Thirteenth Eurographics Workshop on Rendering (2002)
P. Debevec and S. Gibson (Editors)

Introducing Vision-Realistic Rendering

Brian A. Barsky*†‡ Adam W. Bargteil* Daniel D. Garcia* Stanley A. Klein†‡

*Computer Science Division †School of Optometry ‡Bioengineering Graduate Group
University of California, Berkeley, California, 94720-1776, USA
http://www.cs.berkeley.edu/optical/vrr

Abstract

*We introduce the concept of **vision-realistic rendering**–the generation of images that incorporate characteristics of a particular individual's optical system. We then describe a pipeline for creating vision-realistic 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 image. We discuss applications of vision-realistic rendering in computer graphics as well as in optometry and ophthalmology and note that our method is a post-process and can handle simple camera models as a special case.*

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Display algorithms, Viewing algorithms

1. Introduction

We introduce a new concept which we call *vision-realistic rendering*–the generation of images that incorporate characteristics of a particular individual's optical system. We describe a pipeline to achieve vision-realistic rendering and show 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.

There are two distinct impacts of this research, one from the point of view of optometry and ophthalmology, and the other from the perspective of computer graphics. Our technique enables the generation of vision-realistic images and animations that demonstrate specific defects in how a person sees. Such images could be shown to an individual's optometrist or ophthalmologist 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 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-op and post-op, our technique could be used to convey to doctors what the vision of a patient is like before and after surgery (Figure 5). In addition, by using modeled or simulated wavefront measurements, this approach could

Figure 1: *Vision-realistic image simulating vision based on actual wavefront data from a patient with keratoconus.*

provide accurate and revealing medical visualizations of predicted visual acuity and of simulated vision. 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.

There are also interesting applications of our technique in the context of computer graphics and computer animation. For example, vision-realistic rendering could enhance the realism of a first-person view. As a special case, this approach can be used as a post-process to simulate camera effects such as depth of field. Note that the depth map can be manipulated to achieve non-photorealistic focusing effects, such as keeping a range of depths all in perfect focus. This aspect of our system provides powerful control that is not available when the camera model is incorporated in the renderer.

1.1. Previous and Related Work

The first synthetic images with depth of field were computed by Potmesil and Chakravarty[28] who convolved images with depth-based blur filters. However, they ignored issues relating to occlusion, which Shinya[31] subsequently addressed using a ray distribution buffer. Rokita[29] 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[30]. Stochastic sampling techniques were used to generate images with depth of field as well as motion blur by Cook et al.[6], Dippe and Wold[7], and Lee et al.[16]. More recently, Kolb et al.[15] described a more complete camera lens model that addresses both the geometry and radiometry of image formation. Isaksen et al.[13] modeled depth of field effects using dynamically reparameterized light fields. 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 occlusion in an ad hoc manner.

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[35], considered the problem of viewing synthetic images on a CRT and derived post-processing techniques for improved display. Spencer et al.[32] investigated image-based techniques of adding simple ocular and camera effects such as glare, bloom, and lenticular halo. Bolin and Meyer[2] used a perceptually-based sampling algorithm to monitor images as they are being rendered for artifacts that require a change in rendering technique. Tumblin and Rushmeier[34], Chiu et al.[5], Ferweda et al.[8], Ward et al.[36], and Pattanaik et al.[24] 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. Meyer and Greenberg[22] 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[21] 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.[26] 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[18, 25] have been done by artist renditions and physical means, not by computer graphics. For example, Fine and Rubin[9, 10] simulating 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.[3, 4] created a ray tracing algorithm and computer model for evaluation of optical performance. Maguire et al.[19, 20] employed these techniques to analyze post-surgical corneas using their optical bench software. Greivenkamp[12] created a sophisticated model which included the Stiles-Crawford effect[23], 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[11] developed the *CWhatUC* system, which uses a reconstructed corneal shape to blur 2D images to simulate an individual's visual acuity. Since the blurring is done in 2D image space, depth effects are not modeled. In a similar vein, Barsky et al.[1] extended the work of Kolb et al.[15] by tracing rays through a corneal surface extracted from actual patient data instead of a system of camera lenses. Both these approaches only consider effects due to the cornea, whereas we are able to include effects of the entire optical system.

