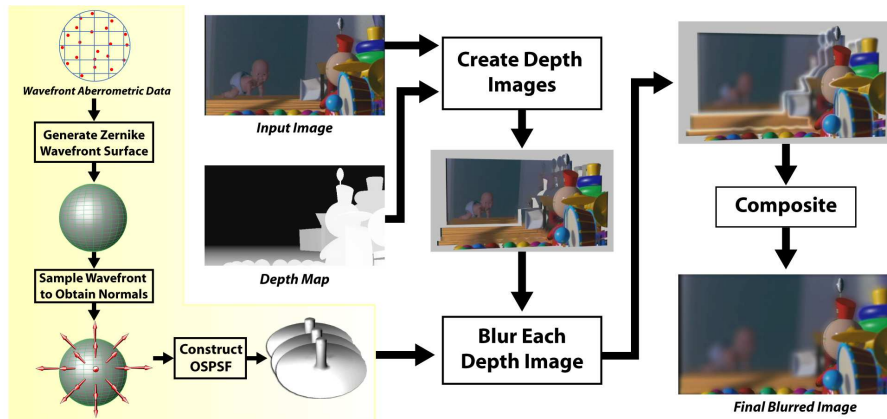
Fig. 2. Overview of the vision-realistic rendering algorithm.

4 Shack-Hartmann Device

The Shack-Hartmann Sensor [46] (Figure 3) is a device that precisely measures the wavefront aberrations, or imperfections, of a subject's eye [47]. It is believed that this is the most effective instrument for the measurement of human eye aberrations [48]. A low-power 1 mm laser beam is directed at the retina of the eye by means of a half-silvered mirror, as in Figure 4.

The retinal image of that laser now serves as a point source of light. From its reflection, a wavefront emanates and then moves towards the front of the eye. The wavewfront passes through the eye's internal optical structures, past the pupil, and eventually out of the eye. The wavefront then goes through a Shack-Hartmann lenslet array to focus the wavefront onto a CCD array, which records it.

The output from the Shack-Hartmann sensor is an image of bright points where each lenslet has focused the wavefront. Image processing algorithms are applied to determine the position of each image blur centroid to sub-pixel resolution and also to compute the deviation from where the centroid would be in for an ideal wavefront. The local slope of the wavefront is determined by the lateral offset of the focal point from the center of the lenslet. Phase information is then derived from the slope [49]. Figures 5 and 6 show the Shack-Hartmann output for eyes with and without aberrations. Figure 7 illustrates the actual output of a Shack-Hartmann sensor for a sample refractive surgery patient.

5 Algorithm

Figure 2 provides an overview of the vision-realistic rendering algorithm, comprising three major components.

5.1 Constructing the Object Space Point Spread Function

A Point Spread Function (PSF) plots the distribution of light energy on the image plane based on light that has emanated from a point source and has passed through an optical

Fig. 3. Measuring the specific vision characteristics of a subject using a Shack-Hartmann wavefront aberrometry device.

system. Thus it can be used as an image space convolution kernel.

We introduce the object space point spread function (OSPSF), which is similar to the usual image space point spread function, as described above, except that it is defined in object space and thus it varies with depth. The OSPSF is a continuous function of depth; however, we discretize it, thereby defining a sequence of depth point spread functions (DPSF) at some chosen depths.

Since human blur discrimination is nonlinear in distance but approximately linear in diopters (a unit measured in inverse meters), the depths are chosen with a constant dioptric spacing ΔD and they range from the nearest depth of interest to the farthest. A theoretical value of ΔD can be obtained from the relation $\theta = p\Delta D$, where θ is the minimum subtended angle of resolution and p is the pupil size in meters. For a human with 20/20 visual acuity, θ is 1 min of arc; that is, $\theta = 2.91 \times 10^{-4}$ [50, 51].

The DPSFs are histograms of rays cast normal to the wavefront (Figure 8). To compute these functions (Figure 9), we first place a grid with constant angular spacing at each of the chosen depths and initialize counters in each grid cell to zero. Then we iteratively choose a point on the wavefront, calculate the normal direction, and cast a ray in this direction. As the ray passes through each grid, the cell it intersects has its counter incremented. This entire process is quite fast and millions of rays may be cast in a few minutes. Finally, we normalize the histogram so that its sum is unity.

In general, wavefront aberrations are measured with the subject's eye focused at
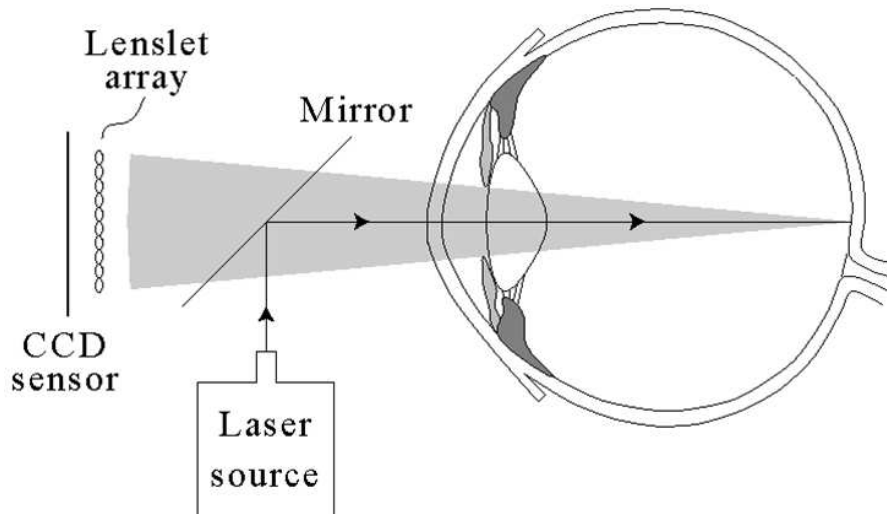
Fig. 4. A side view of a Hartmann-Shack device. A laser projects a spot on the back of the cornea. This spot serves as a point light source, originating a wavefront out of the eye. This wavefront passes through a lattice of small lenslets which focus the wavefront onto a CCD sensor.

infinity. However, it is important to be able to shift focus for vision-realistic rendering. Recent research results in optometry [52] showed that aberrations change significantly with accommodation. When aberrometric data is available for the eye focused at the depth that will be used in the final image, our algorithm exploits that wavefront measurement.

In the situation where such data is not available, then we aassume that the aberrations are independent of accommodation. We can then re-index the DPSFs, which is equivalent to shifting the OSPSF in the depth dimension. Note that this may require the computation of DPSFs at negative distances.

