

SIGGRAPH Asia 2008

Scattering

Course notes

Diego Gutierrez*

Universidad de Zaragoza

Srinivasa G. Narasimhan†

Carnegie Mellon University

Henrik Wann Jensen†

University of California San Diego

Wojciech Jarosz§

University of California San Diego



*diegog@unizar.es

†henrik@cs.ucsd.edu

‡srinivas@cs.cmu.edu

§wjarosz@cs.ucsd.edu

1 About the Lecturers

Diego Gutierrez

University of Zaragoza

<http://giga.cps.unizar.es/~diegog>

Diego Gutierrez is an Associate Professor at the University of Zaragoza, in Spain, where he received his PhD in Computer Science. His areas of expertise include physically based global illumination (specializing in participating media), perception and novel image processing techniques. He's currently Chair of the SIGGRAPH Asia Sketches & Posters programme, and was Papers Chair for ACM Graphite in 2006. He's served on many other international conference committees, including SIGGRAPH (Sketches), Eurographics (Papers) or SIGGRAPH Asia (Courses).

Henrik Wann Jensen

University of California San Diego

<http://graphics.ucsd.edu/~henrik>

Henrik Wann Jensen is an Associate Professor at the University of California San Diego. His contributions to computer graphics include the photon mapping algorithm for global illumination, and the first technique for efficiently simulating subsurface scattering in translucent materials. In 2004, Professor Jensen received an Academy Award (Technical Achievement Award) from the Academy of Motion Picture Arts and Sciences for pioneering research in rendering translucent materials. He also became a Sloan Fellow and he was selected as one of the top 10 scientists in 2004 by Popular Science magazine.

Srinivasa G. Narasimham

Carnegie Mellon University

<http://www.cs.cmu.edu/~srinivas/>

Srinivasa Narasimhan is an Assistant Professor in the School of Computer Science at Carnegie Mellon University (since 2004). He obtained his masters and doctoral degrees in Computer Science from Columbia University in 2000 and 2004 respectively. His research interests are at the intersection of computer vision, computer graphics and optics. He received a Best Paper award in IEEE CVPR 2000 and a CAREER award from NSF in 2007.

Wojciech Jarosz

University of California San Diego

<http://graphics.ucsd.edu/~wjarosz>

Wojciech Jarosz is currently a post-doc at the University of California San Diego. His main research interest is production-quality global illumination techniques, specifically for participating media, and his current list of publications include several SIGGRAPH papers on those topics. He received his B.S. in Computer Science from University of Illinois, Urbana-Champaign and both his M.S. and PhD in Computer Science from UC San Diego.

2 Course Syllabus

09:00	<i>Gutierrez</i>	Welcome and Introduction	(15 min)
09:15	<i>Jarosz</i>	Rendering Scattering Media	(40 min)
09:55	<i>Narasimhan</i>	Real-Time Rendering	(20 min)
10:15	<i>Jensen</i>	Scattering Materials	(30 min)
10:45	 <i>Break</i>	(15 min)
11:00	<i>Gutierrez</i>	Inelastic Scattering	(15 min)
11:15	<i>Narasimhan</i>	Underwater Imaging	(20 min)
11:35	<i>Narasimhan</i>	Scattering and Vision	(25 min)
12:00	<i>Jensen</i>	Acquisition and Measurement	(30 min)
12:30	<i>All</i>	Wrap up and Discussion	(15 min)

3 Course Abstract

Computer graphics and computer vision techniques deal with acquiring, interpreting and presenting the rich visual world around us. These are exciting multi-disciplinary fields of research with a wide spectrum of applications that can impact our daily lives. However, most of the computer generated imagery today represents scenes with clear atmospheres, neglecting light scattering effects. Analogously, most computer vision systems have not enjoyed success when deployed in uncontrolled outdoor environments. Nevertheless, scattering is a fundamental aspect of light transport in a wide range of applications, whether simulating it or interpreting it, from medical imaging to driving simulators or underwater imagery.

In this course we address the challenges that arise when faced with light scattering both in a computer graphics and a computer vision context. Both fields have seen great advances over the past few years; however, most of the existing algorithms still assume that light emitted by a source or reflected off a surface reaches the sensor unaltered. From a computer graphics perspective, this is due mainly to the complex interactions that occur and the high computational costs of simulating them. In computer vision, scattering has traditionally been considered as noise that one should ideally get rid of. Scattering effects are one fundamental hurdle that must be overcome to significantly extend and enhance current state-of-the-art graphics and vision techniques and achieve successful impact in a wide range of domains. Given the increasing overlapping of both fields, including new and hot research fields such as computational photography, we believe the course will be useful for both communities and everybody with research interests at the intersection of the two. We hope to increase awareness about this area and open up new research directions. Topics discussed include appearance modeling, underwater imagery, vision in bad weather and measurement techniques.

This course is intended for people involved in computer graphics, computer vision or related fields such as computational photography. It will provide a good understating of scattering phenomena, state-of-the-art techniques to simulate it and treat it in computer graphics and computer vision contexts, and a wide range of applications. More precisely, it is targeted to delegates interested in general fields such as computer graphics, computer vision or computational photography, or in more particular applications such as medical imaging, oceanography, driving simulators, game developers




Figure 1: The theory of light is described by a series of increasingly complete optical models, where each successive model is able to account for more optical phenomena. In computer graphics, we often restrict ourselves to the simplest model, ray optics.

4 Introduction

Our current understanding of the behavior of light relies on a progression of increasingly complete yet complicated models of light. These are (see Figure 1): ray optics, wave optics, electromagnetic optics, and quantum optics [14]. Computer graphics typically relies on the simplest of these models, ray optics (also called geometric optics). This model makes several simplifying assumptions about the behavior of light that limit the types of phenomena that can be simulated. In essence, in this model light can only be emitted, reflected, and transmitted. Additionally, light is assumed to travel in straight lines and at infinite speed. This means that effects explained by the higher-level models cannot (easily) be incorporated into our simulations. In ray optics, effects such as diffraction and interference (wave optics), polarization and dispersion (electromagnetic optics), and fluorescence and phosphorescence (quantum optics) are completely ignored. In spite of these limitations, we are still able to correctly simulate a wide range of physical phenomena.

In most computer graphics applications, assumptions are made about the properties of the scattering media in order to more easily derive expressions about the behavior of light. In particular, we assume that the participating media can be modeled as a collection of microscopic particles. Since the particles are microscopic and randomly positioned, we do not represent each individual particle in the lighting simulation. Instead, we consider the aggregate probabilistic behavior as light travels through the medium. Moreover, these particles are assumed to be spaced far apart relative to the size of an individual particle. This assumption implies that as a photon travels through the medium and interacts at a particle, this interaction is *statistically independent* from the outcome of subsequent interaction events (Figure 2).

5 Real-time Rendering Techniques for Participating Media

Most of computer generated imagery today in video games, movies, and scientific simulations are of scenes on clear days or nights. Volumetric effects such as the beautiful fog rolling down the hills, the bluish haze of mountains, the eerie night fog or mist reminiscent of Hitchcockian movies, the splendor and brilliance of underwater effects, the light streaming through clouds or the sun rising over the ocean provide pure artistic and entertainment value. They are used in movies and paintings to portray different moods, used in photographs to provide realism, and used even for training in safety and hazardous situations. In the




Figure 2: We treat participating media as a collection of microscopic scattering particles. When lights travels through the medium, a change of radiance, $(\vec{\omega} \cdot \nabla)L(\mathbf{x}, \vec{\omega})$, may occur as the photons interact with the particles.

absence of such effects, current attempts at renderings appear unnatural and cartoonish. Thus, it becomes critical to render these effects accurately for achieving photo-realism.



Figure 3: **Real time rendering of participating media:** (a)-(b) Real time rendering of homogenous media using an analytic single-scattering model of light transport. (a) Clear scene. (b) Scene with fog added. (c)-(d) Fast rendering of smooth, non-homogenous and dynamic media. (c) Clear day scene. (d) Scene with fog added.

Rendering of participating media requires solving complex light transport equations. Brute-force simulations of light transport based on Monte Carlo and finite element simulation can be prohibitively slow (taking CPU-days or even weeks). However, a variety of applications spanning entertainment (video games), medicine (surgery) and autonomous navigation require real-time or interactive performance. Thus, a recent research thrust has been to make smoothness assumptions, either on the media [2, 18] or the lighting [17] in order to achieve interactive rates. Sun et al [18] have implemented their model in modern programmable graphics hardware using a few small numerical lookup tables stored as texture maps, and achieved truly real time performance. Gupta et al [2] present a technique for fast rendering of non-homogenous, as well as dynamic media, by representing the density and intensity fields in a low-dimensional basis. Sloan et al [17] use the concept of *pre-computed radiance transfer* to allow for real-time changes in the environment lighting. Snapshots from a couple of representative papers are shown in Figure 3.

6 Measuring Scattering Properties of Participating Media

The appearance of participating media is governed by their optical properties, which must be input to a rendering algorithm to generate realistic images. Even with the most accurate rendering algorithms, the image quality is often limited by the accuracy of these input parameters. Narasimhan et al presented a simple device and technique [9] for robustly estimating the scattering properties of a broad class of participating media such as juices, beverages, sugar/salt crystals, and suspended impurities (ocean water) from a single

HDR photograph. They measured the scattering parameters of forty commonly found materials. The results are compiled into a freely available database which can be immediately used by the computer graphics community to render realistic images of arbitrary concentrations of the material with multiple scattering (Figure 4), as well as create realistic images of combinations of the original materials. This technique can be used to design portable devices as well, that can be used for in-situ measurements of impurity levels in natural water bodies (oceans, lakes, rivers) for environmental monitoring.




Figure 4: **Measuring scattering properties of participating media:** Renderings of a scene with four liquids in their diluted states (b) and their natural high density states (c). The corresponding input HDR images are shown in (a).

Debevec et al presented a technique [6] for capturing time-varying volumetric data of participating media. In their technique, a laser sheet is swept repeatedly through the volume, and the scattered light is imaged using a high-speed camera. Each sweep of the laser provides a near-simultaneous volume of density values. They demonstrated rendered animations under changing viewpoint and illumination, making use of measured values for the scattering phase function and albedo. An example rendering is shown in Figure 5.

7 Computer Vision in Bad Weather

In recent years, computer vision has seen significant advances in the core areas of image sensing and interpretation. This success has resulted in great demand for vision techniques in application domains ranging from intelligent transportation and security to oceanography (underwater imaging), to astronomy (telescope and satellite imaging), to even biology and medical systems (microscopic and medical imaging). However, unfortunately, most computer vision systems have not enjoyed success when deployed in uncontrolled outdoor environments.

Images of outdoor scenes captured in bad weather suffer from poor contrast. Under bad weather conditions, the light reaching a camera is severely scattered by the atmosphere. The resulting decay in contrast varies across the scene and is exponential in the depths of scene points. Therefore, traditional space invariant image processing techniques are not sufficient to remove weather effects from images. In such scenarios, physics-based models that describe the image-formation process in bad weather conditions can be used to restore contrast of images [12, 10, 11, 16]. Moreover, changes in intensities of scene points under different weather conditions can be used as cues to compute scene structure (see Figures 6 and 7). These techniques



Figure 5: **Measuring scattering properties of dynamic participating media:** Measured media can be rendered under a variety of lighting conditions. (a) A captured smoke volume rendered with two spotlights of different colors. (c) A smoke volume rendered with environmental illumination. Images taken from [6].

can be applied to gray scale, RGB color, multispectral and even IR images. More recently, Fattal et al [1] presented a new method for dehazing the image and estimating the scene structure, given only a single input image, as compared to previous techniques which required multiple images. Sample results are shown in Figure 9. This technique has been used for applications such as image refocusing and novel view synthesis.




Figure 6: **Contrast restoration using scattering-based image formation model:** (a) Image of a traffic scene on a foggy day. (b) The defogged (contrast restored) image. (c) Depth map computed from two foggy images. . Experiments with videos of a traffic scene on a foggy day. Images taken from [12].

Multiple Scattering: Virtually all the previous methods in image processing and computer vision, for removing weather effects from images, assume single scattering of light by particles in the atmosphere. In reality, multiple scattering effects are significant. A common manifestation of multiple scattering is the appearance of glows around light sources in bad weather. Modeling multiple scattering is critical to understanding the complex effects of weather on images, and hence essential for improving the performance of outdoor vision systems. Narasimhan et al [13] develop a new physics-based model for the multiple scattering of light rays as they travel from a source to an observer. This model is valid for various weather conditions including fog, haze, mist and rain. Using this model, the shapes and depths of sources in the scene can be recovered. In addition, the weather condition and the visibility of the atmosphere can be estimated, thus making a camera observing a distant source as a *visual weather meter*.

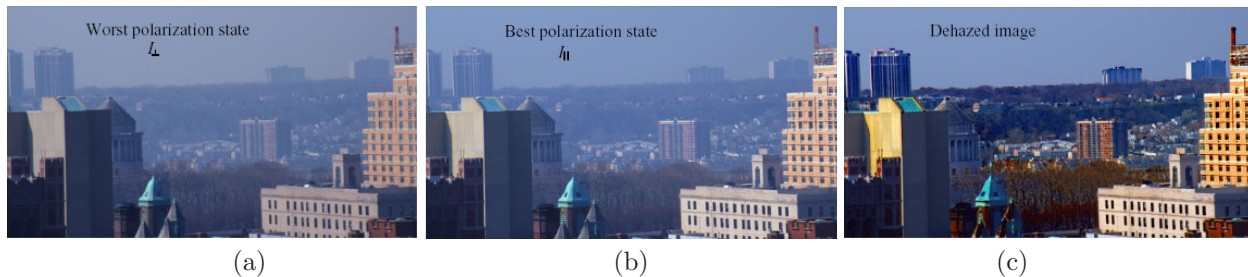


Figure 7: **Contrast restoration using polarization imaging:** Images of the (a) perpendicular and (b) parallel polarization components. The parallel component has the best image contrast that optics alone can give, but in this case it is only slightly better than the contrast in the image of the worst polarization state. (c) The dehazed image has much better contrast and color than what optical filtering alone yields, especially in the distant regions of the scene. Images taken from [16].



Figure 8: Dehazing based on a single input image and the corresponding depth estimate. Images taken from [1].

8 A bio-optical model of ocean water

Simulating the in-water ocean light field is a daunting task. Ocean waters are one of the richest participating media, where light interacts not only with water molecules, but with suspended particles and organic matter as well. The concentration of each constituent greatly affects these interactions, resulting in very different hues. Inelastic scattering events such as fluorescence or Raman scattering imply energy transfers that are usually neglected in the simulations. A bio-optical model of ocean waters is presented in [5], along with a method to obtain the in-water light field based on radiative transfer theory. The bio-optical model of the ocean uses published data from oceanography studies. The method builds on [4], and provides a link between the *inherent optical properties* that define the medium and its *apparent optical properties*, which describe how it looks. For inelastic scattering, all frequency changes at higher and lower energy values are taken into account, based on the spectral quantum efficiency function of the medium. Areas of application for this research span from underwater imagery to remote sensing; the resolution method is general enough to be usable in any type of participating medium simulation. Figure ?? shows some results varying the concentration of some of the components of the model: chlorophyll, minerals and detritus and yellow matter.



Figure 9: Resulting pictures varying the chlorophyll concentration C , the minerals and detritus turbidity α_d at 400nm and the CDOM turbidity α_y at 440nm.

9 Underwater Imaging

Poor visibility conditions due to murky water, bad weather, dust and smoke severely impede the performance of vision systems. A variety of passive methods, most notably based on polarization analysis have been used to restore scene contrast under moderate visibility by digital postprocessing [15]. However, such methods are ineffective when the quality of acquired images is poor to begin with. Active vision systems control light transport even before the image is formed, and hence result in superior performance [8, 3, 7].

Most of the underwater active vision systems are based on using structured lighting. Traditionally, structured light methods assume that the scene and the sources are immersed in pure air and that light is neither scattered nor absorbed. Recently, however, structured lighting has found growing application in underwater and aerial imaging, where scattering effects cannot be ignored. A variety of structured light techniques, such as confocal imaging, structured light striping, polarized light striping have been used to recover scene albedos, depths and normals reliably, even in murky underwater conditions. These results have been used to restore the appearances of scenes as if captured in clear air. A compilation of results is shown in Figure 10.




Figure 10: **Contrast restoration of underwater images:** (a) Contrast restoration using passive polarization based technique. Image taken from [15]. (b) and (c) Recovering contrast using active illumination, such as polarized structured light striping. Images taken from [3, 8].

References

- [1] R. FATTAL, *Single image dehazing*, ACM Transactions on Graphics (SIGGRAPH 2008), 27.
- [2] M. GUPTA AND S. G. NARASIMHAN, *Legendre fluids: A unified framework for analytic reduced space modeling and rendering of participating media*, in Eurographics/ ACM SIGGRAPH Symposium on Computer Animation (2007), August 2007.
- [3] M. GUPTA, S. G. NARASIMHAN, AND Y. Y. SCHECHNER, *On controlling light transport in poor visibility environments*, in Proceedings IEEE CVPR, June 2008.
- [4] D. GUTIERREZ, A. MUNOZ, O. ANSON, AND F. SERON, *Non-linear volume photon mapping*, Rendering Techniques (Eurographics Symposium on Rendering), pp. 291–300.
- [5] D. GUTIERREZ, F. SERON, O. ANSON, AND A. MUNOZ, *Visualizing underwater ocean optics*, Computer Graphics Forum (EUROGRAPHICS 2008), 27, pp. 547–556.
- [6] T. HAWKINS, P. EINARSSON, AND P. DEBEVEC, *Acquisition of time-varying participating media*, ACM Transactions on Graphics (SIGGRAPH 2005), 24 (2005).
- [7] M. LEVOY, B. CHEN, V. VAISH, M. HOROWITZ, I. McDOWALL, AND M. BOLAS, *Synthetic aperture confocal imaging*, ACM Transactions on Graphics (SIGGRAPH 2004), 23, pp. 825–834.
- [8] S. NARASIMHAN AND S. NAYAR, *Structured light methods for underwater imaging: light stripe scanning and photometric stereo*, OCEANS, 2005. Proceedings of MTS/IEEE, (2005), pp. 2610–2617 Vol. 3.
- [9] S. G. NARASIMHAN, M. GUPTA, C. DONNER, R. RAMAMOORTHY, S. K. NAYAR, AND H. W. JENSEN, *Acquiring scattering properties of participating media by dilution*, ACM Transactions on Graphics (SIGGRAPH 2006), 25 (2006).


- [10] S. G. NARASIMHAN AND S. NAYAR, *Interactive deweathering of an image using physical models*, in IEEE IEEE Workshop on Color and Photometric Methods in Computer Vision, In Conjunction with ICCV, October 2003.
- [11] S. G. NARASIMHAN AND S. K. NAYAR, *Vision and the atmosphere*, IJCV, 48 (2002), pp. 233–254.
- [12] S. G. NARASIMHAN AND S. K. NAYAR, *Contrast restoration of weather degraded images*, IEEE PAMI, 25 (2003), pp. 713 – 724.
- [13] S. G. NARASIMHAN AND S. K. NAYAR, *Shedding light on the weather*, in Proceedings of the 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, June 2003, pp. 665 – 672.
- [14] B. E. A. SALEH AND M. C. TEICH, *Fundamentals of Photonics*, Wiley-Interscience, 2 ed., March 2007.
- [15] Y. Y. SCHECHNER AND N. KARPEL, *Recovery of underwater visibility and structure by polarization analysis*, Oceanic Engineering, IEEE Journal of, 30 (July 2005), pp. 570–587.
- [16] Y. Y. SCHECHNER, S. G. NARASIMHAN, AND S. K. NAYAR, *Instant dehazing of images using polarization*, in Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, June 2001, pp. 325 – 332.
- [17] P.-P. SLOAN, J. KAUTZ, AND J. SNYDER, *Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments*, ACM Transactions on Graphics (SIGGRAPH 2002), 21 (2002).
- [18] B. SUN, R. RAMAMOORTHY, S. G. NARASIMHAN, AND S. K. NAYAR, *A practical analytic single scattering model for real-time rendering*, ACM Transactions on Graphics (SIGGRAPH 2005), (2005).

RENDERING PARTICIPATING MEDIA

WOJCIECH JAROSZ
UNIVERSITY OF CALIFORNIA, SAN DIEGO

COURSE: SCATTERING
SIGGRAPH ASIA 2008

FOG



<http://mattmosher.com/>

CLOUDS & CRESPUCULAR RAYS



3

ANTELOPE CANYON, AZ.



4

AERIAL (ATMOSPHERIC) PERSPECTIVE




Wojciech Jarosz



Henrik Wann Jensen

5


LEONARDO DA VINCI (1480)



“Thus, if one is to be five times as distant, make it five times bluer.”
—*Treatise on Painting*, Leonardo Da Vinci, pp 295, circa 1480.

6

NEBULA



T.A.Rector (NOAO/AURA/NSF) and the Hubble Heritage Team (STScI/AURA/NASA)


7

OUTLINE

- Theoretical background
- Rendering general participating media fast:
 - Radiance Caching for Participating Media
 - Volumetric photon mapping
 - The traditional approach
 - The beam radiance estimate


8

PARTICIPATING MEDIA




9

SCENE WITH MEDIUM



10


VOLUME RENDERING EQN



$$L(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) dt + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega})$$

11

VOLUME RENDERING EQN




$$L(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) dt + \boxed{T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega})}$$

Transmittance: $T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) = \exp \left(- \int_0^s \sigma_t(\mathbf{x} + t\vec{\omega}) dt \right)$

12


VOLUME RENDERING EQN



$$L(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) dt + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega})$$

13

VOLUME RENDERING EQN




$$L(\mathbf{x}, \vec{\omega}) = \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) dt + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega})$$

$$L_i(\mathbf{x}_t, \vec{\omega}) = \int_{\Omega_{4\pi}} p(\mathbf{x}_t, \vec{\omega}, \vec{\omega}_t) L(\mathbf{x}_t, \vec{\omega}_t) d\omega_t$$

14

VOLUME RENDERING EQN



$$L(\mathbf{x}, \vec{\omega}) = \int_0^s \boxed{T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t)} L_i(\mathbf{x}_t, \vec{\omega}) dt + T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega})$$

15

AVAILABLE TECHNIQUES

Finite Element

- Zonal Method [Rushmeier and Torrance 87; Bhate and Tokuta 92]
- Diffusion [Stam 95]

- Requires discretization

Monte Carlo


- Path tracing [Kajiya and Herzen 84; Kajiya 86; Lafortune and Willems 96]
- Metropolis [Pauly, Kollig, and Keller 00]
- Path Integration [Premože 03]

- Slow convergence / noisy results.

Biased Monte Carlo

16

VOLUMETRIC PHOTON MAPPING



17

OUTLINE

- Theoretical background
- Rendering general participating media fast:
 - Radiance Caching for Participating Media
 - Volumetric photon mapping
 - The traditional approach
 - The beam radiance estimate

18

RELATED WORK


Global Illumination

Caching:

- “A Ray Tracing Solution for Diffuse Interreflection.” Ward et al. 1988.
- “Irradiance Gradients.” Ward and Heckbert. 1992.
- “Radiance Caching for Efficient Global Illumination Computation.” Křivánek et al. '05

19

INDIRECT ILLUMINATION



Direct Illumination




Indirect Illumination

20

IRRADIANCE CACHING


Ward et al. '88



21

IRRADIANCE GRADIENTS


Ward and Heckbert '92



22

IRRADIANCE GRADIENTS

Ward and Heckbert '92




22


IRRADIANCE GRADIENTS

Ward and Heckbert '92

Without Gradients



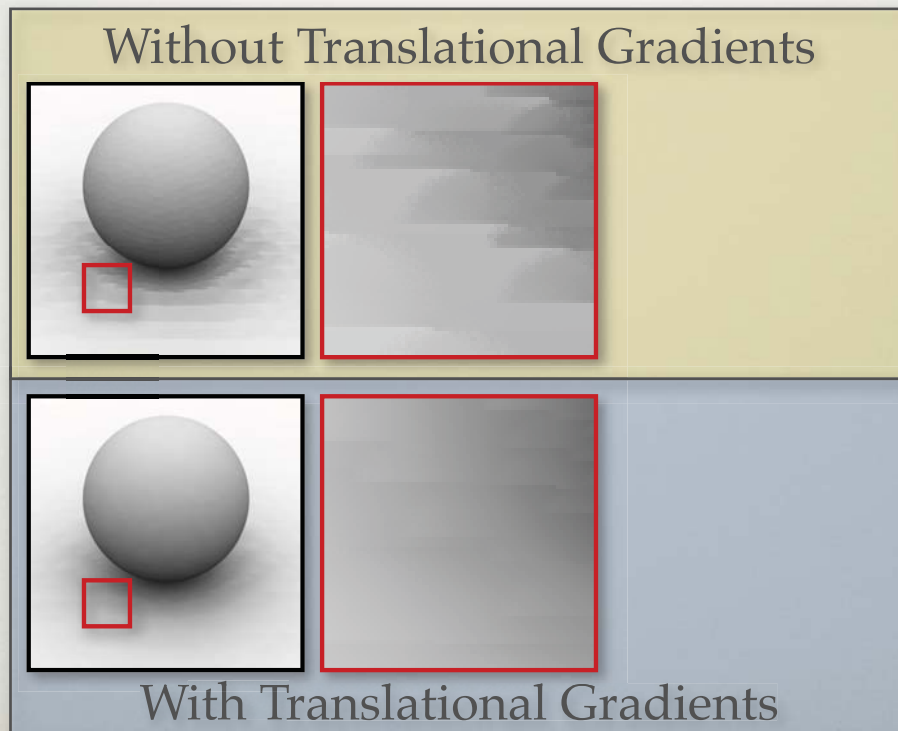
With Gradients



23

IRRADIANCE GRADIENTS


Ward and Heckbert '92



24

IRRADIANCE GRADIENTS

Ward and Heckbert '92



25

OBSERVATIONS



Smooth in large portions of the image


26

GOALS

- Exploit this property by caching lighting within participating media.
- Develop an efficient but general rendering algorithm which can handle:
 - single, multiple, anisotropic scattering
 - heterogeneous media
 - production quality


27

RADIANCE CACHING IN PARTICIPATING MEDIA



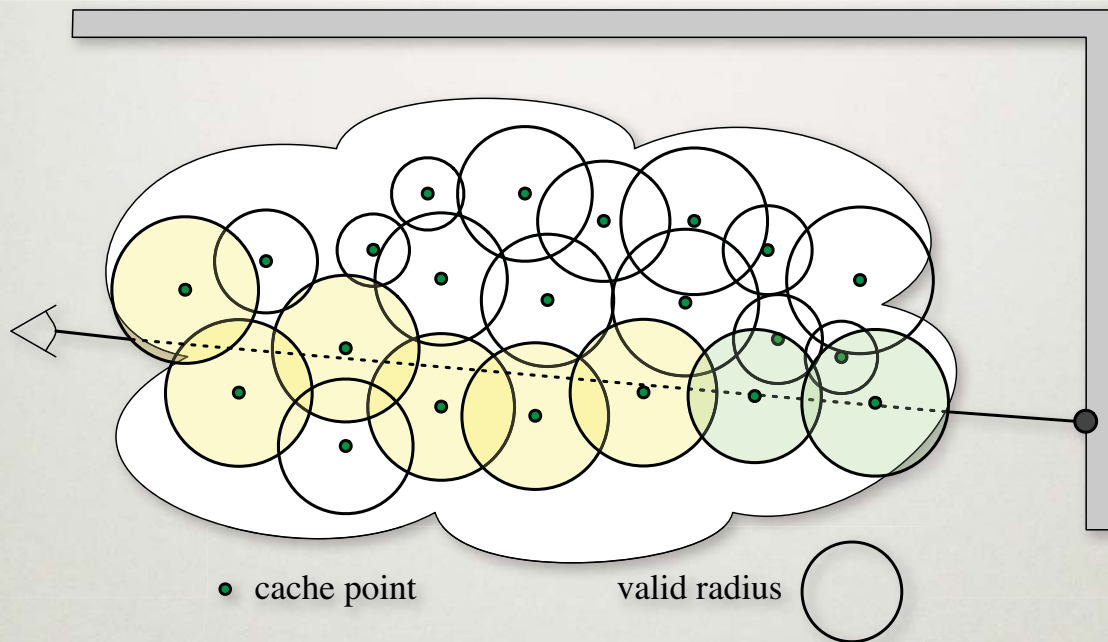
28

RADIANCE CACHING IN PARTICIPATING MEDIA



29

RADIANCE CACHING IN PARTICIPATING MEDIA



30

CHALLENGES

- What should the cache points store?
- Where to place cache points to minimize visible error?
- How to interpolate cache points accurately?

31

APPROACH

- Cache inscattered radiance:

$$L(\mathbf{e}, \vec{\omega}) = \int_0^s T_r(\mathbf{e} \leftrightarrow \mathbf{x}) \sigma_s(\mathbf{x}) L_i(\mathbf{x}, \vec{\omega}) d\mathbf{x} + T_r(\mathbf{e} \leftrightarrow \mathbf{x}') L(\mathbf{x}', \vec{\omega})$$

APPROACH

- Cache inscattered radiance:

$$L(\mathbf{e}, \vec{\omega}) = \int_0^s T_r(\mathbf{e} \leftrightarrow \mathbf{x}) \sigma_s(\mathbf{x}) L_i(\mathbf{x}, \vec{\omega}) d\mathbf{x} + T_r(\mathbf{e} \leftrightarrow \mathbf{x}') L(\mathbf{x}', \vec{\omega})$$

- Compute gradients due to translation

APPROACH

- Cache inscattered radiance:

$$L(\mathbf{e}, \vec{\omega}) = \int_0^s T_r(\mathbf{e} \leftrightarrow \mathbf{x}) \sigma_s(\mathbf{x}) L_i(\mathbf{x}, \vec{\omega}) d\mathbf{x} + T_r(\mathbf{e} \leftrightarrow \mathbf{x}') L(\mathbf{x}', \vec{\omega})$$

- Compute gradients due to translation
- Use gradients to:
 - Estimate valid radius within which it's OK to extrapolate
 - Provide high quality interpolation

32

RADIANCE COMPUTATION


- In order to make gradient derivations more convenient:
 - Split computation into single and multiple scattering components:

$$L_i = L_s + L_m$$

- How do we compute L_s and L_m ?
- How do we compute ∇L_s and ∇L_m ?

33


SINGLE SCATTERING



$$L_s(\mathbf{x}, \vec{\omega}) = \int_A p(\vec{\omega}, \mathbf{x}' \rightarrow \mathbf{x}) L_r(\mathbf{x}' \rightarrow \mathbf{x}) V(\mathbf{x}' \leftrightarrow \mathbf{x}) H(\mathbf{x}' \rightarrow \mathbf{x}) d\mathbf{x}'$$

34

SINGLE SCATTERING



$$L_s(\mathbf{x}, \vec{\omega}) = \int_A p(\vec{\omega}, \mathbf{x}' \rightarrow \mathbf{x}) L_r(\mathbf{x}' \rightarrow \mathbf{x}) V(\mathbf{x}' \leftrightarrow \mathbf{x}) H(\mathbf{x}' \rightarrow \mathbf{x}) d\mathbf{x}'$$

$$\nabla L_s(\mathbf{x}, \vec{\omega}) = \int_A (\nabla p) L_r V H + p(\nabla L_r) V H + p L_r V (\nabla H) d\mathbf{x}'$$

34

CACHE STORAGE

- Cached points store:
 - 3D position
 - Value (inscattered radiance)
 - Gradient
 - Valid Radius

35

CACHE STORAGE

Isotropic Media

- Cached points store:
 - 3D position
 - Value
 - Gradient
 - Valid Radius
- } ▶ inscattered radiance is a scalar

36



CACHE STORAGE

Anisotropic Media

- Cached points store:
 - 3D position
 - Value
 - Gradient
 - Valid Radius
- } ▶ inscattered radiance is a spherical function
▶ projected onto SH

37

VALID RADIUS




38

RESULTS

- All results rendered:
 - at 1K horizontal resolution
 - with up to 16 samples per pixel
 - on a 2.4 GHz Intel Core 2 Duo (one core)

39


RESULTS



1.4 minutes

40



RESULTS



3.6 minutes

41

RESULTS



19 minutes


42

RESULTS



43


RESULTS



5.8 minutes


44

RESULTS



45

RESULTS



20 minutes

46

RESULTS



contrast enhanced

47

QUESTIONS?


48

OUTLINE

- Theoretical background
- Rendering general participating media fast:
 - Radiance Caching for Participating Media
 - Volumetric photon mapping
 - The traditional approach
 - The beam radiance estimate


49

VOLUMETRIC PHOTON MAPPING




50

VOLUMETRIC PHOTON MAPPING



51


VOLUME RENDERING EQ.



$$L(\mathbf{x}, \vec{\omega}) = T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + \int_0^s T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) dt$$

52

RAY MARCHING



$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + \left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) \Delta_t \right)$$

53

VOLUMETRIC PHOTON MAPPING

Conventional Radiance Estimate

$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + \left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) \Delta_t \right)$$

54

VOLUMETRIC PHOTON MAPPING

Conventional Radiance Estimate


$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + \left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) \Delta_t \right)$$

$$L_i(\mathbf{x}_t, \vec{\omega}) \approx \sum_{p=1}^n \frac{p(\mathbf{x}_t, \vec{\omega}, \vec{\omega}_p) \Delta \Phi_p}{\frac{4}{3} \pi r^3}$$

54

VOLUMETRIC PHOTON MAPPING


Conventional Radiance Estimate



55

VOLUMETRIC PHOTON MAPPING


Conventional Radiance Estimate



56

VOLUMETRIC PHOTON MAPPING


Conventional Radiance Estimate



57

VOLUMETRIC PHOTON MAPPING


Conventional Radiance Estimate



58

VOLUMetric PHOTON MAPPING


Conventional Radiance Estimate



59


VOLUMETRIC PHOTON MAPPING

Conventional Radiance Estimate



60


DRAWBACKS



61

DRAWBACKS

- Radiance estimation is expensive
- Requires range search in photon map
- Performed numerous times per ray



61

DRAWBACKS


Large Step-size



62

DRAWBACKS

Large Step-size




62

DRAWBACKS

Large Step-size



Very Small Step-size



62


GOAL

- Render high-quality, noise-free images using photon mapping, *faster*.

GOAL


- Render high-quality, noise-free images using photon mapping, *faster*.
- Eliminate ray marching by finding *all* photons which contribute to the *entire* length of a ray.

ADAPTIVE?



64


ADAPTIVE?



64

PRIMAL VS. DUAL


Primal




65

PRIMAL VS. DUAL

Primal




Dual




65

PRIMAL VS. DUAL

Primal



Dual




allow radius to vary:
adaptive kernel method

65

VOLUMETRIC PHOTON MAPPING


Beam Radiance Estimate



66

VOLUMETRIC PHOTON MAPPING


Beam Radiance Estimate



67

VOLUMETRIC PHOTON MAPPING


Beam Radiance Estimate



68

VOLUMETRIC PHOTON MAPPING






Beam Radiance Estimate



69

ADAPTIVE RADIUS COMPARISON

Conv. Estimate Conv. Estimate Beam Estimate






70

ADAPTIVE RADIUS COMPARISON

Conv. Estimate

Conv. Estimate

Beam Estimate



(∞)

(6:38)

(6:22)

70

ALGORITHM

- 1) Shoot photons from light sources.
- 2) Construct a balanced kD-tree for the photons.
- 3) Assign a radius for each photon (*photon-discs*).
- 4) Create acceleration structure over photon-discs.
- 5) Render:
 - For each ray through the medium, accumulate all *photon-discs* that intersect ray.

ALGORITHM

SAME AS REGULAR PHOTON MAPPING

- 1) Shoot photons from light sources.
- 2) Construct a balanced kD-tree for the photons.
- 3) Assign a radius for each photon (*photon-discs*).
- 4) Create acceleration structure over photon-discs.
- 5) Render:
 - For each ray through the medium, accumulate all *photon-discs* that intersect ray.

71

ALGORITHM


SAME AS REGULAR PHOTON MAPPING

- 1) Shoot photons from light sources.
- 2) Construct a balanced kD-tree for the photons.
- 3) Assign a radius for each photon (*photon-discs*).
- 4) Create acceleration structure over photon-discs.
- 5) Render:
 - For each ray through the medium, accumulate all *photon-discs* that intersect ray.

BEAM GATHERING

71


ALGORITHM



1) Shoot photons from light sources.

72


ALGORITHM



2) Construct a balanced kD-tree for the photons.

73


ALGORITHM



2) Construct a balanced kD-tree for the photons.

74


ALGORITHM



3) Assign a radius for each photon (*photon-discs*).
Adaptive: perform k-NN search at each photon

75

ALGORITHM





4) Create a bounding-box hierarchy over photon-discs

76

ALGORITHM


kD-tree



4) Create a bounding-box hierarchy over photon-discs

77


ALGORITHM



- 4) Create a bounding-box hierarchy over photon-discs
reuse hierarchical structure of kD-tree

77

ALGORITHM




- 5) Render: For each ray through the medium, accumulate all photon-discs that intersect ray.

78

VOLUMETRIC PHOTON MAPPING

Conventional Radiance Estimate

$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + \left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) L_i(\mathbf{x}_t, \vec{\omega}) \Delta_t \right)$$




79

VOLUMETRIC PHOTON MAPPING

Conventional Radiance Estimate

$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + \left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) \sigma_s(\mathbf{x}_t) \boxed{L_i(\mathbf{x}_t, \vec{\omega})} \Delta_t \right)$$




79

VOLUMETRIC PHOTON MAPPING

Conventional Radiance Estimate

$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) +$$

~~$$\left(\sum_{t=0}^{S-1} T_r(\mathbf{x} \leftrightarrow \mathbf{x}_t) L_s(\mathbf{x}_t, \vec{\omega}) \Delta_t \right)$$~~




80

VOLUMETRIC PHOTON MAPPING

Beam Radiance Estimate

$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) +$$

$$L_b(\mathbf{x}, \vec{\omega})$$




81

VOLUMETRIC PHOTON MAPPING

Beam Radiance Estimate

$$L(\mathbf{x}, \vec{\omega}) \approx T_r(\mathbf{x} \leftrightarrow \mathbf{x}_s) L(\mathbf{x}_s, \vec{\omega}) + L_b(\mathbf{x}, \vec{\omega})$$

$$L_b(\mathbf{x}, \vec{\omega}) = \frac{1}{N} \sum_{i=1}^N K_i T_r(\mathbf{x} \leftrightarrow \mathbf{x}_i) \sigma_s(\mathbf{x}_i) p(\mathbf{x}_i, \vec{\omega}, \vec{\omega}_i) \alpha_i$$



81

RESULTS

- 1K horizontal resolution
- 2.4 GHz Core 2 Duo (using one Core)
- Comparing *identical* photon maps

SMOKY CORNELL BOX

Conv. Estimate




Beam Estimate




83

SMOKY CORNELL BOX


Conv. Estimate



Beam Estimate



(4:03)



(3:35)

83

LIGHTHOUSE

Beam Estimate



Conventional Estimate

84

LIGHTHOUSE

Beam Estimate



Conventional Estimate

(1:05)



(1:12)

84

CARS ON FOGGY STREET

Beam Estimate




Conventional Estimate

85

CARS ON FOGGY STREET

Beam Estimate

(1:53)



Conventional Estimate

(2:02)

85

SUMMARY

- Radiance caching for participating media
- Volumetric photon mapping
 - Which should you use?

QUESTIONS?

Acquiring Scattering Properties of Participating Media by Dilution

Srinivasa Narasimhan

Craig Donner

Ravi Ramamoorthi

Mohit Gupta

Henrik Wann Jensen

Shree Nayar

Carnegie Mellon
University

UC San Diego

Columbia
University

SIGGRAPH 2006

Sponsors: NSF, ONR, Sloan

Scattering in Participating Media



Accurate Rendering of Media Critical for Realism

Significant Progress in Volumetric Rendering



[Jensen et al, 01]



[Donner, 03]




[Jensen, 02]




[Fedkiw et al, 01]

Accuracy Limited by the Input Medium Parameters

Measuring Surface Reflectance ['95 - '06]



BRDF [Matusik et al, 03]




BTF [Dana et al, 97]



Time varying BTF [Gu et al, 06]

Measurement Work in Graphics



Attenuation of
Laser Beams

[Hawkins et al, 05]

No Scattering



Directional Scattering
using a Mirror

[Hawkins et al, 05]

No Attenuation




Diffusion-based
BSSRDF Measurement

[Jensen et al, 01;
Joshi et al, 05]

High Scattering

Measurement Work in Graphics



Attenuation of
Laser Beams

[Hawkins et al, 05]

No Scattering



Directional Scattering
using a Mirror

[Hawkins et al, 05]

No Attenuation




Diffusion-based
BSSRDF Measurement

[Jensen et al, 01;
Joshi et al, 05]


High Scattering

One Simple Setup
Robust Parameter Estimation from a Photograph
Cover Entire Parameter Space

Our Measurement Setup



Our Measurement Setup



Dimensions of Tank: 25cm x 30 cm x 30 cm

Problem: Multiple Scattering

- Causes significant **Blurring** of Incident Light
- Inverse Estimation is **Ill-conditioned** and **Not Unique**

[Ishimaru 75,97; McCormick et al., 79-83]



Photo of Milk in Setup

Problem: Multiple Scattering

- Causes significant Blurring of Incident Light
- Inverse Estimation is Ill-conditioned and Not Unique

[Ishimaru 75,97; McCormick et al., 79-83]

- **Key Idea: Avoid Multiple Scattering**

At “low” concentrations:

- Single Scattering dominant
- Multiple Scattering negligible

[Ishimaru 97; Narasimhan et al 99-03]




Photo of Milk in Setup





Increasing Milk Concentrations

So...dilute media “sufficiently” with water to simplify light transport.