2. Methods

2.1. Shack-Hartmann device

The Shack-Hartmann Sensor[27] (Figure 2) is a device that precisely measures the wavefront aberrations, or imperfections, of a subject's eye[33]. It is believed that this is the

Figure 2: *Measuring the specific vision characteristics of a subject using a Shack-Hartmann wavefront aberrometry device.*

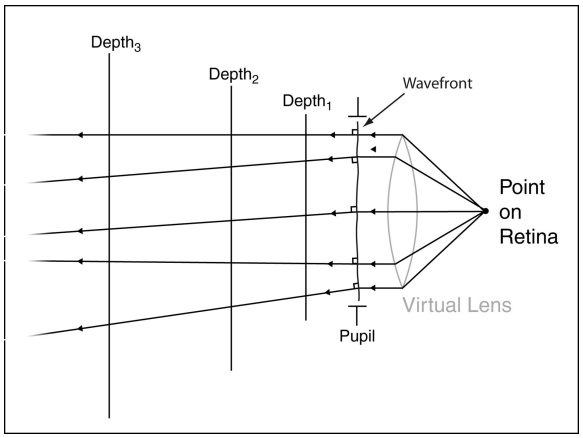

Figure 3: *A simplified view: Rays are cast from a point light source on the retina and pass through a virtual lens, thereby creating the measured wavefront. This wavefront is sampled and rays are cast normal to it. The DPSFs are determined by intersecting these rays at a sequence of depths.*

most effective instrument for the measurement of human eye aberration[17]. A low-power 1 mm laser beam is directed at the retina of the eye by means of a half-silvered mirror. The retinal image of that laser now serves as a point source of light for a wavefront that 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 image 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 these centroids to sub-pixel resolution and also to compute the deviation from where they would ideally be. 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[14].

The limited number of lenslets provides only a sparse sampling of the overall wavefront; thus, a Zernike-polynomial surface is fit to these samples. This provides a continuous surface, enabling us to sample the wavefront at a much higher rate.

2.2. Object Space Point Spread Function

We introduce the object space point spread function (OSPSF), which is similar to the usual image space point spread function, 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.

The DPSFs are histograms of rays cast normal to the wavefront. To compute these functions (Figure 3), 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 infinity. However, it is important to be able to shift focus for vision-realistic rendering. This is achieved by reindexing 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.

2.3. Blurring

Our vision-realistic blurring algorithm can be summarized as follows: (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 image.

We begin with the depth information for the image, and create 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. However, we have

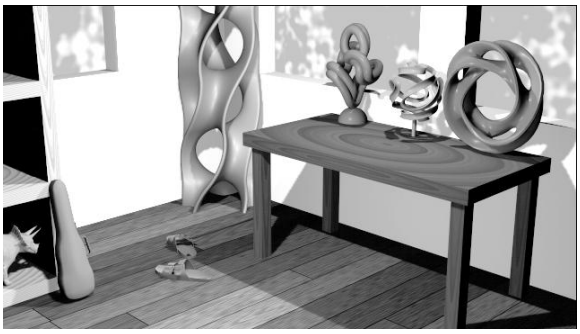

Figure 4: *Simulation of vision of an aberration-free model eye.*

found that in practice, the following technique works well: For each depth, d, those pixels from the original image that are within $\Delta D/2$ diopters of d are copied to the depth image. Since occluded pixels are likely to be similar to neighboring pixels, partial occlusion is handled by convolving the depth image with a Gaussian and copying pixels from the result to pixels occluded in the depth image (i.e., pixels nearer than $d + \Delta D/2$). Although we believe that this approach works fairly well, the observant viewer will notice some subtle artifacts in our example images.