We further assume the OSPSF is independent of the image plane location. In optics, this is called the "isoplanatic" assumption and is the basis for being able to perform convolutions across the visual field. For human vision, this assumption is valid for at least several degrees around the fixation direction.

5.2 Fitting a Wavefront Surface to Aberrometry Data

The output of the Shack-Hartmann device comprises a ray orientation (normal vector) at each lenslet. Current devices yield only 50 to 200 such vectors. To generate the millions of samples necessary to calculate the OSPSF (see Section 5.1 above), we first generate a smooth mathematical surface representation of the wavefront from this sparse data. Our wavefront surface is a fifth degree polynomial bivariate surface defined as a height field whose domain is the pupil plane. This surface is determined by a least squares fit to the Shack-Hartmann data.

We use a particular polynomial form which was developed in 1934 [54] by the Dutch mathematician and physicist Frits Zernike who was awarded the Nobel Prize
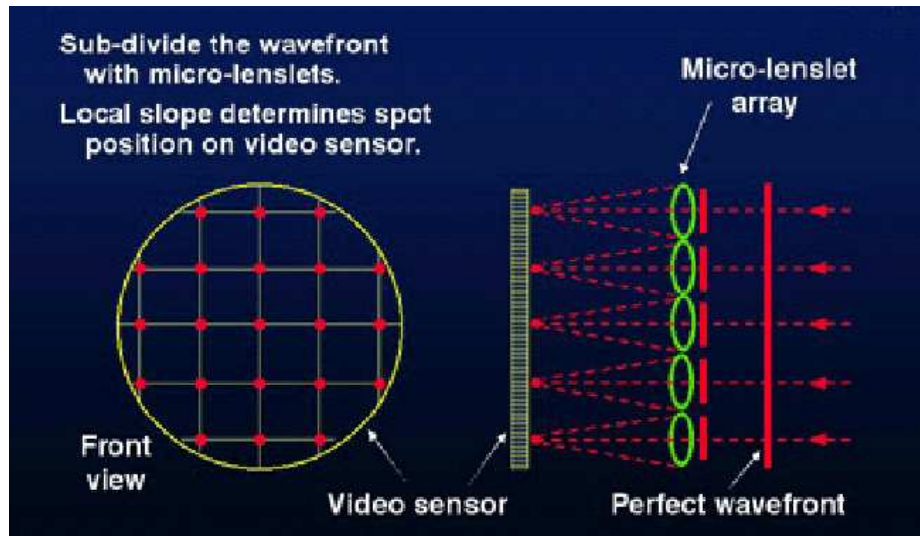
Fig. 5. Hartmann-Shack sensors measuring a perfect eye with no aberrations. Image courtesy of Larry Thibos [53].
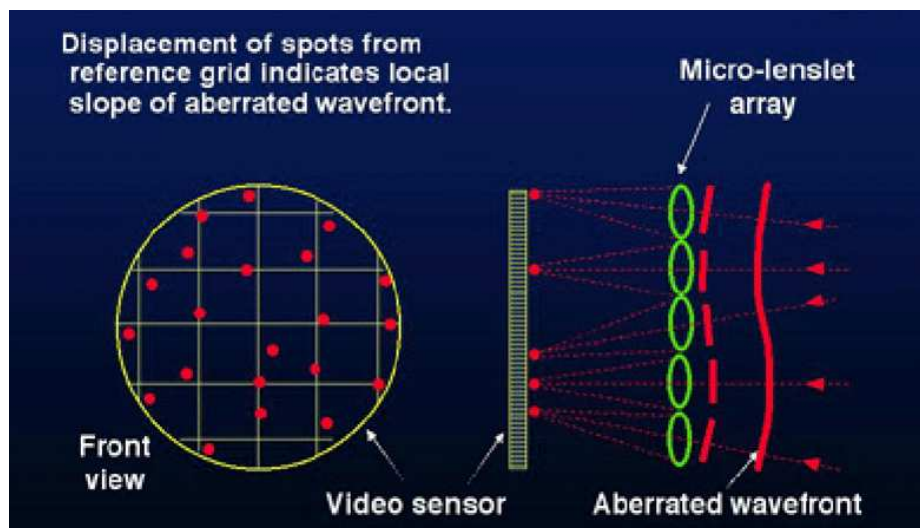

Fig. 6. Hartmann-Shack sensors measuring a normal eye with some aberrations. Image courtesy of Larry Thibos [53].

in Physics 1953 for discovering the phase contrast phenomenon; for a discussion of Zernkie polynomials realted to the optical aberrations of eyes, the reader is referred to [53]. Zernike polynomials are derived from the orthogonalization of the Taylor series. The resulting polynomial basis corresponds to orthogonal wavefront aberrations. The coefficients $Z_{m,n}$ weighting each polynomial have easily derived relations with meaningful parameters in optics. The index m refers to the aberration type, while n distin-
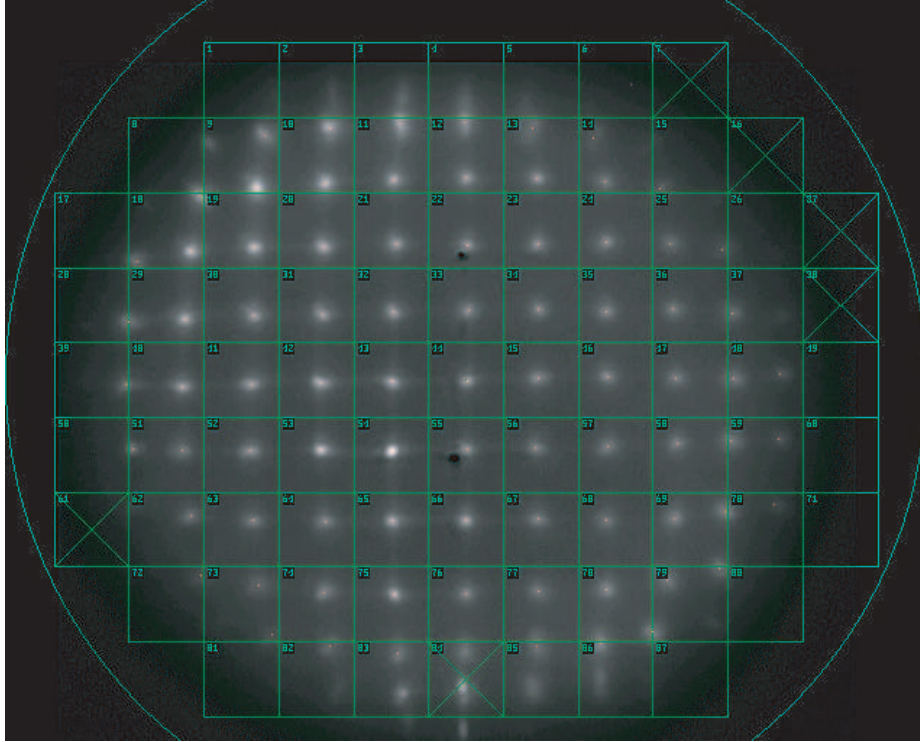
Fig. 7. Hartmann-Shack output for a sample eye. The green overlay lattice is registered to correspond to each lenslet in the array. Image courtesy of David Williams.