Single Scattering Ray Geometry




Physical Apparatus



Captured Ortho-Photo

- Range of Scattering Angles: [0 deg, 175 deg]
- Range of Path-lengths: [125 mm, 610mm]
- All Path-length and Angle Combinations

Single Scattering Model and Estimation




- Image Formation Model:

$$E(x, y, z) = \int \frac{I_0}{d_1^2} \exp(-\sigma d_1) \bullet \beta P(\alpha, g) \bullet \exp(-\sigma d_2) dz$$

Annotations for the equation above:

- σ : Scattering Coefficient
- $P(\alpha, g)$: Phase function
- σ : Extinction Coefficient

Single Scattering Model and Estimation



- Image Formation Model:

$$E(x, y) = \int \frac{I_0}{d_1^2} \exp(-\sigma d_1) \bullet \beta P(\alpha, g) \bullet \exp(-\sigma d_2) dz$$


Annotations for the equation above:

- σ : Scattering Coefficient
- $P(\alpha, g)$: Phase function
- σ : Extinction Coefficient

- Parameter Estimation: Nelder-Meade Search in Matlab

$$\arg \min_{\sigma, \beta, g} \| E^{\text{measured}}(x, y) - E^{\text{model}}(x, y) \|$$

Single Scattering Model and Estimation



- Image Formation Model:

$$E(x, y) = \int \frac{I_0}{d_1^2} \exp(-\sigma d_1) \cdot \beta P(\alpha, g) \cdot \exp(-\sigma d_2) dz$$


Scattering Coefficient
Phase function
Extinction Coefficient

- Parameter Estimation: Nelder-Meade Search in Matlab


$$\arg \min_{\sigma, \beta, g} \| E^{\text{measured}}(x, y) - E^{\text{model}}(x, y) \|$$

**Unique
and
Robust!**

How Much to Dilute?



Single
Scattering



Multiple
Scattering

Database of 40 Common Materials

- **Alcoholic Beverages** – 3 wines, 3 beers...
- **Coffees** – black, with cream, cappuccino,...
- **Milks** – chocolate, whole, 2% fat, vitamin A & D,...
- **Juices** – grape, apple, cranberry,...
- **Soft-drinks** – coke, pepsi, lemonade...
- **Cleaning supplies** – detergents, shampoos,...
- **Powders and Crystals** – sugar, salt, tang,...
- **Pacific Ocean Water** – bay, different depths,...



Sample Photographs: Highly Scattering Media



Pink Lemonade
Powder



Regular
Chocolate Milk



Low Fat Choc
Milk



Ruby Grapefruit
Juice



Salt Powder



Espresso
Coffee



Cappuccino
Powder



Orange
Powder



Low Fat
Milk



Regular
Milk

Sample Photographs: Highly Absorbing Media




Merlot
Wine



Coke




Grape Juice




Yuengling
Beer




Pacific Ocean
Water




Chardonnay
Wine




Era
Detergent



Strawberry
Shampoo



Lemon Tea
Powder



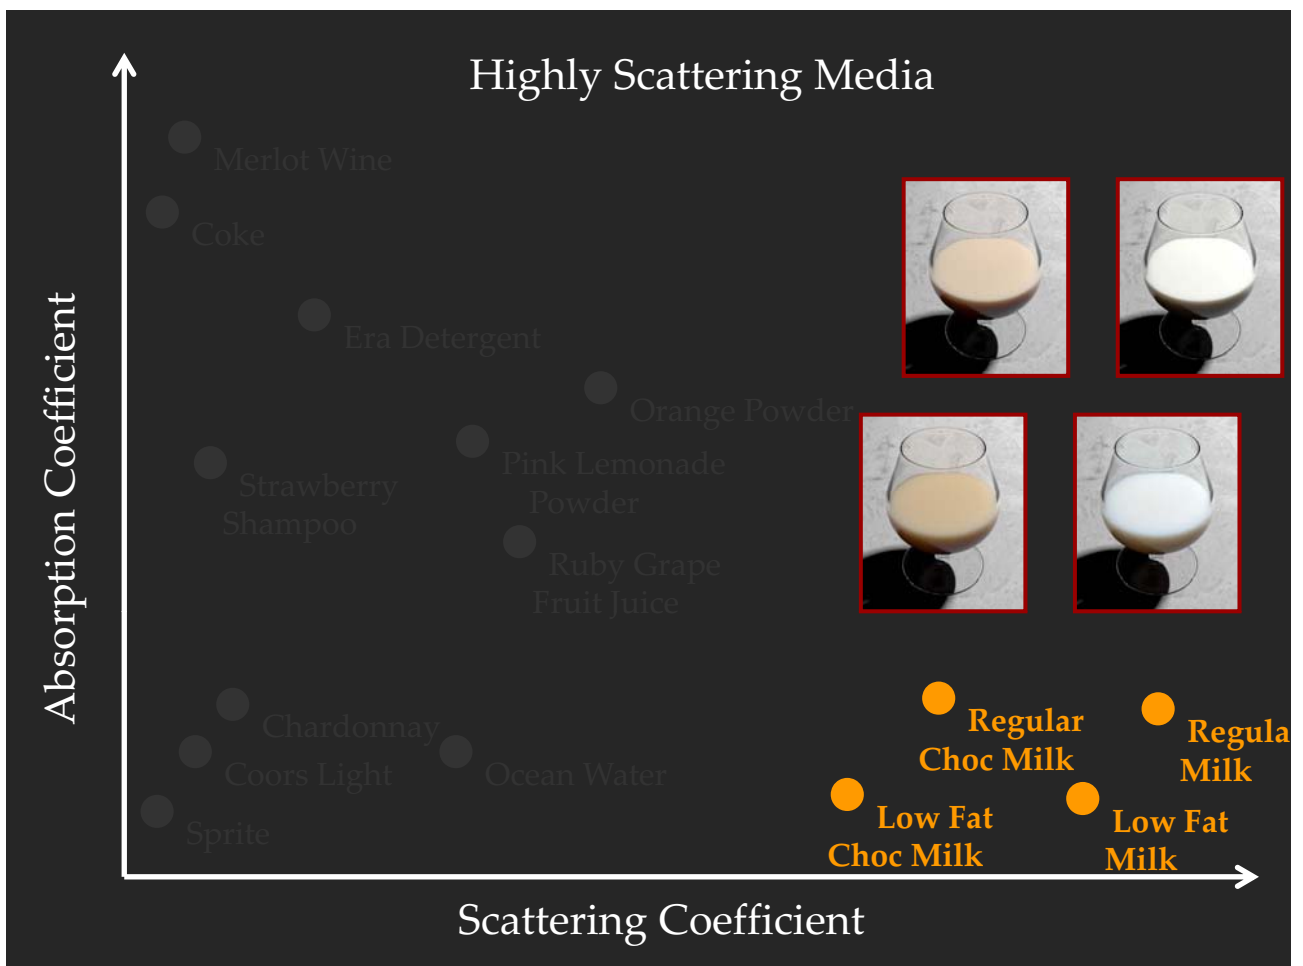
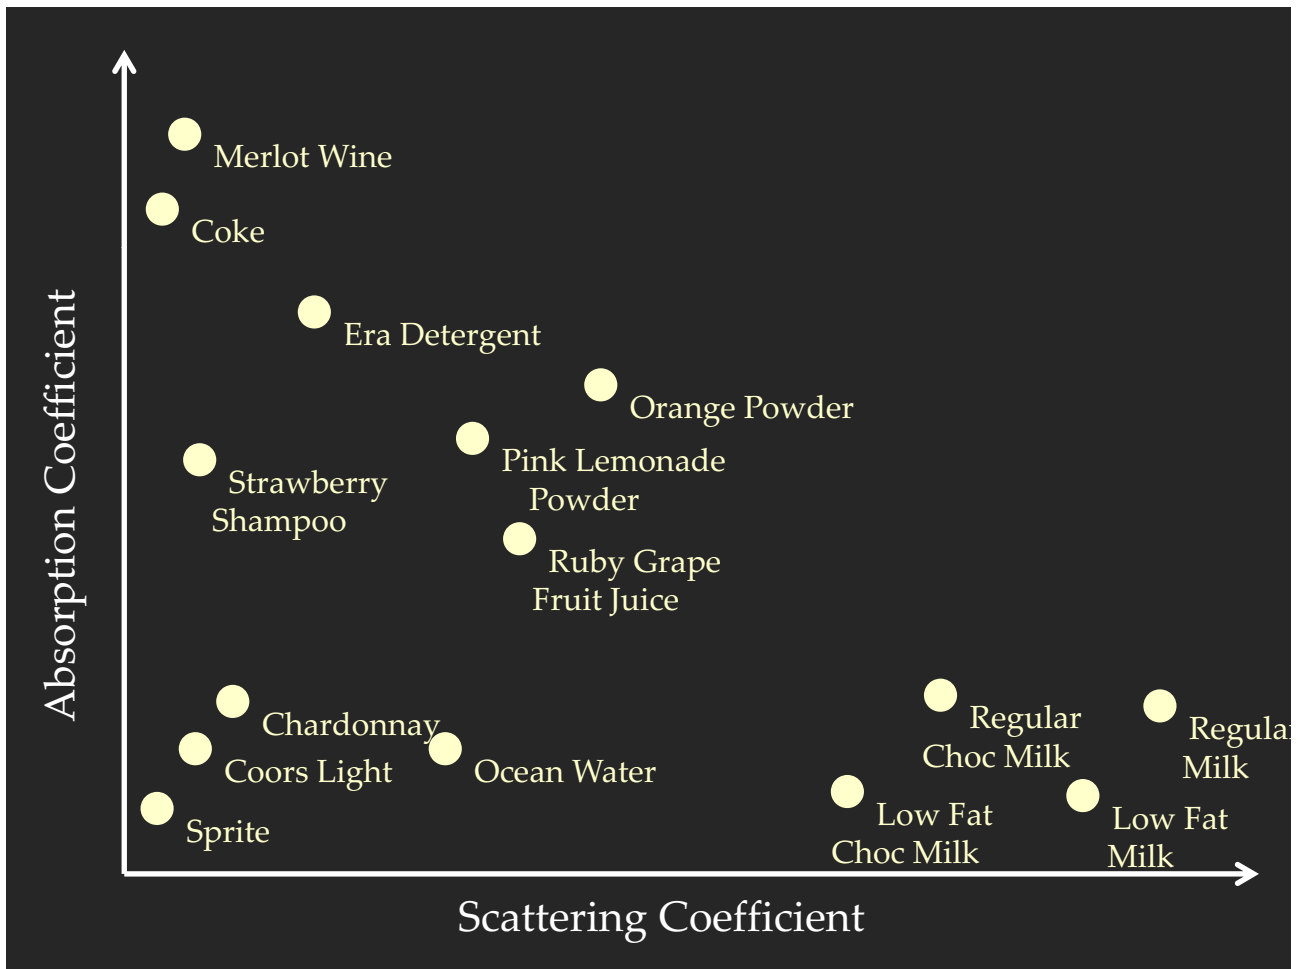
Brown
Sugar

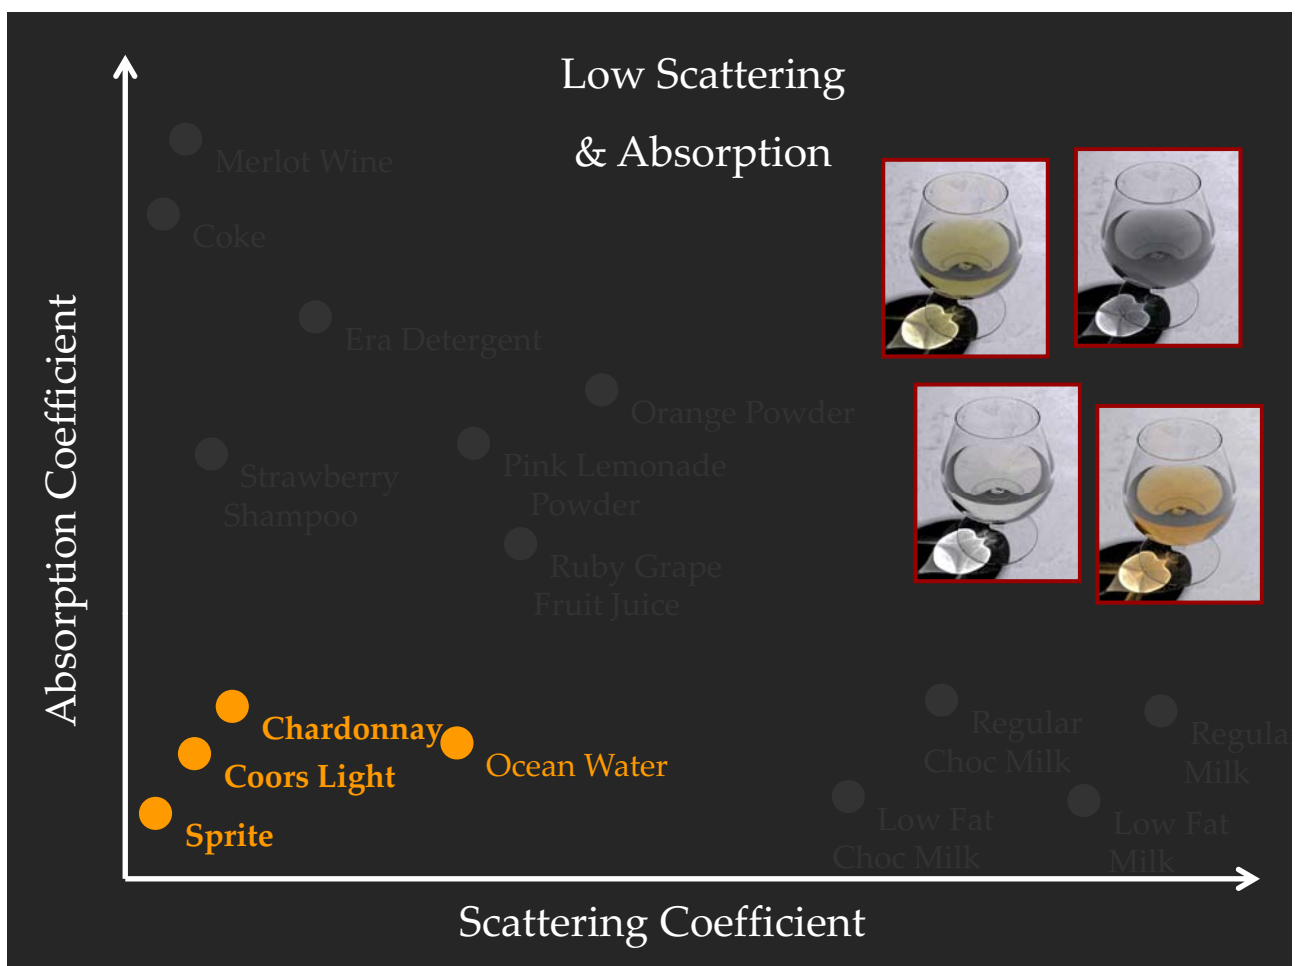
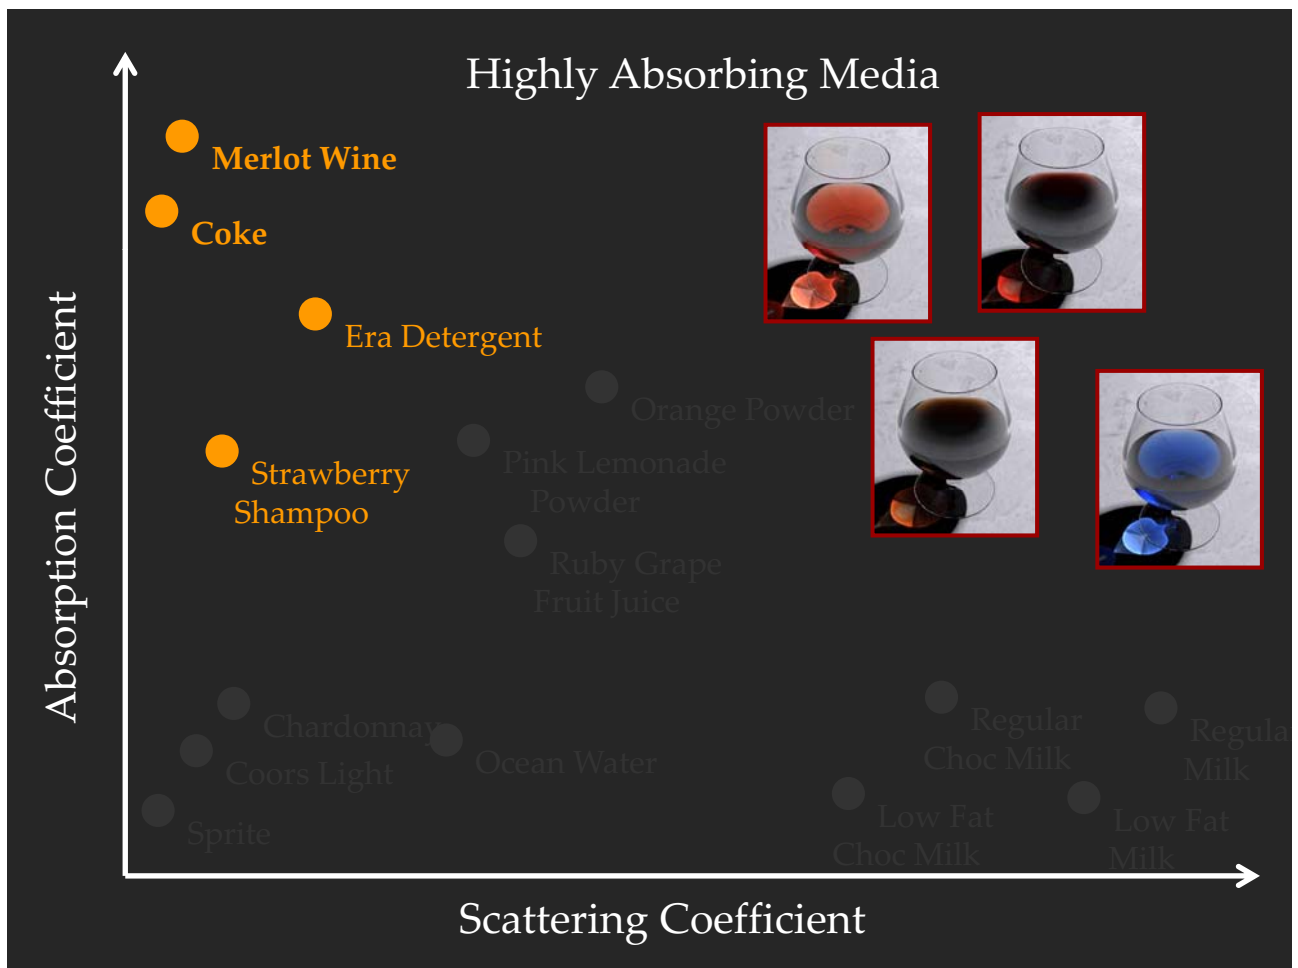
Sample Parameters: Highly Scattering Media

Medium	Volume		Extinction Coefficient (σ) ($\times 10^{-2} \text{ mm}^{-1}$)	Scattering Coefficient (β) ($\times 10^{-2} \text{ mm}^{-1}$)	Average Cosine (g)
Low-Fat Milk	16 ml	R	0.9126	0.9124	0.932
		G	1.0748	1.0744	0.902
		B	1.2500	1.2492	0.859
Regular Milk	15 ml	R	1.1874	1.1873	0.750
		G	1.3296	1.3293	0.714
		B	1.4602	1.4589	0.681
Regular Chocolate Milk	16 ml	R	0.7359	0.7352	0.862
		G	0.9172	0.9142	0.838
		B	1.0688	1.0588	0.806

Sample Parameters: Highly Absorbing Media

Medium	Volume		Extinction Coefficient (σ) ($\times 10^{-2} \text{ mm}^{-1}$)	Scattering Coefficient (β) ($\times 10^{-2} \text{ mm}^{-1}$)	Average Cosine (g)
Yuengling Beer	2900 ml	R	0.1535	0.0495	0.969
		G	0.3322	0.0521	0.969
		B	0.7452	0.0597	0.975
Merlot Wine	1500 ml	R	0.7639	0.0053	0.974
		G	1.6429	0.0000	-
		B	1.9196	0.0000	-
Era Detergent	2300 ml	R	0.7987	0.0553	0.949
		G	0.5746	0.0586	0.950
		B	0.2849	0.0906	0.971





Absorption Coefficient

Moderate Scattering & Absorption

Merlot Wine

Coke

Era Detergent

Orange Powder

Pink Lemonade
Powder

Ruby Grape
Fruit Juice

Strawberry
Shampoo

Chardonnay

Coors Light

Sprite

Regular

Choc Milk

Low Fat

Choc Milk

Regular

Milk


Low Fat

Milk

Scattering Coefficient


Material Name	Volume	Extinction Coefficient (σ) ($\times 10^{-2} \text{ mm}^{-1}$)			Scattering Coefficient (β) ($\times 10^{-2} \text{ mm}^{-1}$)			Average Cosine (g)			% RMS Error
		R	G	B	R	G	B	R	G	B	
Milk (lowfat)	16ml	0.9126	1.0748	1.2500	0.9124	1.0744	1.2492	0.932	0.902	0.859	0.95
Milk (reduced)	18ml	1.0750	1.2213	1.3941	1.0748	1.2209	1.3931	0.819	0.797	0.746	1.27
Milk (regular)	15ml	1.1874	1.3296	1.4602	1.1873	1.3293	1.4589	0.750	0.714	0.681	1.56
Coffee (espresso)	8ml	0.4376	0.5115	0.6048	0.2707	0.2828	0.2970	0.907	0.896	0.880	1.90
Coffee (mint mocha)	6ml	0.1900	0.2600	0.3500	0.0916	0.1081	0.1460	0.910	0.907	0.914	2.00
Soy Milk (lowfat)	16ml	0.1419	0.1625	0.2740	0.1418	0.1620	0.2715	0.850	0.853	0.842	1.75
Soymilk (regular)	12ml	0.2434	0.2719	0.4597	0.2433	0.2714	0.4563	0.873	0.858	0.832	1.68
Chocolate Milk (lowfat)	10ml	0.4282	0.5014	0.5791	0.4277	0.4998	0.5723	0.934	0.927	0.916	1.04
Chocolate Milk (regular)	16ml	0.7359	0.9172	1.0688	0.7352	0.9142	1.0588	0.862	0.838	0.806	2.19
Soda (coke)	1600ml	0.7143	1.1688	1.7169	0.0177	0.0208	0.0000	0.965	0.972	—	4.86
Soda (pepsi)	1600ml	0.6433	0.9990	1.4420	0.0058	0.0141	0.0000	0.926	0.979	—	2.92
Soda (sprite)	15000ml	0.1299	0.1283	0.1395	0.0069	0.0089	0.0089	0.943	0.953	0.952	3.22
Sports Gatorade	1500ml	0.4009	0.4185	0.4324	0.2392	0.2927	0.3745	0.933	0.933	0.935	3.42
Wine (chardonnay)	3300ml	0.1577	0.1748	0.3512	0.0030	0.0047	0.0069	0.914	0.958	0.975	5.10
Wine (white zinfandel)	3300ml	0.1763	0.2370	0.2913	0.0031	0.0048	0.0066	0.919	0.943	0.972	5.49
Wine (merlot)	1500ml	0.7639	1.6429	1.9196	0.0053	0.0000	0.0000	0.974	—	—	4.56
Beer (budweiser)	2900ml	0.1486	0.3210	0.7360	0.0037	0.0069	0.0074	0.917	0.956	0.982	5.61
Beer (coorslight)	1000ml	0.0295	0.0663	0.1521	0.0027	0.0055	0.0000	0.918	0.966	—	4.89
Beer (yuengling)	2900ml	0.1535	0.3322	0.7452	0.0495	0.0521	0.0597	0.969	0.969	0.975	4.48
Detergent (Clorox)	1200ml	0.1600	0.2500	0.3300	0.1425	0.1723	0.1928	0.912	0.905	0.892	1.99
Detergent (Era)	2300ml	0.7987	0.5746	0.2849	0.0553	0.0586	0.0906	0.949	0.950	0.971	4.17
Apple Juice	1800ml	0.1215	0.2101	0.4407	0.0201	0.0243	0.0323	0.947	0.949	0.945	4.92
Cranberry Juice	1500ml	0.2700	0.6300	0.8300	0.0128	0.0155	0.0196	0.947	0.951	0.974	4.60
Grape Juice	1200ml	0.5500	1.2500	1.5300	0.0072	0.0000	0.0000	0.961	—	—	5.19
Ruby Grapefruit Juice	240ml	0.2513	0.3517	0.4305	0.1617	0.1606	0.1669	0.929	0.929	0.931	2.68
White Grapefruit Juice	160ml	0.3609	0.3800	0.5632	0.3513	0.3669	0.5237	0.548	0.545	0.565	2.84
Shampoo (balancing)	300ml	0.0288	0.0710	0.0952	0.0104	0.0114	0.0147	0.910	0.905	0.920	4.86
Shampoo (strawberry)	300ml	0.0217	0.0788	0.1022	0.0028	0.0032	0.0033	0.927	0.935	0.994	2.47
Head & Shoulders	240ml	0.3674	0.4527	0.5211	0.2791	0.2890	0.3086	0.911	0.896	0.884	1.91
Lemon Tea Powder	5tsp	0.3400	0.5800	0.8800	0.0798	0.0898	0.1073	0.946	0.946	0.949	2.83
Orange Powder	4tbsp	0.3377	0.5573	1.0122	0.1928	0.2132	0.2259	0.919	0.918	0.922	2.25
Pink Lemonade Powder	5tbsp	0.2400	0.3700	0.4500	0.1235	0.1334	0.1305	0.902	0.902	0.904	1.02
Cappuccino Powder	0.25tsp	0.2574	0.3536	0.4840	0.0654	0.0882	0.1568	0.849	0.843	0.926	0.67
Salt Powder	1.75cup	0.7600	0.8685	0.9363	0.2485	0.2822	0.3216	0.802	0.793	0.821	1.34
Sugar Powder	5cup	0.0795	0.1759	0.2780	0.0145	0.0162	0.0202	0.921	0.919	0.931	1.80
Suisse Mocha Powder	0.5tsp	0.5098	0.6476	0.7944	0.3223	0.3583	0.4148	0.907	0.894	0.888	1.33
Mission Bay Surface Water (1-2 hours)		3.3623	3.2929	3.2193	0.2415	0.2762	0.3256	0.842	0.865	0.912	2.48
Pacific Ocean Surface Water (1 hour)		3.3645	3.3158	3.2428	0.1800	0.1834	0.2281	0.902	0.825	0.914	2.57
Mission Bay 10ft deep Water (30 min)		3.4063	3.3410	3.2810	0.0990	0.1274	0.1875	0.726	0.820	0.921	5.10
Mission Bay 10ft deep Water (8 hours)		3.3997	3.3457	3.2928	0.1018	0.1033	0.1611	0.929	0.910	0.945	5.13

Experimental Validation: Fits to Measurements




Orange Powder


Sample Fits: Highly Absorbing Media




Era Detergent



Pacific Ocean Surface Water




Yuengling Beer



Merlot Wine

Sample Fits: Highly Scattering Media



Renderings with the “Kitchen” Environment Map



[Debevec et al]



Merlot Wine



Chardonnay Wine



Yuengling Beer



Coca-Cola



Yuengling Beer



Milk

Renderings with a
Single Directional Light Source



Yuengling Beer



Coca-cola



Chardonnay Wine



Orange Powder



Strawberry Shampoo



Era Detergent

Blending Parameters



75% Espresso

+



25% Milk

=



Light Coffee

Blending Parameters




50% Wine

+



50% Milk

=



?

Transitions between Media



Merlot wine

Wine → Water → Milk → Espresso

Concentrations at which Parameters Measured




Real Concentrations



Summary




Simple Apparatus



Robust Parameter Estimation



Render Any Concentration



Mixing Media

<http://graphics.cs.cmu.edu>

Rendering Scattering Materials

Simulating the Appearance of Natural Materials

Henrik Wann Jensen

Computer Science and Engineering

University of California, San Diego



The Natural World



Mist



A Sunset



The Green Flash



A Halo



Human Skin



Simulating Natural Materials



Simulating the appearance of scattering media and materials


Appearance Modeling



Film




Computer Games



Medicin

Radiative Transport




$$(\vec{\omega} \cdot \nabla)L(x, \vec{\omega}) = -\sigma_t L(x, \vec{\omega}) + \sigma_s \int_{4\pi} p(\vec{\omega}, \vec{\omega}') L(x, \vec{\omega}') d\vec{\omega}' + s(x, \vec{\omega})$$

Marble Bust



Monte Carlo photon tracing, rendering time = days

Photon Mapping



[Jensen and Christensen, SIGGRAPH 1998]

Wet Materials



[Jensen, Legakis, and Dorsey -- EGWR 1999]

Rising Smoke Simulation




[Fedkiw, Stam, and Jensen -- SIGGRAPH 2001]

Campfire Simulation



[Nguyen, Fedkiw, and Jensen -- SIGGRAPH 2002]

BRDF




$$\text{BRDF: } f_r(x, \vec{\omega}_i, \vec{\omega}_o) = \frac{dL(x, \vec{\omega}_o)}{dE(x, \vec{\omega}_i)}$$

Marble Bust (BRDF)



BSSRDF



$$\text{BSSRDF: } S(x_i, \vec{\omega}_i, x_o, \vec{\omega}_o) = \frac{dL(x_o, \vec{\omega}_o)}{d\Phi(x_i, \vec{\omega}_i)}$$

Radiative Transport

$$(\vec{\omega} \cdot \nabla)L(x, \vec{\omega}) = -\sigma_t L(x, \vec{\omega}) + \sigma_s \int_{4\pi} p(\vec{\omega}, \vec{\omega}') L(x, \vec{\omega}') d\vec{\omega}' + s(x, \vec{\omega})$$

- Costly for highly scattering materials
- Highly scattering materials are common

Human Skin



The Diffusion Approximation

$$L(x, \vec{\omega}) = L_u(x, \vec{\omega}) + L_d(x, \vec{\omega})$$

$$L_d(x, \vec{\omega}) \approx F_t(x) + \frac{3}{4}\pi \vec{E}(x) \cdot \vec{\omega}$$

The diffusion equation

$$\frac{1}{3\sigma'_t} \nabla^2 F_t(x) = \sigma_a F_t(x) - S_0(x) + \frac{1}{\sigma'_t} \nabla \cdot \vec{S}_1(x)$$

The BSSRDF Model

Multiple scattering:

$$dM_o(x_o) = d\Phi_i(x_i) \frac{\alpha'}{4\pi} \left\{ C_1 \frac{e^{-\sigma_{tr}d_r}}{d_r^2} + C_2 \frac{e^{-\sigma_{tr}d_v}}{d_v^2} \right\}$$

where


$$C_1 = z_r \left(\sigma_{tr} + \frac{1}{d_r} \right) \quad \text{and} \quad C_2 = z_v \left(\sigma_{tr} + \frac{1}{d_v} \right)$$

Single scattering:


$$\begin{aligned} L_o^{(1)}(x_o, \vec{\omega}_o) &= \sigma_s(x_o) \int_{2\pi} F p(\vec{\omega}'_i \cdot \vec{\omega}'_o) \int_0^\infty e^{-\sigma_{tc}s} L_i(x_i, \vec{\omega}_i) ds d\vec{\omega}_i \\ &= \int_A \int_\Omega S^{(1)}(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) L_i(x_i, \vec{\omega}_i) (\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i dA(x_i) \end{aligned}$$

[Jensen, Marschner, Levoy, and Hanrahan - SIGGRAPH 2001]

The BSSRDF Model



Multiple Scattering



Single Scattering

Marble Bust (BSSRDF)



Marble: MCRT versus BSSRDF



MCRT



BSSRDF

Marble: BRDF versus BSSRDF



BRDF



BSSRDF

A Face (BSSRDF Model)




Modeled by Steven Stahlberg

A Face (BRDF Approximation)



Modeled by Steven Stahlberg

Virtual Skin




[Swerdlow, National Geographic 2002]

Shrek




[Jensen and Buhler, SIGGRAPH 2002]

Structure of Skin



Scattering in a Layer




Multipole Diffusion

$$R(r) = \sum_{i=-n}^n \frac{\alpha' z_{r,i} (1 + \sigma_{tr} d_{r,i}) e^{-\sigma_{tr} d_{r,i}}}{4\pi d_{r,i}^3} - \frac{\alpha' z_{v,i} (1 + \sigma_{tr} d_{v,i}) e^{-\sigma_{tr} d_{v,i}}}{4\pi d_{v,i}^3}$$


$$T(r) = \sum_{i=-n}^n \frac{\alpha' (d - z_{r,i}) (1 + \sigma_{tr} d_{r,i}) e^{-\sigma_{tr} d_{r,i}}}{4\pi d_{r,i}^3} - \frac{\alpha' (d - z_{v,i}) (1 + \sigma_{tr} d_{v,i}) e^{-\sigma_{tr} d_{v,i}}}{4\pi d_{v,i}^3}$$

[Donner and Jensen, SIGGRAPH 2005]


Scattering in Two Layers



Scattering in Two Layers



Scattering in Two Layers



Interlayer Scattering

$$T_{12} = T_1 * T_2 + T_1 * R_2 * R_1 * T + 2 + \\ T_1 * R_2 * R_1 * R_2 * R_1 * T_2 + \dots$$

$$\mathcal{T}_{12} = \mathcal{T}_1 \mathcal{T}_2 + \mathcal{T}_1 \mathcal{R}_2 \mathcal{R}_1 \mathcal{T}_2 + \mathcal{T}_1 \mathcal{R}_2 \mathcal{R}_1 \mathcal{R}_2 \mathcal{R}_1 \mathcal{T}_2 + \dots \\ = \frac{\mathcal{T}_1 \mathcal{T}_2}{1 - \mathcal{R}_2 \mathcal{R}_1}$$

Multilayer Diffusion Model


$$R(r) = \sum_{i=-n}^n \frac{\alpha' z_{r,i} (1 + \sigma_{tr} d_{r,i}) e^{-\sigma_{tr} d_{r,i}}}{4\pi d_{r,i}^3} - \frac{\alpha' z_{v,i} (1 + \sigma_{tr} d_{v,i}) e^{-\sigma_{tr} d_{v,i}}}{4\pi d_{v,i}^3}$$

$$T(r) = \sum_{i=-n}^n \frac{\alpha' (d - z_{r,i}) (1 + \sigma_{tr} d_{r,i}) e^{-\sigma_{tr} d_{r,i}}}{4\pi d_{r,i}^3} - \\ \frac{\alpha' (d - z_{v,i}) (1 + \sigma_{tr} d_{v,i}) e^{-\sigma_{tr} d_{v,i}}}{4\pi d_{v,i}^3}$$

$$\mathcal{T}_{12} = \mathcal{T}_1 \mathcal{T}_2 + \mathcal{T}_1 \mathcal{R}_2 \mathcal{R}_1 \mathcal{T}_2 + \mathcal{T}_1 \mathcal{R}_2 \mathcal{R}_1 \mathcal{R}_2 \mathcal{R}_1 \mathcal{T}_2 + \dots \\ = \frac{\mathcal{T}_1 \mathcal{T}_2}{1 - \mathcal{R}_2 \mathcal{R}_1}$$

[Donner and Jensen, JOSA 2006]

A 3 Layer Skin Model



Epidermis

Dermis

Bloody Dermis

[Skin parameters from Tuchin 2000, 3D Model XYZRGB]

Face (3 Layer Skin Model)




Face (3 Layer Skin Model)



Face (Single Layer Model)



Spectral BSSRDF for Human Skin




[Donner and Jensen, EGSR 2006]

Spectral BSSRDF for Human Skin




- Few parameters
- Subsurface scattering
- Surface reflections
- Surface texture


A Two-Layer Skin Model



A Two-Layer Skin Model



Chromophore Spectra



Layer Absorption

Epidermis absorption:

$$\sigma_a^{epi}(\lambda) = C_m(\beta_m \sigma_a^{em}(\lambda) + (1 - \beta_m) \sigma_a^{pm}(\lambda)) + (1 - C_m) \sigma_a^{baseline}$$

Melanin type $\beta_m \in [0, 1]$ and concentration $C_m \in [0, 1]$

Dermis absorption:

$$\sigma_a^{derm}(\lambda) = C_h(\gamma \sigma_a^{oxy}(\lambda) + (1 - \gamma) \sigma_a^{deoxy}(\lambda)) + (1 - C_h) \sigma_a^{baseline}$$

Hemoglobin oxygenation $\gamma = 0.7$ and concentration $C_h \in [0, 1]$

Layer Scattering

Epidemis scattering:

$$\sigma_s^{epi}(\lambda) = 14.74 \cdot \lambda^{-0.22} + 2.2 \cdot 10^{11} \cdot \lambda^{-4}$$

Dermis scattering:

$$\sigma_s^{derm}(\lambda) = 29.48 \cdot \lambda^{-0.22} + 4.4 \cdot 10^{11} \cdot \lambda^{-4}$$

Surface Scattering

🌀 Torrance-Sparrow BRDF

$$f_{r,TS} = \rho_s \frac{F(x, \vec{\omega}_o, \vec{\omega}_i) \cdot D(x, \vec{\omega}_o, \vec{\omega}_i, \sigma) \cdot G(x, \vec{\omega}_o, \vec{\omega}_i)}{4(\vec{\omega}_i \cdot \vec{n})(\vec{\omega}_o \cdot \vec{n})}$$

$$\sigma \approx 0.3$$

Oiliness: ρ_s


Total surface reflectance: ρ_{dr}

$$\rho_{dr}(x, \vec{\omega}_i) = \int_{2\pi} f_{r,TS}(r, \vec{\omega}_o, \vec{\omega}_i)(\vec{\omega}_i \cdot \vec{n}) d\vec{\omega}_o$$

Modulate BSSRDF by $(1 - \rho_{dr})$




Melanin Concentration



Melanin type $\beta_m = 0.7$, hemoglobin concentration $C_h = 0.5\%$

Melanin Type



Melanin concentration $C_m = 50\%$, hemoglobin concentration $C_h = 1\%$


Hemoglobin Concentration



$C_h = 0.1\%$ $C_h = 1\%$ $C_h = 5\%$ $C_h = 10\%$ $C_h = 20\%$


Melanin concentration $C_m = 1\%$ and type $\beta_m = 0.5$

Skin Oiliness





$\rho_s = 0$ $\rho_s = 0.25$ $\rho_s = 0.5$ $\rho_s = 0.75$ $\rho_s = 1$

Skin Matrix



Skin Types

Renderings



Photographs




Caucasian

Asian

African

Skin Texture



Modulate BSSRDF using albedo map normalized by the skin reflectance

Skin Pale



Skin Tan




Skin Ear Closeup



Skin Normal Lighting




Reflectance Prediction



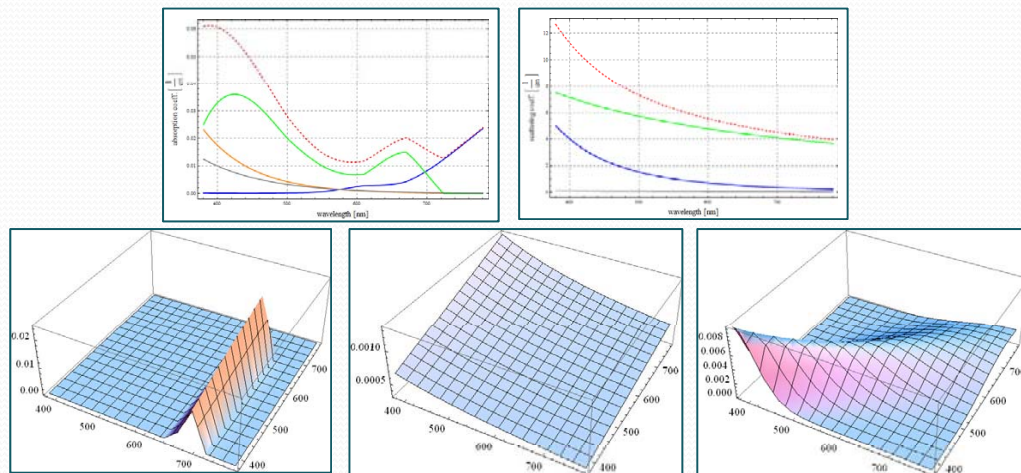
Visualizing Underwater Ocean Optics

Diego Gutierrez
Universidad de Zaragoza

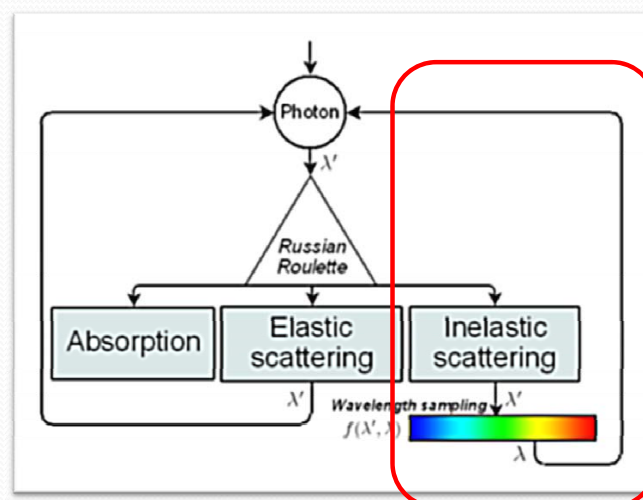
- Ocean water is one of the richest participating media
- The dissolved and particulate matter in the water are optically influential constituents.