Once we have the depth images, we convolve them with the corresponding DPSFs to create a set of blurred depth images. Finally, we composite these blurred depth images into a single, vision-realistic image. This step is performed from far to near, following the alpha channel compositing rules.

3. Sample Images

Figures 1, 4 and 5, are vision-realistic renderings of a room scene. These images demonstrate the vision of actual individuals based on their measured data. The field of view of the image is roughly 46° and the pupil size is rather large at 5.7 mm.

Figure 1 uses data measured from the left eye of female patient KS who has the eye condition *keratoconus*. This image shows the distortion due to the *irregular astigmatism* that is associated with this complex-shaped cornea.

Figure 5 uses both pre-op and post-op data from the right eye of male patient DB who has undergone LASIK vision correction surgery. Pre-op vision is simulated in Figure 5(a) while Figure 5(b) simulates post-op vision. The pre-op image shows the characteristic extreme blur pattern of the highly myopic patients who tend to be prime candidates for this surgery. Although the vision has been improved by the surgery, it is still not as good as the aberration-free model eye (Figure 4). Validation of such images is an area of future work for us.

Figure 6 shows several frames from a *rack focus* applied

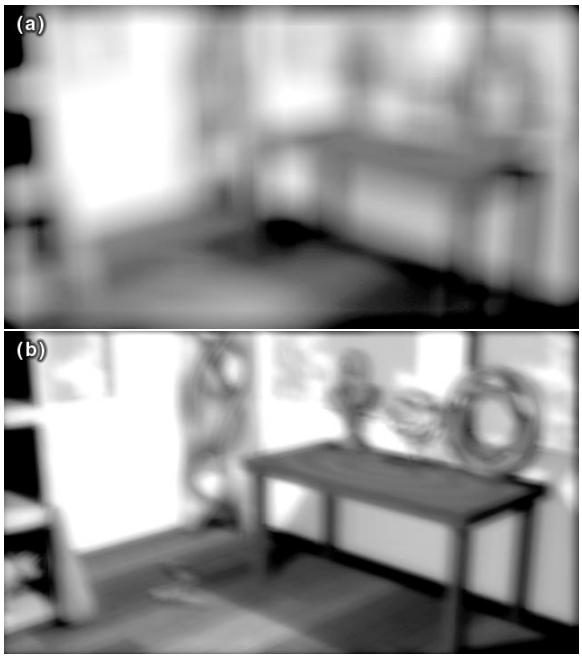

Figure 5: *Simulation of vision of LASIK patient DB based on (a) Pre-op and (b) Post-op data.*

to a shot from Pixar's short film *Tin Toy*. The original short film was rendered without depth of field. We add depth of field to this shot to draw the audience's attention away from the Tin Toy to the baby as he enters the scene. The depth information for this short film no longer exists; thus, we hand-generated the depth map. We deduced a 100 mm focal length lens which corresponds to a 23° field of view. In these frames, we set the aperture to f/2.8.

Table 1 provides computation times, number of DPSFs computed, number of wavefront samples (rays cast), and the maximum width of the DPSFs for the figures in this paper. Not all the DPSFs are used during blurring; the number that were actually used is given in parentheses.

4. Conclusion

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. We presented methods for achieving vision-realistic rendering and demonstrated those methods on sample images. Applications of vision-realistic rendering in computer graphics as well as in optometry and ophthalmology were presented.