guishes between individual aberrations within a harmonic. For a given index m, n ranges from $-m$ to m in steps of two. Specifically, $\mathbf{Z_{0,0}}$ is displacement, $\mathbf{Z_{1,1}}$ is horizontal tilt, $\mathbf{Z_{1,-1}}$ is vertical tilt, $\mathbf{Z_{2,0}}$ is average power, $\mathbf{Z_{2,2}}$ is horizontal cylinder, $\mathbf{Z_{2,-2}}$ is oblique cylinder, $\mathbf{Z_{3,n}}$ are four terms ($n = -3, -1, 1, 3$) related to coma, and $\mathbf{Z_{4,n}}$ are five terms ($n = -4, -2, 0, 2, 4$) related to spherical aberration.

5.3 Rendering Steps

Given the input image and its associated depth map, and the OSPSF, the vision-realistic rendering algorithm comprises three steps: (1) create a set of depth images, (2) blur each depth image, and (3) composite the blurred depth images to form a single vision-realistic rendered image.

Create Depth Images. Using the depth information, the image is separated into a set of disjoint images, one at each of the depths chosen in the preceding section. Ideally, the image at depth d would be rendered with the near clipping plane set to $d + \Delta D/2$ and the far clipping plane set to $d - \Delta D/2$. Unfortunately, this is not possible because we are using previously rendered images and depth maps. Complicated texture synthesis algorithms would be overkill here, since the results will be blurred anyway. The following technique is simple, fast, and works well in practice: For each depth, d, those pixels
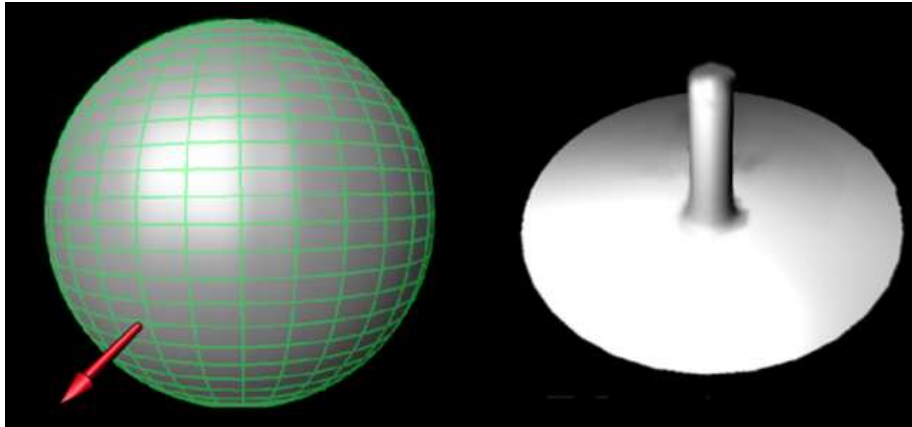
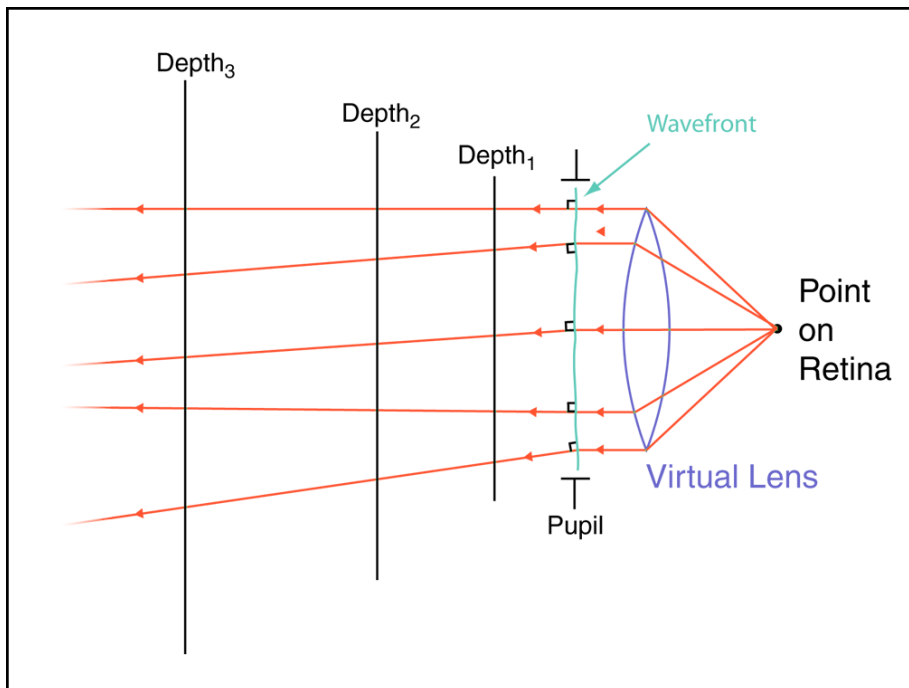**Fig. 8.** Each depth point spread function (DPSF) is a histogram of rays cast normal to the wavefront.

Fig. 9. A simplified view: Rays are cast from a point light source on the retina and pass through a virtual lens, thereby creating the measured wavefront. This wavefront is sampled and rays are cast normal to it. The DPSFs are determined by intersecting these rays at a sequence of depths.

from the original image that are within $\Delta D/2$ diopters of d are copied to the depth image. We handle partial occlusion by the techniques described in Section 6 and in more detail by Barsky et al. [7] [8].

Blur each Depth Image. Once we have the depth images, we do a pairwise convolution: Each depth image is convolved with its corresponding DPSF, thereby producing a set of blurred depth images.

Composite. Finally, we composite these blurred depth images into a single, vision-realistic rendered image. This step is performed from far to near, using alpha-blending following alpha channel compositing rules.