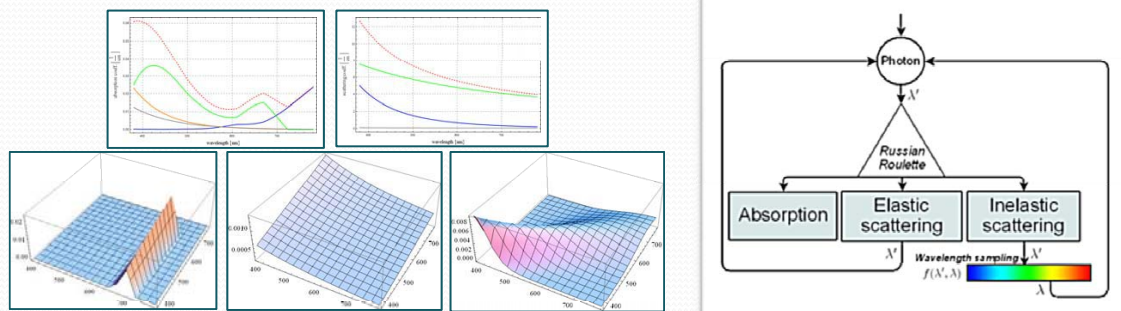
- Two steps:
 - Bio-optical model for ocean waters



- Two steps:
 - Extension of the photon mapping algorithm



- Two contributions:
 - Bio-optical model for ocean waters
 - Extension of the photon mapping algorithm: Anti-Stokes inelastic scattering



What is a participating medium?



What is a participating medium?



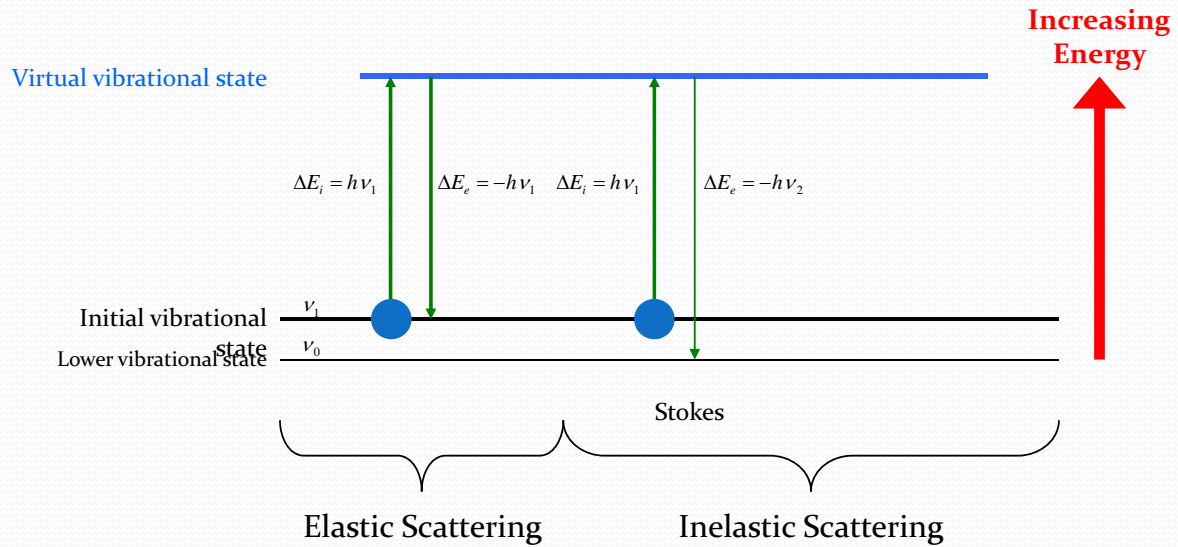


Image by Perez et al.

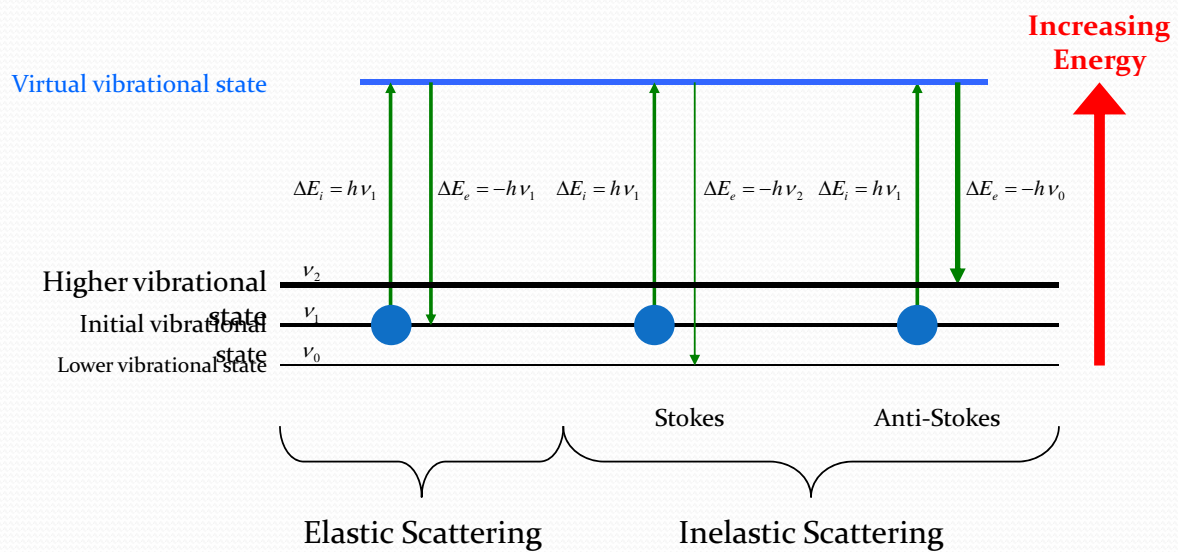
Scattering



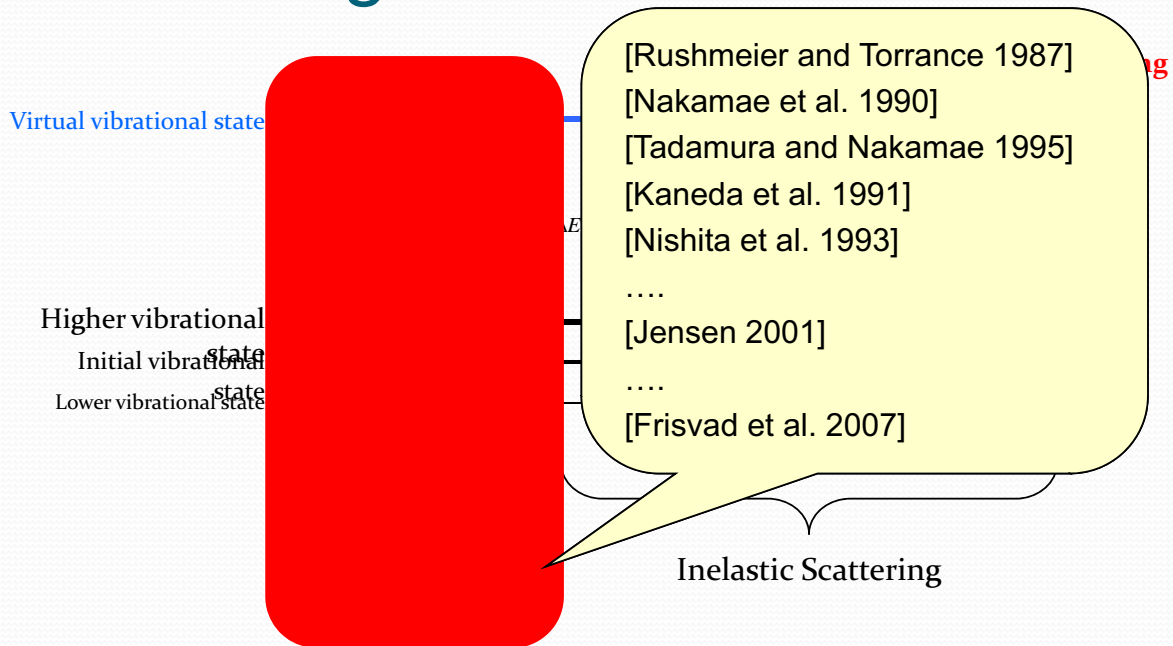
Scattering




Scattering




Scattering



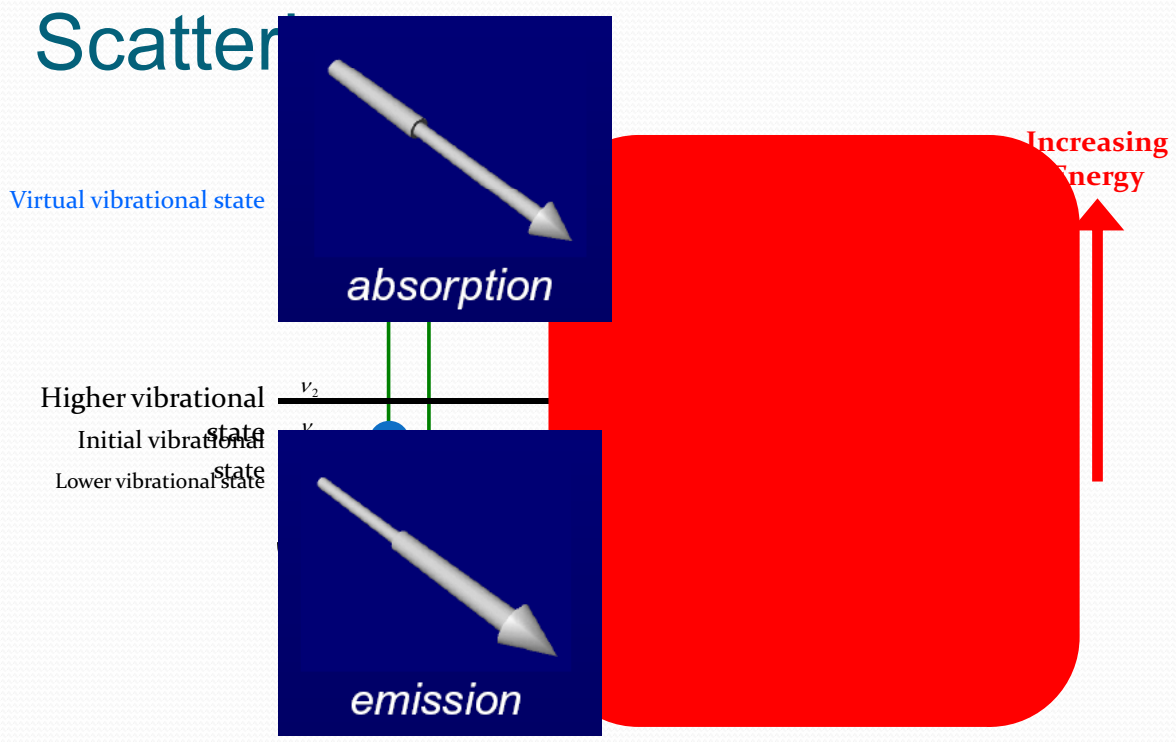
Scattering



Scattering



Scattering





Radiative Transfer Theory



Pure water.....Raman scattering (S, aS)
 Phytonplankton.....Fluorescence (S)
 Detritus
 Yellow matter (CDOM).....Fluorescence (S, aS)

Radiative Transfer Equation

$$\frac{\partial L(\lambda, \vec{\omega}_o)}{\partial x} = \underbrace{\alpha(\lambda)}_{\text{Emission}} L_e(\lambda, \vec{\omega}_o) - \underbrace{\kappa(\lambda)}_{\text{Extinction}} L(\lambda, \vec{\omega}_o) + \underbrace{\sigma(\lambda)}_{\text{Elastic Scattering}} \int_{\Omega} \underbrace{p(\lambda, \vec{\omega}_i, \vec{\omega}_o)}_{\text{Scattering}} L(\lambda, \vec{\omega}_i) d\vec{\omega}_i$$

Optical properties determined by water constituents

Full Radiative Transfer Equation

$$\frac{\partial L(\lambda, \vec{\omega}_o)}{\partial x} = \underbrace{\alpha(\lambda) L_e(\lambda, \vec{\omega}_o)}_{\text{Emission}} - \underbrace{\kappa(\lambda) L(\lambda, \vec{\omega}_o)}_{\text{Extinction}} + \underbrace{\sigma(\lambda) \int_{\Omega} p(\lambda, \vec{\omega}_i, \vec{\omega}_o) L(\lambda, \vec{\omega}_i) d\vec{\omega}_i}_{\text{Elastic Scattering}} + \underbrace{\int_{\Omega} \int_W \left\{ \sigma(\lambda', \lambda) p(\lambda', \lambda, \vec{\omega}_i, \vec{\omega}_o) L(\lambda', \vec{\omega}_i) \right\} d\lambda' d\vec{\omega}_i}_{\text{Inelastic Scattering}}$$


Full Radiative Transfer Equation

Inelastic phase function
Wavelength redistribution function


$$+ \int_{\Omega} \int_W \left\{ \sigma(\lambda', \lambda) p(\lambda', \lambda, \vec{\omega}_i, \vec{\omega}_o) L(\lambda', \vec{\omega}_i) \right\} d\lambda' d\vec{\omega}_i$$

$\rightarrow \alpha(\lambda') f(\lambda', \lambda)$ Inelastic Scattering



Full Radiative Transfer Equation



Full Radiative Transfer Equation



Full Radiative Transfer Equation



• Medium IOPs

- Absorption coefficient
- Scattering coefficient
- Phase function
- Extinction coefficient
- Inelastic absorption coefficient
- Inelastic phase function
- Wavelength redistribution function

$$\alpha(\lambda)$$

$$\sigma(\lambda)$$

$$p(\lambda, \theta)$$

$$\kappa(\lambda)$$

$$\alpha_I(\lambda')$$

$$p_I(\lambda', \lambda, \theta)$$

$$f_I(\lambda', \lambda)$$




- Medium = f(Constituent)

$\alpha(\lambda)$	$= \alpha_d(\lambda) + \alpha_p(\lambda) + \alpha_w(\lambda) + \alpha_y(\lambda)$
$\sigma(\lambda)$	$= \sigma_w(\lambda) + \sigma_d(\lambda) + \sigma_p(\lambda)$
$p(\lambda, \theta)$	$= \frac{\sigma_w(\lambda)p_w(\lambda, \theta) + \sigma_d(\lambda)p_d(\lambda, \theta) + \sigma_p(\lambda)p_p(\lambda, \theta)}{\sigma(\lambda)}$
$\kappa(\lambda)$	$= \alpha(\lambda) + \sigma(\lambda)$
$\alpha_I(\lambda')$	$= \alpha_p(\lambda') + \alpha_w(\lambda') + \alpha_y(\lambda')$
$p_I(\lambda', \lambda, \theta)$	$= \frac{\alpha_p(\lambda')p_p(\lambda', \lambda, \theta) + \alpha_w(\lambda')p_w(\lambda', \lambda, \theta) + \alpha_y(\lambda')p_y(\lambda', \lambda, \theta)}{\alpha_I(\lambda')}$
$f_I(\lambda', \lambda)$	$= \frac{\alpha_p(\lambda')f_p(\lambda', \lambda) + \alpha_w(\lambda')f_w(\lambda', \lambda) + \alpha_y(\lambda')f_y(\lambda', \lambda)}{\alpha_I(\lambda')}$


Absorption

- Water
 - Smith and Baker, Pope and Fry




Absorption

- CDOM
- Bricaud, Morel and Prieur




Absorption

- Phytoplankton
- Sathyendranath, Lazzara and Prieur



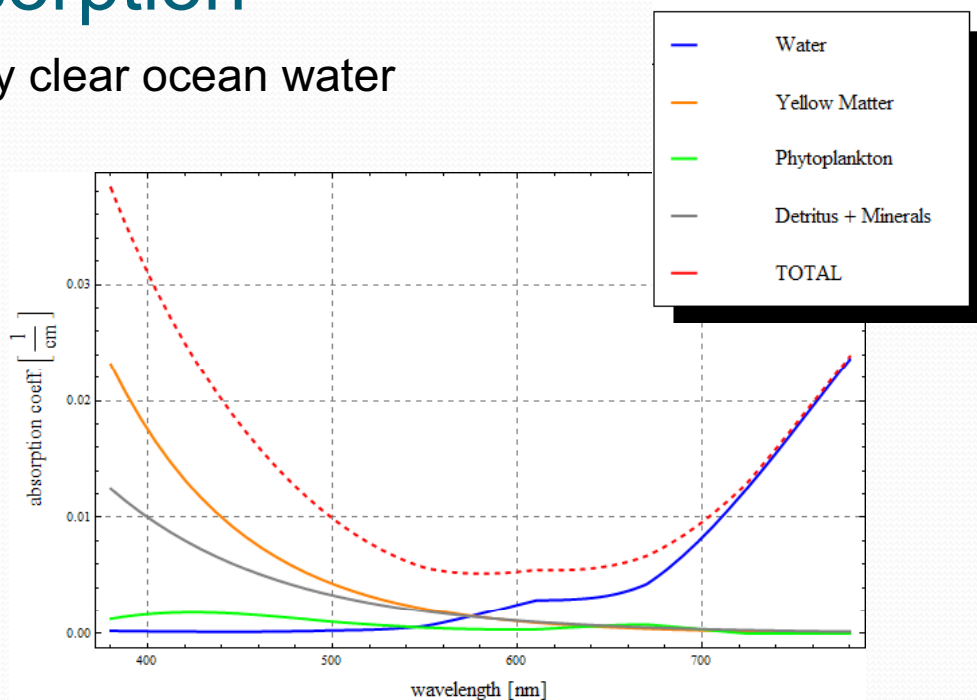
Absorption

- Organic Detritus and minerals
 - Roesler, Perry and Carder




Absorption

- Very clear ocean water




Absorption

- Very muddy ocean water




Absorption

- Ocean water with moderate phytoplankton concentration




Elastic Scattering

- Water
 - Einstein and Smoluchowski




Elastic Scattering

- Phytoplankton
 - Gordon and Morel




Elastic Scattering

- Organic Detritus and minerals
 - Stramski et al.




Elastic Scattering

- Moderately clear ocean water




Elastic Scattering

- Muddy ocean water with high concentration of phytoplankton




Inelastic Scattering

- Fluorescence by Phytoplankton




Inelastic Scattering

- Fluorescence by CDOM



Inelastic Scattering

- Raman by pure water





• Constituents = f(Parameters of the model)

$$\alpha_d(\lambda) = \alpha_d(\lambda_0) e^{-S_d(\lambda-\lambda_0)}$$

$$\alpha_p(\lambda) = C \alpha_p^*(\lambda)$$

λ	[nm]	380	440	500	550	610	670	720	780
α_{dr}	[cm^{-1}]	0.00022	0.000145	0.000257	0.000638	0.00289	0.0043	0.01169	0.0236
α_{pr}^*	[$\text{m}^2 \cdot \text{mg}^{-1}$]	0.025	0.035	0.02	0.01	0.007	0.015	0.001	0.0001

$$\alpha_y(\lambda) = \alpha_y(\lambda_0) e^{-S_y(\lambda-\lambda_0)}$$

$$\sigma_w(\lambda) = 16.06 \beta_w(\lambda_0, 90^\circ) \left(\frac{\lambda_0}{\lambda} \right)^{4.32}$$

λ	[nm]	380	440	500	550	610	670	720	780
debris	σ_{dr}	0.045	0.0375	0.0325	0.03	0.0285	0.0275	0.027	0.027
minerals	σ_{mr}	0.0675	0.0525	0.05	0.045	0.04	0.036	0.034	0.032
total	σ_d	0.1125	0.09	0.0825	0.075	0.0685	0.0635	0.061	0.059

$$\sigma_p(\lambda) = \left(\frac{550}{\lambda} \right) 0.30 C^{0.62}$$

$$p_w(\theta) = 0.06225 \left(1 + 0.835 \cos^2 \theta \right)$$

$$p_d(\theta, \zeta, g_f, g_b) = \zeta p_{HG}(\theta, g_f) + (1 - \zeta) p_{HG}(\theta, g_b)$$

$$p_{HG}(\theta, g) = \frac{1 - g^2}{(1 + g^2 - 2g \cos \theta)^{3/2}} \quad p_p = 1/\pi$$

$$f_w(\lambda', \lambda) = \frac{10^7 \sum_{j=1}^4 A_j \frac{1}{\Delta \tilde{\nu}_j} \exp \left\{ - \frac{[10^7 (\frac{1}{\lambda'} - \frac{1}{\lambda}) - \tilde{\nu}_j]^2}{\Delta \tilde{\nu}_j^2} \right\}}{\lambda'^2 \sqrt{\frac{\pi}{4 \ln 2} \sum_{j=1}^4 A_j}}$$

$$f_y(\lambda', \lambda) = A_0(\lambda') \exp \left\{ - \left(\frac{\frac{1}{\lambda} - \frac{A_1}{\lambda'} - B_1}{0.6 \left(\frac{A_2}{\lambda^2} + B_2 \right)} \right)^2 \right\} \frac{\lambda'}{\lambda}$$

$$f_p(\lambda', \lambda) = \eta_p(\lambda', \lambda) \frac{\lambda'}{\lambda} \equiv \Gamma_p g_p(\lambda') h_p(\lambda) \frac{\lambda'}{\lambda}$$

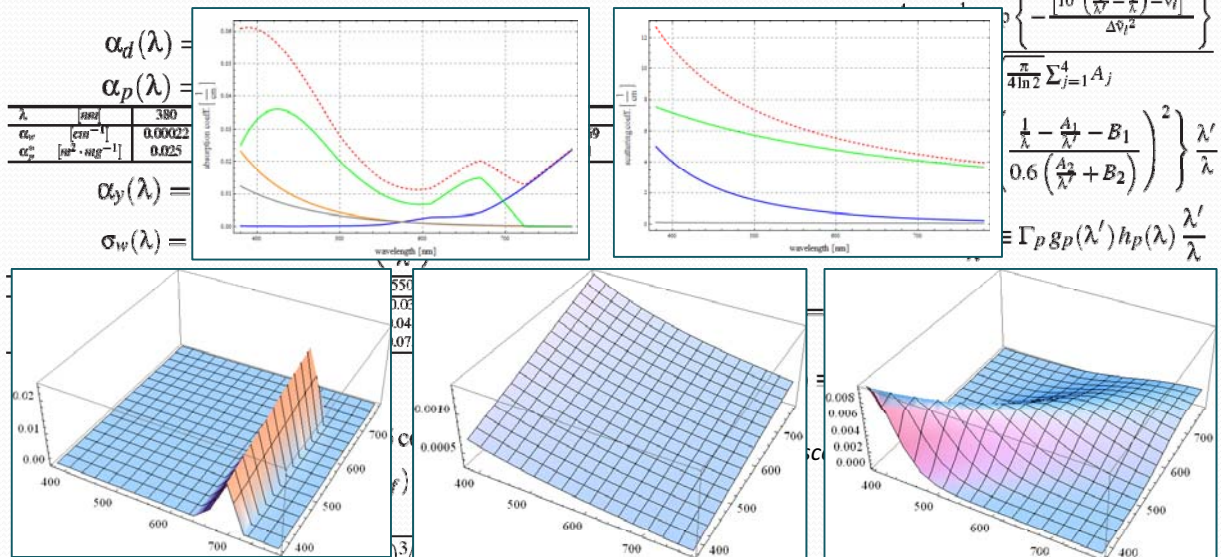
$$h_p(\lambda) = \frac{1}{\sqrt{2\pi} \lambda_\sigma} \exp \left\{ - \frac{(\lambda - \lambda_0)^2}{2(\lambda_\sigma)^2} \right\}$$

$$g_p(\lambda') = \begin{cases} 1 & \text{if } 370 \leq \lambda' \leq 690 \text{ nm} \\ 0 & \text{otherwise} \end{cases}$$

Fluorescence \rightarrow isotropic phase function









• Constituents = f(Parameters of the model)





- We identify 3 main parameters of the model
 - Chlorophyll concentration
 - Detritus and minerals turbidity
 - Yellow matter turbidity
- Most effect on resulting light field.
- Very intuitive!!!



Full Radiative Transfer Equation

$$\frac{\partial L(\lambda, \vec{\omega})}{\partial x} = \text{Emission} - \text{Extinction} + \int_{\Omega} \int_{W} \text{Elastic Scattering} + \int_{\Omega} \int_{W} \text{Inelastic Scattering} \lambda' d\vec{\omega}_i$$

BUT... HOW DO WE SOLVE THIS EQUATION?!?!?




- Algorithm based on photon mapping for participating media





- Photon tracing
 - Photon shooting.




Photon tracing

- Photon carries:
 - Flux
 - Wavelength information



- Photon tracing
 - Photon advances importance-sampled optical distance.

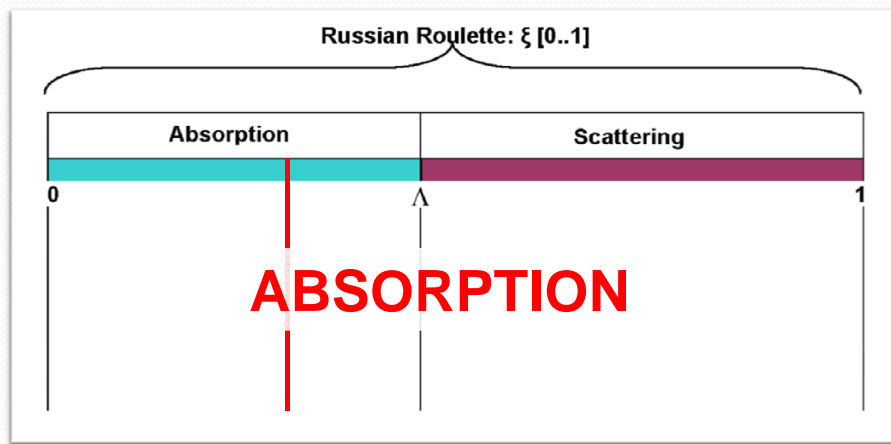

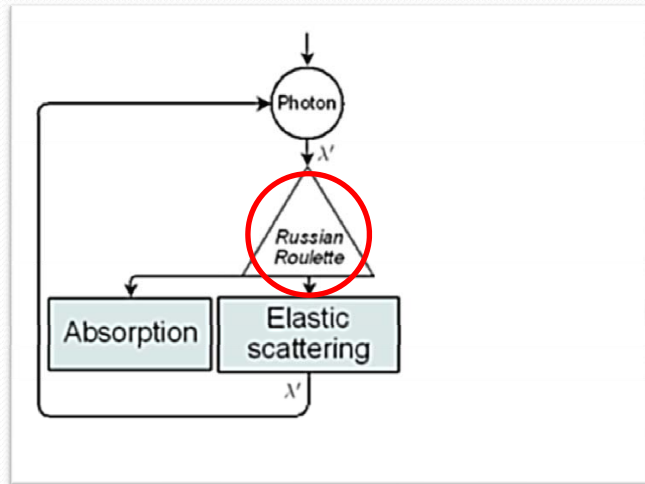

$$d = \frac{\log \xi}{\kappa(\lambda)}$$



Photon tracing


Ray marching

Radiance estimate






Absorption



Elastic scattering



Photon tracing

Ray marching

Radiance estimate



Photon tracing

Ray marching

Radiance estimate

Russian Roulette: $\xi \in [0..1]$


Absorption

Scattering


SCATTERING



Absorption



Elastic scattering



Photon tracing

Ray marching


Radiance estimate




Photon tracing

Ray marching

Radiance estimate





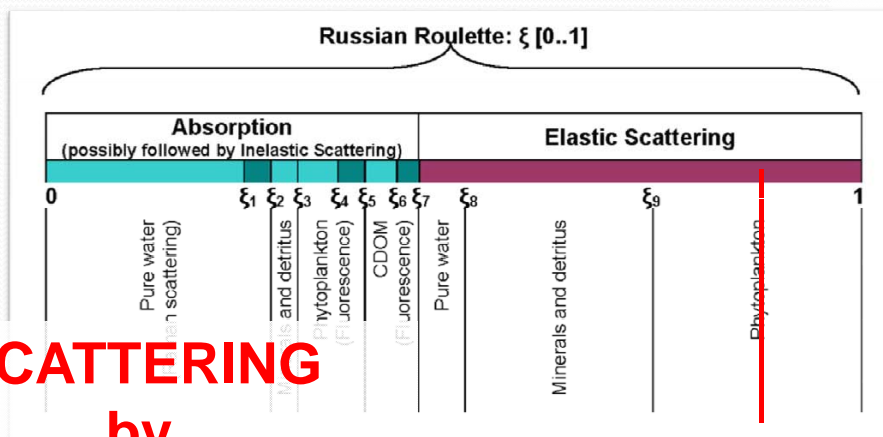
Photon tracing

Ray marching
Radiance estimate

- Russian Roulette
 - Albedo is assigned to each possible interaction.

Photon tracing

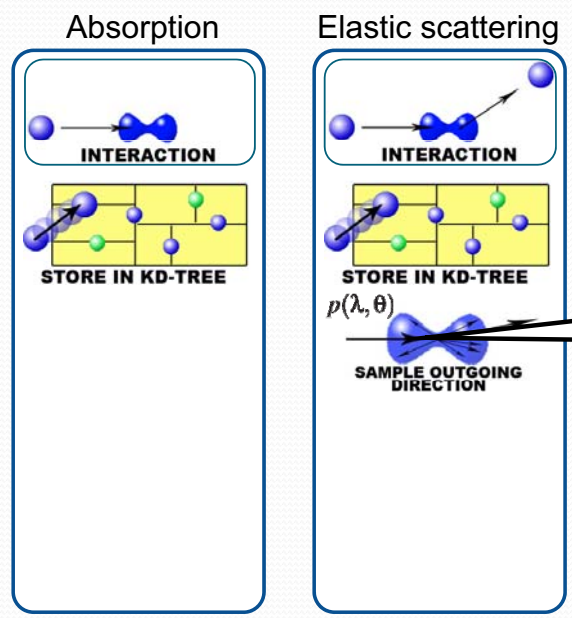
Ray marching
Radiance estimate






SCATTERING
by
PHYTOPLANKTON





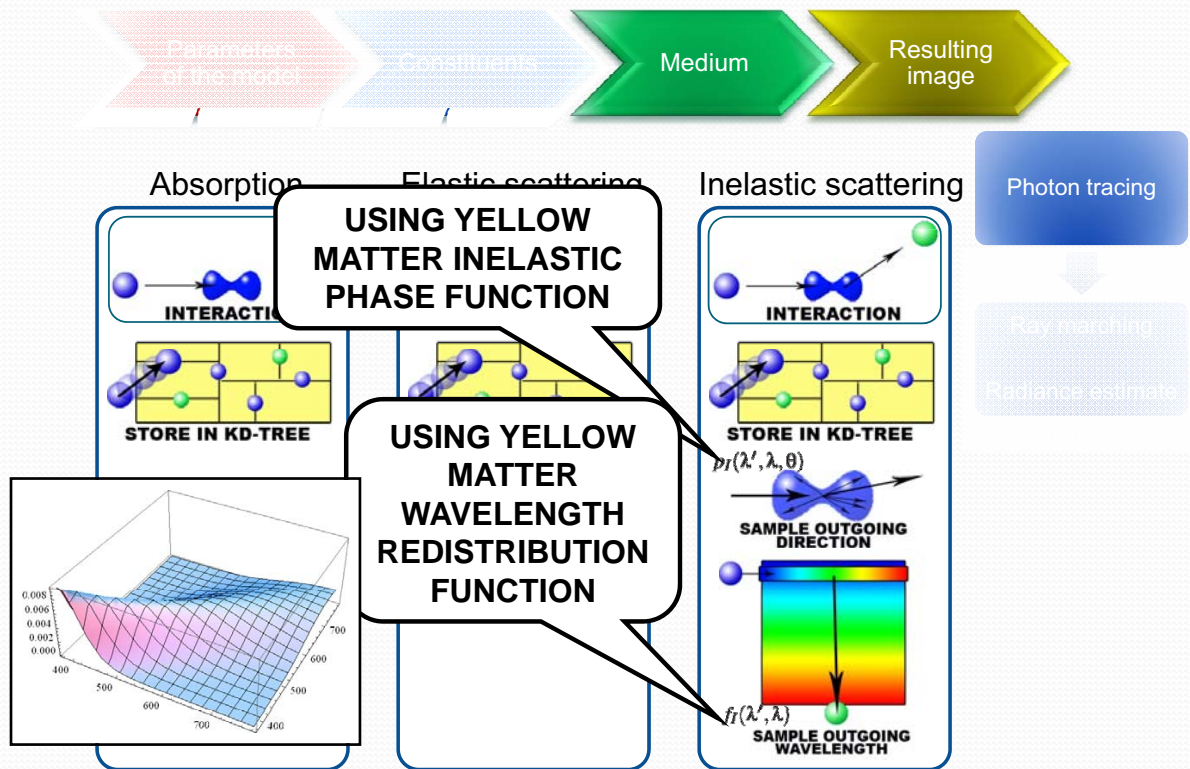
• Russian Roulette



USING
PHYTOPLANKTON
PHASE FUNCTION




**INELASTIC
SCATTERING
by
YELLOW MATTER**





- Ray marching




Photon tracing

Ray marching

Radiance estimate


- Radiance estimate

- Single scattering



$$\sum_{l=1}^N \sum_{i=1}^{N_{\lambda}} \left\{ L_l(\lambda'_i, \vec{w}_l) p_l(\lambda'_i, \lambda, \vec{w}_l, \vec{w}_o) \alpha_l(\lambda'_i) f_l(\lambda'_i, \lambda) \Delta x \right\}$$

- Multiple scattering

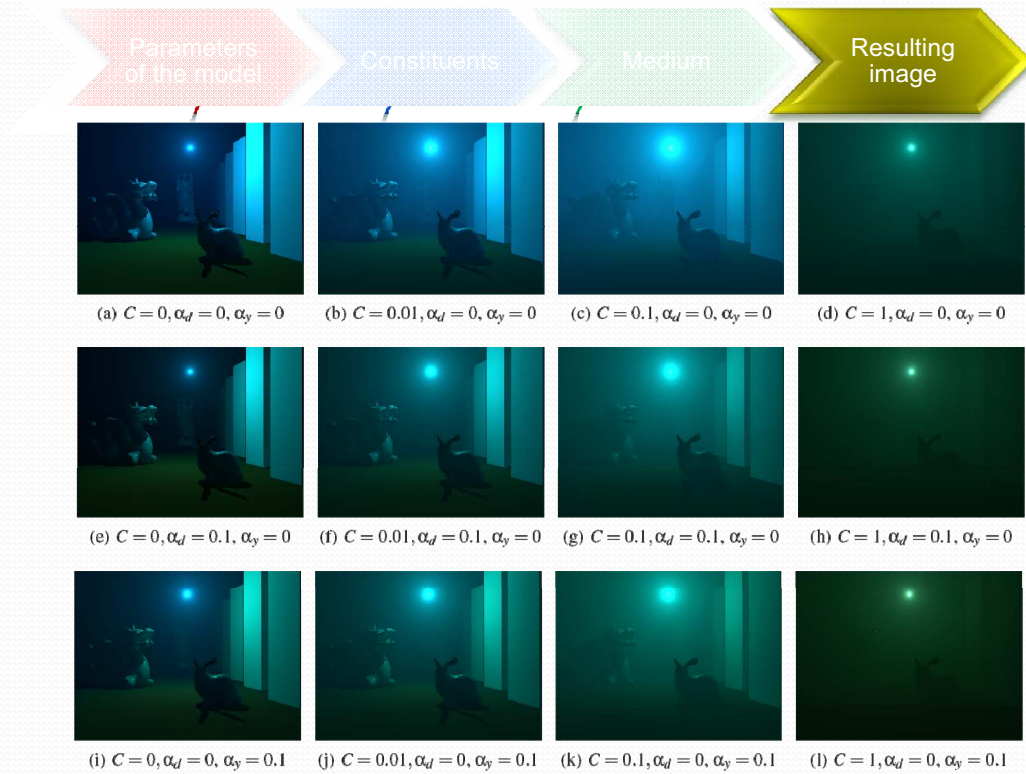


$$\sum_{k=1}^N \left\{ p_l(\lambda'_k, \lambda, \vec{w}_k, \vec{w}_o) f_l(\lambda'_k, \lambda) \frac{\Delta \Phi_k}{\frac{4}{3} \pi r^2} \right\}$$

Photon tracing

Ray marching

Radiance estimate



• Phitoplankton




$C = 0$



$C = 0.01$



$C = 0.1$



$C = 1$

- CDOM




$\alpha_y=0$



$\alpha_y=0.1$

- Detritus and minerals




$\alpha_d=0$




$\alpha_d=0.1$



- Different rendered waters
 - Comparison with [Frisvad et al. 2007]



- Light field





- The effect of inelastic scattering



Without inelastic scattering



With inelastic scattering

Future work

- Make it real time!
- Include bio-luminiscence (medium emission) in the model
- Consider new constituents and new inelastic events (such as Compton, Dragg...)
- Make the algorithm faster without losing physical accuracy (see next paper)

Structured Light in Scattering Media

Srinivasa Narasimhan

Sanjeev Koppal



Robotics Institute
Carnegie Mellon University

Shree Nayar


Bo Sun



Computer Science
Columbia University

Sponsor : ONR

Natural illumination in Scattering Media



[Narasimhan and Nayar, 99 - 03, Schechner et al, 01, 04]

Active illumination in Scattering Media




[Levoy et al., Narasimhan-Nayar, Kocak-Caimi, Jaffe et al., Schechner et al., Negahdaripour et al.]

Floodlighting is Bad in Scattering Media



Structured Light Critical for Good Visibility

Light Strip on Barbed Wire Finding Scattering Media



Light Striping Model in Scattering Media

- Irradiance due to Medium:

$$E_{medium} = L_0 e^{-\beta(x+y)} \frac{\beta}{4\pi} P(\alpha) g \cos \alpha$$

Extinction coefficient
Phase Function


- Irradiance due to Surface:

$$E_{surface} = L_0 R e^{-\beta(D_s + D_v)}$$

Radiance



- Final Image Irradiance:

$$E = E_{surface} \delta(x = D_s) + E_{medium} \delta(x < D_s)$$



Light Striping Algorithm in Scattering Media

Surface Intersection from Brightness Profile:




3D by Triangulation or Temporal Analysis : Same as in clear air.

Medium (β, g) from Fall-off : $\min \| E_{medium} - L_0 e^{-\beta(x+y)} \frac{\beta}{4\pi} (1 + g \cos \alpha) \|_2$


“Clear-Air” Scene Appearance: $L_0 R = E_{surface} e^{+\beta(D_s + D_v)}$

Experimental Setup



Calibration technique similar in spirit to [\[Grossberg-Nayar 01\]](#)

Experimental Setup and Calibration



- No Refractive index and location of glass or medium
- No explicit calibration of camera and projector
- Similar in spirit to [Levoy-Hanrahan 96, Grossberg-Nayar 01]

VIDEO





Floodlit Image



Computed Appearance



How to Place the Camera and Projector?




In Scattering Media:

Still a problem ☹️

Solved if Light Plane is visible 😊

Smoke and Mirrors



Planar Mirror seen
through Dilute Milk



Light Striping of Mirrors
(Dark Intersections)

Reconstruct surfaces with any BRDF if light plane visible

[Discussions with Marc Levoy]

Photometric Stereo in Clear Air

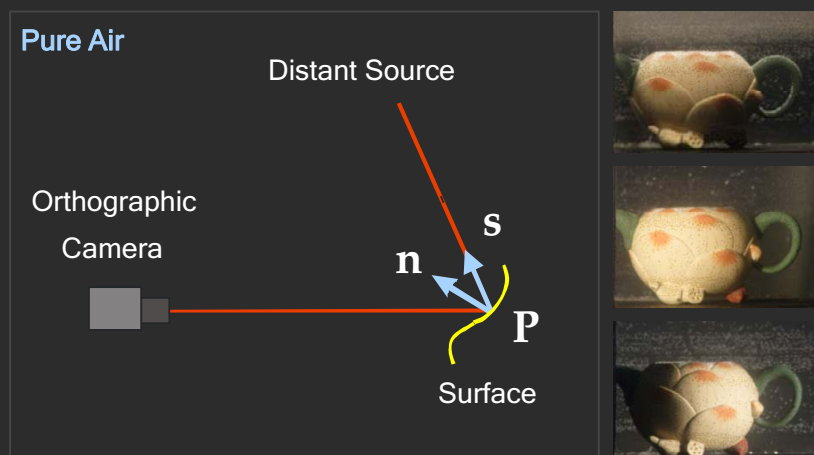


Image Irradiance:

$$E_{\text{surface}} = L_0 \rho \mathbf{n} \cdot \mathbf{s}$$

Surface normal
Albedo
Source direction

Three images
required.