Figure	# DPSFs	# Samples	DPSF Width	DPSF Computation Time	Convolution Time	Total Time
1	26 (11)	1 million	63	0:26	1:46	2:12
4	26 (11)	1 million	1	0:05	0:35	0:40
5(a)	26 (11)	1 million	127	0:26	2:03	2:29
5(b)	26 (11)	1 million	127	0:27	2:02	2:29
6(a)	136 (133)	1 million	33	0:24	19:30	19:54
6(b)	136 (133)	1 million	25	0:23	6:08	6:31
6(c)	136 (133)	1 million	41	0:24	8:35	8:59

Table 1: *Number of DPSFs computed, number of wavefront samples, maximum DPSF width, time to compute the DPSFs, time to perform convolution, and total time for vision-realistic rendering on a Pentium 4 running at 2.4GHz with 1 GB of RAM.*

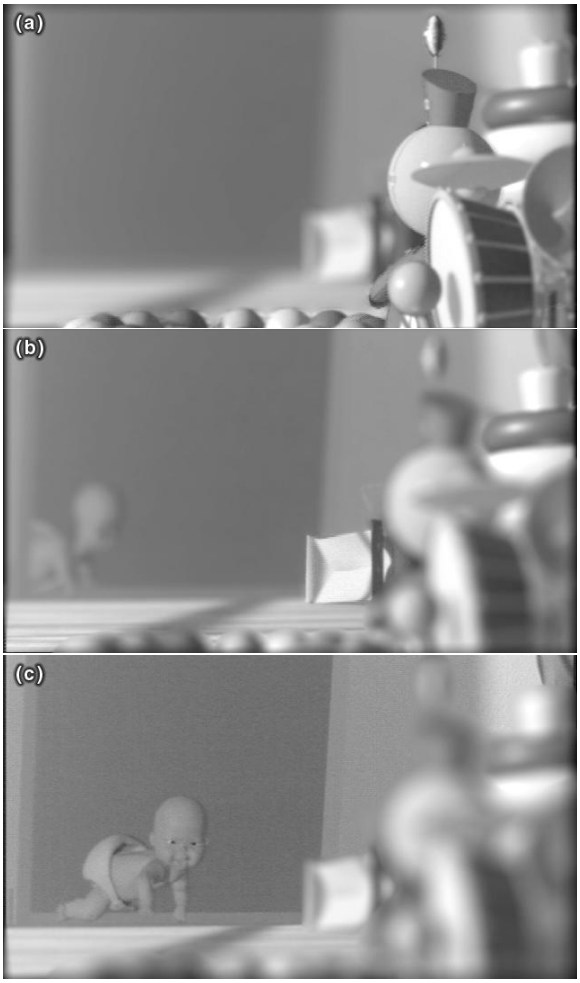

Figure 6: *Frames from a* ***rack focus sequence using a shot from Pixar's Tin Toy as input.***

5. Acknowledgments

The authors would like to thank the following University of California, Berkeley students: Viet Dinh Nguyen for helping to create the figures, Gabriel N. Garza for writing the DPSF generation software, and Vidya Iyer and Clark Zhe Li for creating the Tin Toy depth map. We are also grateful to Ian Cox, Ph.D., and Michele Lagana, O.D., at Bausch and Lomb for providing the Shack-Hartmann data used in this paper. Finally, we wish to thank Edwin E. Catmull, Christine Freeman, and Randy Nelson of Pixar Animation Studios for providing the Tin Toy shot and Prof. Carlo H. Séquin of the University of California, Berkeley for the sculptures used in our room scene.

This work was supported in part by the National Science Foundation grant number ASC-9720252, "Visualization and Simulation in Scientific Computing for the Cornea," and grant number CDA-9726362, "Virtual Environments for Telesurgery and Surgical, Training: Efficient Computation, Visualization, and Interaction."

References

1. Brian A. Barsky, Billy P. Chen, Alexander C. Berg, Maxence Moutet, Daniel D. Garcia, and Stanley A. Klein. Incorporating camera models, ocular models, and actual patient eye data for photo-realistic and vision-realistic rendering. Submitted for publication, 2001. 2

2. Mark R. Bolin and Gary W. Meyer. A perceptually based adaptive sampling algorithm. In Michael F. Cohen, editor, *ACM SIGGRAPH 1998 Conference Proceedings*, pages 299–309, Orlando, July 19–24 1998. SIGGRAPH, ACM. 2