6   Elimination of Artifacts due to Occlusion and Discretization

Although processing in image space allows an increase in speed, the images may have artifacts introduced. This can occur in two ways, which we refer to as *occlusion* and *discretization* [7] [8]. The occlusion problem arises because there is scene geometry that is missing. This results from the finite aperture of the lens, which allows more of the scene to be visible than would be seen through an infinitesimal pinhole. Thus, without additional input, the colors from parts of the scene that are behind objects would have to be approximately reconstructed using the border colors of visible objects.

The discretization problem occurs from separating the image by depth. At adjacent pixels in different sub-images, the calculation of depth of field is complicated. This arises because these adjacent pixels may or may not correspond to the same object. An artifacts can be introduced into the image when a single object straddles two sub-images and the sub-images are blurred. The artifact arises when the far pixel is averaged with neighboring colors behind the near pixel that do not match the far pixel's color. The neighboring colors are often black, which is the default background color. Consequently, a black blurred band occurs at the intersection of the object with the separation of the sub-images that it spans, as can be seen in Figure 10.

7   Object Identification as a Solution for Image Space Artifacts

To eliminate these band artifacts that arise when an object is separated into multiple discrete sub-images, the algorithm attempts to identify entire objects within the image. This eliminates the artifact by avoiding the separation of objects across sub-images. Instead, when a large object straddles several sub-images, each sub-image will include the entire object instead of only a part of that object. Consequently, the object will have minimal artifacts due to blurring.

We will now consider two approaches for object identification to properly blur the scene; these techniques are described in more detail by Barsky et al. [7] [8]. Our first approach uses the depth difference of adjacent pixels to identify objects. In our second approach, the Canny Edge Detection [55] algorithm is applied to draw borders between objects and hence identify them.

7.1   Edge Detection Technique for Object Identification

Our first method for identifying objects begins by using a variant of the Canny Edge Detection algorithm [55]. The Canny algorithm takes as input an intensity map for the

Fig. 10. Black bands appear at the locations where the sub-images are separated.
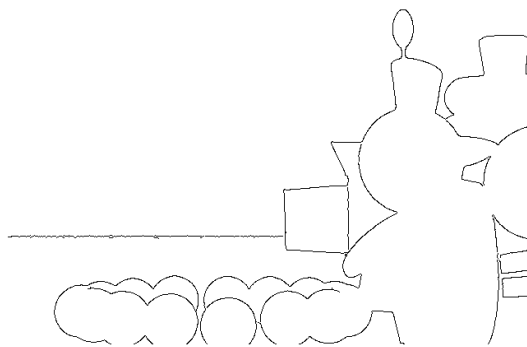

Fig. 11. Using depth map information as input, the edge detection algorithm identifies where object edges lie in the image.

image, and it convolves the intensity map with the first derivative of a Gaussian function. The algorithm then marks pixels in the resulting array whose magnitude exceeds a specified upper threshold. These marked pixels are grouped into edge curves based on the assumption that neighboring marked pixels that have consistent orientations belong to the same edge.

Our technique uses a depth map as the intensity map as input to this Edge Detection algorithm. Figure 11 shows the result of edge detection on the example depth map. Using this variant of the Canny algorithm to segment a scene into distinct objects avoids inadequacies that are common to traditional edge detection methods. In particular, using depth information avoids erroneous detection of edges that correspond to the surface markings and shadows of objects in the scene.

Starting with the region formed by the boundary pixels in the current sub-image, the algorithm extends that region until it is bounded by previously detected edges. Specifically, extending the region involves taking the union of the line segments that begin within the original region and do not intersect the detected edge segments; this is illustrated in Figure 12.

The result of the Canny Edge Detection method to eliminate these artifacts is demonstrated in Figures 13 and 14, focused on the Tin Toy in the foreground, and on the baby in the background, respectively.
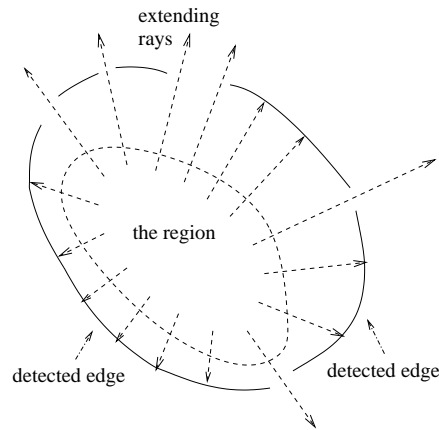
Fig. 12. Extending the region involves taking the union of the line segments that begin within the original region and do not intersect the detected edges.

Fig. 13. Artifacts eliminated by the Edge Detection technique for image focused on the Tin Toy in the foreground.

7.2 Adjacent Pixel Difference Technique for Object Identification

The second technique for including points from objects that span several sub-images assumes a surface with a given order of continuity. As input to the algorithm, we select the order of continuity, denoted C^n, of the surface. In addition, the bound on the nth derivative of depth with respect to the image plane coordinates is selected such that adjacent pixels within the bound correspond to the same object. Since image space is a discrete representation of continuous geometry, we use the difference as the discretized counterpart of the derivative. Figure 15 illustrates a first degree difference map for an arbitrary image.

The algorithm assigns an object identifier to each pixel and then groups together those pixels that share an object identifier. Once all objects are located, it is straightforward to determine whether the neighboring colors should be obtained from objects in front of, at, or behind, the current sub-image.

In Section 7.1, Figures 13 and 14 demonstrated the results of the Canny Edge Detection technique. This eliminated the artifacts illustrated in Figure 10 and generated a

Fig. 14. Artifacts eliminated by the Edge Detection technique for image focused on the baby in the background.

3	1	4	1
5	9	2	6
5	3	5	8
9	7	9	3

2	−3	3
−4	7	−4
2	−2	−3
2	−2	6

Fig. 15. An example of a first degree difference map (right) resulting from applying a horizontal difference to the first 16 digits of π (left).

correctly blurred image. We will now use the Adjacent Pixel Difference technique to generate a similar artifact-free blurred image, which are shown in Figures 16 and 17, focused on the Tin Toy in the foreground, and on the baby in the background, respectively.

8 Sample Images

Figures 1, 19, 20, and 21 are vision-realistic renderings of a room scene. Figure 18 is a simulation that models ideal vision and Figures 1, 19, 20, and 21 are simulations of the vision of actual individuals based on their measured data. Notice that the nature of the blur is different in each image. The field of view of the image is approximately 46° and the pupil size is rather large at 5.7 mm.

For Figure 18, we constructed an OSPSF from a planar wavefront to yield a simulation of vision for an aberration-free model eye.

The simulation of vision shown in Figure 19 is based on the data from the left eye of male patient GG who has astigmatism. Note how the blur is most pronounced in one direction (in this case horizontal), which is symptomatic of astigmatism.