[Woodham 80, Horn 86]

Photometric Stereo in Scattering Media




Image Irradiance:

$$E_{surface} = L_0 \rho e^{-\beta D_s (1 + \cos \alpha)} \mathbf{n} \cdot \mathbf{s}$$


Optical Thickness

$$E_{medium} = L_0 (1 + g \cos \alpha) e^{-\beta D_s} (1 - e^{-\beta D_s \cos \alpha})$$

Phase Function

$$E = E_{surface} + E_{medium}$$

Photometric Stereo in Scattering Media




5 Parameter Non-linear Optimization (4 per pixel, 1 global) :

$$\min \| E - (E_{surface} + E_{medium}) \|$$

Five Non-degenerate Sources are Necessary and Sufficient

Simulations: Error Histograms



Zero error with zero noise.

Robust estimation with 5% uniform noise.

Experiments: Teapot in Pure Water




Experiments: Teapot in Dilute Milk




Low Contrast, Flat Appearance

Results: Traditional Photometric Stereo



3D Shape from Normals


Too Flat



Albedos

Scattering effects present

Results: Our Five-Source Algorithm




3D Shape from Normals



Albedos

Results: Depth from Photometric Stereo



3D Shape from Normals



Depth map

Impossible using traditional method


Milk Concentration	3 ml	4 ml	5 ml	6 ml	12 ml	15 ml
% RMS Error	2.0	2.5	3.0	3.3	5.8	6.3

Summary


- Structured light improves visibility



- Physics of scattering crucial



- Surprising results possible because of scattering




Good Vision in Bad Weather

Srinivasa Narasimhan


Robotics Institute
Carnegie Mellon University

Support:
NSF, DARPA

Traditional Computer Vision Problems




Shape from Shading



Reflectance Estimation




Illumination Estimation



Stereo and Motion

The Fundamental Assumption in Vision



Vision and the Weather



Haze



Mist



Fog



Rain

Effects of Bad Weather



Haze



Mist



Fog




Rain

Image Processing does not Suffice

How often do we see Bad Weather?



Manhattan, Every Hour, 9 Months



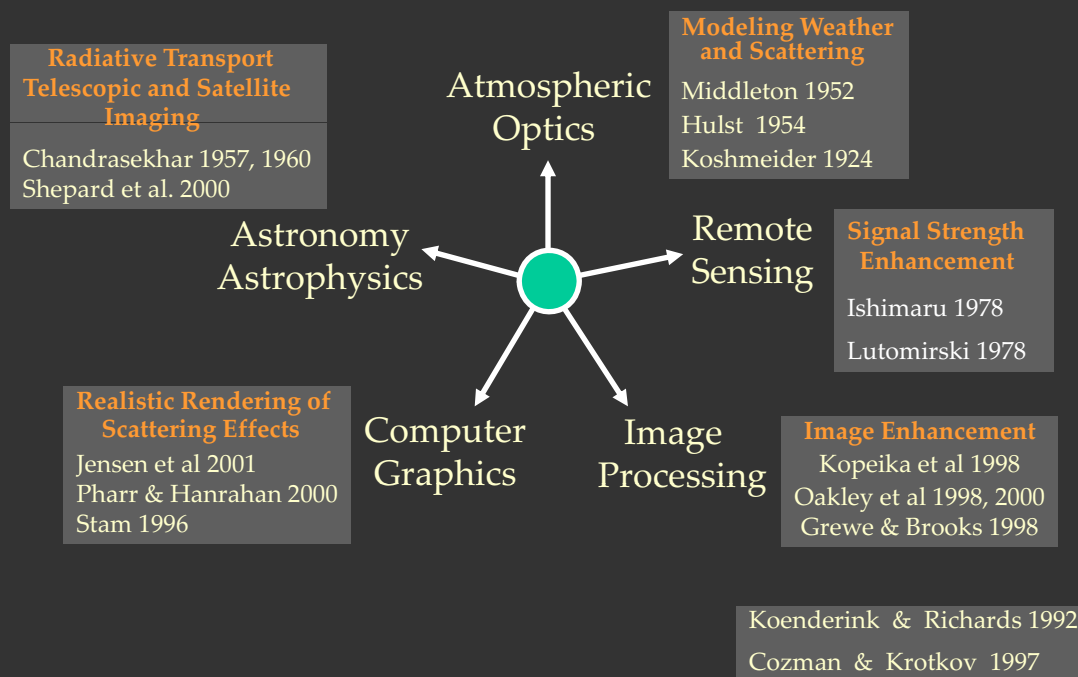
(Narasimhan et. al, ECCV 2002)

Weather Conditions and Particles

CONDITION	PARTICLE TYPE	RADIUS (μm)	CONCENTRATION(cm^{-3})
AIR	Molecule	10^{-4}	10^{19}
HAZE	Aerosol	$10^{-2} - 1$	$10^3 - 10$
FOG	Water Droplet	$1 - 10$	$100 - 10$
CLOUD	Water Droplet	$1 - 10$	$300 - 10$
RAIN	Water Drop	$10^2 - 10^4$	$10^{-2} - 10^{-5}$

(Mie 1908, McCartney 1975)


The Study of the Atmosphere



Computer Vision :

What can Images Reveal about the Scene and the Atmosphere ?

Attenuation Model



Brightness at Distance d :

$$E(d) = E_0 e^{-\beta d}$$

(Bouguer's Law, 1729)

Scattering Coefficient

Direct Transmission

Attenuation of Diverging Beams:

$$E(d) = \frac{I_0 e^{-\beta d}}{d^2}$$

(Allard's Law, 1876)

Optical Thickness:

$$T = \beta d$$

Scaled Depth


BUT ... Distant objects appear Bright !



Mountains

Airlight Model

(Koschmeider, 1924)



Brightness due to a Path of Length d :

$$E(d) = E_{\infty} (1 - e^{-\beta d})$$

Horizon Brightness

Depth of Sources from Attenuation




Clear Day





Clear Night



Foggy Night



Computed Source Locations



Structure from Airlight



Mountain Range



Foggy Day Image



Computed Depth Map



3D Structure




Urban Scene



Foggy Day Image




Computed Depth Map




3D Structure

Attenuation + Airlight




Dichromatic Atmospheric Scattering

(Narasimhan et al, IJCV 2002)



$$E = p \hat{D} + q \hat{A}$$


Scattering Coefficient (β) vs. Wavelength (λ)



$d = 0.2$ km, viewing direction is 5 degrees to horizontal plane

Dichromatic Atmospheric Scattering

(Narasimhan et al, IJCV 2002)



$$E = p \hat{D} + q \hat{A}$$

Horizon
Brightness

Surface
Radiance


$$p = E_{\infty} e^{-\beta d} R$$

$$q = E_{\infty} (1 - e^{-\beta d})$$



Dichromatic Planes



Different Visibilities



Dichromatic Planes





Weather	Densities	Sky	Mean Angular Error
Fog	5	Overcast	0.58
Mist	5	Overcast	1.25
Rain	5	Overcast	1.13
Dense Haze	5	Overcast	1.47
Mild Haze	5	Sunny	3.61

Color Constraints using Dichromatic Model

Visibility 1



Visibility 2



Two Images give Scaled Depth:



$$(\beta_1 - \beta_2) d = \ln \frac{E_{\infty 1}}{E_{\infty 2}} - \ln \frac{p_1}{p_2}$$

Color Constraints using Dichromatic Model

Visibility 1




Visibility 2



Two Images give Scaled Depth:

$$(\beta_1 - \beta_2) d = \ln \frac{E_{\infty 1}}{E_{\infty 2}} - \ln \frac{p_1}{p_2}$$

Color Cube Boundary Algorithm



Minimum Time to Collision

$$\text{Min } \alpha = \frac{\tilde{\beta} d}{(\beta_2 - \beta_1) d}$$

First Collision with Color Boundary

Clear Day from Hazy Days

Unknown Hazy Conditions



Time: 3:00 PM



Time: 3:30 PM

Deweathering

Computed Clear Day Image



(Narasimhan et. al, IJCV 2002)


Gray World: Contrast Restoration and Structure



Mild Fog, 5 PM



Dense Fog, 5:30 PM




Computed Depth Map
Quantized – 20 levels



Contrast Restored Image

(Narasimhan et. al, PAMI 2003)



Defogging Videos



Foggy Video




Defogged Video




Histogram Equalized Video

Haze and Polarization



Trivial case



Yoav Schechner, Srinivasa Narasimhan, Shree Nayar

Instant Dehazing Using Polarization



Polarized Image (50 deg)



Polarized Image (140 deg)



Computed Depth Map



Dehazed Image

Instant Dehazing Using Polarization

Polarized Image (50 deg)



Polarized Image (140 deg)




Computed Depth Map



Dehazed Image



Multiple Scattering in the Atmosphere



Glows of Light Sources




Mist



Fog



How are Glows formed in Images?



Multiple Scattering Equation :


$$\cos \theta \frac{\partial E}{\partial T} + E(\theta, \theta) - \frac{1 - \cos^2 \theta}{T} \frac{\partial E}{\partial \theta} + \frac{1}{4\pi} \int P(\theta, \theta') E(T, \theta') d\theta' d\phi'$$

Phase Function Glow Intensity

(Chandrasekhar 1960)

Axially Symmetric Phase Functions

(Ishimaru 1978)



Legendre Polynomial Expansion:

$$P(\cos \alpha) = \sum_{m=0}^{\infty} [(2m+1) q^m] L_m(\cos \alpha)$$

Forward Scattering Parameter

Legendre Polynomial

Analytic Multiple Scattering Solution

Glow Intensity:

$$E(T, \theta) = \sum_{m=0}^{\infty} (g_m(T) + g_{m+1}(T)) L_m(\theta)$$

Legendre Polynomial

Exponential Coefficients:

$$g_m(T) = E_0 \exp \left[-(m+1) \ln T - \frac{2m+1}{m} (1 - q^{m-1}) T \right]$$

Optical Thickness


Phase Function Parameter

Radiant Intensity of Source


(Narasimhan et. al, CVPR 2003)

Highlights of the Solution


- Absorbing and Purely Scattering Media
- Isotropic and Anisotropic Phase Functions
- Small Number of Coefficients (m) :



What are the Shapes of the Glows?



Mild Weather ($T = 1.2$)



Dense Weather ($T = 4$)

Validation: Multiple Scattering in Milk

Image acquired
With No Milk



Increasing Milk Concentrations




Original Milk
Images




Rendered Milk
Images

[1% – 3% error]




Model Fit Accuracy: Experiments with Milk

Low Milk Concentration



High Milk Concentration




Number of Milk Concentrations : 15


Model Fitting Error : [1 % to 3 %]

Diffusion Fitting Error : [20 % to 50 %]

Model Fit Accuracy: Monte Carlo Simulations



Low Concentration




Higher Concentration

Effect of Source Visibility



Observed Milk Images

Increasing Milk Concentrations



What do we get from Glows?


Inverse RTE :

$$\arg \min_{T, q} \left\| E(T, \mu) - \sum_{m=0}^{\infty} (g_m(T, q) + g_{m+1}(T, q)) L_m(\mu) \right\|$$

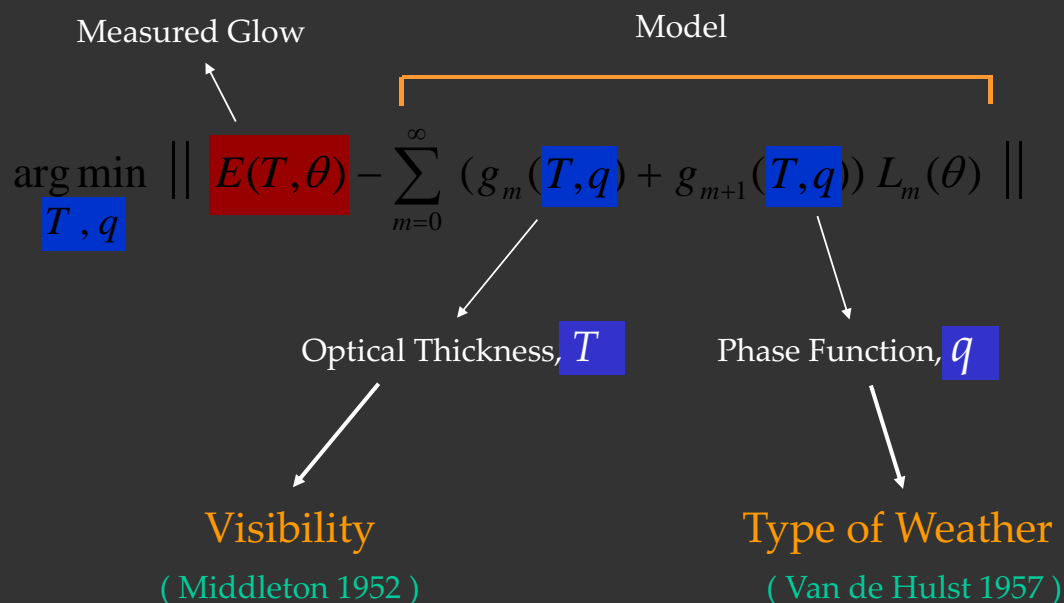
Optical Thickness
Phase Function Parameter

Meteorological Visibility: (Middleton 1952)


$$V \approx \frac{3.912}{T} R$$



What do we get from Glows?




Camera-based Weather Station




Images of a Light Source

Computed
Visibilities (miles)





Computed
Type of Weather



Active Visibility Meters




Camera-based Weather Station




Images of a Light Source

Computed
Visibilities (miles)




Computed
Type of Weather



Rendering Glows using Convolution

Original Image



Rendered Images




Increasing Fog →



Rendering Glows using Convolution

Original Image



Rendered Images



————— Increasing Fog —————>



Single versus Multiple Scattering



Original Image



Single Scattering



Multiple Scattering (Mild Condition)



Multiple Scattering (Dense Condition)

Summary

- No Escape from Bad Weather.
- Image Processing does Not Suffice.
- Bad Weather can be Good for Vision.

Non-uniform and Dynamic Weather



Non-uniform Fog



Rain Drops and Rain Streaks



Snow Flakes and Snow Streaks

The Beauty of the Atmosphere



Clear Blue Skies



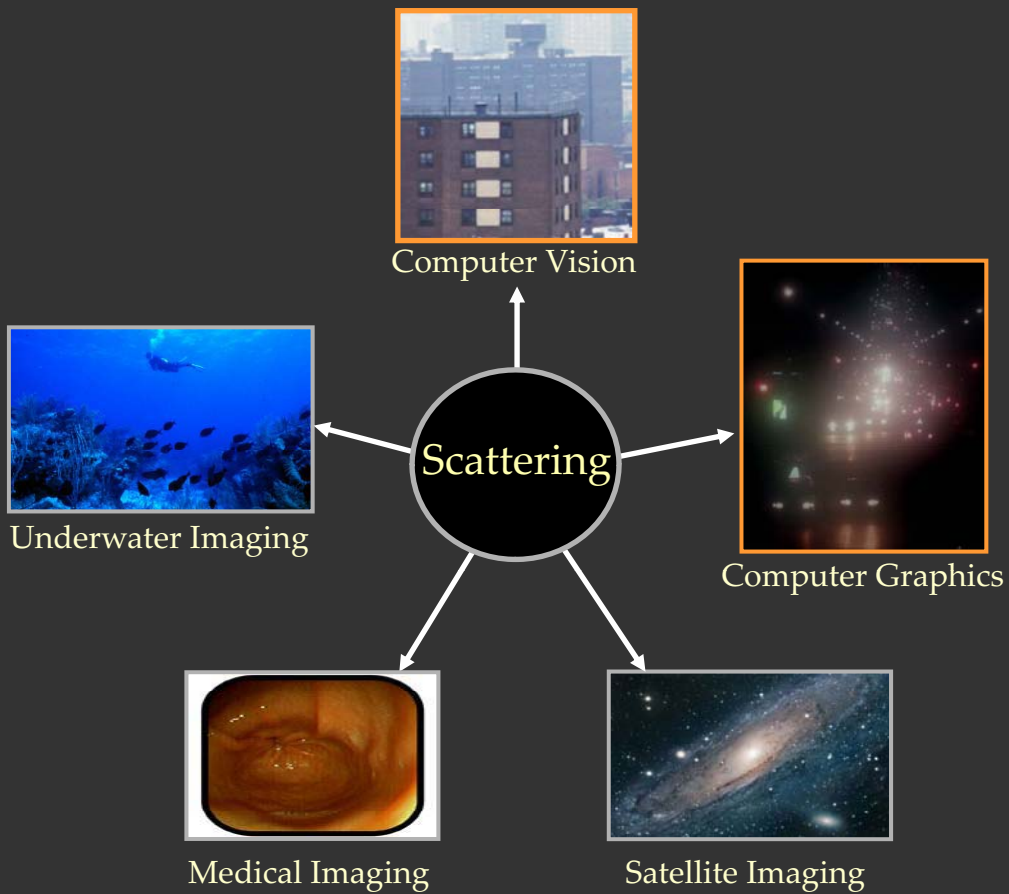
Colors of Sunset and Sunrise



Clouds



Rainbow



Acknowledgements



Shree Nayar



Yoav Schechner



Ravi Ramamoorthi




Jan Koenderink




Support: NSF, DARPA

THANK YOU!


Acquiring Scattering Properties



A High-Dynamic Range Image




Monte Carlo Verification



Measured Materials


Material	σ'_s [mm^{-1}]			σ_a [mm^{-1}]		
	R	G	B	R	G	B
Apple	2.29	2.39	1.97	0.0030	0.0034	0.046
Chicken1	0.15	0.21	0.38	0.015	0.077	0.19
Chicken2	0.19	0.25	0.32	0.018	0.088	0.20
Cream	7.38	5.47	3.15	0.0002	0.0028	0.0163
Ketchup	0.18	0.07	0.03	0.061	0.97	1.45
Marble	2.19	2.62	3.00	0.0021	0.0041	0.0071
Potato	0.68	0.70	0.55	0.0024	0.0090	0.12
Skimmilk	0.70	1.22	1.90	0.0014	0.0025	0.0142
Skin1	0.74	0.88	1.01	0.032	0.17	0.48
Skin2	1.09	1.59	1.79	0.013	0.070	0.145
Spectralon	11.6	20.4	14.9	0.00	0.00	0.00
Wholemilk	2.55	3.21	3.77	0.0011	0.0024	0.014

Measuring Scattering Anisotropy




[Joshi, Donner, and Jensen - Optics Letters 2006]

Measuring Scattering Anisotropy




Measurement via Dilution




[Narasimhan, Gupta, Donner, Ramamoorthi, Nayar, and Jensen - SIGGRAPH 2006]


Measurement via Dilution




Measurement via Dilution



Measurement via Dilution



Skin Acquisition




[Weirich et al., SIGGRAPH 2006]



Translucency Acquisition




Skin Rendering




Hair Measurements




Black Hair



Blond Hair




Synthetic Hair



Tilted Cuticle Scales



Hair Model



Real Hair




Real Black Hair




Real Brown Hair




The Final Measurement




Caustics




Hair Glints



Hair Scattering



Hair Model



Real Hair



Photograph



Gray 1997

Photograph



Gray 1997

Photograph vs. Kajiya-Kay



Photograph



Kajiya-Kay

Photograph vs. New Model




Photograph




New Model

Components of the New Model



Components of the New Model



Glints



Different Lighting Directions



Lorenz-Mie Theory

$$s_1(\theta) = \sum_{n=1}^{\infty} \frac{2n+1}{n(n+1)} (a_n \pi_n(\cos \theta) + b_n \tau_n(\cos \theta)) ,$$
$$s_2(\theta) = \sum_{n=1}^{\infty} \frac{2n+1}{n(n+1)} (b_n \pi_n(\cos \theta) + a_n \tau_n(\cos \theta)) .$$

[Revall, Christensen, and Jensen - SIGGRAPH 2007]

Milk Components



Computed scattering properties of milk components

Milk



Milk with varying fat concentration

Ice Dragons



Pure



Compacted




White

Green Icebergs




10 am




7 pm


Ocean Water



Atlantic



Mediterranean



Baltic



North Sea

What's Next?



A Practical Analytic Single Scattering Model for Real Time Rendering*

Bo Sun
Columbia University

Ravi Ramamoorthi
Columbia University

Srinivasa G. Narasimhan
Carnegie Mellon University

Shree K. Nayar
Columbia University

Abstract

We consider real-time rendering of scenes in participating media, capturing the effects of light scattering in fog, mist and haze. While a number of sophisticated approaches based on Monte Carlo and finite element simulation have been developed, those methods do not work at interactive rates. The most common real-time methods are essentially simple variants of the OpenGL fog model. While easy to use and specify, that model excludes many important qualitative effects like glows around light sources, the impact of volumetric scattering on the appearance of surfaces such as the diffusing of glossy highlights, and the appearance under complex lighting such as environment maps. In this paper, we present an alternative physically based approach that captures these effects while maintaining real-time performance and the ease-of-use of the OpenGL fog model. Our method is based on an explicit analytic integration of the single scattering light transport equations for an isotropic point light source in a homogeneous participating medium. We can implement the model in modern programmable graphics hardware using a few small numerical lookup tables stored as texture maps. Our model can also be easily adapted to generate the appearances of materials with arbitrary BRDFs, environment map lighting, and precomputed radiance transfer methods, in the presence of participating media. Hence, our techniques can be widely used in real-time rendering.

1 Introduction

Many real-time rendering applications like games or interactive simulations seek to incorporate atmospheric effects such as mist, fog and haze. These participating media lead to a number of qualitative effects not present in clear-day conditions (compare figure 1a with our result in figure 1c). For instance, there are often glows around light sources because of scattering. The shading on objects is also softer, with specular highlights diffused out, dark regions brightened and shadows softer. It is critical to capture these effects to create realistic renderings of scenes in participating media.

In computer graphics, the approaches for capturing these effects represent two ends in the spectrum of speed and quality. For high-quality rendering, a number of Monte Carlo and finite element techniques have been proposed. These methods can model very general volumetric phenomena and scattering effects. However, they are slow, usually taking hours to render a single image. Significant gains in efficiency can generally be obtained only by substantial precomputation, and specializing to very specific types of scenes.

At the other extreme, perhaps the most common approach for interactive rendering is to use the OpenGL fog model, which simply blends the fog color with the object color, based on the distance of the viewer (figure 1b). The fog model captures the attenuation of surface radiance with distance in participating media. This model is also popular because of its simplicity—implementation requires almost no modification to the scene description, and the user need only specify one parameter, β , corresponding to the scattering coefficient of the medium (density of fog). However, many qualitative



(a) Clear day



(b) OpenGL fog



(c) Our model

Figure 1: Rendered images of a scene with 66,454 texture-mapped triangles and 4 point lights. The insets show an image for another view of the vase, with highlights from all 4 sources, to amplify shading differences. (a) Standard OpenGL rendering (without fog), (b) OpenGL fog which captures attenuation with distance and blending with fog color, and (c) Our real-time model, that includes the glows around light sources, and changes to surface shading such as dimming of diffuse radiance (floor and wall), brightening of dark regions (back side of pillars and vases) and dimming and diffusing of specular highlights (inset). All the visual effects in this complex scene are rendered by our method at about 20 frames per second.

*e-mail: {bosun,ravir,nayar}@cs.columbia.edu; srinivas@cs.cmu.edu




Figure 2: Diagrams showing three cases of how light travels to the viewer through the participating medium. In (a) light travels in a straight line and directly reaches the surface and the viewer. This is essentially what previous interactive models such as OpenGL fog compute. In (b), in addition to what happens in (a), airlight scatters to the viewer and produces effects like glows around the light source. In (c), in addition to what happens in (b), airlight also scatters to the surface and gets reflected, leading to effects such as the diffusing out of specular highlights and brightening of darker regions. In image (d), reflected rays from the surface also scatter to the viewer.

effects are missing, such as the glows around light sources, the effect of scattering on object shading, and the ability to incorporate complex lighting effects like environment maps.

In this paper, we take a significant step towards improving the realism of rendered images with participating media (figure 1c), while maintaining the real-time performance and ease of use of the OpenGL fog model. Our model can be implemented as a simple vertex or pixel shader (pseudocode in figure 13), allowing it to be easily added to almost any interactive application. The method can also be applied with complex lighting, allowing environment mapping and precomputed radiance transfer to be used interactively with participating media for the first time (figures 15 and 16).

Figure 2 illustrates three important visual effects due to light transport in scattering media. In this discussion, and this paper, we assume single scattering (i.e. that light scatters at most once in the medium), which is a common approximation in volumetric scattering and can be shown to be accurate in many common situations such as thin fog. Figure 2a corresponds to *direct transmission* of light from the source or surfaces to the viewer. We can simply attenuate the clear-day radiance values based on the distance (optical thickness). This simple approach is essentially what interactive models like OpenGL fog implement. Figure 2b also includes the glows around light sources, commonly referred to as *airlight* [Koschmeider 1924]. Glows occur because light reaches the viewer from different directions due to scattering in the atmosphere. Figure 2c further includes the effect of airlight on the outgoing *surface radiance*, leading to effects such as the spreading out of specular highlights and softening of shadows. These are important effects, usually neglected in previous interactive methods. Our model renders all of the effects in figure 2c in real-time.

Figure 2d illustrates the case where the surface radiance is single scattered in addition to being attenuated, before reaching the view point. On one hand, the attenuation decreases the brightness of the radiance at the surface according to the distance of the surface from the viewer. On the other hand, the single scattering results in slight brightening and blurring of this surface radiance. Implementing the latter effect requires a depth-dependent convolution. In this paper, we will only consider attenuation of surface radiance, and we will set aside a more thorough investigation of the latter effect for future work¹. The specific technical contributions of this paper are:

Explicit Compact Formula for Single Scattering: The common approach to using single scattering is to numerically integrate brightness contributions while marching along the viewing ray. However, this approach is too slow for interactive applications, which require an explicit formula such as the OpenGL fog model. One of the main contributions of this paper is the derivation of an explicit compact formula for the single scattering from an isotropic point source in a homogeneous participating medium, by analytically integrating the single scattering equations. This *airlight model* (section 3) allows us to simulate effects like the glows around light sources (figure 2b). We can also use the model to calculate the effects of scattering on the surface shading (figure 2c). These calculations are very expensive even for numerical integration, because we

¹Single scattering from different surface points in the scene can partially compensate for the loss of brightness due to attenuation. Neglecting this can produce consistently darker images, especially for indoor scenes.

must consider incident airlight from the entire visible hemisphere. However, they can be directly implemented using our explicit *surface radiance model* (section 4).

Implementation on Programmable Graphics Hardware:

We speculate that an explicit formula for the single scattering integrals has previously not been derived because of the complexity of the calculations involved. In this paper, we reduce these difficult integrals to a combination of analytic functions that depend only on the physical parameters of the problem, and a few lookups of tabulated 2D functions, that have no simple analytic form, but are smooth and purely numerical—independent of the physical parameters. The numerical functions can be precomputed and stored as 2D texture maps, and the entire analytic computation and table lookups can be implemented in simple pixel or vertex shaders in modern programmable graphics hardware (section 5).

Extensions to Complex Lighting and BRDFs: It is also possible to extend our airlight and surface radiance models to incorporate more complex illumination models and material properties (section 6). Mathematically, we derive a point-spread function (PSF) to represent the glow around a light source. We can convolve an environment map with this PSF to get the appearance of a foggy scene under natural lighting. We can also use a frequency domain spherical harmonic representation to enable rendering with arbitrary BRDFs, and add in shadows and interreflections with precomputed radiance transfer methods. This approach enables methods such as environment mapping and precomputed radiance transfer to be used with volumetric scattering effects for the first time.

Our goal is to achieve interactive rendering of participating media. To enable this, and derive an explicit compact expression that can be implemented in real-time, we make a number of assumptions—isotropic point light sources, homogeneous media, the single scattering approximation, and no cast or volumetric shadows (shadows can, however, be added using precomputed light transport methods). More complex and general scattering effects are certainly desirable in many situations, but are not possible to obtain at real-time rates for general scenes. On the other hand, our method captures most of the important visual effects of scattering, while being very simple to add to any interactive application.

2 Related Work

The literature on simulating volumetric effects is large, going back to [Blinn 1982], and we only discuss important representative papers. Most techniques are based on numerical or analytic approximations to the radiative transfer equation [Chandrasekhar 1960]. Monte Carlo ray tracing methods were adapted by computer graphics researchers to render impressive effects including multiple scattering and non-homogeneous media [Kajiya and Herzen 1984; Max 1994; Jensen 2001]. However, such methods can take hours to render a single image. To speed up rendering, numerical methods that only simulate single scattering have also been proposed [Pattanaik and Mudur 1993; Nakamae et al. 1990; Sakas 1990; Rushmeier and Torrance 1987]. However, they still require significant running times, and are not suitable for interactive applications.

Hardware-accelerated numerical methods: A number of recent hardware-accelerated techniques can significantly decrease the running times of numerical simulations, although they are still usually not fast enough for many interactive applications such as games. Dobashi et al. [2002] describe a multi-pass rendering technique that numerically integrates the single scattering equations, using graphics hardware to accumulate the results at a number of planes in the scene, similar to volume rendering. Harris and Lastra [2001] render clouds by including a forward scattering term in addition to single scattering. Note that their method is geared toward the case when the viewer is far from the clouds, and they apply a different and slower approach when the viewer and scene are immersed inside the medium, as is the scenario in our work.

These methods are intended to apply to specific phenomena like the sky or clouds [Dobashi et al. 2002; Riley et al. 2004; Harris and Lastra 2001]. This allows them to make use of complex tabular volume specifications, precomputed lighting solutions or multipass rendering techniques to produce effects including inhomogeneous media and simple heuristics for multiple scattering. They allow for viewpoint, and in a few cases interactive lighting variation, but usually fix the medium properties and scene specification.

In contrast, our technique, while focusing on homogeneous media and single scattering, can be encapsulated in a simple shader for general scenes, and allows for real time variation of the viewpoint, lighting, scattering properties of the medium, and even scene geometry and reflectance. Another major benefit of our method is that it addresses the effects of scattering on surface shading (figure 2c) and complex lighting like environment maps. These effects are not included in previous methods because they are difficult to numerically simulate efficiently, requiring an integration over all incident scattered lighting directions at each surface point.

Analytically based methods: The diffusion approximation for optically thick media was applied to subsurface scattering [Stam 1995; Jensen et al. 2001]. An analytic form for the single scattering term was also derived by Hanrahan and Krueger [1993]. However, the problem we are solving is very different from that of subsurface scattering, where the light sources and viewer are outside the medium. In our case, both the sources and viewer are immersed *inside* the medium. Also, unlike in the case of diffusion, we are interested in strongly directional effects like glows around sources.

Analytic expressions for airlight with directional light sources, based on the derivation by Koschmeider [1924], are used frequently for rendering skies [Preetham et al. 1999; Hoffman and Preetham 2003; Narasimhan and Nayar 2002]. However, our focus is different. We wish to derive an analytic model with “near-field” point sources, which is a significantly more complex lighting situation as compared to distant lighting (collimated beams).

Analytic expressions for the glows around point light sources *inside* homogeneous media have also been derived [Max. 1986; Biri et al. 2004; Narasimhan and Nayar 2003]. Therefore, those methods could be used to render glows in real time. However, it is not clear how to extend them to a complete real-time rendering system that also considers the effects of airlight on surface shading, or handles complex environment map lighting. Furthermore, their derivations involve approximations that are not feasible in several common rendering scenarios. For instance, the model derived by Max [1986] does not take into account attenuation. Biri et al. [2004] use a polynomial approximation to single scattering which results in inaccurate glows along viewing directions near the source. The multiple scattering model in [Narasimhan and Nayar 2003] is not strictly valid when objects are present in the medium, especially near the sources (as is generally true in most common scenes), or for optically thin media. Further, the integration required for surface radiance cannot be computed analytically or simulated numerically at interactive rates.

3 The Airlight Model

In this section, we will derive an explicit model for the single scattered radiance at a viewer, due to an isotropic point light source, assuming that both the viewer and the source are immersed in a homogeneous scattering medium. Consider the scenario illustrated in

s, v, p	Subscripts for Source, Viewer, surface Point
γ	Angle between light source and viewing ray
D_{sv}	Distance between source and viewer
D_{vp}	Distance between viewer and closest surface point
D_{sp}	Distance between source and surface point
T_{sv}	Optical thickness between source, viewer (βD_{sv})
T_{vp}	Optical thickness between viewer, surface point (βD_{vp})
T_{sp}	Optical thickness between source, surface point (βD_{sp})
β	Scattering coefficient of the participating medium
α	Angle of scattering
x	Distance along the ray from viewer (integration variable)
d	Distance of single scattering from light source
I_0	Radiant intensity of point light source
f_r	BRDF of surface

Figure 3: Notation used in our derivations.




Figure 4: Diagram showing how light is scattered once and travels from a point light source to the viewer.

figure 4 (the notations used are indicated in figure 3). The point light source has a radiant intensity I_0 and is at a distance D_{sv} from the view point, making an angle γ with the viewing direction. The radiance, L , is composed of the direct transmission, L_d , and the single scattered radiance or airlight, L_a ,

$$L = L_d + L_a. \quad (1)$$

The direct term L_d simply attenuates the incident radiance from a point source (I_0/D_{sv}^2) by an exponential corresponding to the distance between source and viewer, and the scattering coefficient² β ,

$$L_d(\gamma, D_{sv}, \beta) = \frac{I_0}{D_{sv}^2} e^{-\beta D_{sv}} \cdot \delta(\gamma), \quad (2)$$

where the delta function indicates that for direct transmission, we receive radiance only from the direction of the source (no glows).

3.1 The Airlight Integral

We focus most of our attention on the airlight L_a . The standard expression [Nishita and Nakamae 1987] is given by an integral along the viewing direction,

$$L_a(\gamma, D_{sv}, D_{vp}, \beta) = \int_0^{D_{vp}} \beta k(\alpha) \cdot \frac{I_0 \cdot e^{-\beta d}}{d^2} \cdot e^{-\beta x} dx, \quad (3)$$

where D_{vp} is the distance to the closest surface point along the viewing ray or infinity if there are no objects in that direction, and $k(\alpha)$ is the particle phase function. The exponential attenuation corresponds to the total path length traveled, $d + x$. The two parameters d and angle α in the integrand depend on x . In particular, d is

²When there is light absorption in addition to scattering, β is called the extinction coefficient and is given by the sum of the scattering and absorption coefficients. In this paper, we simply refer to β as the scattering coefficient, and it is straightforward to include absorption in our models.




Figure 5: 3D plot of special function $F(u, v)$ in the range of $0 \leq u \leq 10$ and $0 \leq v \leq \frac{\pi}{2}$. The plot shows that the function is well-behaved and smooth and can therefore be precomputed as a 2D table. As expected from the definition in equation 10, the function decreases as u increases, and increases as v increases. The maximum value in the plot above therefore occurs at $(u = 0, v = \frac{\pi}{2})$. Also note from equation 10, that for $u = 0$, there is no attenuation so the function is linear in v .

given by the cosine rule as

$$d = \sqrt{D_{sv}^2 + x^2 - 2xD_{sv}\cos\gamma}. \quad (4)$$

Let us now substitute equation 4 into equation 3. For now, we also assume the phase function $k(\alpha)$ is isotropic and normalized to $1/4\pi$ (our approach can also be generalized to arbitrary phase functions—see appendix D on CDROM). In this case,

$$L_a(\gamma, D_{sv}, D_{vp}, \beta) = \frac{\beta I_0}{4\pi} \int_0^{D_{vp}} \frac{e^{-\beta\sqrt{D_{sv}^2 + x^2 - 2xD_{sv}\cos\gamma}}}{D_{sv}^2 + x^2 - 2xD_{sv}\cos\gamma} \cdot e^{-\beta x} dx. \quad (5)$$

We refer to this equation as the *airlight single scattering integral* and next focus on simplifying it further to derive an explicit form.

3.2 Solution to the Airlight Integral

We take a hybrid approach to solve equation 5. The key result is that this integral can be factorized into two expressions—(a) an **analytic expression** that depends on the physical parameters of the scene and (b) a two-dimensional numerically tabulated function that is **independent** of the physical parameters. Essentially, this factorization enables us to evaluate the integral in equation 5 analytically. A high-level sketch of the derivation is given below and detailed simplifications are included in appendix A.

STEP 1. Reducing the dimensions of the integral: Since the integral in equation 5 depends on 4 parameters, our first step is to apply a series of substitutions that reduce the dependency of the integrand to only one parameter. For this, we first write the expressions in terms of optical thicknesses $T_* = \beta D_*$ and $t = \beta x$. In most cases, this eliminates the separate dependence on both β and the distance parameters, somewhat reducing the complexity, and giving us a simpler intuition regarding the expression's behavior. Then, we combine the dependence on T_{sv} and γ by making the substitution $z = t - T_{sv}\cos\gamma$, to obtain

$$L_a(\gamma, T_{sv}, T_{vp}, \beta) = \frac{\beta^2 I_0}{4\pi} e^{-T_{sv}\cos\gamma} \int_{-T_{sv}\cos\gamma}^{T_{vp}-T_{sv}\cos\gamma} \frac{e^{-z - \sqrt{z^2 + T_{sv}^2 \sin^2 \gamma}}}{T_{sv}^2 \sin^2 \gamma + z^2} dz. \quad (6)$$

Now, the integrand really depends on only one physical parameter $T_{sv} \sin \gamma$, beginning to make the computation tractable.

It is possible to further simplify equation 6, as described in appendix A. To encapsulate the dependence on the physical parameters of the problem, we define the following two auxiliary expressions, corresponding respectively to the normalization term outside the integrand, and the single physical parameter in the integrand,




Figure 6: Accuracy of the airlight model. The plots show the error (versus numerically integrating equation 5) as a function of the resolution for the 2D tables for $F(u, v)$. We report the fractional error, normalizing by the total airlight over the hemisphere. The error for each resolution is averaged over 40000 parameter values of β , D_{sv} , D_{vp} and γ . Bilinear (red) and nearest neighbor (green) interpolation is used to interpolate $F(u, v)$ at non-grid locations of the indices (u, v) . The plots clearly indicate the high accuracy of our compact formula, and that a 64×64 table for $F(u, v)$ suffices for a maximum error of less than 2%.




Figure 7: Comparison of the airlight model with a standard Monte Carlo simulation that includes multiple scattering. The plots show the relative RMS error between the two methods for the case of isotropic phase function. [Left] The low RMS errors show that our model is physically accurate (less than 4% error) for optically thin media ($T_{sv} \leq 2$). [Right] From this plot, it is evident that multiple scattering becomes more important as optical thickness increases. However, the actual errors grow slowly and are still low for a wide range of optical thicknesses ($T_{sv} < 10$). It is also interesting to note that for very high optical thicknesses ($T_{sv} > 20$), attenuation dominates over scattering and once again the RMS errors decrease.

$T_{sv} \sin \gamma$:

$$A_0(T_{sv}, \gamma, \beta) = \frac{\beta^2 I_0 e^{-T_{sv}\cos\gamma}}{2\pi T_{sv} \sin \gamma} \quad (7)$$

$$A_1(T_{sv}, \gamma) = T_{sv} \sin \gamma. \quad (8)$$

It is then possible to derive, as shown in appendix A, that

$$L_a = A_0(T_{sv}, \gamma, \beta) \int_{\gamma/2}^{\frac{\pi}{4} + \frac{1}{2} \arctan \frac{T_{vp}-T_{sv}\cos\gamma}{T_{sv}\sin\gamma}} \exp[-A_1(T_{sv}, \gamma) \tan \xi] d\xi. \quad (9)$$

Although equation 9 might seem complicated, it is really in a simplified form. We already have simple analytic expressions for A_0 and A_1 . Further, the function A_1 is a numerical constant as far as the integration is concerned.

STEP 2. Evaluating the integral using a special function: To encapsulate the key concepts in the integrand of equation 9, we define the special function,

$$F(u, v) = \int_0^v \exp[-u \tan \xi] d\xi. \quad (10)$$



Figure 8: The images show glows around three identical point light sources (street lamps) at different distances from the viewer. From left to right, we show three different values of the scattering coefficient β ($\beta = 0, 0.01, 0.04$). Larger values of β correspond to larger optical thicknesses T_{sv} . We clearly see the effect of greater glows for larger β . Also, the radiance from farther light sources is attenuated more in each individual image, resulting in smaller glows for the farther lights. In the fourth (rightmost) image, we show a different view with $\beta = 0.04$, where all the light sources are approximately equidistant, with the result that they have similar glows. (The shading on the surfaces is computed using the surface radiance model in section 4.)

Unfortunately, there exists no simple analytic expression for $F(u, v)$. However, the function is a well behaved 2D function as shown in figure 5. Therefore, we can simply store it numerically as a 2D table. This is really no different from defining functions like sines and cosines in terms of lookup tables. In practice, we will use texture mapping in graphics hardware to access this 2D table. Note that $F(u, v)$ is purely numerical (independent of the physical parameters of the problem), and thus needs to be precomputed only once.

Finally, we can obtain for $L_a(\gamma, T_{sv}, T_{vp}, \beta)$,

$$L_a = A_0 \left[F(A_1, \frac{\pi}{4} + \frac{1}{2} \arctan \frac{T_{vp} - T_{sv} \cos \gamma}{T_{sv} \sin \gamma}) - F(A_1, \frac{\gamma}{2}) \right], \quad (11)$$

where we have omitted the parameters for L_a, A_0 and A_1 for brevity.