3. Jon J. Camp, Leo J. Maguire, Bruce M. Cameron, and Richard A. Robb. A computer model for the evaluation

of the effect of corneal topography on optical performance. *Am. J. Ophthalmol.*, 109:379–386, 1990. 2

4. Jon. J. Camp, Leo J. Maguire, and Richard A. Robb. An efficient ray tracing algorithm for modeling visual performance from corneal topography. In *First Conference on Visualization in Biomedical Computing*, pages 279–285, Atlanta, GA, May 22–25 1990. The Institute of Electrical and Electronics Engineers, Inc., IEEE Computer Society Press. 2

5. K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman. Spatially non-uniform scaling functions for high contrast images. In *Proceedings of Graphics Interface '93*, pages 245–254, Toronto, May 1993. Canadian Information Processing Society. 2

6. Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. In Hank Christiansen, editor, *ACM SIGGRAPH 1984 Conference Proceedings*, pages 137–145, Minneapolis, July 23–27 1984. SIGGRAPH, ACM. 2

7. Mark A. Z. Dippe and Erling H. Wold. Antialiasing through stochastic sampling. In Brian A. Barsky, editor, *ACM SIGGRAPH 1985 Conference Proceedings*, pages 69–78, San Francisco, July 22–26 1985. SIGGRAPH, ACM. 2

8. James A. Ferwerda, Sumanta N. Pattanaik, Peter Shirley, and Donald P. Greenberg. A model of visual adaptation for realistic image synthesis. In Holly Rushmeier, editor, *ACM SIGGRAPH 1996 Conference Proceedings*, pages 249–258, New Orleans, August 4–9 1996. SIGGRAPH, ACM. 2

9. Elisabeth M. Fine and Gary S. Rubin. Effects of cataract and scotoma on visual acuity: A simulation study. *Optometry and Vision Science*, 76(7):468–473, July 1999. 2

10. Elisabeth M. Fine and Gary S. Rubin. The effects of simulated cataract on reading with normal vision and simulated central scotoma. *Vision Research*, 39(25):4274–4285, 1999. 2

11. Daniel D. Garcia, Brian A. Barsky, and Stanley A. Klein. CWhatUC: A visual acuity simulator. In *Proceedings of Ophthalmic Technologies VIII, SPIE International Symposium on Biomedical Optics*, pages 290–298, San Jose, CA, January 24–30 1998. SPIE. 2

12. J. E. Greivenkamp, J. Schweigerling, J. M. Miller, and M. D. Mellinger. Visual acuity modeling using optical raytracing of schematic eyes. *Am. J. Ophthalmol.*, 120:227–240, 1995. 2

13. Aaron Isaksen, Leonard McMillan, and Steven J. Gortler. Dynamically reparameterized light fields. In Kurt Akeley, editor, *Proceedings of ACM SIGGRAPH 2000*, pages 297–306, New Orleans, July 23–28 2000. SIGGRAPH, ACM. 2

14. Stanley A. Klein. Optimal corneal ablation for eyes with arbitrary Hartmann-Shack aberrations. *J. Opt. Soc. Am. A*, 15(9):2580–2588, 1998. 3

15. Craig Kolb, Don Mitchell, and Pat Hanrahan. A realistic camera model for computer graphics. In Robert Cook, editor, *ACM SIGGRAPH 1995 Conference Proceedings*, pages 317–324, Los Angeles, August 6–11 1995. SIGGRAPH, ACM. 2

16. Mark E. Lee, Richard A. Redner, and Samuel P. Uselton. Statistically optimized sampling for distributed ray tracing. In Brian A. Barsky, editor, *ACM SIGGRAPH 1985 Conference Proceedings*, pages 61–67, San Francisco, July 22–26 1985. SIGGRAPH, ACM. 2

17. Junzhong Liang. *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, Heidelberg, Germany, December 1991. 3

18. Jeffrey Lubin. A visual discrimination model for imaging system design and evaluation. In Eli Peli, editor, *Vision Models for Target Detection and Recognition*, volume 2, pages 245–357. World Scientific Publishing Co., Inc., Washington, DC, 1995. 2