Next, we show vision-realistic rendered images based on pre- and post-operative data of patients who have undergone LASIK vision correction surgery. Specifically, the vision for the right eye of male patient DB is simulated in Figure 20, and then Figure

Fig. 16. Artifacts eliminated by the Adjacent Pixel Difference technique for image focused on the Tin Toy in the foreground.

Fig. 17. Artifacts eliminated by the Adjacent Pixel Difference technique for image focused on baby in the background.

Fig. 18. Simulation of vision of an aberration-free model eye.

21 simulates the vision of the left eye of male patient DR. For each patient, the pre-operative vision is simulated in the top image while the lower image simulates the post-

Fig. 19. Simulation of vision of astigmatic patient GG.

operative vision. The images demonstrating pre-operative vision show the characteristic extreme blur pattern of the highly myopic (near-sighted) patients who tend to be the prime candidates for this surgery. Although, in both cases, the vision has been improved by the surgery, it is still not as good as the aberration-free model eye. Furthermore, the simulated result of the surgery for patient DB is slightly inferior to that depicted for patient DR. However, note that the patient (DB) with the inferior surgical result had significantly inferior pre-operative vision compared to that of patient DR.

Figure 1 is computed based on data measured from the left eye of female patient KS who has the eye condition known as *keratoconus*. This image shows the distortion of objects that is caused by the complex, irregular shape of the keratoconic cornea. Note how the nature of these visual artifacts is distinct from what would generally be a smooth blur that occurs in more pedestrian vision problems such as myopia (see Figure 21(a)). This distinction is often not understood by clinicians when it is articulated by keratoconic patients. We hope our techniques could be used in optometry and oph-thalmology for the education of students and residents as well as for the continuing education of clinicians.

Our approach can also be applied to photographs (with associated depth maps) of real scenes, not only to synthetic images. For example, in Figure 22, the top image is a photograph showing the Campanille at U.C. Berkeley with San Francicsco's Golden Gate Bridge in the background, with both the tower and bridge are in sharp focus. Constructing an OSPSF with the point of focus at the Campanille and then applying our algorithm yields an image with the background blurred, as shown in bottom image.

To consider the computational requirements of our technique, note that it comprises three parts: fitting the wavefront surface, construction of the OSPSF and the rendering step. The computation time for the surface fitting is negligible. The time to compute the OSPSF depends on the number of wavefront samples. For the images in this paper, the computation of the OSPSF, using one million samples, was less than half a minute. The rendering step is dominated by the FFTs performed for convolution (our kernels
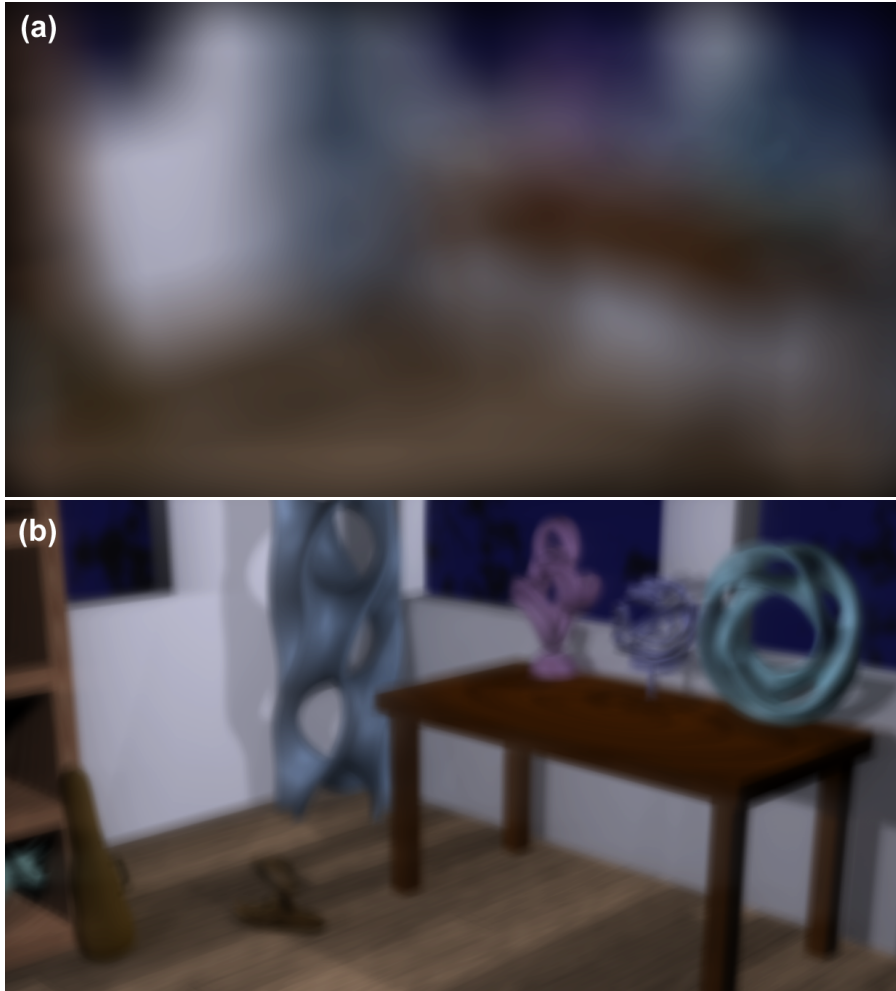
Fig. 20. Simulation of vision of LASIK patient DB based on (a) Pre-operative and (b) Post-operative data.

are generally too large to convolve in the spatial domain). Thus, the computation time for the rendering step is dependent on the image size and the number of non-empty depth images. The room scene has a resolution of 1280 X 720 and each image took about 8 minutes to render, using 11 depth images, on a Pentium 4 running at 2.4 GHz, using Matlab. This time could be significantly reduced by converting to a C or C++ implementation with a standard FFT library.