In the important special case of $T_{vp} = \infty$, corresponding to no objects along the viewing ray, we get $L_a(\gamma, T_{sv}, \infty, \beta)$ as

$$L_a = A_0(T_{sv}, \gamma, \beta) \left[F(A_1(T_{sv}, \gamma), \frac{\pi}{2}) - F(A_1(T_{sv}, \gamma), \frac{\gamma}{2}) \right]. \quad (12)$$

In summary, we have reduced the computation of a seemingly complex single scattering integral in equation 5 into a combination of an analytic function computation that depends on the physical parameters of the problem and a lookup of a pre-computed 2D smooth function that is independent of the physical parameters of the problem. In the rest of the paper, we will demonstrate several extensions and applications of our model.

3.3 Accuracy of the Airlight Model

We first investigate the accuracy of our analytic model as compared to numerically integrating equation 5. Figure 6 shows plots of the mean error in L_a as a function of the resolution of the 2D numerical table for the special function $F(u, v)$. We use interpolation to evaluate $F(u, v)$ at non-grid locations for the indices (u, v) (bilinear and nearest neighbor interpolations are shown in figure 6). For each resolution, the error computed is averaged over 40000 sets of parameter values for $\beta, D_{sv}, D_{vp}, \gamma$. The error bars in the figure show the standard deviation of the errors. The plots indicate that even a low resolution 64×64 table suffices to compute $F(u, v)$ accurately, with a maximum error of less than 2%. As expected, bilinear interpolation performs better, but, for faster rendering, one can use nearest neighbor interpolation with only a small loss in accuracy.

We also validate the accuracy of the single scattering assumption in our airlight model. Figure 7 shows the relative RMS errors between glows around light sources computed using our model and a standard volumetric Monte Carlo approach that takes into account multiple scattering as well. The Monte Carlo simulation took approximately two hours to compute each glow, whereas our explicit model runs in real-time. The comparison was conducted for optical thicknesses over a wide range $T_{sv} \in (0.25, 25)$ and $T_{vp} \in (0.5, 50)$, which covers almost all real situations. As expected, for optically thin media ($T_{sv} \leq 2$), our model is very accurate (less than 4% relative RMS error). Interestingly, even for greater optical thicknesses ($T_{sv} > 2$), the error only increases slowly. Thus, our single scattering model may be used as a viable approximation for most common

real-time rendering scenarios, such as games.

3.4 Visual Effects of the Airlight Model

The dependence of the model on the viewing direction γ and the distance of the source from the observer D_{sv} , predicts visual effects like the glows around light sources and the fading of distant objects. As discussed above, these effects are physically accurate for thin fog (low β and T), and qualitatively reasonable in other cases. In figure 8, we also see how these glows change as a function of the medium properties (the scattering coefficient β) and distance to the sources. As β increases, we go from no glow ($\beta = T = 0$) to a significant glow due to scattering. The differences in the 3 light sources should also be observed. The farther lights are attenuated more, and we perceive this effect in the form of reduced glows around more distant sources. The final (rightmost) image in figure 8 shows a different viewpoint, where the sources are at approximately the same distance, and the glows therefore look the same.

4 The Surface Radiance Model

In this section, we discuss the effects of airlight on the outgoing surface radiance. Consider the illustration in figure 9, where an isotropic point light source s illuminates a surface point p . We will calculate the reflected radiance at the surface. To get the actual appearance at the viewer, we need to attenuate by $\exp[-T_{vp}]$ as usual, where T_{vp} is the optical thickness between viewer and surface point.

The reflected radiance L_p is the sum of contributions, $L_{p,d}$ and $L_{p,a}$, due to direct transmission from the source, and single scattered airlight from the source respectively,

$$L_p = L_{p,d} + L_{p,a}. \quad (13)$$

The direct transmission corresponds to the standard surface reflectance equation, only with an attenuation of $\exp[-T_{sp}]$ added because of the medium, where T_{sp} is the optical thickness between the source and the surface point:

$$L_{p,d} = \frac{I_0 e^{-T_{sp}}}{D_{sp}^2} f_r(\theta_s, \phi_s, \theta_v, \phi_v) \cos \theta_s, \quad (14)$$

where f_r is the BRDF, (θ_s, ϕ_s) is the direction to the source, and therefore also the incident direction, and (θ_v, ϕ_v) is the viewing direction. All angles are measured with respect to the surface normal, in the local coordinate frame of the surface.

4.1 The Surface Radiance Integral

On the other hand, the single-scattered radiance $L_{p,a}$ is more complicated, involving an integral of the airlight (L_a from equation 12) over all incident directions,

$$L_{p,a} = \int_{\Omega_{2\pi}} L_a(\gamma(\theta_s, \omega_i), T_{sp}, \infty, \beta) f_r(\theta_i, \phi_i, \theta_v, \phi_v) \cos \theta_i d\omega_i. \quad (15)$$




Figure 9: Diagram showing how light travels from a point light source to a surface point and gets reflected towards the viewer by the surface point.

Consider the parameters of L_a in the integrand of the above equation. The angle γ' in this case is the angle³ between the incident direction ω_i and the source direction (θ_s, ϕ_s) . Note that for isotropic BRDFs, we can always rotate the coordinate system so $\phi_s = 0$, allowing us to write $\gamma'(\theta_s, \omega_i)$. Finally, L_a also depends on the optical thickness between the source and the surface point T_{sp} (instead of between source and viewer in equation 12).

We refer to equation 15 as the *surface radiance single scattering integral*, analogous to the airlight single scattering integral in equation 5, and next focus on deriving an explicit compact form.

4.2 Solution to the Surface Radiance Integral for Lambertian and Phong BRDFs

First consider the Lambertian case, so the BRDF is a constant k_d . The integral will then depend only on the parameters of L_a , i.e. γ' , T_{sp} and β . Of these, the dependency on β is primarily a normalization factor and does not affect the integrand. The angle γ' is a function of the source direction θ_s , and the integration variable ω_i . Hence, the integrand depends on only two physical variables, T_{sp} and θ_s . Thus, as in the previous section, we can define a special two-dimensional numerical function $G_0(T_{sp}, \theta_s)$.

For the Phong BRDF, we employ the reparameterization method in [Ramamoorthi and Hanrahan 2002], measuring angles from the reflection of the viewing direction about the surface normal, rather than the surface normal itself. To indicate this, we denote by θ'_s the angle the source makes with respect to this reflected direction. Upon making this transformation, it can be shown that the Phong BRDF is mathematically analogous to the Lambertian case. To allow for the Phong exponent n , we define the 2D function G_n instead of G_0 . These functions are well-defined and smooth as shown by the plots in figure 10. The details of these calculations are in appendix B, and the formula for G_n is

$$G_n(T_{sp}, \theta'_s) = \int_{\Omega_{2\pi}} \frac{e^{-T_{sp} \cos \gamma'}}{\sin \gamma'} \left[F(A_1, \frac{\pi}{2}) - F(A_1, \frac{\gamma'}{2}) \right] \cos^n \theta'_i d\omega_i, \quad (16)$$

where γ' and $A_1(T_{sp}, \gamma')$ are functions of θ'_s and ω_i , i.e. $\gamma'(\theta'_s, \omega_i)$.

The final shading formula, considering both direct transmission and single scattering is then given by (derivation in appendix B):

$$L_p = I_0 k_d \left[\frac{e^{-T_{sp}}}{D_{sp}^2} \cos \theta_s + \beta^2 \frac{G_0(T_{sp}, \theta_s)}{2\pi T_{sp}} \right] + \quad (17)$$

$$I_0 k_s \left[\frac{e^{-T_{sp}}}{D_{sp}^2} \cos^n \theta'_s + \beta^2 \frac{G_n(T_{sp}, \theta'_s)}{2\pi T_{sp}} \right]. \quad (18)$$

As in the airlight model derivation, we have reduced the computation of surface radiance due to single scattering to a few analytic function evaluations and a few 2D table lookups.

³We use the prime on γ' to make a technical distinction from the angle γ between the source and viewer used in the previous section. An explicit trigonometric formula for γ' is given in appendix B.




Figure 10: 3D plots of functions G_0 and G_n for $n = 20$ in the range of $0 \leq T_{sp} \leq 10$ and $0 \leq \theta_s \leq \frac{\pi}{2}$. The plots show that both functions are well-defined and smooth and can therefore be precomputed as 2D tables. The functions reach their peak values for $\theta_s = T_{sp} = 0$, decaying with increases in both parameters. The decay is faster for the peakier G_{20} on the right.




Figure 11: Influence of scattering on Lambertian surface radiance. In the foggy image (b), created using our surface radiance model, we see a dimming due to attenuation and diffusing of shading (note the brightening of darker areas compared to the clear day image in (a)). These effects are perhaps more apparent in (c), where we also include airlight from the source.




Figure 12: Influence of scattering on specular surface radiance (top row has Phong exponent 10, bottom has Phong exponent 20). In the foggy images (b), we see a dimming and diffusing of the specular highlight compared to the clear-day image in (a). Note also the overall loss in color saturation and contrast, especially in (c). These are important visual effects, usually missing in previous interactive techniques.

4.3 Visual Effects of the Surface Radiance Model

To illustrate the different qualitative effects we see on surfaces due to scattering, we rendered spheres with Lambertian (figure 11) and Phong BRDFs (figure 12) using our model above. The columns are from left to right (a) no scattering, (b) the effects of scattering on surface shading, and (c) combining this with the airlight effects directly from the source. For the Lambertian sphere in figure 11, we see a dimming due to attenuation of light through the scattering medium, and the diffusing of shading leading to a brightening of darker shadow regions. In the specular case, we see a dimming and diffusing out of the specular highlight due to scattering, combined with an overall reduction in color saturation and contrast. These are important qualitative shading effects that add to the realism of scene appearance in scattering media.

```

frag2app fmain(
    float4 objPos : TEXCOORD3,           // 2D texture coords
    ...
    uniform samplerRECT F,               // 2D special functions
    uniform samplerRECT G0,
    uniform samplerRECT Gn)
{
    frag2app OUT;                        // output radiance
    // Set up and calculate Tsv, γ, Dsv, Tvp, θs and θ's

    /***** Compute La from equation 11 *****/
    A0 = (β * I0 * exp[-Tsv * cos γ]) / (2π * Dsv * sin γ); // equation 7
    A1 = Tsv * sin γ; // equation 8
    v = π/4 + (1/2) arctan[(Tvp - Tsv * cos γ) / (Tsv * sin γ)];
    // v is one of texture coords
    f1 = texRECT(F, float2(A1, v)); // 2D texture lookup
    f2 = texRECT(F, float2(A1, γ/2));
    airlight = A0 * (f1 - f2); // equation 11

    /***** Diffuse surface radiance from equation 17 *****/
    d1 = kd * exp[-Tsp] * cos θs * I0 / (Dsp * Dsv);
    d2 = (kd * I0 * β * β) / (2π * Tsp) * texRECT(G0, float2(Tsp, θs));
    diffuse = d1 + d2;

    /***** Specular surface radiance from equation 18 *****/
    s1 = ks * exp[-Tsp] * cosn θs * I0 / (Dsp * Dsv);
    s2 = (ks * I0 * β * β) / (2π * Tsp) * texRECT(Gn, float2(Tsp, θ's));
    specular = s1 + s2;

    /***** Final Color (equation 19) *****/
    OUT.color = airlight + (diffuse + specular) * exp[-Tvp];
    return OUT;
}

```

Figure 13: Pseudocode for the Cg fragment shader that implements our combined model for airlight and surface radiance.

5 The Complete Model and its Hardware Implementation

While the mathematical derivations in the previous two sections are somewhat involved, the actual implementation is straightforward. Our model provides an explicit form that can be easily implemented in modern programmable graphics hardware. This requires minimal changes to both the original rendering code and scene description, and can therefore be easily integrated into other existing real-time rendering methods. Indeed, the user need only specify the coefficient β of the medium, as in standard OpenGL fog, and use the shader corresponding to our model (pseudocode is in figure 13).

To compute the final appearance, we sum up the attenuated reflected radiance from the surface and the airlight from the source,

$$L = e^{-T_{vp}} L_p + L_a. \quad (19)$$

L_a is the airlight and is given by equation 11. L_p is the exitant radiance at the surface and is given by equations 17 and 18. **We only need to compute a few simple analytic terms and do 4 texture lookups for each vertex or pixel, two for special function F, and one each for G_0 and G_n** (these texture lookups correspond to the texRECT function call in the pseudocode of figure 13). Clearly, these computations can be done by modern programmable graphics cards interactively in a single rendering pass.

In practice, we implement the model using Cg in the fragment shader of an NVidia Geforce 6800 graphics card. The special functions F , G_0 and G_n are precomputed and tabulated as 64×64 floating point textures. Since these textures are precomputed only once, we minimize frequent data transfer between the graphics card and main memory.

The rendering speed depends on a variety of variables, and rendering time is linear in the number of light sources. As shown in the video, we are able to achieve real-time rates even for fairly

complex scenes with several light sources. As an example, we rendered the scene shown in figure 1, with 39,999 vertices and 66,454 triangles. We simulated the scattering effects from 4 light sources and achieved about 20 fps using the graphics hardware mentioned above. The model for the scene was obtained from <http://hdri.cgtechniques.com>.

6 Complex BRDFs and Lighting

So far, we have considered arbitrarily located point lights, and simple Lambertian and Phong BRDFs, showing how an explicit expression can be derived and implemented. Rendering time is linear in the number of lights. In this section, we show how these ideas can be extended to efficiently handle complex BRDFs and environment map lighting using convolution, if we are willing to make particular simplifying assumptions. We first introduce the notion of a point-spread function (PSF) for the radiance or glow from a point source due to single scattering. This is similar in spirit to the PSFs derived by Narasimhan and Nayar [2003] and Premoze et al. [2004] in the context of multiple scattering for offline rendering. We will then discuss a number of applications including

- Rendering arbitrary BRDFs with point light sources, by convolving the BRDF with this PSF, as shown in figure 14. This approach can be used if we are willing to precompute a tabular BRDF representation, instead of using a simple explicit formula, as for Lambertian and Phong BRDFs.
- Convolving an environment map with the PSF to efficiently handle very complex lighting (with possibly thousands of lights, corresponding to the pixels of an environment map). This convolution is possible if we assume that all light sources are equally far away, as in a distant environment map. This enables us to obtain the characteristic glows and blurriness around light sources on foggy days, as shown in figure 15.
- Integrating volumetric scattering into precomputed radiance transfer methods that include complex lighting, realistic materials, cast shadows and interreflections (figure 16). The idea of convolving with the point-spread function can be applied to almost any technique that uses environment maps, enabling environment mapping and precomputed radiance transfer to be used in participating media for the first time.

Throughout the section, we apply the signal-processing results of Ramamoorthi and Hanrahan [2001] and Basri and Jacobs [2003] to efficiently compute the convolutions in the frequency domain using spherical harmonics.

6.1 Airlight Point Spread Function (PSF)

In section 3, we determined the radiance $L_a(\gamma, D_{sv}, D_{vp}, \beta)$ from a point source reaching a viewer, due to single scattering. If we fix the distance to the source D_{sv} , the integrating distance D_{vp} , and the scattering coefficient β of the medium, the radiance becomes a function only of the angle γ . We normalize this function by I_0/D_{sv}^2 to account for the intensity of the source, and define the PSF as

$$\text{PSF}(\gamma)_{D_{sv}, D_{vp}, \beta} = \frac{D_{sv}^2 L(\gamma, D_{sv}, D_{vp}, \beta)}{I_0}. \quad (20)$$

Since the PSF is mostly applied for surface shading, we will generally set $D_{vp} = \infty$, as in section 4.

6.2 Empirical PSF factorization for Speedup

The PSF defined above still depends on the parameters of the medium such as the coefficient β . So, changing these parameters changes the PSF and requires us to redo any convolutions. However, we have observed empirically that the PSF above can be factored into **a purely angular component that is independent of the medium parameters and an amplitude component that depends on the medium parameters. This factorization enables us**




Figure 14: [Left] A teapot rendered using the measured blue metallic BRDF. [Right] The teapot as it appears in a scattering medium. The brightening of darker regions, and softening of shading, is clearly visible.

to change the medium parameters interactively without having to re-compute the PSF or redo any convolutions. Specifically,

$$\text{NPSF}(\gamma) = \frac{F(\sin \gamma, \frac{\pi}{2}) - F(\sin \gamma, \frac{\gamma}{2})}{2\pi \sin \gamma \cdot e^{(\cos \gamma - 1)}} \quad (21)$$

$$\text{PSF}(\gamma)_{D_{sv}, \infty, \beta} = T_{sv} e^{-T_{sv}} \cdot \text{NPSF}(\gamma), \quad (22)$$

where NPSF has only angular dependence, independent of other physical parameters. In appendix C, we derive and validate this approximation, showing plots that indicate there is little noticeable numerical error.

6.3 Rendering with arbitrary BRDFs

We can use convolution with the PSF to render with arbitrary tabulated BRDFs, such as measured reflectance. For each outgoing direction, we tabulate the BRDF as a function over the sphere of incident directions. A new effective BRDF can be obtained for that outgoing direction by convolving this function with the PSF.

Mathematically, we first write the (isotropic) BRDF in terms of spherical harmonic coefficients for each outgoing angle as

$$\rho^{\text{orig}}(\theta_i, \theta_o, \phi) = \sum_{l,m} \rho_{lm}^{\text{orig}}(\theta_o) Y_{lm}(\theta_i, \phi), \quad (23)$$

where ρ_{lm}^{orig} are the coefficients, and Y_{lm} is the spherical harmonic. To perform the convolution [Ramamoorthi and Hanrahan 2001], we multiply the coefficients ρ_{lm}^{orig} of the original BRDF by the corresponding coefficients of the point-spread function⁴ PSF_l ,

$$\rho_{lm}^{\text{eff}}(\theta_o) = \sqrt{\frac{4\pi}{2l+1}} \text{PSF}_l \rho_{lm}^{\text{orig}}(\theta_o). \quad (24)$$

Then, we can use the effective BRDF to compute the reflected radiance due to airlight, and the original BRDF for the reflected radiance due to direct transmission. Thus, standard rendering algorithms can be executed with only slight modification and at virtually no additional cost. Note however, that while our previous formulae for Lambertian and Phong models required no precomputations, the convolution approach requires precomputation of the spherical harmonic coefficients for a collection of outgoing angles.

Figure 14 shows images rendered with the Blue metallic BRDF measured by Matusik et al. [2003]. In the left image, we simply render a tabular description of the BRDF without scattering. In the right image, we use the formula above to compute a new effective tabulated BRDF, including the effects of airlight. The brightening of darker regions owing to scattering is clearly visible on the right.

6.4 Rendering with Environment Maps

Our point spread function can be applied directly to environment maps, with the effects of scattering obtained by convolving the environment map with the PSF. To use a single PSF for all sources, we must assume that the lighting is made of small equidistant light

⁴Since the PSF is radially symmetric, depending only on γ , only spherical harmonic coefficients with $m = 0$ are nonzero.



Figure 15: [Top] Grace cathedral environment map with no scattering. [Middle] The environment map is convolved with the airlight single scattering PSF to create a foggy/misty appearance. Notice the glows around the light sources, the blurring of the sources and the brightening of dark areas. [Bottom] A scene illuminated by the environment map without scattering (left) and with scattering (right). Notice the spreading of the highlights and brightening of the spheres due to scattering.

sources (fixed D_{sv}). This is a good approximation when the size of the objects is small compared to the distance to the environment⁵.

We first consider simply looking at the environment, where we would like to see the glows around the bright light sources, to create the effects of foggy or misty appearance. To achieve this effect, we simply need to convolve the environment map with the PSF,

$$L_{lm}^{\text{convolved}} = \sqrt{\frac{4\pi}{2l+1}} \text{PSF}_l L_{lm}^{\text{original}}. \quad (25)$$

Furthermore, similar to equation 1, we can simply use a combination of the original attenuated environment map $L^{\text{attenuated}}$ (for direct transmission, and corresponds to L_d in equation 1) and the convolved version $L^{\text{convolved}}$ above (for airlight, and corresponds to L_a in equation 1) to compute the surface shading,

$$L^{\text{final}} = L^{\text{attenuated}} + L^{\text{convolved}} \quad (26)$$

$$L^{\text{attenuated}} = L^{\text{original}} e^{-T_{sv}}. \quad (27)$$

Figure 15 shows results obtained by convolving the Grace Cathedral environment map [Debevec 1998] with the single scattering PSF. The blurring of light sources and the overall increase in brightness in dark regions can be immediately seen. Below that, we compare the appearances of spheres rendered illuminated by this environment map with and without scattering. Notice the spreading of highlights and the brightening of the objects.

⁵Note that while this assumption is similar to standard environment mapping, our PSF requires us to specify a finite (but possibly large) D_{sv} .




Figure 16: [Left] A scene rendered using precomputed radiance transport, to capture the complex shadows and specular surface BRDFs. [Middle] The same scene rendered as though it was immersed in a scattering medium, with the observer close to the object. Notice the blurring and fogging of the environment in the background. In the insets, we see a number of shading changes, such as the brightening of dark regions in the face because of the scattering of light, the attenuation due to dimming and diffusing of specular highlights on the base, and the softening of shadows on the plane. [Right] The same scene including effects of airlight between viewer and object (as if the viewer were far away, seeing the Buddha through fog). Note the loss in contrast and saturation.

6.5 Precomputed Radiance Transfer

The traditional environment map rendering techniques do not take shadows or interreflections into account. Precomputed radiance transport methods [Sloan et al. 2002] compute the visibility in an off-line manner, followed by interactive rendering. To add participating media, we only need to convolve the lighting (environment map) with our PSF and use the result as input to the existing precomputed radiance transfer techniques. To demonstrate this, we used the technique of [Wang et al. 2004; Liu et al. 2004], which handles non-diffuse objects under all-frequency environment illumination using a separable BRDF approximation.

We show the result of the Happy Buddha model rendered with the Ashikhmin-Shirley BRDF [2000] in figure 16. The left image is the standard result with no scattering. In the middle image, we show a view of the Buddha, where we include the effect of airlight from the environment on surface appearance, but there is no attenuation or scattering between the viewer and object itself (as if the observer were very close to the Buddha). We clearly see the foggy appearance of the background and the glows or airlight due to the light sources. On the face of the Buddha, we see a brightening of darker regions, along with a dimming and diffusing of specular highlights. A similar effect is seen on the base, where the dimming and diffusing of highlights reduces saturation and contrast. Finally, the shadows on the plane are blurred out, with a considerable softening and loss of detail. In the right image, there is also scattering or airlight between the object and the viewer (as if the observer were far away and seeing the Buddha through fog). This leads to a further loss of detail and contrast, so that the original glossy appearance of the object is essentially lost.

7 Conclusions and Future Work

We have presented a simple method to add the effects of participating media to interactive applications. Our approach can be easily implemented in programmable graphics hardware and leads to a number of new effects in the real-time domain, such as interactive rendering with glows around light sources, the effects of scattering on surface shading, environment maps, and precomputed light transport. The key insight is a new analytic model for integrating the light transport equations assuming single scattering, which can also be extended to predict the impact of scattering or airlight on

the inherent appearance of surfaces.

More broadly, this paper indicates the power of using explicit formulae to simulate difficult effects like volumetric scattering, speeding up such a process by many orders of magnitude. We do sacrifice some generality, considering only isotropic point light sources, single scattering, homogeneous media, and excluding most cast and volumetric shadowing, but believe this is a worthwhile tradeoff to enable a simple technique that achieves real-time rates.

Future work can follow many avenues. For instance, we can attempt to extend our theoretical model to consider non-isotropic light sources (like spotlights) and inhomogeneous media. Our preliminary work in this area indicates that some of these generalizations, while relatively simple in standard surface calculations, are rather non-trivial for volumetric media. However, we believe that many common cases can be addressed by extensions of the basic methods proposed in this paper. In general, we believe that analytic models of difficult to simulate volumetric phenomena are critical to achieving efficient renderings for real-time applications.

Acknowledgements: We thank Rui Wang, John Tran and David Luebke for the precomputed radiance transfer code, and Simon Premoze for the Monte Carlo simulation code. This research was supported by a Columbia University Presidential Fellowship, an ONR Grant #N00014-05-1-0188 (A Physical Approach to Underwater Optical Imaging), an NSF Grant #0305322 (Real-Time Rendering and Visualization of Complex Scenes), an NSF CAREER award #0446916 (Mathematical and Computational Fundamentals of Visual Appearance), and equipment donations from Intel and NVIDIA.

References

- ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic phong model. *Journal of Graphics Tools* 5, 2, 25–32.
- BASRI, R., AND JACOBS, D. W. 2003. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 2, 218–233.
- BIRI, V., MICHELIN, S., AND ARQUES, D. 2004. Real-time single scattering with shadows. In *In review* http://igm.univ-mlv.fr/~biri/indexCA_en.html.
- BLINN, J. 1982. Light reflection functions for simulation of clouds and dusty surfaces. In *SIGGRAPH* 82, 21–29.
- CHANDRASEKHAR, S. 1960. *Radiative Transfer*. Oxford Univ. Press.
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH* 98, 189–198.

- DOBASHI, Y., YAMAMOTO, T., AND NISHITA, T. 2002. Interactive rendering of atmospheric scattering effects using graphics hardware. In *Graphics Hardware Workshop 02*, 99–109.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *SIGGRAPH 93*, 165–174.
- HARRIS, M., AND LASTRA, A. 2001. Real-time cloud rendering. In *Eurographics 2001*, 76–84.
- HOFFMAN, N., AND PREETHAM, A. J. 2003. Real-time light-atmosphere interactions for outdoor scenes. *Graphics programming methods*, 337–352.
- JENSEN, H., MARSCHNER, S., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *SIGGRAPH 01*, 511–518.
- JENSEN, H. W. 2001. *Realistic Image Synthesis Using Photon Mapping*. AK Peters.
- KAJIYA, J., AND HERZEN, B. 1984. Ray tracing volume densities. In *SIGGRAPH 84*, 165–174.
- KOSCHMEIDER, H. 1924. Theorie der horizontalen sichtweite. *Beitr. Phys. freien Atm.*, 12.
- LIU, X., SLOAN, P.-P. J., SHUM, H.-Y., AND SNYDER, J. 2004. All-frequency precomputed radiance transfer for glossy objects. In *EuroGraphics Symposium on Rendering 04*, 337–344.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics (SIGGRAPH 03)* 22, 3, 759–769.
- MAX, N. L. 1986. Atmospheric illumination and shadows. In *SIGGRAPH 86*, 117–124.
- MAX, N. 1994. Efficient light propagation for multiple anisotropic volume scattering. In *Eurographics Rendering Workshop 94*, 87–104.
- NAKAMAE, E., KANEDA, K., OKAMOTO, T., AND NISHITA, T. 1990. A lighting model aiming at drive simulators. In *SIGGRAPH 90*, 395–404.
- NARASIMHAN, S., AND NAYAR, S. 2002. Vision and the atmosphere. *IJCV* 48, 3 (August), 233–254.
- NARASIMHAN, S., AND NAYAR, S. 2003. Shedding light on the weather. In *CVPR 03*, 665–672.
- NISHITA, T., AND NAKAMAE, E. 1987. A shading model for atmospheric scattering considering luminous intensity distribution of light sources. In *SIGGRAPH*, 303–310.
- PATTANAIK, S., AND MUDUR, S. 1993. Computation of global illumination in a participating medium by monte carlo simulation. *Journal of Visualization and Computer Animation* 4, 3, 133–152.
- PREETHAM, A. J., SHIRLEY, P., AND SMITS, B. 1999. A practical analytic model for daylight. In *SIGGRAPH*, 91–100.
- PREMOZE, S., ASHIKHMIN, M., TESSENDORF, J., RAMAMOORTHY, R., AND NAYAR, S. 2004. Practical rendering of multiple scattering effects in participating media. In *EuroGraphics Symposium on Rendering 04*, 363–374.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, 117–128.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2002. Frequency space environment map rendering. *ACM Transactions on Graphics (SIGGRAPH 02)* 21, 3, 517–526.
- RILEY, K., EBERT, D., KRAUS, M., TESSENDORF, J., AND HANSEN, C. 2004. Efficient rendering of atmospheric phenomena. In *EuroGraphics Symposium on Rendering 2004*, 375–386.
- RUSHMEIER, H., AND TORRANCE, K. 1987. The zonal method for calculating light intensities in the presence of a participating medium. In *SIGGRAPH 87*, 293–302.
- SAKAS, G. 1990. Fast rendering of arbitrary distributed volume densities. In *Eurographics 90*, 519–530.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (SIGGRAPH 02)* 21, 3, 527–536.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop 95*, 41–50.
- WANG, R., TRAN, J., AND LUEBKE, D. 2004. All-frequency relighting of non-diffuse objects using separable BRDF approximation. In *EuroGraphics Symposium on Rendering 2004*, 345–354.

Appendix A: Solution to Airlight Integral We start the derivation from equation 5.

$$L_a = \frac{\beta I_0}{4\pi} \int_0^{D_{vp}} \frac{e^{-\beta \sqrt{D_{sv}^2 + x^2} - 2xD_{sv} \cos \gamma}}{D_{sv}^2 + x^2 - 2xD_{sv} \cos \gamma} \cdot e^{-\beta x} dx \quad (28)$$




Figure 17: [Left] Plot of $PSF(\gamma)_{D_{sv}, \infty, \beta}$ normalized by $T_{sv} e^{-T_{sv}}$ for different optical thicknesses T_{sv} ranging from 0.1 to 3.1. After normalization, the PSF depends on γ and is largely independent of T_{sv} . This implies that we can factor it into a purely angular component and an amplitude component depending on T_{sv} . [Right] The average and standard deviation of the absolute error of the empirical PSF. While the error increases for smaller angles γ , it remains well below 0.05.

$$\begin{aligned} & \text{---} > \text{substitute } T_s = \beta D_s \text{ and } t = \beta x \\ & = \frac{\beta^2 I_0}{4\pi} \int_0^{T_{vp}} \frac{e^{-\sqrt{T_{sv}^2 + t^2} - 2tT_{sv} \cos \gamma}}{T_{sv}^2 + t^2 - 2tT_{sv} \cos \gamma} \cdot e^{-t} dt \end{aligned} \quad (29)$$

$$\begin{aligned} & \text{---} > \text{substitute } z = t - T_{sv} \cos \gamma \\ & = \frac{\beta^2 I_0 e^{-T_{sv} \cos \gamma}}{4\pi} \int_{-T_{sv} \cos \gamma}^{T_{vp} - T_{sv} \cos \gamma} \frac{e^{-\sqrt{z^2 + T_{sv}^2 \sin^2 \gamma}}}{z^2 + T_{sv}^2 \sin^2 \gamma} \cdot e^{-z} dz \end{aligned} \quad (30)$$

$$\begin{aligned} & \text{---} > \text{substitute } z = T_{sv} \sin \gamma \tan \eta \\ & = \frac{\beta^2 I_0 e^{-T_{sv} \cos \gamma}}{4\pi T_{sv} \sin \gamma} \int_{\gamma - \frac{\pi}{2}}^{\arctan \frac{T_{vp} - T_{sv} \cos \gamma}{T_{sv} \sin \gamma}} e^{-T_{sv} \sin \gamma \frac{1 + \sin \eta}{\cos \eta}} d\eta \end{aligned} \quad (31)$$

$$\begin{aligned} & \text{---} > \text{substitute } \eta = 2\xi - \frac{\pi}{2} \\ & = \frac{\beta^2 I_0 e^{-T_{sv} \cos \gamma}}{2\pi T_{sv} \sin \gamma} \int_{\gamma/2}^{\frac{\pi}{4} + \frac{1}{2} \arctan \frac{T_{vp} - T_{sv} \cos \gamma}{T_{sv} \sin \gamma}} \exp[-T_{sv} \sin \gamma \tan \xi] d\xi, \end{aligned} \quad (32)$$

from which we obtain equation 9.

Appendix B: Formula for Lambertian and Phong BRDFs

Here, we derive the expression for Lambertian and Phong BRDFs. We first consider the Lambertian BRDF, beginning with equation 15. Note that in the derivation below, γ' is given from trigonometry by $\gamma'(\theta_s, \omega_i) = \cos \theta_i \cos \theta_s + \sin \theta_i \sin \theta_s \cos \phi_i$.

$$\begin{aligned} L_{p,a} &= \int_{\Omega_{2\pi}} L_a(\gamma'(\theta_s, \omega_i), T_{sp}, \infty, \beta) f_r(\theta_i, \phi_i, \theta_v, \phi_v) \cos \theta_i d\omega_i \\ & \text{---} > \text{substitute equation 12 for } L_a \text{ and a constant } k_d \text{ for } f_r \\ & = \int_{\Omega_{2\pi}} A_0(T_{sp}, \gamma, \beta) \left[F(A_1(T_{sp}, \gamma'), \frac{\pi}{2}) - F(A_1(T_{sp}, \gamma'), \frac{\gamma'}{2}) \right] k_d \cos \theta_i d\omega_i \\ & \text{---} > \text{substitute equation 7 for } A_0 \text{ and take constants out of integration} \\ & = \frac{\beta^2 I_0 k_d}{2\pi T_{sp}} \int_{\Omega_{2\pi}} \frac{e^{-T_{sp} \cos \gamma'}}{\sin \gamma'} \left[F(A_1(T_{sp}, \gamma'), \frac{\pi}{2}) - F(A_1(T_{sp}, \gamma'), \frac{\gamma'}{2}) \right] \cos \theta_i d\omega_i \\ & = \frac{\beta^2 I_0 k_d}{2\pi T_{sp}} G_0(T_{sp}, \theta_s). \end{aligned} \quad (33)$$

For the Phong BRDF after reparameterization, instead of $k_d \cos \theta_i$, we will obtain $k_s \cos^n \theta_i$, where n is the Phong exponent. This can be handled exactly as above, simply replacing G_0 with G_n .

Appendix C: Empirical PSF factorization The empirical PSF factorization is inspired by the observation that after being normalized by $T_{sv} e^{-T_{sv}}$, the PSF becomes essentially independent of the medium physical parameters (optical thickness) and largely depends on angle γ as shown in figure 17 (left). This implies we can factor the PSF into a purely angular component and an amplitude component that depends on the medium parameters. We define the angular component $NPSF(\gamma)$ as the $PSF(\gamma)_{T_{sv}=1}$ normalized by $T_{sv} e^{-T_{sv}}$ and define the amplitude component as the normalization factor $T_{sv} e^{-T_{sv}}$. Then, the PSF can be expressed using these two terms as in equation 22. The absolute approximation error is plotted in figure 17 (right) for 11 different optical thickness ranging from 0.1 to 3.1.

Acquiring Scattering Properties of Participating Media by Dilution

Srinivasa G. Narasimhan¹, Mohit Gupta¹, Craig Donner², Ravi Ramamoorthi³, Shree K. Nayar³, Henrik Wann Jensen²

¹Carnegie Mellon University* ²University of California, San Diego ³Columbia University




Figure 1: (a) Photographs of our simple setup consisting of a glass tank and a bulb, filled with diluted participating media (from top, MERLOT, CHARDONNAY, YUENGLING beer and milk). The colors of the bulb and the glow around it illustrate the scattering and absorption properties in these media. At low concentrations, single scattering of light is dominant while multiple scattering of light is negligible. From a single HDR photograph, we robustly estimate all the scattering properties of the medium. Once these properties are estimated, a standard volumetric Monte Carlo technique can be used to create renderings at any concentration and with multiple scattering, as shown in (b) and (c). While the colors are only slightly visible in the diluted setting in (b), notice the bright colors of the liquids - deep red and golden-yellow wines, soft white milk, and orange-red beer - in their natural concentrations. Notice, also the differences in the caustics and the strong interreflections of milk onto other liquids.

Abstract

The visual world around us displays a rich set of volumetric effects due to participating media. The appearance of these media is governed by several physical properties such as particle densities, shapes and sizes, which must be input (directly or indirectly) to a rendering algorithm to generate realistic images. While there has been significant progress in developing rendering techniques (for instance, volumetric Monte Carlo methods and analytic approximations), there are very few methods that measure or estimate these properties for media that are of relevance to computer graphics. In this paper, we present a simple device and technique for robustly estimating the properties of a broad class of participating media that can be either (a) diluted in water such as juices, beverages, paints and cleaning supplies, or (b) dissolved in water such as powders and sugar/salt crystals, or (c) suspended in water such as

impurities. The key idea is to dilute the concentrations of the media so that single scattering effects dominate and multiple scattering becomes negligible, leading to a simple and robust estimation algorithm. Furthermore, unlike previous approaches that require complicated or separate measurement setups for different types or properties of media, our method and setup can be used to measure media with a complete range of absorption and scattering properties from a single HDR photograph. Once the parameters of the diluted medium are estimated, a volumetric Monte Carlo technique may be used to create renderings of any medium concentration and with multiple scattering. We have measured the scattering parameters of forty commonly found materials, that can be immediately used by the computer graphics community. We can also create realistic images of combinations or mixtures of the original measured materials, thus giving the user a wide flexibility in making realistic images of participating media.

*e-mail: srinivas@cs.cmu.edu

1 Introduction

Very often in our daily lives, we see participating media such as fluids (juices, beverages, milks) and underwater impurities (natural ocean, river and lake waters). The propagation of light through these media results in a broad range of effects, including softer appearance of milk, coloring of wines and juices, the transformation of appearances when liquids are mixed (coffee with milk, and cocktails), the brilliant caustics from glasses containing these liquids, and low visibility in underwater situations. These effects inherently depend on several physical properties of the media such as

scattering nature, sizes, shapes, and densities of particles [Hulst 1957; Chandrasekhar 1960]. Rendering these effects accurately is critical to achieving photo-realism in computer graphics.

In the past few years, there has been a considerable effort towards developing efficient and accurate rendering algorithms for participating media, based on Monte Carlo simulation and analytic approximations. All these algorithms and models contain parameters (scattering coefficient, absorption coefficient, phase function) that directly or indirectly represent the physical properties of the medium. In order to faithfully render the effects of any participating medium, the right parameters must be input. Given the progress in developing rendering algorithms, the quality of images is now often limited by the quality of these input parameters. Since there has so far been relatively little work in measuring or estimating scattering properties of media relevant to computer graphics, the parameters are currently often set in an ad-hoc manner.

This situation is similar in some ways to that of standard surface rendering. In that case, global illumination algorithms have progressed to the point of creating almost photo-realistic images, leaving the realism limited by the quality of the reflectance models, and leading to much recent effort on measuring BRDFs. [Marschner 1998; Dana et al. 1997; Matusik et al. 2003]. However, existing methods for directly measuring physical properties for media usually require very expensive equipment, such as the particle sizing apparatus used in colloidal chemistry [Finsy and Joosten 1991; Jaeger et al. 1991], resulting in little usable data for graphics.

Earlier efforts to estimate scattering properties from images of media have often yielded ill-conditioned and non-unique results, because of the difficulties of solving the inverse light transport problem. The reasoning for the ill-conditioning of the inverse problem is mainly due to multiple scattering, which blurs the incident light field and results in significant loss of information [McCormick 1981; McCormick 1985; Antyufeev 2000]. This is analogous to the ill-conditioning of BRDF estimation under complex illumination [Ramamoorthi and Hanrahan 2001]. In this paper, we take a completely different approach. The key idea is to estimate properties of media by acquiring the data in a state where multiple scattering effects are negligible. Instead, the data is acquired when single scattering (which does not degrade the incident light significantly) is the dominant effect. This is achieved by diluting the material to low concentrations.

We present a simple and inexpensive experimental setup, along with a robust and accurate technique for measuring the scattering properties of a broad class of participating media that can be either (a) diluted in water such as juices, beverages, paints and cleaning supplies, or (b) suspended in natural waters such as impurities and organisms, or even (c) dissolved in water such as powders and sugar or salt crystals. These media collectively have a wide range of scattering and absorption properties. We first derive a simple image formation model for single scattering of light in our setup. Through extensive simulations of both our model and ground truth (with multiple scattering), we then determine the space of concentrations and scattering properties of media for which single scattering is dominant. Within this regime of valid concentrations, we conduct simulations to demonstrate that our estimation technique uniquely solves the inverse single scattering light transport problem. Finally, we present a simple experimental procedure to determine the best concentration (dilution) for any material despite no prior knowledge of its scattering properties.

We have used our approach to create a dataset of scattering parameters for **forty commonly found materials**, which can be directly used for computer graphics rendering. Once the scattering parameters have been estimated, they can be used to render realistic images of arbitrary concentrations of the material with multiple scattering, using a standard physically based volumetric rendering algorithm. Figure 1 shows two renderings of a scene with four

Medium Property	Notation
Concentration or Volume Fraction	C
Scattering Coefficient (mm^{-1})	β
Absorption Coefficient (mm^{-1})	κ
Extinction Coefficient (mm^{-1})	$\sigma = \beta + \kappa$
Single Scattering Albedo	$\omega = \beta / \sigma$
Scattering Angle	θ
Henyey-Greenstein (H-G) Parameter	g
H-G Phase Function	$P(g, \theta) = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g \cos \theta)^{3/2}}$

Figure 2: The different scattering properties of a participating medium and their notations used in this paper. Light transport equations are usually written in terms of three parameters σ , β and g . We estimate these parameters for participating media based on single scattering.

liquids in their natural high density states and their diluted states. The scattering parameters of each material were computed using a single HDR photograph of our setup. Notice the bright saturated colors obtained despite the murky appearance of the diluted states. We can also create realistic images of mixtures of the original measured materials, thus giving the user a wide flexibility in creating realistic images of participating media.

2 Related Work

Figure 2 shows the most common properties of participating media including the scattering and absorption coefficients, and the phase function (angular scattering distribution represented by the Henyey-Greenstein (H-G) model [Henyey and Greenstein 1941]). The scattering and absorption coefficients are proportional to the concentration or volume fraction of the particulate medium. We will briefly review some of the representative works on the direct measurement and indirect estimation of these parameters.