19. Leo J. Maguire, Jon J. Camp, and Richard A. Robb. 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 1992*, pages 601 – 609, Chapel Hill, NC, October 13 – 16 1992. SPIE — The International Society for Optical Engineering, SPIE. 2

20. Leo J. Maguire, Ralph W. Zabel, Paula Parker, and Richard L. Lindstrom. Topography and raytracing analysis of patients with excellent visual acuity 3 months after excimer laser photorefractive keratectomy for myopia. *Refract. Corneal Surg.*, 8:122–128, March/April 1991. 2

21. Gary W. Meyer. Image synthesis and color vision. In David F. Rogers and Rae A. Earnshaw, editors, *Computer Graphics Techniques*, pages 45–77. Springer Verlag, New York, 1990. 2

22. Gary W. Meyer and Donald P. Greenberg. Color-defective vision and computer graphics displays. *IEEE Computer Graphics and Applications*, 8(5):28–40, September 1988. 2

23. Parry Moon and Domina Eberle Spencer. On the Stiles-Crawford effect. *J. Opt. Soc. Am.*, 34(6):319–329, 1944. 2

24. Sumanta N. Pattanaik, James A. Ferwerda, Mark D. Fairchild, and Donald P. Greenberg. A multiscale model of adaptation and spatial vision for realistic image display. In Michael F. Cohen, editor, *ACM SIGGRAPH 1998 Conference Proceedings*, pages 287–298, Orlando, 1998. SIGGRAPH, ACM. 2

25. Eli Peli. Test of a model of foveal vision by using simulations. *Journal of Optical Society of America*, 13(6):1131–1138, June 1996. 2

26. Fabio Pellacini and Donald P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In Kury Akeley, editor, *ACM SIGGRAPH 2000 Conference Proceedings*, pages 55–64, New Orleans, July 23-28 2000. SIGGRAPH, ACM. 2

27. Ben C. Platt and Roland V. Shack. Lenticular Hartmann-screen. Newsletter 5, 15, Optical Science Center, University of Arizona, 1971. 2

28. Michael Potmesil and Indranil Chakravarty. Synthetic image generation with a lens and aperture camera model. *ACM Transactions on Graphics*, 1(2):85–108, April 1982. (Original version in ACM SIGGRAPH 1981 Conference Proceedings, Aug. 1981, pp. 297-305.). 2

29. Przemyslaw Rokita. Fast generation of depth-of-field effects in computer graphics. *Computers & Graphics*, 17(5):593–595, September 1993. 2

30. Przemyslaw Rokita. Generating depth-of-field effects in virtual reality applications. *IEEE Computer Graphics and Applications*, 16(2):18–21, March 1996. 2

31. M. Shinya. Post-filtering for depth of field simulation with ray distribution buffer. In *Proceedings of Graphics Interface '94*, pages 59–66, Banff, Alberta, May 1994. Canadian Information Processing Society. 2

32. Greg Spencer, Peter Shirley, Kurt Zimmerman, and Donald P. Greenberg. Physically-based glare effects for digital images. In Robert Cook, editor, *ACM SIGGRAPH 1995 Conference Proceedings*, pages 325–334, Los Angeles, August 6–11 1995. SIGGRAPH, ACM. 2

33. Larry N. Thibos. Principles of Hartmann-Shack aberrometry. In V. Lakshminarayanan, editor, *Vision Science and its Applications*, volume 15, pages 163–169. OSA-TOPS, Washington, DC, 2000. 2

34. Jack Tumblin and Holly E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993. 2

35. Steven D. Upstill. *The Realistic Presentation of Synthetic Images: Image Processing in Computer Graphics*. PhD thesis, Computer Science Division, University of California, Berkeley, California, August 1985. 2

36. Gregory Ward-Larson, Holly Rushmeier, and Christine D. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, Oct. – Dec. 1997. 2