9 Validation

An important area of future work is validation, and will involve the establishment of psychophysical experiments. Nonetheless, some preliminary experiments are possible
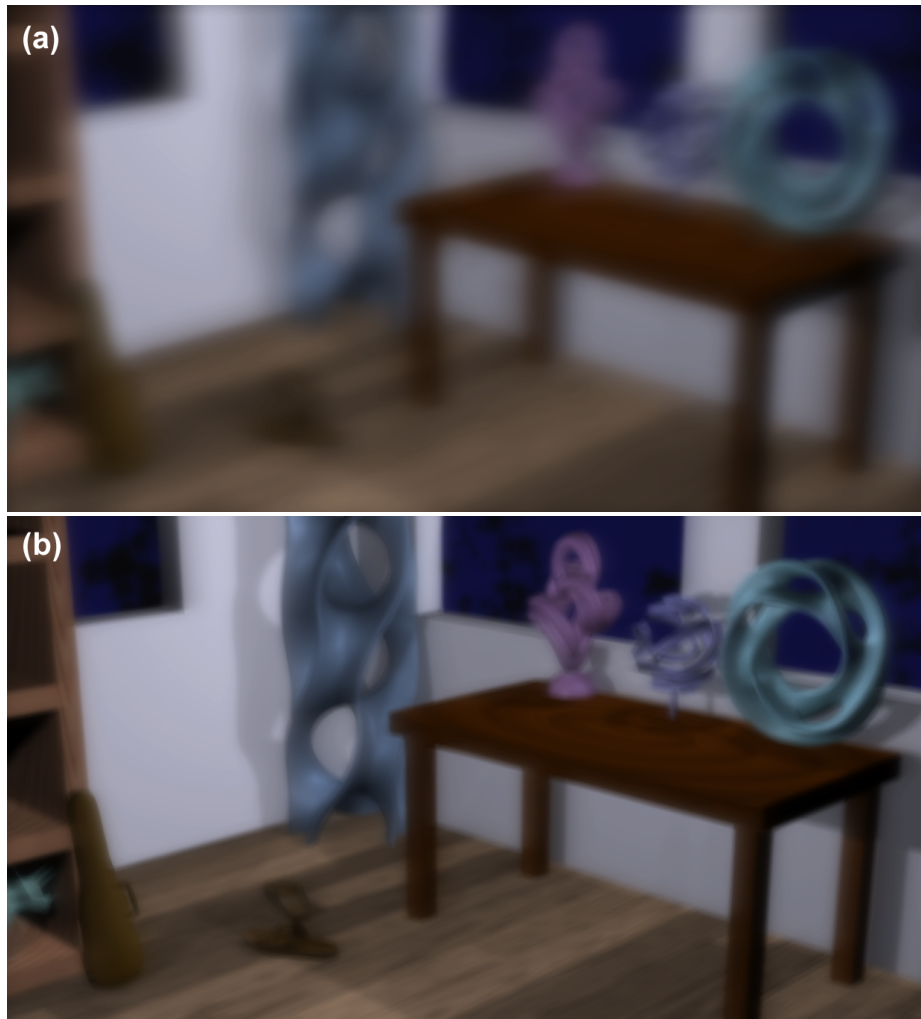
Fig. 21. Simulation of vision of LASIK patient DR based on (a) Pre-operative and (b) Post-operative data.

immediately, and our initial results have been positive. First, patients who have unilateral vision problems can view our simulations of the vision in their pathological eye using their contralateral eye, thereby evaluating the fidelity of the simulation. Second, consider patients who have vision conditions such as myopia, hyperopia, and astigmatism, that are completely corrected by spectacles or contact lenses. More precisely, in optometry terms, they might have 20/20 BSCVA (best spectacle corrected visual acuity). Such patients could validate the quality of the depiction of their vision in vision-realistic rendered images simply by viewing them while wearing their corrective eyewear. Third, the visual anomalies present in keratoconus are different from those in more common conditions such as myopia, and this distinction is indeed borne out in our example im-

Fig. 22. Original photograph with both Campanille and Golden Gate Bridge in focus (top) and output image with background blurred (bottom).

ages. Specifically, keratoconus can cause the appearance of diplopia (double-vision) whereas myopia usually engenders a smooth blur around edges. Indeed, exactly this distinction can be observed upon close examination of our sample images. Fourth, severe astigmatism causes more blur in one direction than in the orthogonal direction, and this is exactly what is depicted in our sample image of astigmatism. Fifth, our simulations of the vision of patients with more myopia are more blurred than those of patients with less myopia.

10   Conclusions and Future Work

We introduced the concept of *vision-realistic rendering* – the computer generation of synthetic images that incorporate the characteristics of a particular individual's entire optical system. This paper took the first steps toward this goal, by developing a method for simulating the scanned foveal image from wavefront data of actual human subjects, and demonstrated those methods on sample images. First, a subject's optical system is measured by a Shack-Hartmann wavefront aberrometry device. This device outputs a measured wavefront which is sampled to calculate an object space point spread function (OSPSF). The OSPSF is then used to blur input images. This blurring is accomplished by creating a set of depth images, convolving them with the OSPSF, and finally compositing to form a vision-realistic rendered image. Applications of vision-realistic rendering in computer graphics as well as in optometry and ophthalmology were discussed.

The problem of vision-realistic rendering is by no means solved. Like early work on photo-realistic rendering, our method contains several simplifying assumptions and other limitations. There is much interesting research ahead.

The first limitations are those stemming from the method of measurement. The Shack-Hartmann device, although capable of measuring a wide variety of aberrations, does not take into account light scattering due to such conditions as cataracts. The wavefront measurements can have some error, and fitting the Zernike polynomial surface to the wavefront data can introduce more. However, since the wavefronts from even pathological eyes tend to be continuous, smooth interpolation of the Shack-Hartmann data should not produce any significant errors. Consequently, any errors that are introduced should be small and, furthermore, such small errors would imperceptible in final images that have been discretized into pixels.

Strictly speaking, the pupil size used for vision-realistic rendering should be the same as the pupil size when the measurements are taken. However, the error introduced in using only part of the wavefront (smaller pupil) or extrapolating the wavefront (larger pupil) should be quite small. We have made use of three assumptions commonly used in the study of human physiological optics: isoplanarity, independence of accommodation, and off-axis aberrations being dominated by on-axis aberrations. Although we have argued that these assumptions are reasonable and provide a good first-order approximation, a more complete model would remove at least the first two.

As discussed in Section 5.1, we have assumed "independence of accommodation" since aberrometric measurements with the eye focused at the depth is not usually available. However, this is not a limitation of our algorithm. Our algorithm can exploit wave-

front data where the eye is focused at the depth that will be used in the final image, when such a measurement is made.

We currently do not take chromatic aberration into account, but again that is not a limitation of our algorithm. Since the data we acquire is from a laser, it is monochromatic. However, some research optometric colleagues have acquired polychromatic data and will be sharing it with us. It is again interesting that recent research in optometry by Marcos [56] has shown that except for the low order aberrations, most aberrations are fairly constant over a range of wavelengths.