Estimation based on analytic approximations to light transport. Surprisingly, little work has been done in computer graphics on the measurement of scattering properties of media. A recent work is that of [Jensen et al. 2001], on the diffusion model for subsurface scattering. They present a measurement of a number of translucent materials. However, the diffusion approximation assumes multiple scattering for optically dense media, so that only a limited amount of information on the scattering parameters can be estimated. For instance, this approximation is independent of the phase function of the medium, and therefore this important property cannot be estimated. Furthermore, the diffusion is a poor approximation when scattering is comparable to absorption [Prahl 1988]. The analytic multiple scattering model presented in [Narasimhan and Nayar 2003] has also been used to estimate properties of only purely scattering media (visibility and type of weather such as fog and mist). Our focus is somewhat different in considering fluids like juices or beverages, instead of subsurface scattering in translucent solids like marble and skin, or weather conditions such as fog. Nevertheless, our approach is valid for media with the entire range of absorbing and scattering properties, significantly extending the class of measurable media for graphics.

Most recently, Hawkins et. al., [2005] measure the extinction coefficient of optically thin smoke from the exponential attenuation of a laser beam in a tank. They also use a separate mirror setup to directly measure the phase function (see below). In contrast, our setup uses divergent beams from a simple bulb to include more light in the volume (than a single laser beam) for robust measurements, and requires only a single photograph to measure all scattering properties shown in Figure 2.

Numerical solution to inverse light transport: In cases where there are no analytic solutions to light transport, several works have taken a numerical approach to estimate scattering properties [McCormick 1996; Antyufeev 2000]. However, it is widely

known, that inverse problems in radiative transfer that take into account multiple scattering are ill-conditioned and require regularizing assumptions to obtain reliable estimates. See the reports and critiques by McCormick et al [1981; 1985]. Furthermore, the computational complexity of such inverse estimation techniques make it hard for measuring large sets of media for computer graphics or vision applications. Our focus here is on estimating scattering properties of media that can be measured in a state where multiple scattering is negligible.

The observation that single scattering is dominant for optically thin media has been made by [Hawkins et al. 2005; Sun et al. 2005]. We exploit this observation and apply the single scattering model for the first time to a large class of materials which exhibit significant multiple scattering in their natural states of existence. We also determine the exact range of optical thicknesses for which single scattering is dominant for media with arbitrary scattering properties, and propose an experimental procedure to ensure the dominance of single scattering in real data.

Goniophotometry is often used to directly measure the phase function. Here, several detectors measure radiance in different directions after being scattered by a very small volume of the medium. [Fuchs and Jaffe 2002] use thin laser light sheet microscopy for detecting and localizing microorganisms in ocean waters. [Boss and Pegau 2001; Oishi 1990] investigate the relationship of light scattering at a single angle and the extinction coefficient using specialized receivers and transmitters. However, all these techniques assume that there is no attenuation of light through the sample and require expensive devices with precise alignment of detectors and transmitters. In contrast, our setup is extremely simple (consisting of a glass tank and an off the shelf bulb), and our technique robustly estimates all properties from only a single photograph, thus making it inexpensive and easy to measure a large number of participating media.

3 Single Scattering in Dilute Media

Our approach is to measure media in a state where single scattering is dominant and multiple scattering is negligible. This is achieved by diluting the otherwise optically thick media, such as fluids, in water. The process of dilution does not usually corrupt the inherent scattering properties of media¹ since the scattering and absorption of pure water itself is negligible for very small distances (less than 50 cm) [Sullivan 1963]. We begin by presenting our acquisition setup and an image formation model for single scattered light transport within the measurement volume. We will then present extensive simulations of this model and compare with traditional Monte-Carlo approaches that include multiple scattering, to derive a valid space of scattering parameters over which single scattering is dominant. Based on this simulation, we design a simple experimental procedure to choose the best concentration for any particular medium. Later, we will describe our algorithm to estimate the scattering parameters using our image formation model.

3.1 Acquisition Setup

The measurement apparatus, shown in Figure 3, consists of a $25 \times 30 \times 30 \text{ cm}^3$ tank that is filled with the diluted scattering medium. The depth of the tank is large enough to ensure the scattering angles are adequately covered (0 to 175 degrees). The volume of the tank is designed to be large enough to dilute concentrated media such as milk. Two sides of the tank are constructed using anti-reflection glass and the other sides using diffuse black coated acrylic. A small frosted (diffuse) glass bulb fixed to a side

¹When crystals are dissolved in water, they may exhibit different scattering properties due to ionization.




Figure 3: Two views of the apparatus used to measure scattering properties of water-soluble media. A glass tank with rectangular cross-section is fitted with a small light bulb. The glass is anti-reflection coated. Different volumes of participating media are diluted with water in the tank, to simulate different concentrations. A camera views the front face of the tank at normal incidence to avoid refractions at the medium-glass-air boundaries.




Figure 4: A volume filled with a homogeneous participating medium and illuminated by an isotropic point light source. A camera views the front face of the volume at normal incidence. The path of one single-scattered ray as it travels from the source to the camera is shown. This ray is first attenuated in intensity over a distance d , is then scattered at an angle $\pi - \theta$, and finally, is attenuated again over a distance z , before reaching the camera. The irradiances due to all the rays that scatter into a viewing direction must be integrated to obtain the final camera irradiance.

of the tank illuminates the medium. A Canon EOS-20D 12-bit 3504x2336 pixel digital camera with a zoom lens is placed five meters away from the tank and observes a face of the tank at normal incidence. The field of view occupied by the tank in the image is three degrees and is therefore approximately orthographic. Orthographic projection avoids the need for modeling refractions of light rays at the medium-glass-air interfaces. In all our experiments, about 25 different exposures (1/500s to 10s) were used to acquire HDR images.

3.2 Image Formation Model

Although the basic principles of single scattering are well known, the exact nature of the image formation model depends on the geometry of the volume and the locations of the source and the camera. Figure 4 illustrates the illumination and measurement geometry based on our acquisition setup. For simplicity, we will assume that the medium is illuminated by an isotropic point light source (later we extend the analysis to area sources) of intensity I_0 that is located at the coordinates $(0, B, H)$.

Consider the path of one single-scattered light ray (thick ray in Figure 4) in the medium as it travels from the source to the camera. This ray is first exponentially attenuated in intensity for a distance d . At location $U(x, y, z)$, depending on the phase function P , a fraction of the light intensity is scattered at an angle $\pi - \theta$. Finally, the ray is attenuated again for a distance z , before it reaches the camera. Mathematically, the irradiance at the camera produced by

this ray is written as [Sun et al. 2005],

$$E(x, y, z) = \frac{I_0}{d^2} \cdot e^{-\sigma d} \cdot \beta P(g, \pi - \theta) \cdot e^{-\sigma z}.$$

$$d = \sqrt{x^2 + (y - H)^2 + (z - B)^2}, \quad \cos \theta = (z - B) / d \quad (1)$$

Here, $P(g, \pi - \theta)$ is the Henyey-Greenstein (H-G) phase function, and β and σ are the scattering and extinction coefficients (Figure 2). Then, the total irradiance E at pixel (x, y) in the camera is obtained by integrating intensities due to all rays that are scattered at various angles along the pixel's line of sight (Z-direction),

$$E(x, y) = \int_0^{2B} E(x, y, z) dz$$

$$= \beta \int_0^{2B} \frac{I_0 e^{-\sigma(z + \sqrt{x^2 + (y-H)^2 + (z-B)^2})}}{x^2 + (y-H)^2 + (z-B)^2} P(g, \pi - \theta) dz. \quad (2)$$

The above equation relates the camera irradiances as a function of the three medium parameters, σ , β and g . Although obtaining an analytic (closed-form) solution to the above integral is hard [Sun et al. 2005], it is straightforward to evaluate it numerically.

3.3 Space of valid medium parameters

Different materials have their own natural densities and scattering properties, which are all unknown before experimentation. So, how do we know if single scattering is dominant at a particular concentration for a given material? Note that the scattering β , absorption κ and extinction σ , coefficients are proportional to the concentration (fraction of volume diluted in water) of the medium. Thus, we performed exhaustive simulations to derive the complete space of parameters for which the above image formation model is accurate². For ground truth, we simulated the irradiances obtained using multiple scattering for the same set of parameter values, using a standard volumetric Monte Carlo technique. Figure 5 shows a plot of the differences between energies captured by the single scattering and multiple scattering simulations for a set of parameter values. From the RMS errors in the plot, we can define the upper bounds on the parameters κ and $\sigma = \beta + \kappa$ as those for which the energy differences between our model and the ground truth are less than five percent. For example, the valid domain where single scattering is dominant, is approximately $\sigma < 0.04$ for $\kappa < 0.004$.

3.4 How to choose the best concentration?

Based on the simulations, we present an experimental method to determine the best concentration for our measurements. Figure 6 shows images acquired of different concentrations of milk and MERLOT. Which among these images should we use to measure the scattering properties? Several heuristics may be used to decide on a particular concentration. For instance, the extent of blurring of the light source provides us a good clue to determine whether multiple scattering is significant (rightmost image in Figure 6). A better heuristic is to compute an approximation to the extinction coefficient σ from the attenuated brightness of the light source. Under single scattering, the radiance in the direction of the source (distance d) can be approximated using exponential attenuation as:

$$E(0) \approx \left(\frac{I_0}{d^2} \right) e^{-\hat{\sigma} d}, \quad (3)$$

²This extends the simulations in [Sun et al. 2005], where a small part of the possible parameter space (pure isotropic scattering) was considered.




Figure 5: Plot showing the differences between irradiances obtained by simulating single scattering and multiple scattering (ground truth) models, for a large space of parameter values σ and $\kappa = \sigma - \beta$. An upper bound on the differences of, say, 5%, can be used to define the range of parameters for which single scattering is a valid approximation. From the plot, the valid range is approximately $\sigma < 0.04$ for $\kappa < 0.004$.

where $\hat{\sigma}$ is an estimate of the extinction coefficient σ . In the absence of multiple scattering, this estimate is closer to the true value of σ (and varies linearly with concentration), whereas, in the presence of multiple scattering, this estimate is called *diffuse or reduced attenuation coefficient* [Ishimaru 1978] and is usually much lesser than σ . Thus, we can determine whether the concentration can be used for measurement by observing the plot (Figure 7 of $\hat{\sigma}$ versus the volume fraction of the medium diluted with water). Figure 7 shows that after a certain amount of milk is added to water, the $\hat{\sigma}$ no longer remains linear with concentration (dashed line), and must not be used for measurements. For a purely absorbing liquid like wine (MERLOT), the plot is completely linear and any image that has the best signal-to-noise ratio may be used. Similarly, the plot shows that coke scatters, albeit weakly, and ESPRESSO coffee scatters light strongly. We use this simple procedure to try several concentrations and observe where the linearity in the plot fails to determine the best concentration. As a further test, we check if the estimated parameters from this concentration lie within the valid space of parameters simulated above.




Figure 6: Images illustrating different degrees of scattering and absorption. [Top row] Images of milk at various concentrations. Since milk is a highly scattering liquid, we observe an increase in blurring with increasing concentration. [Bottom Row] Images of red wine at various concentrations. Red wine is a highly absorbing liquid, showing only a saturation of the bulb color with increasing concentration, and no blurring. The highlighted images are chosen for estimating the parameters.




Figure 7: Plot of extinction coefficient estimate $\hat{\sigma}$ as a function of the volume of the media diluted in water in the measurement apparatus. The plots are linear when multiple scattering is negligible and single scattering is dominant. As the concentrations of media (and hence multiple scattering) increase, the estimated $\hat{\sigma}$ is less than the true extinction coefficient σ . For a highly scattering medium such as milk, the linearity fails at very low concentrations, while for an absorbing medium such as MERLOT, the linearity is always preserved.

4 Estimating Medium Properties based on Single Scattering

In this section, we present a non-linear minimization algorithm to estimate the properties of the medium (σ , β and g), from the measured image irradiances $E(x, y)$ (see Equation (2)). We then demonstrate the accuracy of the algorithm through extensive simulations.

4.1 Formulating the Error Function

The error at each pixel is written as the difference between the measured irradiance $E(x, y)$ and the irradiance predicted by the model in equation 2,

$$\mathcal{F}(x, y) = E(x, y) - RHS(x, y). \quad (4)$$

Here $RHS(x, y)$ is the numerically evaluated right hand side integral in the model of equation 2. Then, the parameters σ , β and g can be estimated by computing the global minimum of the sum of squares of the errors of all the pixels, as,

$$\min_{\beta, \sigma, g} \sum_y \sum_x \mathcal{F}^2(x, y). \quad (5)$$

The above function essentially requires a 3-parameter search. However, note that the parameter β is a global scale factor. Thus, we can eliminate β by defining a normalized error function as,

$$\mathcal{F}_{norm}(x, y) = \frac{E(x, y)}{\max_{x, y} E(x, y)} - \frac{RHS(x, y)}{\max_{x, y} RHS(x, y)}. \quad (6)$$

Now, instead of requiring a 3-parameter search, the above problem can be reduced to a 2-parameter search that minimizes the normalized objective function to estimate σ and g :

$$\min_{\sigma, g} \sum_y \sum_x \mathcal{F}_{norm}^2(x, y). \quad (7)$$

Then, the scale factor β can be recovered using the original function \mathcal{F} . To compute the global minimum, we use Nelder-Mead search implemented by the MatlabTM function "fminsearch".




Figure 8: Plot showing the errors in reconstruction of the single scattering parameters σ and $q = |g|$, where $-1 < g < 1$, compared to ground truth values. The low errors indicate the accuracy of our estimation technique. The maximum of the errors for positive or negative g is shown.

4.2 Estimation Accuracy using Simulations

Fortunately, since the space of the possible parameters is small (see Section 3.3), exhaustive simulation of the above algorithm is possible. We only show the correctness of the estimated parameters σ and g , using Equation (7). The estimation of the scale factor β then follows trivially. Gaussian noise of unit standard deviation was added in all our simulations. The non-linear search was initialized randomly for both the parameters σ and g . The plot in Figure 8 shows the error in the estimated parameters as compared to ground truth values. In all the cases, the estimation errors were less than 0.0001%, and the number of iterations required for convergence was less than 100. Since the numerical evaluation of the integral is very fast, the time for convergence is usually of the order of a few minutes. This demonstrates that the inverse estimation is fast and results in unique and correct parameters.

4.3 Implementation Issues

We present two issues that need careful implementation for our algorithm to be successful on real images.

Calibrating the area source: Our method does not rely on isotropic point sources but requires only a calibrated divergent source to take advantage of the different phase angles measured in the same view and hence, any off-the-shelf bulb suffices. For our real setup, we have implemented a spherical diffuse area source. To compute the irradiance at any point P within the tank, we sample (using roughly 10×10 samples) the hemisphere of the bulb that is visible to that point P . The non-uniform directional intensities and intensity fall-off were calibrated carefully by using a light meter at discrete 3D locations within the tank. The camera also measures a pure water image (without any scattering or absorption) to give the image irradiance of each source element (sample). This irradiance along with the fall-off value and the pixel solid angle is used to determine the intensity without scattering.

Instabilities in the H-G phase function for highly absorbing media: The H-G phase function was designed for scattering media and is not defined for purely absorbing media. However, for highly absorbing media, the scattering coefficient β is very low and the average cosine $g \approx 1$ since rays only pass straight through, much like highly forward scattering media. Even though this was not a problem in simulations, the instability for $g > 0.95$ can be high in real experiments. For this special case, we simply use a truncated legendre polynomial expansion of the H-G phase function as $P(g, \theta) = \sum_i (2i + 1) g^i L_i(\theta)$, and truncate to less than 100 terms. As an undesirable byproduct the fits may show some "ringing" at the tail of the phase function. However, this truncated function still fits higher brightness well and thus does not affect appearance strongly. Despite this instability, the H-G phase function is flexible enough to model the scattering behavior of all our materials.




Figure 9: Captured photographs of a variety of water-soluble media illustrating different degrees of scattering and absorption. For highly scattering media such as milk, chocolate milk and espresso, we observe a significant blur around the bulb. For highly absorbing media such as grape juice, there is very little scattering. All the images have wide dynamic range of intensities and hence, we have tone-mapped them for illustration. Please see supplementary material for more images.

5 Actual Measurements and Validation

Using our approach, we have measured the scattering properties of a broad class of **forty** commonly found participating media that can be either (a) diluted in water such as juices (for example, apple, strawberry, orange), beverages (for example, coffee, soft drinks, milks, wines, beers), cleaning supplies (detergents), or (b) suspended in natural waters such as impurities and organisms, or even (c) dissolved in water such as powders and sugar, salt crystals. In addition to liquids available at the usual supermarkets, we have also collected four samples from different locations and depths in the Pacific ocean. We then present detailed validation by showing that our parameters extrapolate correctly to higher concentrations as well, where multiple scattering is prominent.

A subset of nine photographs of the diluted set of liquids contained in the glass tank is shown in Figure 9, similar to the four in Figure 1. Together, these include representative types of media such as highly scattering, highly absorbing and moderate levels of absorption and scattering. The images show a high dynamic range of brightness and are enhanced to show the scattering effects. The set of scattering parameters for all the media is shown in Table 1. The extinction (σ) and scattering (β) coefficients are given for each of the three color channels, red, green and blue. The phase function parameter g is also shown for the three color channels. Note that all the extinction and scattering coefficients are less than 0.04 in accordance with our simulations in Section 3.3. Also, as expected, in all cases, the scattering coefficient does not increase with wavelength.

5.1 Fits to Measured Brightness Profiles

We demonstrate the accuracy of our technique by reconstructing the photographs using the estimated parameters. Although we considered the brightness at all pixels in the captured photographs, for illustration purposes we show only the profile of intensity values in the direction that is radially outward from the source. Figure 10

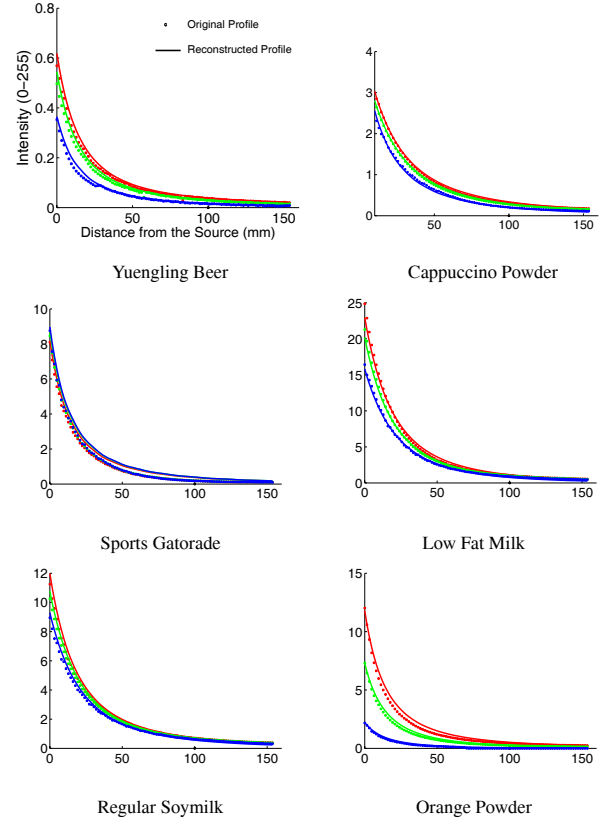


Figure 10: Fits obtained using the estimated parameters as compared against the corresponding measured brightness profiles in the captured photographs. The brightness profile is measured radially outward from the source in the image. The red, green and blue plots correspond to the three color channels of the camera. The match between the estimated and measured data demonstrates the accuracy of the estimation technique. The fits for six (out of 40) representative materials with varying degrees of absorption and scattering are shown. Please see the supplementary material for more plots.

shows the good fits obtained using the estimated parameters compared against the measured profiles for a subset of six materials of varying degrees of scattering and absorption properties (please review supplementary document for plots of other materials). When there is no scattering (pure absorption), fitting a scattering model can induce some “ringing” effect in the dark tail end of the profile. We can detect this special case and use the attenuation model to compute the absorption coefficient ($\kappa = \sigma$).

5.2 Extrapolation to higher concentrations

The extinction and scattering coefficients are proportional to the concentration of the medium. Thus, if β_1 and σ_1 are estimated at concentration c_1 , then the coefficients β_2 and σ_2 at another concentration c_2 can be extrapolated using:

$$\beta_2 = \beta_1 \left(\frac{c_2}{c_1} \right), \quad \sigma_2 = \sigma_1 \left(\frac{c_2}{c_1} \right). \quad (8)$$

Note, however, that g is independent of the medium concentration. While we estimate the parameters from lower concentrations, it is important to ensure that the parameters can be scaled to any concentration (where multiple scattering cannot be ignored) to produce accurate scattering effects. We show an example validation using fits obtained in comparison to the measured brightness profiles of chocolate milk at various concentrations. Figure 11 shows the fits

Material Name	Volume	Extinction Coefficient (σ) ($\times 10^{-2} \text{ mm}^{-1}$)			Scattering Coefficient (β) ($\times 10^{-2} \text{ mm}^{-1}$)			Average Cosine (g)			% RMS Error
		R	G	B	R	G	B	R	G	B	
Milk (lowfat)	16ml	0.9126	1.0748	1.2500	0.9124	1.0744	1.2492	0.932	0.902	0.859	0.95
Milk (reduced)	18ml	1.0750	1.2213	1.3941	1.0748	1.2209	1.3931	0.819	0.797	0.746	1.27
Milk (regular)	15ml	1.1874	1.3296	1.4602	1.1873	1.3293	1.4589	0.750	0.714	0.681	1.56
Coffee (espresso)	8ml	0.4376	0.5115	0.6048	0.2707	0.2828	0.2970	0.907	0.896	0.880	1.90
Coffee (mint mocha)	6ml	0.1900	0.2600	0.3500	0.0916	0.1081	0.1460	0.910	0.907	0.914	2.00
Soy Milk (lowfat)	16ml	0.1419	0.1625	0.2740	0.1418	0.1620	0.2715	0.850	0.853	0.842	1.75
Soymilk (regular)	12ml	0.2434	0.2719	0.4597	0.2433	0.2714	0.4563	0.873	0.858	0.832	1.68
Chocolate Milk (lowfat)	10ml	0.4282	0.5014	0.5791	0.4277	0.4998	0.5723	0.934	0.927	0.916	1.04
Chocolate Milk (regular)	16ml	0.7359	0.9172	1.0688	0.7352	0.9142	1.0588	0.862	0.838	0.806	2.19
Soda (coke)	1600ml	0.7143	1.1688	1.7169	0.0177	0.0208	0.0000	0.965	0.972	—	4.86
Soda (pepsi)	1600ml	0.6433	0.9990	1.4420	0.0058	0.0141	0.0000	0.926	0.979	—	2.92
Soda (sprite)	15000ml	0.1299	0.1283	0.1395	0.0069	0.0089	0.0089	0.943	0.953	0.952	3.22
Sports Gatorade	1500ml	0.4009	0.4185	0.4324	0.2392	0.2927	0.3745	0.933	0.933	0.935	3.42
Wine (chardonnay)	3300ml	0.1577	0.1748	0.3512	0.0030	0.0047	0.0069	0.914	0.958	0.975	5.10
Wine (white zinfandel)	3300ml	0.1763	0.2370	0.2913	0.0031	0.0048	0.0066	0.919	0.943	0.972	5.49
Wine (merlot)	1500ml	0.7639	1.6429	1.9196	0.0053	0.0000	0.0000	0.974	—	—	4.56
Beer (budweiser)	2900ml	0.1486	0.3210	0.7360	0.0037	0.0069	0.0074	0.917	0.956	0.982	5.61
Beer (coorslight)	1000ml	0.0295	0.0663	0.1521	0.0027	0.0055	0.0000	0.918	0.966	—	4.89
Beer (yuengling)	2900ml	0.1535	0.3322	0.7452	0.0495	0.0521	0.0597	0.969	0.969	0.975	4.48
Detergent (Clorox)	1200ml	0.1600	0.2500	0.3300	0.1425	0.1723	0.1928	0.912	0.905	0.892	1.99
Detergent (Era)	2300ml	0.7987	0.5746	0.2849	0.0553	0.0586	0.0906	0.949	0.950	0.971	4.17
Apple Juice	1800ml	0.1215	0.2101	0.4407	0.0201	0.0243	0.0323	0.947	0.949	0.945	4.92
Cranberry Juice	1500ml	0.2700	0.6300	0.8300	0.0128	0.0155	0.0196	0.947	0.951	0.974	4.60
Grape Juice	1200ml	0.5500	1.2500	1.5300	0.0072	0.0000	0.0000	0.961	—	—	5.19
Ruby Grapefruit Juice	240ml	0.2513	0.3517	0.4305	0.1617	0.1606	0.1669	0.929	0.929	0.931	2.68
White Grapefruit Juice	160ml	0.3609	0.3800	0.5632	0.3513	0.3669	0.5237	0.548	0.545	0.565	2.84
Shampoo (balancing)	300ml	0.0288	0.0710	0.0952	0.0104	0.0114	0.0147	0.910	0.905	0.920	4.86
Shampoo (strawberry)	300ml	0.0217	0.0788	0.1022	0.0028	0.0032	0.0033	0.927	0.935	0.994	2.47
Head & Shoulders	240ml	0.3674	0.4527	0.5211	0.2791	0.2890	0.3086	0.911	0.896	0.884	1.91
Lemon Tea Powder	5tsp	0.3400	0.5800	0.8800	0.0798	0.0898	0.1073	0.946	0.946	0.949	2.83
Orange Powder	4tbsp	0.3377	0.5573	1.0122	0.1928	0.2132	0.2259	0.919	0.918	0.922	2.25
Pink Lemonade Powder	5tbsp	0.2400	0.3700	0.4500	0.1235	0.1334	0.1305	0.902	0.902	0.904	1.02
Cappuccino Powder	0.25tsp	0.2574	0.3536	0.4840	0.0654	0.0882	0.1568	0.849	0.843	0.926	0.67
Salt Powder	1.75cup	0.7600	0.8685	0.9363	0.2485	0.2822	0.3216	0.802	0.793	0.821	1.34
Sugar Powder	5cup	0.0795	0.1759	0.2780	0.0145	0.0162	0.0202	0.921	0.919	0.931	1.80
Suisse Mocha Powder	0.5tsp	0.5098	0.6476	0.7944	0.3223	0.3583	0.4148	0.907	0.894	0.888	1.33
Mission Bay Surface Water (1-2 hours)		3.3623	3.2929	3.2193	0.2415	0.2762	0.3256	0.842	0.865	0.912	2.48
Pacific Ocean Surface Water (1 hour)		3.3645	3.3158	3.2428	0.1800	0.1834	0.2281	0.902	0.825	0.914	2.57
Mission Bay 10ft deep Water (30 min)		3.4063	3.3410	3.2810	0.0990	0.1274	0.1875	0.726	0.820	0.921	5.10
Mission Bay 10ft deep Water (8 hours)		3.3997	3.3457	3.2928	0.1018	0.1033	0.1611	0.929	0.910	0.945	5.13

Table 1: Scattering properties for 40 different water-soluble materials estimated using our technique. The second column lists the volumes V of the materials dissolved in $23 - V$ litres of water to achieve the desired levels of dilution where single scattering is dominant. These parameters can be proportionately scaled to any other volume V_n , using a scale factor of V_n/V . The percentage RMS errors (obtained over all pixels) quantify the accuracy of fits achieved with the estimated parameters to the measured intensity profiles. Errors for all the highly scattering media are less than 3%. For low-scattering materials, the total intensity of profiles is relatively low, thus making the estimation more sensitive to noise. Even for such low-scattering media, the errors are less than 5 – 6%. The last four rows are the parameters for various ocean water samples at their original concentrations. The **time elapsed** between the collection of samples and the image acquisition is listed in the parentheses. Since the suspended particles in ocean water settle down with time, we observe a small decrease in scattering coefficients in the sample for which 8 hours had been elapsed as compared to the one which was imaged just 30 minutes after collection. Note that all the extinction and scattering coefficients are less than 0.04 in accordance with our simulations in Section 3.3. As expected, the scattering coefficients do not decrease with wavelength. The scattering albedos (ratio of scattering coefficients to the extinction coefficients) is much higher for the scattering media (milk, coffee, orange powder) as compared to the absorbing ones (coke, wine). For materials that have $\beta = 0$, the phase function parameter g is undefined. As seen from the values of g which are closer to 1, several media are predominantly forward scattering. The parameters for the milks match those in [Jensen et al. 2001] up to a scale factor (due to the different fat contents in the milks used), providing further support for our estimation.

in this validation experiment. First, we estimate the parameters from the photograph of only 8ml of chocolate milk diluted in water, where single scattering is dominant. In (a), we show the fit obtained compared against the measured intensity profile. However, for higher concentrations of 50ml, 100ml and 150ml, multiple scattering cannot be ignored. For these cases, we scaled the coefficients (σ and β) by factors of $\{50/8, 100/8, 150/8\}$ (see Equation

8) and use them in a standard volumetric Monte Carlo renderer that includes multiple scattering. The plots in (b) - (d) demonstrate the strong fits obtained. This demonstrates that our parameters are robust enough to be extrapolated to higher concentrations. In fact, we will show renderings of most of the liquids at their natural concentrations (Section 6) despite measuring the parameters at significantly dilute states.

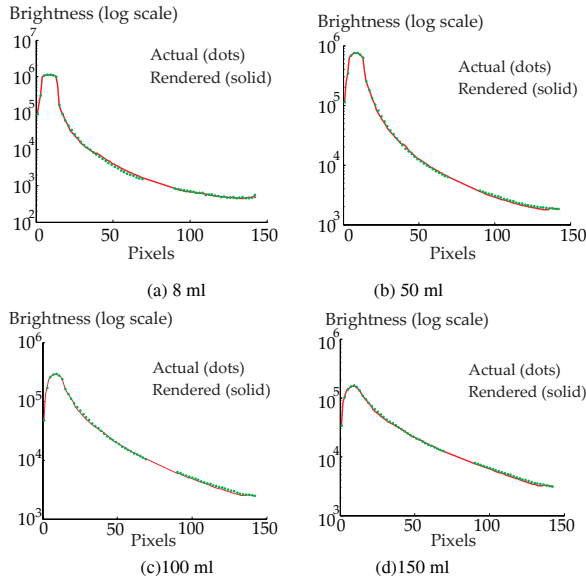


Figure 11: Extrapolation of parameters to higher concentrations with multiple scattering. (a) 8 ml of chocolate milk is diluted in water and the parameters are estimated using the measured brightness profile. (b) - (d) The parameters estimated in (a) are scaled to higher concentrations (50ml, 100ml and 150ml) where multiple scattering cannot be ignored. Plots show a good fit between the brightness profile obtained by extrapolating our estimated parameters with a Monte Carlo renderer, and the ground truth measurements. The fits are shown in logarithmic scale.



Figure 12: Rendered scenes of liquids in a cognac glass under complex lighting. The KITCHEN environment map [Debevec 1998] was used for the lighting. The natural colors, shading and caustics indicate the high accuracy of our parameters.




Figure 13: Rendered scenes of liquids and powders in a cognac glass illuminated with a single directional white light source. The bright caustics show the colors transmitted through the media.

6 Example Volumetric Renderings

The scattering properties estimated in this work can be input to any volumetric rendering algorithm to create visual effects of participating media. Here, we chose brute-force volumetric Monte-Carlo path tracing since it can be used to render arbitrary materials³. We use photon mapping for rendering caustics. For display purposes, we have applied a tone-mapping operator [Ward-Larson et al. 1997] to the renderings. Indices of refraction (IOR) of these media are also important for rendering. In initial experiments, we found the IOR to be between 1.33 (water) and 1.42 (milk) and varying linearly with concentrations, by using location of total internal reflection from the top of the water surface in the tank. In current renderings, we have simply used an IOR proportionate to the medium concentrations between 1.33 and 1.42, since this does not alter the visual appearance of the liquid drastically. We wish to perform thorough experiments in the future.

Figure 12 shows a mosaic of images of liquids rendered in their natural concentrations, partially filled in a cognac glass and illuminated by the “Kitchen Environment Map” [Debevec 1998]. These include two different types of wine (deep red MERLOT and golden-yellow CHARDONNAY), dark brown coffee ESPRESSO, and the golden-orange YUENGLING beer. Notice the color differences between MERLOT (no scattering) and ESPRESSO (moderate scattering) even though both of them are dark liquids. Observe that while beer and CHARDONNAY are very clear liquids, coffee is noticeably more opaque. Similarly, Figure 13 shows a mosaic of predominantly bright colored liquids such as the deep

³Under-sampling of path-traces can cause speckle noise seen in the renderings, and is not an artifact of our estimation.



Figure 14: *Effect of changing concentrations of a highly absorbing (MERLOT) and a highly scattering (milk) liquid. In the case of wine, notice that while the color gradually becomes deep red, the liquid remains clear; due to the lack of scattering. In the case of milk, however, we see a quick transition from a murky appearance to a soft white appearance, due to the high scattering albedo of milk.*

blue ERA detergent, the reddish strawberry shampoo, and powders dissolved in water such as the “pinkish” strawberry lemonade and orange powders. These images are illuminated only by a strong directional source to illustrate the bright caustics whose colorings are primarily due to absorption. We also present different types of novel visual effects obtained by changing or blending the parameters of different media to create realistic images of dilutions and mixtures of the original measured materials.

Effect of changing concentrations: Figure 14 illustrates the effects of changing concentrations of media in water. The top row shows a transition from pure water to MERLOT, obtained by scaling parameters of wine as in Equation 8. Notice the changes in caustics and the gradual deepening of the red color of the liquid. Note that as the transition occurs, the liquid remains clear even though the color changes; this is due to the pure absorbing nature of wine, as depicted by our parameters. The bottom row shows the effect of changing milk concentration in water. Since milk is a highly scattering medium, as expected, the appearance quickly changes from murky whitish water to soft and thick white milk. This is because the scattering albedo β/σ is high and the phase function parameter g is such that a significant amount of light diffuses into different directions.

Blending parameters for mixtures of media: For example, what are the properties of a mixture of ESPRESSO and milk, or otherwise known as *light coffee*? Consider a medium containing a mixture of two types of media, A and B . The properties of the individual media are denoted with the subscripts A and B . The scattering coefficient of the mixture is obtained by a weighted average,

$$\beta_{\text{mix}} = \frac{V_A \beta_A + V_B \beta_B}{V_A + V_B}. \quad (9)$$

The absorption and extinction coefficients are similarly defined.

Unlike above where we just changed the scattering and absorption coefficients, here a new phase function parameter must be defined for the mixture as the weighted average [Key 2005],

$$g_{\text{mix}} = \frac{g_A \beta_A + g_B \beta_B}{\beta_{\text{mix}}}. \quad (10)$$

These equations can be used to render mixtures of participating media or morph from one medium into another. Figure 15 shows mixing of different proportions of milk and wine. The second example shows a more common mixing of milk and coffee. Such mixing between materials, for the first time, gives a user the flexibility to create novel renderings of participating media.

7 Conclusion

Rendering the rich visual effects of participating media, like fluids or underwater impurities, requires precise measurements of their scattering properties. In this paper, we have developed a simple device and method for accurately estimating the scattering properties of a variety of media that can be diluted in water. Our approach only requires a single high dynamic range photograph. By diluting the medium, we work in the single scattering regime, where the inverse light transport problem is well conditioned—however, we can later render at arbitrary concentrations and even mix materials. We have presented a database of scattering parameters for 40 commonly found materials. This database is the first of its kind, and enables computer graphics practitioners to accurately render a wide variety of participating media, rather than having to set parameters in an ad-hoc fashion. In the future, we would like to improve this work by investigating different phase functions and measuring indices of refraction more accurately.



Figure 15: Mixing two liquids - milk and coffee (top) and milk and wine (bottom), in different proportions. The wine-milk combination produces a soft pink appearance while the ESPRESSO-milk combination produces soft but brown appearance. (Minor noise due to Monte-Carlo under-sampling.)

Acknowledgements: This work was supported by NSF Grants 0541307, 0305322, 0446916 and 0305399, and an ONR grant N00014-05-1-0188. Ramamoorthi, Jensen and Donner were supported also by Sloan Fellowships. The authors thank Terrence Boulton for ocean water samples, Chao-Kuo Lin for help with experiments, and Simon Premoze for early discussions.

References

- ANTYUFEEV, S. 2000. *Monte Carlo Method for Solving Inverse Problems of Radiative Transfer*. Inverse and Ill-Posed Problems Series, VSP Publishers.
- BOSS, E., AND PEGAU, W. S. 2001. Relationship of light scattering at an angle in the backward direction to the backscattering coefficient. *Applied Optics* 40 (30), 5503–5507.
- CHANDRASEKHAR, S. 1960. *Radiative Transfer*. Oxford University Press.
- DANA, K., NAYAR, S., VAN GINNEKEN, B., AND KOENDERINK, J. 1997. Reflectance and texture of real-world surfaces. In *Proc CVPR*, 151–157.
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Proc. SIGGRAPH* 98, 189–198.
- FINSY, E. G., AND JOOSTEN, J. 1991. Maximum entropy inversion of static light scattering data for the particle size distribution by number and volume. In *Advances in measurements and control of colloidal processes*. Butterworth-Heinemann, Ch. 30.
- FUCHS, E., AND JAFFE, J. S. 2002. Thin laser light sheet microscope for microbial oceanography. *OPTICS EXPRESS* 10 (2), 145–154.
- HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. Acquisition of time-varying participating media. *ACM Trans. on Graphics (SIGGRAPH)* 24, 3, 812–815.
- HENYEY, L., AND GREENSTEIN, J. 1941. Diffuse radiation in the galaxy. vol. 93, 70–83.
- HULST, V. D. 1957. *Light Scattering by small Particles*. John Wiley and Sons.
- ISHIMARU, A. 1978. *Wave Propagation and Scattering in Random Media. Volume 1: Single Scattering and Transport Theory*. Academic Press.
- JAEGER, D., DEMEYERE, H., FINSY, R., SNEYERS, R., VANDERDEELEN, J., VAN-DER-MEEREN, P., AND VAN-LAETHEM, M. 1991. Particle sizing by photon correlation spectroscopy. part i: Monodisperse latices: influence of scattering angle and concentration of dispersed material. In *Part. Syst. Charact.* 8, 179.
- JENSEN, H., MARSCHNER, S., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proc. SIGGRAPH* 01, 511–518.
- KEY, J. R. 2005. Streamer: User's guide. *Tech Report, NOAA/NESDIS, Madison, Wisconsin*.
- MARSCHNER, S. 1998. Inverse rendering for computer graphics. *PhD Thesis, Cornell University*.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. on Graphics (SIGGRAPH)* 22, 3, 759–769.
- MCCORMICK, N. J. 1981. A critique of inverse solutions to slab geometry transport problems. *Prog. Nucl. Energy* 8.
- MCCORMICK, N. J. 1985. Sensitivity of multiple-scattering inverse transport methods to measurement errors. *JOSA A* 2.
- MCCORMICK, N. J. 1996. Analytical transport theory applications in optical oceanography. *Annals of Nuclear Energy* 23, 381–395.
- NARASIMHAN, S. G., AND NAYAR, S. K. 2003. Shedding light on the weather. In *CVPR* 03, 665–672.
- OISHI, T. 1990. Significant relationship between the backward scattering coefficient of sea water and the scatterance at 120 degrees. *Applied Optics* 29 (31), 4658–4665.
- PRAHL, S. A. 1988. Light transport in tissue. *PhD Thesis, University of Texas at Austin*.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. A signal processing framework for inverse rendering. *Proc. SIGGRAPH* 01, 117–128.
- SULLIVAN, S. A. 1963. Experimental study of the absorption in distilled water, artificial sea water, and heavy water in the visible region of the spectrum. *JOSA* 53.
- SUN, B., RAMAMOORTHY, R., NARASIMHAN, S. G., AND NAYAR, S. K. 2005. A practical analytic single scattering model for real time rendering. *ACM Trans. on Graphics (SIGGRAPH)* 24, 3, 1040–1049.
- WARD-LARSON, RUSHMEIER, H., AND PIATKO. 1997. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. on Visualization and Computer Graphics* 3, 4, 291–306.

Non-Linear Volume Photon Mapping

Diego Gutierrez[†], Adolfo Munoz, Oscar Anson and Francisco J. Seron

GIGA, Universidad de Zaragoza, Spain

Abstract

This paper describes a novel extension of the photon mapping algorithm, capable of handling both volume multiple inelastic scattering and curved light paths simultaneously. The extension is based on the Full Radiative Transfer Equation (FRTE) and Fermat's law, and yields physically accurate, high-dynamic data than can be used for image generation or for other simulation purposes, such as driving simulators, underwater vision or lighting studies in architecture. Photons are traced into the participating medium with a varying index of refraction, and their curved trajectories followed (curved paths are the cause of certain atmospheric effects such as mirages or rippling desert images). Every time a photon is absorbed, a Russian roulette algorithm based on the quantum efficiency of the medium determines whether the inelastic scattering event takes place (causing volume fluorescence). The simulation of both underwater and atmospheric effects is shown, providing a global illumination solution without the restrictions of previous approaches.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

Simulation of nature has always been one of the loftiest goals of computer graphics, providing a rich range of visual phenomena. Most of the times, the effect to be reproduced can be faked using a top-down approach, where the final desired result guides the implementation. This usually turns out relatively fast, ad-hoc methods that yield more than acceptable results. However, a physically correct simulation is necessary in certain fields where accuracy is a must. Underwater vision, driving simulators, the military, architectural lighting design etc. are fields where it is not enough to render an image which resembles reality. Predictive algorithms must be developed instead, where the image is the final visualization of the physically correct data generated. A bottom-up approach is then necessary: first, the basic laws of physics that govern the phenomenon need to be described and fed to the rendering system; the phenomenon itself will just be the logical, inevitable output. This approach sacrifices rendering speed in exchange for reliable, physically accurate numerical data that can be used for purposes beyond image generation.