We only compute the aberrations for one point in the fovea, and not for other points in the visual field. However, it is important to note that for computer graphics, the on-axis aberrations are critically important because viewers move their eyes around when viewing a scene. If we had actually included the off-axis aberrations of the eye, then the off-axis parts of the scene would have been improperly blurred for a person who is scanning the scene. The off-axis aberrations are of minor concern even without eye movements since the retinal sampling of cones is sparse in peripheral vision. The image that we are simulating is formed by viewing the entire scene using the on-axis aberrations because we assume that the viewer is scanning the scene.

However, since peripheral vision does makes important contributions to visual appearance, viewers are affected by optical distortions of peripheral vision. Thus, it is of interest to extend this method to properly address the off-axis effects.

References

1. Barsky, B.A., Horn, D.R., Klein, S.A., Pang, J.A., Yu, M.: Camera models and optical systems used in computer graphics: Part I, object based techniques. In: Proceedings of the 2003 International Conference on Computational Science and its Applications (ICCSA'03), Montreal,Second International Workshop on Computer Graphics and Geometric Modeling (CGGM'2003), Springer-Verlag Lecture Notes in Computer Science (LNCS), Berlin/Heidelberg (2003) 246–255

2. Barsky, B.A., Horn, D.R., Klein, S.A., Pang, J.A., Yu, M.: Camera models and optical systems used in computer graphics: Part II, image based techniques. In: Proceedings of the 2003 International Conference on Computational Science and its Applications (ICCSA'03), Montreal,Second International Workshop on Computer Graphics and Geometric Modeling (CGGM'2003), Springer-Verlag Lecture Notes in Computer Science (LNCS), Berlin/Heidelberg (2003) 256–265

3. Potmesil, M., Chakravarty, I.: Synthetic image generation with a lens and aperture camera model. ACM Transactions on Graphics 1 (1982) 85–108 (Original version in ACM SIGGRAPH 1981 Conference Proceedings, Aug. 1981, pp. 297-305.).

4. Shinya, M.: Post-filtering for depth of field simulation with ray distribution buffer. In: Proceedings of Graphics Interface '94, Banff, Alberta, Canadian Information Processing Society (1994) 59–66

5. Rokita, P.: Fast generation of depth-of-field effects in computer graphics. Computers & Graphics 17 (1993) 593–595

6. Rokita, P.: Generating depth-of-field effects in virtual reality applications. IEEE Computer Graphics and Applications 16 (1996) 18–21

7. Barsky, B.A., Tobias, M.J., Horn, D.R., Chu, D.P.: Investigating occlusion and discretization problems in image space blurring techniques. In: First International Conference on Vision, Video, and Graphics, Bath, Peter Hall and Philip Willis (editors). (2003) 97–102

8. Barsky, B.A., Tobias, M.J., Chu, D.P., Horn, D.R.: Elimination of artifacts due to occlusion and discretization problems in image space blurring techniques. Graphical Models (2005)

9. Cook, R.L., Porter, T., Carpenter, L.: Distributed ray tracing. In: ACM SIGGRAPH 1984 Conference Proceedings, Minneapolis, SIGGRAPH, ACM (1984) 137–145

10. Dippe, M.A.Z., Wold, E.H.: Antialiasing through stochastic sampling. In: ACM SIGGRAPH 1985 Conference Proceedings, San Francisco, SIGGRAPH, ACM (1985) 69–78

11. Lee, M.E., Redner, R.A., Uselton, S.P.: Statistically optimized sampling for distributed ray tracing. In: ACM SIGGRAPH 1985 Conference Proceedings, San Francisco, SIGGRAPH, ACM (1985) 61–67

12. Kolb, C., Mitchell, D., Hanrahan, P.: A realistic camera model for computer graphics. In: ACM SIGGRAPH 1995 Conference Proceedings, SIGGRAPH, ACM (1995) 317–324

13. Loos, J., Slusallek, P., Seidel, H.P.: Using wavefront tracing for the visualization and optimization of progressive lenses. In: Computer Graphics Forum (Eurographics '98 Conference Proceedings, Eurographics (1998) 255–266 held in Lisbon, Portugal, 02-04 September 1998.

14. Levoy, M., Hanrahan, P.: Light field rendering. In: ACM SIGGRAPH 1996 Conference Proceedings, New Orleans, SIGGRAPH, ACM (1996) 31–42

15. Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F.: The lumigraph. In: ACM SIGGRAPH 1996 Conference Proceedings, SIGGRAPH, ACM (1996) 43–54

16. Adelson, E., Bergen, J.: Computational Models of Visual Processing. The MIT Press, Cambridge, Mass. (1991)

17. Heidrich, W., Slusallek, P., Seidel, H.P.: An image-based model for realistic lens systems in interactive computer graphics. In: Graphics Interface, Canadian Information Processing Society (1997) 68–75

18. Isaksen, A., McMillan, L., Gortler, S.J.: Dynamically reparameterized light fields. In: Proceedings of ACM SIGGRAPH 2000, New Orleans, SIGGRAPH, ACM (2000) 297–306

19. Upstill, S.D.: The Realistic Presentation of Synthetic Images: Image Processing in Computer Graphics. PhD thesis, Computer Science Division, University of California (1985)

20. Spencer, G., Shirley, P., Zimmerman, K., Greenberg, D.P.: Physically-based glare effects for digital images. In: ACM SIGGRAPH 1995 Conference Proceedings, SIGGRAPH, ACM (1995) 325–334

21. Bolin, M.R., Meyer, G.W.: A perceptually based adaptive sampling algorithm. In: ACM SIGGRAPH 1998 Conference Proceedings, Orlando, SIGGRAPH, ACM (1998) 299–309

22. Tumblin, J., Rushmeier, H.E.: Tone reproduction for realistic images. IEEE Computer Graphics and Applications 13 (1993) 42–48

23. Chiu, K., Herf, M., Shirley, P., Swamy, S., Wang, C., Zimmerman, K.: Spatially non-uniform scaling functions for high contrast images. In: Proceedings of Graphics Interface '93, Toronto, Canadian Information Processing Society (1993) 245–254

24. Ferwerda, J.A., Pattanaik, S.N., Shirley, P., Greenberg, D.P.: A model of visual adaptation for realistic image synthesis. In: ACM SIGGRAPH 1996 Conference Proceedings, New Orleans, SIGGRAPH, ACM (1996) 249–258

25. Ward-Larson, G., Rushmeier, H., Piatko, C.D.: A visibility matching tone reproduction operator for high dynamic range scenes. IEEE Transactions on Visualization and Computer Graphics 3 (1997) 291–306

26. Pattanaik, S.N., Ferwerda, J.A., Fairchild, M.D., Greenberg, D.P.: A multiscale model of adaptation and spatial vision for realistic image display. In: ACM SIGGRAPH 1998 Conference Proceedings, SIGGRAPH, ACM (1998) 287–298

27. Reinhard, E., Stark, M., Shirley, P., Ferwerda, J.: Photographic tone reproduction for digital images. In: ACM SIGGRAPH 2002 Conference Proceedings, SIGGRAPH, ACM (2002) 267–276

28. Meyer, G.W., Greenberg, D.P.: Color-defective vision and computer graphics displays. IEEE Computer Graphics and Applications 8 (1988) 28–40

29. Meyer, G.W.: Image synthesis and color vision. In Rogers, D.F., Earnshaw, R.A., eds.: Computer Graphics Techniques. Springer Verlag, New York (1990) 45–77

30. Pellacini, F., Greenberg, D.P.: Toward a psychophysically-based light reflection model for image synthesis. In: ACM SIGGRAPH 2000 Conference Proceedings, New Orleans, SIGGRAPH, ACM (2000) 55–64

31. Lubin, J.: A visual discrimination model for imaging system design and evaluation. In Peli, E., ed.: Vision Models for Target Detection and Recognition. Volume 2. World Scientific Publishing Co., Inc., Washington, DC (1995) 245–357

32. Peli, E.: Test of a model of foveal vision by using simulations. Journal of Optical Society of America 13 (1996) 1131–1138

33. Fine, E.M., Rubin, G.S.: Effects of cataract and scotoma on visual acuity: A simulation study. Optometry and Vision Science 76 (1999) 468–473

34. Fine, E.M., Rubin, G.S.: The effects of simulated cataract on reading with normal vision and simulated central scotoma. Vision Research 39 (1999) 4274–4285

35. Camp, J.J., Maguire, L.J., Cameron, B.M., Robb, R.A.: A computer model for the evaluation of the effect of corneal topography on optical performance. Am. J. Ophthalmol. 109 (1990) 379–386

36. Camp, J.J., Maguire, L.J., Robb, R.A.: An efficient ray tracing algorithm for modeling visual performance from corneal topography. In: First Conference on Visualization in Biomedical Computing, The Institute of Electrical and Electronics Engineers, Inc., IEEE Computer Society Press (1990) 279–285

37. Maguire, L.J., Zabel, R.W., Parker, P., Lindstrom, R.L.: Topography and raytracing analysis of patients with excellent visual acuity 3 months after excimer laser photorefractive keratectomy for myopia. Refract. Corneal Surg. 8 (1991) 122–128

38. Maguire, L.J., Camp, J.J., Robb, R.A.: Informing interested parties of changes in the optical performance of the cornea caused by keratorefractive surgery — a ray tracing model that tailors presentation of results to fit the level of sophistication of the audience. In: SPIE Vol. 1808 Visualization in Biomedical Computing, The International Society for Optical Engineering, SPIE (1992) 601 – 609

39. Greivenkamp, J.E., Schweigerling, J., Miller, J.M., Mellinger, M.D.: Visual acuity modeling using optical raytracing of schematic eyes. Am. J. Ophthalmol. 120 (1995) 227–240

40. Moon, P., Spencer, D.E.: On the Stiles-Crawford effect. J. Opt. Soc. Am. 34 (1944) 319–329

41. Garcia, D.D., Barsky, B.A., Klein, S.A.: CWhatUC: A visual acuity simulator. In: Proceedings of Ophthalmic Technologies VIII, SPIE International Symposium on Biomedical Optics, SPIE (1998) 290–298

42. Garcia, D.D., Barsky, B.A., Klein, S.A.: The OPTICAL project at UC Berkeley: Simulating visual acuity. Medicine Meets Virtual Reality: 6 (Art, Science, Technology: Healthcare (r)Evolution) (1998)

43. Garcia, D.D.: CWhatUC : Software Tools for Predicting, Visualizing and Simulating Corneal Visual Acuity. PhD thesis, Computer Science Division, University of California, Berkeley, California (2000)

44. Barsky, B.A., Bargteil, A.W., Garcia, D.D., Klein, S.A.: Introducing vision-realistic rendering. In: Eurographics 2002 Poster Proceedings, Eurographics (2002) 1–7 Held in Pisa, Italy 26-28 June 2002.

45. Barsky, B.A.: Vision-realistic rendering: Simulation of the scanned foveal image from wavefront data of human subjects. In: First Symposium on Applied Perception in Graphics and Visualization, co-located with ACM SIGGRAPH, Los Angeles (2004) 73–81

46. Platt, B.C., Shack, R.V.: Lenticular Hartmann-screen. Newsletter 5, 15, Optical Science Center, University of Arizona (1971)

47. Thibos, L.N.: Principles of Hartmann-Shack aberrometry. In: Vision Science and its Applications. Volume 15. Optical Society of America – Trends in Optics and Photonics (2000) 163–169

48. Liang, J.: A New Method to Precisely Measure the Wave Aberrations of the Human Eye with a Hartmann-Shack Wavefront Sensor. PhD thesis, Department of Mathematics, Universität Heidelberg (1991)

49. Klein, S.A.: Optimal corneal ablation for eyes with arbitrary Hartmann-Shack aberrations. J. Opt. Soc. Am. A 15 (1998) 2580–2588

50. Keating, M.P.: Geometric, Physical, and Visual Optics, second edition. Butterworth-Heinemann Ltd. (2002)

51. Bennett, A.G., Rabbetts, R.B.: Clinical Visual Optics. Butterworth-Heinemann Ltd. (1998)

52. He, J.C., Burns, S.A., Marcos, S.: Monochromatic aberrations in the accommodated human eye. Vision Research 40 (2000) 41–48

53. Thibos, L.N., Applegate, R., Schwiegerling, J.T., Webb, R., Members, V.S.T.: Standards for reporting the optical aberrations of eyes. In Lakshminarayanan, V., ed.: Vision Science and its Applications. Volume 35. Optical Society of America – Trends in Optics and Photonics (2000) 232–244

54. Zernike, F.: Beugungstheorie des schneidenverfahrens und seiner verbesserten form, der phasenkontrastmethode. Physica 1 (1934) 689–704

55. Canny, J.F.: A computational approach to edge detection. IEEE Transaction on Pattern Analysis and Machine Intelligence 8 (1986) 679–698

56. Marcos, S., Burns, S.A., Moreno-Barriuso, E., Navarro, R.: A new approach to the study of ocular chromatic aberrations. Vision Research 39 (1999) 4309–4323