Two of the greatest sources of visually appealing phenomena in nature are participating media and a varying index of refraction. Participating media are the cause of such well-known effects such as fog, clouds or blurry underwater vision, whereas a varying index of refraction yields mirages, rippling images, twinkling stars or some spectacular sunsets. Sources of inelastic scattering in ocean waters can greatly affect visibility and alter its color, whereas distortions caused by temperature differences can further alter the perception of things in such environment. Simulating underwater rescue missions, laying submarine data cables or even the correct interpretation of ancient World Heritage sites can benefit from an accurate description of light that includes an ampler range of phenomena.

We present in this paper a physically-based spectral simulation of light, solving the Full Radiative Transfer Equation (FRTE) and applying Fermat's law, which includes multiple inelastic scattering as well as an accurate description of the non-linear paths followed by the light rays in media with a varying index of refraction. It is based on an extension of the volume photon map algorithm presented by Wann Jensen and Christensen [JC98]. The main contributions are a full global illumination solution which supports non-linear light

[†] e-mail: diegog@unizar.es

paths and is free of the restrictions of previous works, and the physically-correct simulation of volume fluorescence in participating media, caused by inelastic scattering, including efficient computation of caustics. Atmospheric effects and underwater imagery are simulated as case studies to demonstrate the algorithm. To our knowledge, there is no previous research in computer graphics literature that models together physically-based inelastic scattering in participating, inhomogeneous media where the index of refraction varies continuously. Related previous works therefore span two different categories: inelastic scattering in participating media and non-linear light propagation.

Rendering participating media is not a new field in computer graphics, and an exhaustive review can be found in [PPS97]. There are two types of scattering events in a participating medium: *elastic* scattering, where no transfer of energy occurs between wavelengths, and *inelastic* scattering, where such energy transfers do occur, from shorter to longer wavelengths. Spectral global illumination algorithms that handle participating media only take into account elastic scattering, with the strategy consisting on decoupling the solutions for each sampled wavelength, then adding them to obtain the final image. No interaction between wavelengths is computed. To the authors' knowledge, the only previous work that simulates volume inelastic scattering in participating media is owed to Cerezo and Seron [CS03], using a discrete ordinate method. Unfortunately their method requires both rectangular meshing of the geometry, as well as an angular and spatial discretization which imposes high memory requirements, thus limiting the complexity of the scenes that can be reproduced (the problem is aggravated when simulating highly anisotropic scattering). They also cannot provide a full solution, failing to render caustics. Surface inelastic scattering works include [Gla95b] or [WTP01], but their methods are not extensible to participating media.

With respect to non-linear ray tracing, the first method to deal with non-straight light paths is owed to Berger et al. [BTL90], refracting the ray according to Snell's law in each of a series of flat homogeneous layers, thus achieving a piece-wise linear approximation of a curved path. This was challenged by Musgrave [Mus90], who develops a purely reflective model where rays follow a parabolic path, following the Kuhlar/Fabri physical model [FFLV82]. A more general approach to non-linear ray tracing is proposed by Gröller [Grö95], although the work does not study the influence of the index of refraction in the curvature of the rays, visualizing mathematical and physical systems instead. In the paper by Stam and Languenou [SL96], the authors use geometrical optics to describe how light bends if the index of refraction of the medium varies continuously. They nevertheless fail to provide a physically-based analytical expression for the index of refraction as a function of temperature and wavelength, and solve the equations only for two specific cases, thus losing generality. Seron et al. [SGGC05] implement a method of curved ray tracing capable of simulating

the inferior mirage and some sunset effects, although they do not attempt to calculate any lighting, deforming pre-lit textures instead. In [HW01] gravitational light bending is visualized according to the theory of general relativity, whereas other relativity- and physics-related papers include the bending caused by neutron stars or black holes [Nem93], so they cannot (nor pretend to) simulate the phenomena described in this paper. Yngve et al. [YOH00] describe a simple method to simulate the bending of light by interpolating a density field, but they need to exaggerate the variation of the index of refraction tenfold for the effect to be visible.

The paper is organized as follows: section 2 provides the physically-based background, with an overview of inelastic scattering, the FRTE and the Fermat's law. In section 3 we describe our extension of the volume photon map algorithm to include inelastic scattering and curved light paths, with sections 4 and 5 providing case studies of underwater imagery and atmospheric effects respectively. The discussion of the results and some additional images are presented in section 6, to finish the paper in section 7 with the conclusions and future work.

2. Physically-based Framework

We now present the physical framework of our work, by first introducing what inelastic scattering is, then deriving the FRTE that needs to be solved to account for it. In order to be able to compute non-linear light paths, we will use Fermat's law to obtain the correct trajectories.

2.1. Inelastic scattering

Inelastic scattering implies an energy transfer from wavelength λ' to λ , with $\lambda' < \lambda$ within the visible spectrum, and gives rise to fluorescence and phosphorescence phenomena. Fluorescence occurs when a molecule absorbs a photon of wavelength λ' (called excitation wavelength), and re-emits it at a longer wavelength λ according to a *fluorescence efficiency function* $P_f(\lambda)$. The time lapse between the two events is 10^{-11} to 10^{-8} seconds, so for computer graphics it can be taken as an instantaneous process. For pure substances, re-emission is isotropic and the wavelength of the re-emitted photons is independent of the different excitation wavelengths, although the intensity of the re-emission does depend on them. Phosphorescence is a similar process, governed by the *phosphorescence efficiency function*, with the main difference being that the re-emitted energy declines with time according to a function $d(t)$.

2.2. Full Radiative Transfer Equation

Usually, participating media algorithms solve the integro-differential Radiative Transfer Equation (RTE), which takes into account emission, absorption and elastic scattering, but

does not yield a solution for inelastic scattering events. Following the notation in [JC98], and reformulating to include wavelength dependencies, the RTE can be written as:

$$\frac{\partial L_\lambda(x, \vec{w})}{\partial x} = \alpha_\lambda(x) L_{e,\lambda}(x, \vec{w}) + \sigma_\lambda(x) L_{i,\lambda}(x, \vec{w}) - \alpha_\lambda(x) L_\lambda(x, \vec{w}) - \sigma_\lambda(x) L_\lambda(x, \vec{w}) \quad (1)$$

where $\frac{\partial L_\lambda(x, \vec{w})}{\partial x}$ represents the variation of radiance L at a point x in the direction \vec{w} , α and σ are the absorption and scattering coefficients, L_e is the emitted radiance and L_i is the in-scattered radiance. Defining the extinction coefficient as $\kappa_\lambda(x) = \alpha_\lambda(x) + \sigma_\lambda(x)$ and integrating $L_{i,\lambda}$ over the sphere Ω we get:

$$\frac{\partial L_\lambda(x, \vec{w})}{\partial x} = \alpha_\lambda(x) L_{e,\lambda}(x, \vec{w}) + \sigma_\lambda(x) \int_\Omega p_\lambda(x, \vec{w}', \vec{w}) L_\lambda(x, \vec{w}') d\vec{w}' - \kappa_\lambda(x) L_\lambda(x, \vec{w}) \quad (2)$$

which is the integro-differential, wavelength-dependent RTE governing the transport of light in participating media, with $p_\lambda(x, \vec{w}', \vec{w})$ being the phase function that defines the re-emission direction. However, this equation does not account for energy transfers between wavelengths, the phenomenon known as inelastic scattering. To be able to compute these inelastic scattering events, we need to develop the RTE equation further, by adding a term that accounts for such energy transfers. This term can be expressed as a double integral over the domains of the solid angle and wavelength:

$$\int_\Omega \int_\lambda \alpha_{\lambda_i}(x) f(x, \lambda_i \rightarrow \lambda) L_{\lambda_i}(x, \vec{w}') \frac{p_\lambda(x, \vec{w}', \vec{w})}{4\pi} d\vec{w}' d\lambda_i \quad (3)$$

where α_{λ_i} is the absorption coefficient for wavelength λ_i (remember there is no inelastic scattering without previous absorption), $f(x, \lambda_i \rightarrow \lambda)$ is the function that governs the efficiency of the energy transfer between wavelengths, defined as the probability of a photon of λ_i being re-emitted at λ . For fluorescence and phosphorescence, this phase function is isotropic [Mob94]. Adding this term to the RTE (equation 2) we obtain the FRTE:

$$\frac{\partial L_\lambda(x, \vec{w})}{\partial x} = \alpha_\lambda(x) L_{e,\lambda}(x, \vec{w}) + \sigma_\lambda(x) \int_\Omega p_\lambda(x, \vec{w}', \vec{w}) L_\lambda(x, \vec{w}') d\vec{w}' - \kappa_\lambda(x) L_\lambda(x, \vec{w}) + \int_\Omega \int_\lambda \alpha_{\lambda_i}(x) f(x, \lambda_i \rightarrow \lambda) L_{\lambda_i}(x, \vec{w}') \frac{p_\lambda(x, \vec{w}', \vec{w})}{4\pi} d\vec{w}' d\lambda_i \quad (4)$$

which is the equation that must be solved to take into account multiple inelastic scattering in participating media, thus being able to render volume fluorescence effects.

2.3. Varying index of refraction in inhomogeneous media

A varying index of refraction n_λ defines an inhomogeneous medium where light travels in curved paths. These curved paths result in a distorted image, with the mirages being probably the best known manifestation of the effect. To be able to simulate this type of inhomogeneous medium, we therefore need to obtain the curved trajectory of light as it traverses it. The direction \vec{w} in equation 4 therefore needs to be recomputed at each differential step, accounting for the changes in n_λ . We obtain this corrected direction at each step by solving Fermat's law, which defines how light traverses one given medium.

The following derivation of Fermat's law uses the work of Gutierrez et al. [GSMA04] and is not meant to be exhaustive. As stated in [Gla95a], *a ray of light, when travelling from one point to another, follows a path that corresponds to a stationary value of the optical path length (OPL)*. The OPL is defined as the index of refraction times the travelled path (or the distance the light would have travelled in a vacuum during the flight time through the material), and in its differential form it can be formulated as $d(OPL) = ndl$, where l is the path travelled by the light ray. The equation shows how light gets bent towards the areas with a greater index of refraction, as Snell's law also predicts for the boundary of two homogeneous media. A stationary value corresponds to a maximum or a minimum in the function, thus the derivative equals zero. We can therefore write:

$$\delta(OPL) = \delta \int_A^B ndl = \int_A^B \delta n dl + \int_A^B n \delta(dl) = \int_A^B \frac{\delta n}{\delta x_i} \delta x_i dl + \int_A^B n \delta(dl) = 0 \quad (5)$$

where x_i are the vector components of l . Considering dx_i as variables and taking increments we get $\delta(dl) = \frac{dx_i}{dl} \delta(dx_i)$. Since light trajectories start and end at the stationary points A and B , we get $\delta x_i(A) = 0$ and $\delta x_i(B) = 0$. Equation 5 then results:

$$\delta L = \int_A^B \left[\frac{\partial n}{\partial x_i} - \frac{d}{dl} \left(n \frac{dx_i}{dl} \right) \right] \delta x_i dl = 0 \quad (6)$$

Since this equation must hold for any value of δx_i , the integrand must equal zero, so we finally come up with the equation that must be solved to obtain the path followed by light while traversing any medium, as a function of the index of refraction at each point:

$$\frac{d}{dl} \left(n \frac{d\vec{r}}{dl} \right) - \nabla n = 0 \Leftrightarrow \frac{d}{dl} \left(n \frac{dx_j}{dl} \right) - \frac{\partial n}{\partial x_j} = 0 \quad (j = 1, 2, 3) \quad (7)$$




Figure 1: Error and rendering time (secs.) as functions of the error tolerance in the Dormand-Prince RK5(4)7M method for a test scene.

where $\vec{r} = x_j$ are the coordinates (x, y, z) of each point. This equation cannot be solved analytically, and thus we must apply a numerical method. We now need to rewrite equation 7 in order to solve it in a more efficient way than the Euler method presented in [GSMA04]:

$$\frac{d^2 x_j}{dl^2} = \frac{1}{n} \left(\frac{\partial n}{\partial x_j} - \frac{dn}{dl} \frac{dx_j}{dl} \right) \quad (8)$$

Doing the change of variable $y_j = \frac{dx_j}{dl}$ we obtain:

$$y_j' = \frac{1}{n} \left(\frac{\partial n}{\partial x_j} - \frac{dn}{dl} y_j \right) \quad (9)$$

where $\frac{dn}{dl} = \frac{dn}{dx_j} \frac{dx_j}{dl}$. The change of variable can also be written as:

$$x_j' = y_j \quad (10)$$

Equations 9 and 10 define a system where x_j represents the position and y_j the velocity at a given point in the trajectory, which can be written in matrix form as:

$$\begin{pmatrix} x_j \\ y_j \end{pmatrix}' = \begin{pmatrix} y_j \\ \frac{1}{n} \left(\frac{\partial n}{\partial x_j} - \frac{dn}{dl} y_j \right) \end{pmatrix} \quad (11)$$

This equation 11 has the form $Y' = f(l, Y)$, which defines an Initial Value Problem with $Y(0) = \alpha$. We solve this problem by applying the embedded Runge-Kutta method RK5(4)7M from the Dormand-Prince family. A detailed description of the method and the error tolerance can be found in [DP80].

We have tested the implementation in a simple scene

where the index of refraction varies according to the equation $n = 1 + ky$, with y representing height, and k varying from -0.1 to 0.1. This distribution of n can be solved analytically, so we can measure the numerical error against the exact solution. Figure 1 shows the error of the Dormand-Prince RK5(4)7M method as the tolerance is reduced, along with the time it takes to reach the solution. As it can be seen, error tolerances in the range of 10^{-8} to 10^{-12} yield good results without much of a time penalty. Error tolerances beyond 10^{-14} start increasing rendering times considerably.

3. Extension of the Volume Photon Mapping Algorithm

Ray tracing techniques involve shooting rays into the scene from the camera and following them to detect hits with the geometry, then shooting shadow rays to the lights to find out direct illumination. With curved light paths this turns out to be highly impractical, though, since finding the ray with the physically-correct curvature which goes from the intersection point to the light is computationally very expensive (or the solution might not even exist). Groeller [Grö95] proposes three solutions: considering shadow rays to follow straight paths, retrieving all lighting information straight from the textures, and finally voxelizing the space and pre-storing the approximated incident directions of light sources for each voxel, by launching rays from the light sources into the scene prior to the render pass. The first two are clearly not physically-based, while the third only approximates the solution with a preprocessing step.

In order to obtain a physically-based solution for multiple inelastic scattering in inhomogeneous media with a varying index of refraction n , we have extended the volume photon mapping algorithm [JC98] to account both for volume fluorescence and the distortions caused by the changing n .

For inelastic scattering, we need to model the possibility of an absorbed photon being re-emitted at a different wavelength. Equation 4 includes a term $f(x, \lambda_i \rightarrow \lambda)$ known as *wavelength redistribution function*, which represents the ef-

efficiency of the energy transfer between wavelengths. It is defined as the quotient between the energy of the emitted wavelength and the energy of the absorbed excitation wavelength, per wavelength unit. Reformulating in terms of photons instead of energy we have the *spectral quantum efficiency function* $\eta(x, \lambda_i \rightarrow \lambda)$, defined as the number of photons emitted at λ per wavelength unit, divided by the number of absorbed photons at λ_i . Both functions are dimensional (nm^{-1}), and are related as follows:

$$f(x, \lambda_i \rightarrow \lambda) = \eta(x, \lambda_i \rightarrow \lambda) \frac{\lambda_i}{\lambda} \quad (12)$$

A related dimensionless function that describes inelastic scattering is the *quantum efficiency* Γ , defined as the total number of photons emitted at all wavelengths divided by the number of photons absorbed at excitation wavelength λ_i . It is related to the spectral quantum efficiency function by the equation:

$$\Gamma(\lambda_i) = \int_{\lambda} \eta(x, \lambda_i \rightarrow \lambda) d\lambda \quad (13)$$

Our extension to the volume photon mapping algorithm includes a) solving Fermat's law to obtain the curved trajectory of each photon if the index of refraction varies (and also for the eye rays shot during the radiance estimate phase), thus being able to overcome the shadow ray problem presented above and to obtain a full solution including effects such as color bleeding and caustics; and b) the inclusion of the quantum efficiency Γ to govern the probability of an inelastic scattering event. As shown in figure 2, once the albedo-based Russian roulette determines that a certain photon has been absorbed by the medium, a second Russian roulette based on the quantum efficiency determines whether an inelastic scattering event takes place, and therefore the photon has to be re-emitted at a different wavelength. This is done by generating a random number $\xi_{in}[0, 1]$ so that:

$$\xi_{in}[0, 1] \rightarrow \begin{cases} \xi_{in} \leq \Gamma & \text{Photon is re-emitted} \\ \xi_{in} > \Gamma & \text{Photon remains absorbed} \end{cases} \quad (14)$$

If re-emitted, the new wavelength must be obtained, for which we must sample the spectral quantum efficiency function $\eta(x, \lambda_i \rightarrow \lambda)$ for the excitation wavelength λ_i . This can be simply done by rejection sampling the function, but to increase efficiency we perform importance sampling using the inverse of its cumulative distribution function (cdf). A random number $\psi[0, 1]$ therefore yields the new wavelength for the re-emitted photon. Steeper areas of the cdf increase the probability of a photon being re-emitted at the corresponding wavelengths.

Figure 2 shows the basic scheme of the algorithm. The

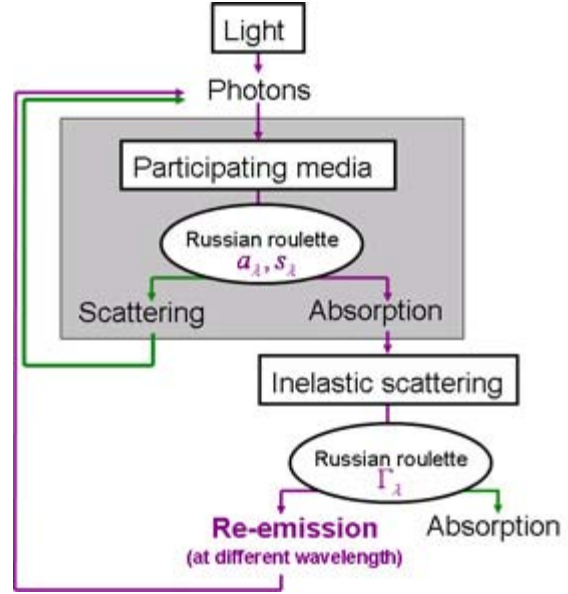


Figure 2: Our extended volume photon mapping algorithm.

sequence of events in the original volume photon mapping by [JC98] is represented inside the grey area.

4. Case Study: Underwater Imagery

We chose deep ocean waters as our first case study, given its rich range of elastic and inelastic scattering phenomena and the fact that it is a medium well studied by oceanographers. Pure seawater absorbs most wavelengths except for blue: the absorption coefficient peaks at 760 nanometers, and reaches a minimum at 430 nm. The phase function p is modelled as the phase function in pure sea water plus the phase function of the scattering by suspended particles, as proposed in [Mob94] ($p = p_w + p_p$). For pure water we use a phase function similar to Rayleigh's:

$$p_w(\theta) = 0.06225(1 + 0.835\cos^2\theta) \quad (15)$$

while the scattering caused by particles is modelled using a Henyey-Greenstein phase function with $g = 0.924$:

$$p_p(\theta, g) = \frac{1 - g^2}{(1 + g^2 - 2g\cos\theta)^{3/2}} \quad (16)$$

It is very common in ocean waters to see a color shift ranging from greenish to very bright green, or even yellowish. These hue shifts are due to the variation in the concentration and type of the suspended microorganisms, mainly phytoplankton, which presents a maximum absorption at 350 nm, rapidly decreasing to almost zero beyond 500 nm. The




Figure 3: Fluorescent ocean water in Cornell rooms. (a), (b) and (c) show varying concentrations of chlorophyll ($0.05\text{mg}/\text{m}^3$, $0.1\text{mg}/\text{m}^3$ and $5\text{mg}/\text{m}^3$ respectively). (d) High concentration of yellow matter ($5\text{mg}/\text{m}^3$).

most important element in the phytoplankton is chlorophyll, which presents spectral absorption peaks in the blue and red ends of the spectrum and is the most important source of volume fluorescence in the waters. For chlorophyll, $\Gamma^c(\lambda_i)$ is wavelength-independent, with values ranging from 0.01 to 0.1 (we use the superscript c for chlorophyll). As with most inelastic scattering event, the re-emission phase function is isotropic.

Another important source of fluorescence is the Color Dissolved Organic Matter (CDOM), also called yellow matter, present in shallow ocean waters and harbors. $\Gamma^y(\lambda_i)$ is also wavelength-independent, with values between 0.005 and 0.025, and re-emission is also isotropic [Haw92].

All the images in the paper have been rendered on a Beowulf system composed of six nodes, each one being a Pentium 4 @ 2.8 GHz. with 1 Gb. of RAM. Figure 3 shows different colorations of ocean water, according to varying chlorophyll and yellow matter concentrations which trigger inelastic scattering events with different probabilities. The images were rendered with 250,000 photons stored in the volume photon map and 200 photons used for the radiance estimate. This high numbers are needed to obtain accurate results, since we use the volume photon map to compute both direct and indirect illumination. Direct illumination in participating media with a varying index of refraction cannot be efficiently computed using ray tracing techniques, as explained at the beginning of section 3. The spectrum was sampled at nine intervals. Below each picture, the resulting absorption and extinction curves (functions of the different concentrations of chlorophyll in the modelled waters) are shown for each case. Image (a) shows little fluorescence (low chlorophyll concentration of $0.05\text{mg}/\text{m}^3$), and the wa-

ters are relatively clear. When chlorophyll concentration increases, fluorescence events become more prominent and the image first gets a milky aspect (b), losing visibility and reaching a characteristic green hue when chlorophyll reaches $5\text{mg}/\text{m}^3$. Image (d) shows fluorescence owed to yellow matter. The absorption function in this case has been modelled after [Mob94]: $a_y(\lambda) = a_y(440)^{-0.014(\lambda-440)}$ where $a_y(440)$ is the empirical absorption at 440 nm. Rendering times for the images were six minutes.

5. Case Study: Atmospheric Phenomena

The images in this section illustrate some of the most relevant effects in nature owed to curved light paths. To achieve physically correct results we have modelled the Earth as a sphere with a radius of 6371 units (one unit equals one kilometer); the atmosphere is another concentric sphere with a thickness of 40 kilometers. Taking the 1976 *USA Standard Atmosphere* (USA76) [USG76], we first obtain a standard temperature and pressure profile of the whole 40 kilometers, with temperature decreasing at an approximate rate of 0.6°C per 100 meters. In order to curve light correctly according to Fermat's law, we need to obtain the wavelength-dependent index of refraction as a function of both the temperature and pressure given by the USA76. To do so, we follow the method described in [GSMA04], by first obtaining density as a function of temperature $T(h)$ and pressure $P(h)$ using the Perfect Gas law $\rho(h) = \frac{P(h)M}{RT(h)}$, where M and R are constants of values $28.93 \cdot 10^{-3} \text{ kg/mol}$ and $8.3145 \text{ J/mol} \cdot \text{K}$ respectively. The Gladstone-Dale law [GD58] relates $n(\lambda, h)$ as a function of both $\rho(h)$ and $n(\lambda)$, given by the expression:

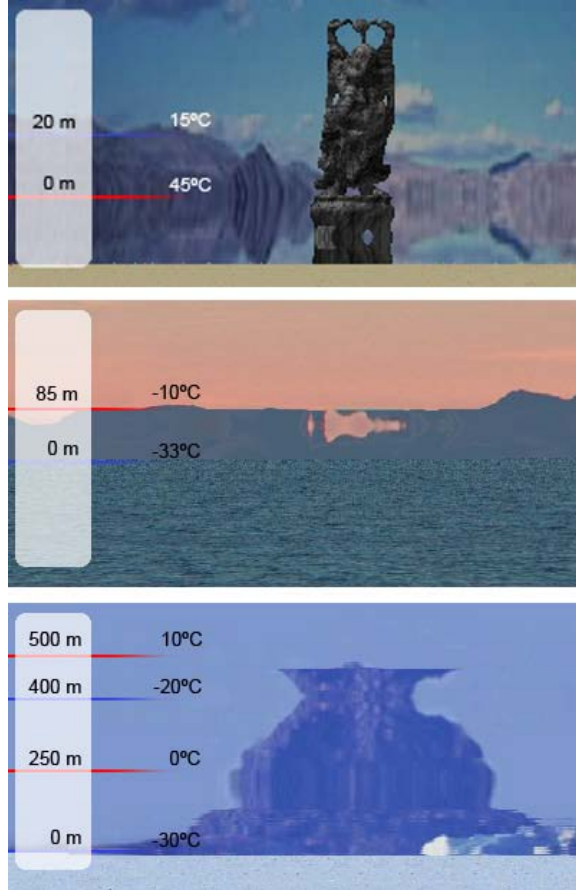


Figure 4: Simulation of several atmospheric phenomena. Top: inferior mirage. Middle: superior mirage. Bottom: Fata Morgana.

$$n(h, \lambda) = \rho(h) \cdot (n(\lambda) - 1) + 1 \quad (17)$$

The only missing function is now $n(\lambda)$, which we obtain from Cauchy's analytical formula [BW02]:

$$n(\lambda) = a \cdot \left(1 + \frac{b}{\lambda^2}\right) + 1 \quad (18)$$

where a and b depend on the medium considered (for air, their values are $a = 29.79 \cdot 10^{-5}$ and $b = 5.67 \cdot 10^{-5}$). Sellmeier [BW02] provides a slightly more elaborated formula, but we have chosen Cauchy's for efficiency reasons.

Combining equations 17 and 18 we finally obtain our profile for $n(\lambda, h)$, which we can alter at will to obtain the desired effects. To interpolate the complete, altered profiles for the whole 40 km, we use Fermi's distribution, as proposed in [VDWGL00].

The camera in the scenes is placed far from the mirages at a specific height for each effect to be seen (they can only appear if the observer's line of vision forms an angle less than one degree with the horizon). The error tolerance in the Dormand-Prince RK5(4)7M method has been set to 10^{-9} , and the spectrum has been sampled in three wavelengths. Figure 4 (top) shows our simulation of an inferior mirage, which occurs when the ground is very hot and heats up the air layers right above it, thus creating a steep temperature gradient (30°C in 20 meters). As a consequence, light rays get bent upwards, and an inverted image of the Happy Buddha and the background appears on the ground. The camera is placed 10 meters above the ground. The image took 14 minutes to render.

Inversion layers are caused by an increase of air temperature with height, reversing the standard behavior where temperature decreases as a function of height. This happens most commonly above cold sea waters, and the light rays get bent downward, giving rise to the superior mirage. Figure 4 (middle) shows our simulation, modelling an inversion layer with a temperature gradient of 23°C . The apparent hole in the mountains is actually formed by the superior inverted image of the real mountains. The camera is placed also 10 meters above the ground, and the image took four minutes and 32 seconds to render. The great decrease in rendering time compared to the inferior mirage is owed to the simpler geometry of the scene, since the far away mountains are textured low-resolution objects.

Maybe less known than the two previous examples, the Fata Morgana occurs as a concatenation of both superior and inferior mirages, and is a much rarer phenomenon. Figure 4 (bottom) shows our simulation with two inversion layers with steep temperature gradients. There is an inferior mirage image across the middle of the mountain plus a superior mirage with the inverted image on top. The shape of the mountain gets greatly distorted; the Fata Morgana has historically tricked arctic expeditions, making them believe they were seeing huge mountains that were just a complicated pattern of upright and inverted images of the real, much lower hill (*Fata Morgana* is in fact the name of a fairly enchantress skilled in the art of changing shape, which she learnt from Merlin the Magician). The camera is placed at 300 meters (for the Fata to be visible it needs to be between the inversion layers), and the rendering time was five minutes.

6. Discussion

The method described has been implemented in Lucifer, our in-house global illumination renderer. It can handle multiple inelastic scattering in inhomogeneous participating media with a varying index of refraction, thus rendering effects such as mirages or fluorescence in ocean waters with full lighting computation. It deals well with strong anisotropy in the phase functions and the effects of backscattering, since no discretizations of the scene must be performed,

and thus the shortcoming of the only previous work on volume fluorescence [CS03] is overcome. It also supports real light sources, with photometric data input specified in the standard CIBSE TM14 format [CIB88]. This is a must for predictive rendering and for generating physically accurate data. The real light sources are sampled so that photons are emitted proportionally to the distribution of the light, given by its photometry.

Spectral images are calculated in high dynamic range, in order to obtain accurate data from the simulations. For tone reproduction purposes we map luminances to the display based on the work by Ward et al. [LRP97] and Pattanaik et al. [PTYG00]. To increase realism during the visualization of the images, an additional operator has been added which simulates the effects of chromatic adaptation in the human eye. This operator is specially important in the realistic depiction of underwater imagery, where the cones in the human eye might undergo a loss of spectral sensitivity after having been exposed to the same wavelength for a long period of time (underwater imagery being usually blue or green mostly). The complete description of such operator can be found in [GSMA04].

As stated in the introduction, the algorithm implemented is general and physically-based. This allows us to use the radiometric and photometric data obtained from the simulations for any purpose other than rendering, such as professional architectural lighting or accurate simulations of deep underwater vision, given the exact description of the luminaire to be used and the water conditions. This accuracy obviously increases rendering times compared to faked, ad-hoc solutions. To improve efficiency, we impose an early light path termination and an adaptive integration step while solving Fermat's law. Choosing the Dormand-Prince RK5(4)7M numerical method over the more standard Euler method has produced speedups of up to 106.4. We have also used a parallel implementation on a six-PC Beowulf system of our non-linear photon mapping algorithm, achieving additional speedups between 4.2 and 4.8.

The non-linear photon mapping implementation allows us to extend several sunset effects similar to the ones simulated in [GSMA04], by including a thin layer of fog between the observer and the sun. The solar disk gets distorted into different shapes, while light is scattered through the layer of fog, thus achieving a "winter sunset" look (figure 5, left and middle). Figure 5 right shows volume caustics generated by a crystal sphere in a fluorescent medium.

Figure 6 shows several renders obtained with Lucifer. All of them are lit by a Philips SW-type[©] luminaire, specified according to the CIBSE TM14 format. The only light source is immersed in the medium, so no caustics from the interaction of sunlight with the surface appear. The medium modelled does not emit light, although adding that to the model is straightforward and would allow us to simulate effects such as bioluminescence in the water. Fluorescence




Figure 5: Sunset effects through a layer of fog. Left: flattened sun. Middle: split sun. Right: Volume caustics in a fluorescent medium.

owed to inelastic scattering is computed according to the varying concentrations of chlorophyll in each image (between 0.01 and $0.1\text{mg}/\text{m}^3$). The volume photon map in all the images contains 500.000 photons, and the radiance estimate used 250. Again, these high numbers are needed since we compute direct lighting with the photon map. The top two images represent a sunken boat along a Happy Buddha in clear, shallow waters (left) or deep underwater with a chlorophyll concentration of $0.05\text{mg}/\text{m}^3$ (right). For the bottom-left image, we have added a volume temperature field that simulates a heat source outside the image as explained in [SGGC05], deriving the index of refraction using the formula $n = 1 + \frac{T_p}{T} (n_o - 1)$ as proposed by Stam and Languenou [SL96]. The distortions caused by the varying index of refraction are visible, similar to the characteristic rippling in a real desert scene. The bottom-middle image uses a smoke-like medium, modelled as a 3D turbulence function, whereas the last to the right shows the effects of a highly anisotropic medium. The images are 400 pixels wide and took between 30 and 40 minutes to render, without any penalty imposed by the anisotropy in the last image.

7. Conclusion and Future Work

We proposed a novel extension of the widely used photon mapping technique, which accounts for multiple inelastic scattering and can provide a full global illumination solution in inhomogeneous media with a varying index of refraction, where light paths are bent. No pre-lit textures are needed in this case, since both direct and indirect lighting is calculated from the photon map. The method is physically-based and yields accurate high-dynamic results that can either be output as an image to a display device (via tone mapping), or used in other fields as raw data. Inelastic scattering is calculated during the photon tracing stage, so the extra cost required is just a second Russian roulette per absorption. The accompanying video shows the feasibility of the approach for animations.

Practically all inelastic scattering effects in the visible range of the spectrum mean a transfer of energy from shorter to longer wavelengths. Nevertheless, the algorithm presented in this work can handle rarer inelastic scattering events where energy gets transferred from longer to shorter wave-




Figure 6: Different images with inelastic scattering in participating media. Top left: very low chlorophyll concentration. Top right: higher concentration yields more inelastic scattering events. Bottom left: distortions caused by a 3D temperature field. Bottom middle: 3D turbulence field simulating smoke. Bottom right: highly anisotropic medium.

lengths (such as a fraction of the Raman scattering that occurs naturally in several solids, liquids and gases [Mob94]), since it does not follow a cascade, one-way scheme from the blue end to the red end of the spectrum. The application of these type of inelastic scattering to computer graphics is probably just marginal, but the data generated can be very useful to physicists or oceanographers. Adding phosphorescence effects could make use of the work by Cammarano and Wann Jensen [CJ02], although a more straightforward approach would be to use the decay function $d(t)$ in each frame. Any number of light sources can be used in the scene, even with different photometric descriptions.

The bottleneck of the algorithm is solving the paths for each photon and eye-ray using Fermat's law. Although the use of a Dormand-Prince method has drastically reduced rendering times by two orders of magnitude, additional work needs to be done to achieve near real-time frame rates. Importance maps could be used for this purpose, although two other promising fields of research lay ahead: the first one is the implementation of the algorithm on GPUs, as proposed by Purcell et al. [PDC*03]. The second would try to take advantage of temporal coherence of light distribution, as presented by Myszkowski et al. [MTAS01].

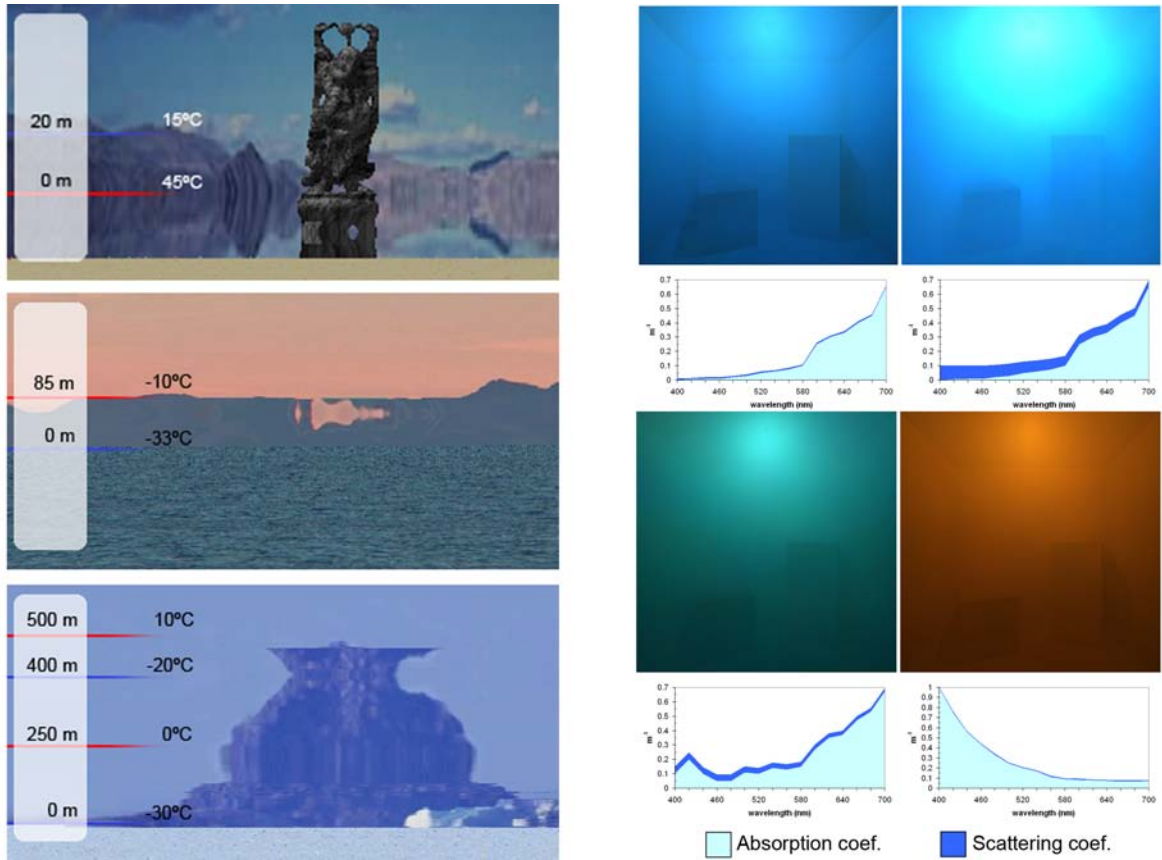
8. Acknowledgements

This research was partly done under the sponsorship of the Spanish Ministry of Education and Research through the project TIN2004-07672-C03-03. The authors would also like to thank Eduardo Jiménez for his initial contribution to this work.

References

- [BTL90] BERGER M., TROUT T., LEVIT N.: Ray tracing mirages. *IEEE Computer Graphics and Applications* 10, 3 (May 1990), 36–41.
- [BW02] BORN M., WOLF E.: *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 2002.
- [CIB88] CIBSE: *Standard File Format for Transfer of Luminaire Photometric Data*. The Chartered Institution of Building Services Engineers, 1988.
- [CJ02] CAMMARANO M., JENSEN H. W.: Time dependent photon mapping. In *Proceedings of the 13th Eurographics workshop on Rendering* (2002), Eurographics Association, pp. 135–144.

- [CS03] CEREZO E., SERON F.: Inelastic scattering in participating media. application to the ocean. In *Proceedings of the Annual Conference of the European Association for Computer Graphics, Eurographics 2003* (2003), pp. CD-ROM.
- [DP80] DORMAND J., PRINCE P.: A family of embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics* 6(1) (1980), 19–26.
- [FFLV82] FABRI E., FIORZO G., LAZZERI F., VIOLINO P.: Mirage in the laboratory. *Am. J. Physics* 50(6) (1982), 517–521.
- [GD58] GLADSTONE J. H., DALE J.: On the influence of temperature on the refraction of light. *Phil. Trans.* 148 (1858), 887.
- [Gla95a] GLASSNER A.: *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Francisco, California, 1995.
- [Gla95b] GLASSNER A. S.: A model for fluorescence and phosphorescence. In *Photorealistic Rendering Techniques* (1995), Sakas P. S. G., Müller S., (Eds.), Eurographics, Springer-Verlag Berlin Heidelberg New York, pp. 60–70.
- [Grö95] GRÖLLER E.: Nonlinear ray tracing: visualizing strange worlds. *The Visual Computer* 11, 5 (1995), 263–276.
- [GSMA04] GUTIERREZ D., SERON F., MUNOZ A., ANSON O.: Chasing the green flash: a global illumination solution for inhomogeneous media. In *Spring Conference on Computer Graphics* (2004), (in cooperation with ACM SIGGRAPH A. P., Eurographics), (Eds.), pp. 95–103.
- [Haw92] HAWES S.: *Quantum Fluorescence Efficiencies of Marine Fulvic and Humic Acids*. PhD thesis, Dept. of Marine Science, Univ. of South Florida, 1992.
- [HW01] HANSON A. J., WEISKOPF D.: Visualizing relativity. siggraph 2001 course 15, 2001.
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *SIGGRAPH 98 Conference Proceedings* (jul 1998), Cohen M., (Ed.), Annual Conference Series, ACM SIGGRAPH, Addison Wesley, pp. 311–320. ISBN 0-89791-999-8.
- [LRP97] LARSON G. W., RUSHMEIER H., PIATKO C.: A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics* 3, 4 (Oct. 1997), 291–306.
- [Mob94] MOBLEY C.: *Light and Water. Radiative Transfer in Natural Waters*. Academic Press, Inc., 1994.
- [MTAS01] MYSZKOWSKI K., TAWARA T., AKAMINE A., SEIDEL H. P.: Perception-guided global illumination solution for animation. In *Computer Graphics Proceedings, Annual Conference Series, 2001 (ACM SIGGRAPH 2001 Proceedings)* (Aug. 2001), pp. 221–230.
- [Mus90] MUSGRAVE F. K.: A note on ray tracing mirages. *IEEE Computer Graphics and Applications* 10, 6 (Nov. 1990), 10–12.
- [Nem93] NEMIROFF R. J.: Visual distortions near a neutron star and black hole. *American Journal of Physics* 61(7) (1993), 619–632.
- [PDC*03] PURCELL T. J., DONNER C., CAMMARANO M., JENSEN J., HANRAHAN P.: Photon mapping on programmable graphics hardware. In *SIGGRAPH/Eurographics Workshop on Graphics Hardware* (2003), Eurographics Association, pp. 041–050.
- [PPS97] PEREZ F., PUEYO X., SILLION F.: Global illumination techniques for the simulation of participating media. In *Proc. of the Eighth Eurographics Workshop on Rendering* (1997), pp. 16–18.
- [PTYG00] PATTANAIK S., TUMBLIN J. E., YEE H., GREENBERG. D. P.: Time-dependent visual adaptation for realistic image display. In *SIGGRAPH 2000, Computer Graphics Proceedings* (2000), Akeley K., (Ed.), Annual Conference Series, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, pp. 47–54.
- [SGGC05] SERON F., GUTIERREZ D., GUTIERREZ G., CEREZO E.: Implementation of a method of curved ray tracing for inhomogeneous atmospheres. *Computers and Graphics* 29(1) (2005).
- [SL96] STAM J., LANGUÉNOU E.: Ray tracing in non-constant media. In *Eurographics Rendering Workshop 1996* (New York City, NY, June 1996), Pueyo X., Schröder P., (Eds.), Eurographics, Springer Wien, pp. 225–234. ISBN 3-211-82883-4.
- [USG76] USGPC: *U.S. Standard Atmosphere*. United State Government Printing Office, Washington, D.C., 1976.
- [VDWGL00] VAN DER WERF S., GUNTHER G., LEHN W.: Novaya zemlya effects and sunsets. *Applied Optics* 42, 3 (2000).
- [WTP01] WILKIE A., TOBLER R., PURGATHOFER W.: Combined rendering of polarization and fluorescence effects. In *Rendering Techniques '01 (Proc. Eurographics Workshop on Rendering 2001)* (2001), Gortler S.J. M. K. e., (Ed.), Eurographics, Springer-Verlag, pp. 197–204.
- [YOH00] YNGVE G. D., O'BRIEN J. F., HODGINS H.: Animating explosions. In *Proceedings of the Computer Graphics Conference 2000 (SIGGRAPH-00)* (New York, July 23–28 2000), Hoffmeyer S., (Ed.), ACM Press, pp. 29–36.



Visualizing Underwater Ocean Optics

Diego Gutierrez, Francisco J. Seron, Adolfo Munoz and Oscar Anson

University of Zaragoza, Spain

Abstract

Simulating the in-water ocean light field is a daunting task. Ocean waters are one of the richest participating media, where light interacts not only with water molecules, but with suspended particles and organic matter as well. The concentration of each constituent greatly affects these interactions, resulting in very different hues. Inelastic scattering events such as fluorescence or Raman scattering imply energy transfers that are usually neglected in the simulations. Our contributions in this paper are a bio-optical model of ocean waters suitable for computer graphics simulations, along with an improved method to obtain an accurate solution of the in-water light field based on radiative transfer theory. The method provides a link between the inherent optical properties that define the medium and its apparent optical properties, which describe how it looks. The bio-optical model of the ocean uses published data from oceanography studies. For inelastic scattering we compute all frequency changes at higher and lower energy values, based on the spectral quantum efficiency function of the medium. The results shown prove the usability of the system as a predictive rendering algorithm. Areas of application for this research span from underwater imagery to remote sensing; the resolution method is general enough to be usable in any type of participating medium simulation.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

Ocean water is arguably the richest participating medium in terms of optical thickness and the number and type of interactions that occur in it. This paper deals with the physically-based rendering of underwater scenes by simulating the in-water light field, based on a compact bio-optical model that takes into account the *dissolved* and *particulate* matter, optically influential constituents of the water. To ensure accuracy, we use published data obtained from a wide range of literature in the field of oceanography. Our model is not restricted to just the visible spectrum and can be adapted to any type of known ocean water in particular, or to any kind of participating medium in general.

Scattering in water is caused by interactions of light at molecular level and with particles [Mob94]. It can be classified in two broad categories: *elastic* or *inelastic* scattering, depending on whether the scattered photon maintains or changes its energy in the process. The inelastic scattering events can be further subclassified according to the nature of the energy transfer: *Stokes* scattering, when a molecule of the medium absorbs the photon and re-emits it with a lower en-

ergy, and *anti-Stokes* scattering, when the re-emitted photon has a higher energy. Both cases are covered by our model. The process implies an energy transfer from wavelength λ' to λ , with λ' being the excitation wavelength and λ the re-emitted wavelength. The former case implies a shift towards longer wavelengths, whereas in the latter the scattered photon has a shorter wavelength. Major forms of elastic events in water include Einstein-Smoluchowski scattering (see Section 3.2), whereas for inelastic events, *Raman* scattering and *fluorescence* are the two most prominent (see Section 3.3).

The presence and concentrations of the constituents in the water determine its optical properties. These optical properties are divided in two classes: *inherent* and *apparent*. The *inherent optical properties* (IOP) only depend on the constituents of the water, whereas the *apparent optical properties* (AOP) are not properties of the aquatic medium itself, although they do depend on its characteristics. Typical IOP are the absorption coefficient, the scattering coefficient or the scattering phase function. Some of the AOP include irradiance reflectance, attenuation coefficients or the average cosines [Pre76]. To obtain the in-water light field,

we rely on the physically based theory of *radiative transfer* [Cha60], which relates the IOP and AOP. More precisely, the link is provided through the Radiative Transfer Equation (RTE) [SCP94], which takes into account emission, absorption and elastic scattering. Unfortunately this equation can not account for the phenomenon known as *inelastic scattering* described previously, which is of significant importance in ocean waters. We consequently expand the RTE by adding an extra term, thus obtaining the Full Radiative Transfer Equation (FRTE) [Gla95] and solving it by using an extended version of the method presented by Gutierrez et al. [GMAS05]:

$$\begin{aligned} \frac{\partial L(\lambda, \vec{\omega}_o)}{\partial x} = & \alpha(\lambda) L_e(\lambda, \vec{\omega}_o) - \kappa(\lambda) L(\lambda, \vec{\omega}_o) \\ & + \sigma(\lambda) \int_{\Omega} p(\lambda, \vec{\omega}_i, \vec{\omega}_o) L(\lambda, \vec{\omega}_i) d\vec{\omega}_i \\ & + \int_{\Omega} \int_W \left\{ \sigma(\lambda', \lambda) p(\lambda', \lambda, \vec{\omega}_i, \vec{\omega}_o) L(\lambda', \vec{\omega}_i) \right\} d\lambda' d\vec{\omega}_i \end{aligned} \quad (1)$$

where L is the radiance and $\vec{\omega}_i$ and $\vec{\omega}_o$ are, respectively, the incoming and outgoing directions of that radiance. α , σ and κ are the absorption, scattering and extinction coefficients respectively. We assume $L_e(\lambda, \vec{\omega}_o)$ to be zero, thus making the medium non-emissive. Note that the last term models the inelastic scattering events and is expressed as a double integral over the domains of the solid angle Ω and wavelength W . Here $p(\lambda', \lambda, \vec{\omega}_i, \vec{\omega}_o)$ is the phase function for inelastic events and $\sigma(\lambda', \lambda)$ is the inelastic scattering function for the energy exchange between λ' and λ . For simplicity, when considering elastic interactions ($\lambda = \lambda'$) parameters λ, λ' are simplified to a single parameter λ . For processes such as fluorescence, where the photons are inelastic scattered to longer wavelengths, the function $\sigma(\lambda', \lambda)$ is usually expressed as:

$$\sigma(\lambda', \lambda) = \alpha(\lambda') f(\lambda', \lambda) \quad (2)$$

where $\alpha(\lambda')$ is the inelastic absorption coefficient and $f(\lambda', \lambda)$ is the *wavelength redistribution function*, which governs the efficiency of the energy transfer between wavelengths. It is defined as the probability of a photon of λ' that inelastically scatters being re-emitted at λ . Therefore, (2) expresses the inelastic scattering as a percentage of the inelastic absorption coefficient. Section 3.3 gives more details on how to model this redistribution function $f(\lambda', \lambda)$.

Our research on water simulation encompasses the fields of both computer graphics and oceanography, and it is free from the restrictions of previous works. The main contributions of this paper are:

- A compact, parameterized bio-optical model of ocean waters which can be used in computer graphics applications.
- A resolution method based on the theory of radiative transfer, which solves the FRTE by handling *all* kinds of inelastic scattering events and modeling both absorption and elastic scattering accurately. This method is based on photon mapping [Jen01].

- A link between the IOP of water and the resulting light field, which in turn defines its AOP, based on radiative transfer theory.

The remainder of this paper is organized as follows: Section 2 presents previous work on the simulation of light transport in water bodies. In Section 3 a comprehensive bio-optical model is developed, whilst section 4 presents our simulation method. The paper ends with the results and conclusions.

2. Related work

The simulation of light transport in participating media usually either relies on Monte-Carlo techniques for ray tracing (Rushmeier and Torrance [RT87]; Nakamae et al. [NKON90]; Tadamura and Nakamae [TN95]) or attempts to solve the RTE, such as the method proposed by Kaneda et al. [KYNN91]. Nishita et al. [NSTN93] display water from outer space modifying this method, but both works only take into account single scattering. In the work of Premoze and Ashikhmin [PA01], no radiance due to scattering is calculated at all, using empirical equations based on experimental data instead. Mobley [Mob94] developed a method to solve the RTE analytically, but it cannot be extended to take into account inelastic scattering. Recently, the Lorenz-Mie theory has been generalized and applied to rendering natural waters by Frisvad, Christensen and Jensen [FCJ07], also neglecting the effects of inelastic scattering. Cerezo and Seron [CS04] also develop a bio-optical model. Whilst the goal of their work is closely related to ours, we overcome here significant shortcomings:

- They use a discrete ordinate method, which requires an angular and spatial discretization of the volume to be rendered. This imposes high memory requirements which seriously limit the complexity of the scenes that can be reproduced.
- In their work, inelastic scattering simulations are limited to fixed re-emissions in the 680 nm. wavelength..
- They cannot provide a full solution to the light transport problem.

Gutierrez et al. [GMAS05] present a method that deals with participating media in which the index of refraction is not homogeneous, while also taking into account the simulation of some inelastic scattering events. They apply their method to the simulation of underwater imagery using a simplified, four-parameter model of ocean waters. In this regard, our paper offers improvement in the following ways:

- Our bio-optical model of ocean waters is more complete, thus making the simulations more accurate.
- They also fail to develop a complete description for the complex inelastic scattering events that occur underwater, and the method is limited to re-emissions at lower energy levels and at fixed wavelengths. In this paper *all* inelastic scattering events can be modeled, including *Anti Stokes* scattering events like *Raman scattering* (see Section 3.3).

- We additionally offer simulations using real data from different seas as a means of visual validation.

3. The Bio-Optical Model

The various constituents of ocean water have a great influence in its optical properties. In order to solve the forward problem in ocean optics, the IOP have to be modeled and used in the FRTE. The values of these IOP can be obtained as the sum of the contributions of pure water and the dissolved particles and particulate matter present in the water, as proposed in [Mob94]. Optically pure water is devoid of any dissolved or suspended matter, and thus there is no scattering or absorption owed to particles or organic material [Mor74]. For *saline* pure water the salt concentration (35 to 39 parts per thousand) does influence the scattering and absorption functions. In particular it absorbs most wavelengths except for blue, with the absorption coefficient peaking at 760 nm, and reaching a minimum at 430 nm.

We develop our bio-optical model from three main IOP, with others like the extinction coefficient or the albedo derived from those three. These IOP are the absorption coefficient (3), the scattering coefficient (4) and the phase function (5), which for the elastic case can be written as (see Table 4 for a more detailed description of the functions used, including both the elastic and inelastic cases):

$$\alpha(\lambda) = \alpha_w(\lambda) + \sum_i \alpha_i(\lambda) \quad (3)$$

$$\sigma(\lambda) = \sigma_w(\lambda) + \sum_i \sigma_i(\lambda) \quad (4)$$

$$p(\lambda, \theta) = \frac{\sigma_w(\lambda)}{\sigma(\lambda)} p_w(\lambda, \theta) + \sum_i \frac{\sigma_i(\lambda)}{\sigma(\lambda)} p_i(\lambda, \theta) \quad (5)$$

where θ is the angle between the incoming $\vec{\omega}_i$ and outgoing $\vec{\omega}_o$ directions, the subscript w stands for the contribution of the pure water (fresh or salty) and the subscript i stands for the constituents in the water body such as biological particles or dissolved substances. We include three types of such constituents in our model, namely CDOM (Colored Dissolved Organic Matter, also known as yellow matter, present mainly in shallow ocean waters and harbors), phytoplankton (microscopic plants rich in chlorophyll) and minerals and organic detritus. The rest of this section will characterize the three main IOP (with elastic and inelastic scattering treated separately) for pure water and the three constituents. The next section will show how radiative transfer theory is applied to simulate the light field (which define the AOP) and render the final images.

3.1. Modeling Absorption

For the spectral absorption function of pure water $\alpha_w(\lambda)$ we rely on the work of Smith and Baker [SB81], whose tabulated values are well known in oceanography studies (shown in Table 1). Following further studies by Pope and

Fry [PF97], we use those values as an upper bound, to account for the fact that the true absorption can be, in fact, lower. The function shows that absorption is more prominent both in the UV and red ends of the spectrum. [PF97] also shows that absorption by salt in oceanic water is negligible. Based on the data by Bricaud, Morel and Prieur [BMP81], we model absorption by CDOM by fitting an exponential curve of the form:

$$\alpha_y(\lambda) = \alpha_y(\lambda_0) e^{-S_y(\lambda - \lambda_0)} \quad (6)$$

where the subscript y denotes the constituent CDOM. λ_0 is a reference wavelength, often chosen to be 440 nm for yellow matter, and S_y is the slope of the semilogarithmic absorption curve [Kir94]. S_y is usually taken to be constant, with a value of 0.014 nm^{-1} , but has been found to vary both geographically and temporally, and is also dependent on the wavelength range over which it is calculated [BMP81]. The values of absorption $\alpha_y(\lambda_0)$ at reference wavelengths also vary in a range between 0.01 m^{-1} to 20 m^{-1} , as a function of turbidity [Kir94].

Phytoplankton absorbs a great amount of visible light, due to its chlorophyll pigment. The absorption function for chlorophyll peaks strongly at 430 nm and 670 nm, being very weak in the mid range of the visible spectrum (thus the more phytoplankton the greener the hue of the water). The concentration of the chlorophyll in the water usually ranges from 0.01 mg/m^3 for open waters to 100 mg/m^3 . The spectral absorption coefficient of the phytoplankton is usually expressed as a function of this concentration C as:

$$\alpha_p(\lambda) = C \alpha_p^*(\lambda) \quad (7)$$

where C can be defined as the concentration of the main pigment chlorophyll- a (Chl $_a$) or as the sum of the concentrations of Chl $_a$ and its degradation products, the pheopigments. α_p^* is the *specific spectral absorption coefficient* (the absorption per unit of concentration) for a particular species of phytoplankton, given in m^2/mg . Typical values for specific absorptions of different species of phytoplankton can be found in the work of Sathyendranath, Lazzara and Prieur [SLP87] (see Table 1). A rough correspondence between chlorophyll concentrations and several oceanic water types is given by Morel [Mor88]. The absorption owed to organic detritus and minerals can be approximated by an exponential function, according to Roesler, Perry and Carder [RPC89]:

$$\alpha_d(\lambda) = \alpha_d(\lambda_0) e^{-S_d(\lambda - \lambda_0)} \quad (8)$$

Here the reference wavelength 400 nm is selected for λ_0 and typical values for the exponent coefficient S_d will be in the range between 0.006 nm^{-1} to 0.014 nm^{-1} , although 0.011 nm^{-1} is the most common value [RPC89]. Further studies confirm that the absorption spectra of minerals and detritus is well described by an exponential function with an average slope S_d of 0.0123 nm^{-1} , with slightly lower values than predicted at wavelengths below 440 nm [BSF*03].

Table 1: Absorption coefficient for a clear water body α_w (after Smith and Baker [SB81]) and specific absorption coefficient for phytoplankton α_p^* (after Sathyendranath, Lazzara and Prieur [SLP87]).

λ	[nm]	380	440	500	550	610	670	720	780
α_w	[cm^{-1}]	0.00022	0.000145	0.000257	0.000638	0.00289	0.0043	0.01169	0.0236
α_p^*	[$m^2 \cdot mg^{-1}$]	0.025	0.035	0.02	0.01	0.007	0.015	0.001	0.0001

3.2. Modeling Elastic Scattering

For the pure water term we use the volume scattering function defined by the Einstein-Smoluchowski theory [Maz02], which models scattering at molecular level as small-scale fluctuations. Whilst usually Rayleigh's scattering is used instead, Einstein-Smoluchowski provides more accurate results, is well defined and imposes no overheads in the simulations. Its scattering coefficient and phase function are given by:

$$\sigma_w(\lambda) = 16.06 \beta_w(\lambda_0, 90^\circ) \left(\frac{\lambda_0}{\lambda} \right)^{4.32} \quad (9)$$

$$p_w(\theta) = 0.06225 \left(1 + 0.835 \cos^2 \theta \right) \quad (10)$$

Typical values for $\beta_w(\lambda_0, 90^\circ)$ for both fresh and saline pure water are given in [Mor74]. These values range from $14.1 \cdot 10^{-4} m^{-1}$ to $134.5 \cdot 10^{-4} m^{-1}$. All the scattering produced by CDOM has inelastic nature and thus will be described in next section.

Gordon and Morel [GM83] found that phytoplankton, even in small concentrations, also contribute to the total elastic scattering in the water. Its contribution is given by:

$$\sigma_p(\lambda) = \left(\frac{550}{\lambda} \right) 0.30 C^{0.62} \quad (11)$$

where the constant 0.30 is selected to fit the data collected from many types of waters. The actual upper bound for this constant has a value of 0.45 [GM83]. The phase function due to phytoplankton is given by an isotropic function ($p_p = 1/\pi$).

The elastic scattering caused by organic detritus and minerals can be modeled based on Mie theory [GSO03]. The Henyey-Greenstein phase function models forward scattering fairly well but fails to reproduce backscattering with the same precision. We found that we can achieve a better fit by using a Two-Terms Henyey-Greenstein phase function (TTHG) [HG41]:

$$p_d(\theta, \zeta, g_f, g_b) = \zeta p_{HG}(\theta, g_f) + (1 - \zeta) p_{HG}(\theta, g_b) \quad (12)$$

where ζ is a weighting function between zero and one. This common way of utilizing this combination defines a forward scattering lobe (first term), plus a backscattering lobe (second term), with $g_f \in [0..1]$ and $g_b \in [-1..0]$. p_{HG} represents a simple Henyey-Greenstein phase function (HG):

$$p_{HG}(\theta, g) = \frac{1 - g^2}{(1 + g^2 - 2g \cos \theta)^{3/2}} \quad (13)$$

The TTHG function not only models backscattering more precisely, but it can describe more complex particle scattering models, improving the fit at large and small angles as well. The shape of each of the two HG functions can be approximated by an ellipsoid, avoiding the relatively expensive exponent in its evaluation. The observation was first introduced by Schlick [BLSS93]. Due to the great variety of particulate matter, the scattering coefficient σ_d can adopt a wide range of values. Table 2 shows typical values of this function (data after Stramski et al. [SBM01]).

3.3. Modeling Inelastic Scattering

For inelastic scattering, we need to model the possibility of an absorbed photon being re-emitted at a different wavelength. (2) includes a term $f(\lambda', \lambda)$ known as *wavelength redistribution function*, which represents the efficiency of the energy transfer between wavelengths. It is defined as the quotient between the energy of the emitted wavelength and the energy of the absorbed excitation wavelength, per wavelength unit. Reformulating in terms of photons instead of energy we have the *spectral quantum efficiency function* $\eta(\lambda', \lambda)$, defined as the ratio between the number of photons emitted at λ per wavelength unit, and the number of absorbed photons at λ' . Both functions are dimensional (nm^{-1}), and are related as follows:

$$f(\lambda', \lambda) = \eta(\lambda', \lambda) \frac{\lambda'}{\lambda} \quad (14)$$

The wavelength redistribution function f , and therefore its associated spectral quantum efficiency function η , can be seen as a re-radiation matrix. A related dimensionless function that describes inelastic scattering is the *quantum yield* $\Gamma(\lambda')$, defined as the total number of photons emitted at all wavelengths divided by the number of photons absorbed at excitation wavelength λ' . It is related to the spectral quantum efficiency function by:

$$\Gamma(\lambda') = \int_w \eta(\lambda', \lambda) d\lambda \quad (15)$$

The three functions $\Gamma(\lambda')$, $f(\lambda', \lambda)$ and $\eta(\lambda', \lambda)$, depend on both the medium and the type of inelastic event. The two inelastic events with more influence in the in-water light field are fluorescence and Raman scattering. Phytoplankton and CDOM are important fluorescence sources, whilst Raman scattering is produced by pure water; minerals and detritus, on the other hand, do not produce any inelastic event.

Table 2: Scattering coefficient for detritus σ_{dt} and minerals σ_m (After Stramski et al. [SBM01]).

	λ	[nm]	380	440	500	550	610	670	720	780
detritus	σ_{dt}	[m^{-1}]	0.045	0.0375	0.0325	0.03	0.0285	0.0275	0.027	0.027
minerals	σ_m	[m^{-1}]	0.0675	0.0525	0.05	0.045	0.04	0.036	0.034	0.032
total	σ_d	[m^{-1}]	0.1125	0.09	0.0825	0.075	0.0685	0.0635	0.061	0.059

3.3.1. Fluorescence

Fluorescence occurs when a molecule absorbs a photon of wavelength λ' , and re-emits it at a longer wavelength λ according to the *fluorescence efficiency function* $\eta_F(\lambda', \lambda)$. For the two main sources of fluorescence (phytoplankton and CDOM), re-emission follows an isotropic phase function. For phytoplankton, the wavelength of the re-emitted photons is independent of the excitation wavelength, although the intensity does show wavelength dependency [Mob94].

It is very common in ocean waters to see a color shift ranging from greenish to very bright green, or even yellowish. These hue shifts are mainly due to the variation in the concentration and type of the suspended microorganisms, specially phytoplankton and its related chlorophyll concentration, which presents an absorption function peaking at 350 nm and rapidly decaying to almost zero beyond 500 nm. Only wavelengths between 370 and 690 nm can trigger fluorescence due to phytoplankton. This can be modeled as a dimensionless function $g_p(\lambda')$ so that:

$$g_p(\lambda') \equiv \begin{cases} 1 & \text{if } 370 \leq \lambda' \leq 690 \text{ nm} \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

The wavelength-independent quantum yield for phytoplankton $\Gamma_p(\lambda')$ ranges from 0.01 to 0.1. Using (14) and (16), the relationship between the wavelength redistribution function $f_p(\lambda', \lambda)$ and the spectral quantum efficiency function $\eta_p(\lambda', \lambda)$ is:

$$f_p(\lambda', \lambda) = \eta_p(\lambda', \lambda) \frac{\lambda'}{\lambda} \equiv \Gamma_p g_p(\lambda') h_p(\lambda) \frac{\lambda'}{\lambda} \quad (17)$$

where $h_p(\lambda)$ is the *fluorescence emission function* per unit wavelength, and can be approximated by a gaussian [Mob94]:

$$h_p(\lambda) = \frac{1}{\sqrt{2\pi}\lambda_\sigma} \exp\left\{-\frac{(\lambda - \lambda_0)^2}{2(\lambda_\sigma)^2}\right\} \quad (18)$$

$\lambda_0 = 685 \text{ nm}$ is the wavelength of maximum emission and $\lambda_\sigma = 10.6 \text{ nm}$ represents the standard deviation. Using (7) and (17) we can now compute the inelastic scattering coefficient owed to phytoplankton $\sigma_p(\lambda', \lambda)$ following (2).

The other important source of fluorescence in water is CDOM. For relatively high concentrations of CDOM, its quantum yield $\Gamma_y(\lambda')$ varies between 0.005 and 0.025. Following the work of Hawes [Haw92] we use the following formula to describe its spectral fluorescence quantum efficiency function:

Table 3: Water constituents and interactions

Constituent	Absorption	Elastic Scat.	Inelastic Scat.
Pure water (w)	Yes	Yes	Raman Scattering
Minerals, detritus (d)	Yes	Yes	No
Phytoplankton (p)	Yes	Yes	Fluorescence
CDOM (y)	Yes	No	Fluorescence

$$f_y(\lambda', \lambda) = A_0(\lambda') \exp\left\{-\left(\frac{\frac{1}{\lambda} - \frac{A_1}{\lambda'} - B_1}{0.6\left(\frac{A_2}{\lambda^2} + B_2\right)}\right)^2\right\} \frac{\lambda'}{\lambda} \quad (19)$$

where A_0, A_1, A_2, B_1 and B_2 are empirical parameters whose values depend on the specific composition of the CDOM and can be found in [Mob94] (see Table 5). A_1 and A_2 are dimensionless, whereas the rest are given in nm^{-1} . Like fluorescence due to phytoplankton, we can use (6) and (19) to compute the inelastic scattering coefficient $\sigma_y(\lambda', \lambda)$ following (2).

Our model can be easily extended to account for phosphorescence phenomena, which are intrinsically similar to fluorescence and are governed by the *phosphorescence efficiency function*. The only difference is that the re-emitted energy declines with time according to a function $d(t)$.

3.3.2. Raman scattering

Raman scattering influences the in-water light field, specially at great depths where sun irradiance becomes zero and only Raman radiance remains. It occurs when vibration and rotation in water molecules exchange energy with incoming photons, re-emitting them with approximately the same wavelength, but allowing for small shifts towards longer or shorter wavelengths. It can also be considered a spontaneous process. To isolate Raman inelastic events from fluorescence and other scattering events, it is usually studied in pure water, filtered several times, so that the second term in (4) becomes zero.

The Raman wavelength redistribution function $f_w(\lambda', \lambda)$ is usually described in terms of a sum of four Gaussian functions [Mob94]:

$$f_w(\lambda', \lambda) = \frac{10^7 \sum_{j=1}^4 A_j \frac{1}{\Delta \tilde{\nu}_j} \exp\left\{-\frac{\left[10^7\left(\frac{1}{\lambda'} - \frac{1}{\lambda}\right) - \tilde{\nu}_j\right]^2}{\Delta \tilde{\nu}_j^2}\right\}}{\lambda'^2 \sqrt{\frac{\pi}{4 \ln 2}} \sum_{j=1}^4 A_j} \quad (20)$$

where $\tilde{\nu}$ is the wavenumber ($\tilde{\nu} = 10^7/\lambda$) given in cm^{-1} . Typical parameter values $A_j, \tilde{\nu}_j$ and $\Delta \tilde{\nu}_j$ for the Raman redistribution function are given by Walrafen [Wal69] and are

shown in Table 5. The inelastic scattering coefficient can now be obtained using α_w and f_w in (2).

4. The simulation method

Having so far developed our bio-optical model, we can now formalize it into a set of parameters and equations to fully simulate the in-water light field. To summarize, the four constituents of the model and their interactions with light are given in Table 3. Table 4 shows how the main functions that define the model are derived from IOP and related functions at constituent level.

Table 4: The main functions of the model

Equations	
$\alpha(\lambda)$	$= \alpha_d(\lambda) + \alpha_p(\lambda) + \alpha_w(\lambda) + \alpha_y(\lambda)$
$\sigma(\lambda)$	$= \sigma_w(\lambda) + \sigma_d(\lambda) + \sigma_p(\lambda)$
$p(\lambda, \theta)$	$= \frac{\sigma_w(\lambda)p_w(\lambda, \theta) + \sigma_d(\lambda)p_d(\lambda, \theta) + \sigma_p(\lambda)p_p(\lambda, \theta)}{\sigma(\lambda)}$
$\kappa(\lambda)$	$= \alpha(\lambda) + \sigma(\lambda)$
$\alpha_I(\lambda')$	$= \alpha_p(\lambda') + \alpha_w(\lambda') + \alpha_y(\lambda')$
$p_I(\lambda', \lambda, \theta)$	$= \frac{\alpha_p(\lambda')p_p(\lambda', \lambda, \theta) + \alpha_w(\lambda')p_w(\lambda', \lambda, \theta) + \alpha_y(\lambda')p_y(\lambda', \lambda, \theta)}{\alpha_I(\lambda')}$
$f_I(\lambda', \lambda)$	$= \frac{\alpha_p(\lambda')f_p(\lambda', \lambda) + \alpha_w(\lambda')f_w(\lambda', \lambda) + \alpha_y(\lambda')f_y(\lambda', \lambda)}{\alpha_I(\lambda')}$

Table 5: Parameters of the model

Parameter	Equations	Simulated values	Units
C	(7) (11)	[0..1.0]	$\frac{mg}{m^3}$
$\alpha_d(400)$	(8)	[0..0.1]	m^{-1}
$\alpha_y(440)$	(6)	[0..0.1]	m^{-1}
S_y	(6)	0.014	nm^{-1}
S_d	(8)	0.011	nm^{-1}
A_0	(19)	$\frac{150}{700}$	nm^{-1}
A_1	(19)	4	-
A_2	(19)	4	-
B_0	(19)	$\frac{1}{450 \cdot 10^{-7}}$	nm^{-1}
B_1	(19)	$\frac{1}{650 \cdot 10^{-7}}$	nm^{-1}
Γ_p	(17)	0.1	-
Γ_y	(19)	0.025	-
$A_i, i = 1..4$	(20)	0.41, 0.39, 0.10, 0.10	-
$\tilde{v}_i, i = 1..4$	(20)	3250, 3425, 3530, 3625	-
$\Delta\tilde{v}_i, i = 1..4$	(20)	210, 175, 140, 140	-

The model allows for easy adjusting of its parameters to simulate different types of water and thus obtain different in-water light fields. As well as minerals and detritus, other particulate components of water can be added from oceanographic studies (although minerals and detritus have the greatest influence in the final appearance of water). Mie theory can again be used to model the scattering by these new particles, and the phase function can be approximated by using a Two Terms Henyey-Greenstein phase function (12). An overview of the most significant parameters of the model, the equations in which they can be found and the corresponding values used for the simulations in this paper can be found in Table 5. Note that for simplicity we have not included the values that are already specified throughout the text during the explanation of the bio-optical model (more

specifically, those included in tables 1 and 2). The first three correspond to the parameters analyzed in Figure 2.

Once we have formalized the model into a set of equations, we rely on radiative transfer theory to obtain a solution for the in-water light field. We solve the Full Radiative Transfer Equation (1) by extending the traditional photon mapping algorithm [Jen01] by taking into account all ten different events specified in Table 3, while allowing for both Stokes or anti-Stokes inelastic scattering. This enhancement is done in both stages: photon tracing and radiance estimation.

During the photon tracing stage in the original photon mapping method [Jen01], a Russian roulette algorithm is triggered at each interaction with the medium, deciding whether the photon is scattered or absorbed. In [GMAS05] the authors add a second Russian roulette which separates absorption from inelastic scattering; in the latter case, a new photon is generated at a different wavelength, but the algorithm considers just a single type of inelastic event with Stokes behavior. No anti-Stokes events are simulated. In contrast, our method uses just a single Russian roulette to choose between ten different kinds of interactions (including three types of inelastic events where the photons may gain or lose energy), and can be easily extended to handle an arbitrary number of different interactions. Finally, we improve the radiance estimation stage over previous methods by adding a term to take into account the contributions from the inelastic scattering events. The next subsections present the algorithm in more detail.

4.1. Stage 1: Photon tracing

We shoot photons from the light sources and let them interact with the geometry and the medium according to its optical distance, which is a function of the extinction coefficient (as in the original photon mapping method). We statistically decide at each interaction which type of event occurs (refer to Table 3) with just a single Russian roulette. At the interactions, photons are stored in a kd-tree as in traditional photon mapping.

The wavelength spectrum is box sampled into N_λ samples, so absorption ($\alpha(\lambda)$) and scattering coefficients ($\sigma(\lambda)$) are implemented as N_λ -dimensional arrays while wavelength redistribution functions ($f(\lambda', \lambda)$) are implemented as $N_\lambda \times N_\lambda$ square matrices. Each of the photons carries information about a portion of flux ($\Delta\Phi$) at a certain sampled wavelength (λ'). Importance sampling is used for computing the optical distance, so $\Delta\Phi$ does not change along the photon tracing stage, while λ' changes for inelastic scattering events.

In order to apply the Russian roulette algorithm, we will define an albedo $\Lambda_j(\lambda)$ for each interaction j as follows:

- If interaction j represents an elastic scattering event, then $\Lambda_j(\lambda) = \frac{\sigma_j(\lambda)}{\kappa(\lambda)}$
- If j represents an absorption interaction that does not

show inelastic scattering (detritus and minerals, basically), then $\Lambda_j(\lambda) = \frac{\alpha_j(\lambda)}{\kappa(\lambda)}$

- For each absorption interaction that could generate inelastic scattering (pure water, phytoplankton and CDOM) we define its *inelastic probability* (χ_j), the probability that an absorption event generates an inelastic scattering event:

$$\chi_j(\lambda') = \int_{\lambda_a}^{\lambda_b} f_I(\lambda', \lambda) d\lambda \approx \sum_{i=1}^{N_\lambda} f_I(\lambda', \lambda_i) \quad (21)$$

where λ_a and λ_b are the lower and upper limits of the simulated wavelengths, and $i \in [1..N_\lambda]$ refer to samples in wavelength domain:

- If interaction j represents the effective inelastic scattering event within the absorption interaction: $\Lambda_j(\lambda) = \frac{\alpha_j(\lambda)}{\kappa(\lambda)} \chi_j(\lambda)$
- If interaction j represents the pure absorption event (no inelastic scattering happening at all): $\Lambda_j(\lambda) = \frac{\alpha_j(\lambda)}{\kappa(\lambda)} (1 - \chi_j(\lambda))$

Thus, at each interaction a random number ξ between 0 and 1 is generated resulting in (between parenthesis, example values of Λ_j at $\lambda = 500nm$ that determine the size of the corresponding interval are included):

- $\xi \in [0, \xi_1) \rightarrow$ absorption by pure water ($2.51 \cdot 10^{-1}$).
- $\xi \in [\xi_1, \xi_2) \rightarrow$ Raman scattering, inelastic scattering by pure water ($1.21 \cdot 10^{-9}$).
- $\xi \in [\xi_2, \xi_3) \rightarrow$ absorption by minerals and detritus ($7.12 \cdot 10^{-2}$).
- $\xi \in [\xi_3, \xi_4) \rightarrow$ absorption by phytoplankton ($4.90 \cdot 10^{-3}$).
- $\xi \in [\xi_4, \xi_5) \rightarrow$ inelastic scattering by phytoplankton ($2.18 \cdot 10^{-3}$).
- $\xi \in [\xi_5, \xi_6) \rightarrow$ absorption by CDOM ($7.83 \cdot 10^{-2}$).
- $\xi \in [\xi_6, \xi_7) \rightarrow$ inelastic scattering by CDOM ($1.21 \cdot 10^{-2}$).
- $\xi \in [\xi_7, \xi_8) \rightarrow$ elastic scattering by pure water ($7.44 \cdot 10^{-3}$).
- $\xi \in [\xi_8, \xi_9) \rightarrow$ elastic scattering by minerals and detritus ($2.94 \cdot 10^{-1}$).
- $\xi \in [\xi_9, 1] \rightarrow$ elastic scattering by phytoplankton ($2.79 \cdot 10^{-1}$).

where $\xi_i(\lambda)$ is given by $\xi_i(\lambda) = \sum_{j=1}^i \Lambda_j(\lambda)$

To compute the new re-emitted wavelength after a inelastic scattering event i , the normalized wavelength redistribution function $\frac{f_i(\lambda', \lambda)}{\chi_i(\lambda')}$ is treated as a probability distribution function (PDF) given the excitation wavelength λ' . To sample it efficiently we first build its normalized cumulative distribution function (CDF) and then inverse importance sample this CDF. Greater values of the PDF for a given wavelength will translate to steeper areas of the CDF, thus increasing the probability of a re-emission at such wavelength. Note that the definition of $f_i(\lambda', \lambda)$ is not limited to the visible spectrum, which might result in re-emissions happening at wavelengths beyond the visible spectrum. However, as $\chi_i(\lambda')$ is limited to the simulated (visible) spectrum, only inelastic interactions within this spectrum are considered. It could happen that a photon inelastically scattered at such

wavelengths suffers a second inelastic scattering event that brings it back to the visible light range. Given the low probability of this chain of events and our computer graphics approach, we assume that a photon beyond the visible spectrum is definitely absorbed. Figure 1 shows a global overview of the algorithm during the photon tracing stage.




Figure 1: Photon tracing algorithm. Inelastic scattering events generate a photon with a different associated wavelength according to the wavelength redistribution function.

4.2. Stage 2: radiance estimate

To estimate radiance we adopt a tradeoff between speed and memory requirements similar to the proposed by Jensen and Christensen [JC98]: we only store photons in the photon map if they have been reflected or transmitted from surfaces, or if they have already been scattered at least once. Thus, we can compute single scattering more efficiently by ray marching through the medium and sampling the light sources by casting shadow rays. Taking into account the wavelength redistribution function for inelastic scattering, a new addend will be added at each step of the ray marching process:

$$\sum_{l=1}^N \sum_{i=1}^{N_\lambda} \left\{ L_l(\lambda'_i, \vec{w}_l) p_l(\lambda'_i, \lambda, \vec{w}_l, \vec{w}_o) \alpha_l(\lambda'_i) f_l(\lambda'_i, \lambda) \Delta x \right\} \quad (22)$$

where $i \in [1..N_\lambda]$ and $l \in [1..N]$ refer to samples in the wavelength and light source domain respectively, \vec{w}_l is the direction to the light with an incoming radiance L_l and Δx represent the ray marching steps.

Multiple scattering will be computed from the photon map, finding in the kd-tree the n photons which are closest to the estimation point by using the typical nearest neighbours algorithm. To account for multiple *inelastic* scattering we modify the radiance estimate expression of [JC98] by including a new term:

$$\sum_{k=1}^n \left\{ p_l(\lambda'_k, \lambda, \vec{w}_k, \vec{w}_o) f_l(\lambda'_k, \lambda) \frac{\Delta \Phi_k}{\frac{4}{3} \pi r^2} \right\} \quad (23)$$

where r is the radius of the sphere that contains the n closest photons, and k represents each of the stored photons.

5. Results

We have used the values from Table 5 for our simulations. In the images produced we only vary the chlorophyll concentration C , minerals and detritus turbidity $\alpha_d(400)$ and



Figure 2: Resulting pictures varying the chlorophyll concentration C , the minerals and detritus turbidity α_d at 400nm and the CDOM turbidity α_y at 440nm.

CDOM turbidity $\alpha_y(440)$. The choice of those three parameters to reduce the dimensionality of the model was based on their greater overall influence on the resulting light field. The photon map contains 400000 photons, with 250 used in the estimation of radiance. Ray-marching depth is set at 200 steps. Each of the images has been rendered in a Dual Xeon Pentium 4 at 2.8GHz with 2GB RAM at 512×384 resolution, casting one ray per pixel, and took approximately 20 minutes to render. This time is roughly independent of the number of parameters of the bio-optical model. In order to reduce these computation times, several optimization techniques could be adopted, like using adaptive ray-marching or radiance caching strategies [JDZJ08]. Additionally, perceptual issues could be taken into account, using just an approximate solution in areas of the image where the error is known to be perceptually negligible [SGA*07].

Energy balances show that on average almost 99% of the energy emitted by the light sources is absorbed after just a few interactions of the photons, with very incremental variation after the fourth interaction and negligible contribution after the fifth. This relatively fast convergence is due to the strong absorption in water. We have therefore limited the number of interactions per photon to five, in order to speed up the simulations. Variations of the parameters C , $\alpha_d(400)$ and $\alpha_y(440)$ yield different probabilities for absorption, elastic and inelastic scattering events, which in turn

affect the in-water light field. The results can be seen in Figure 2, with each of the varying parameters influencing the final light field as follows:

- Chlorophyll concentration (C) affects mainly both elastic and inelastic scattering. The effects of inelastic scattering are mostly masked by the more predominant elastic scattering and absorption, which increases slowly. The third column in Figure 2 shows brighter images than the previous two due to in-scattering. For higher values (fourth column), out-scattering prevails and the images become darker.
- Minerals and detritus turbidity ($\alpha_d(400)$) increases absorption at lower wavelengths, thus reducing the brightness of the scene and the overall blue hue. Scattering is also increased, making the images appear murkier. Figure 2 shows variations of the minerals and detritus turbidity between the first and second rows for direct comparison.
- CDOM turbidity ($\alpha_y(440)$) slightly increases absorption (darker images) and introduces inelastic scattering (change in hue). This can be seen by comparing the first and third rows in Figure 2.

We have undergone a visual validation of our model by rendering different natural waters. Figure 3 shows the resulting underwater images for Atlantic, Mediterranean, Baltic, North Sea and shallow coastal waters rich in CDOM respectively. All the images have been simulated at the same depth



Figure 3: Rendered images of different waters. From left to right: Atlantic, Mediterranean, Baltic, North Sea and shallow coastal waters rich in CDOM. Smaller patches below for comparison purposes by Frisvad et al. [FCJ07] (used with permission).

and are illuminated by the same isotropic point light source. The changes in color are clearly noticeable, from a darker blue in the case of Atlantic water, to the greener hue in the image of the North Sea. The smaller patches below the first four images correspond to the simulations by Frisvad et al. [FCJ07] for the same types of water, and are shown for comparison purposes. Our simulations based on radiative transfer approximately match their simulations based on Lorenz-Mie theory. The differences are mainly owed to two factors: on the one hand, the overall darker tone in our images is due to in-water absorption, whereas [FCJ07] renders the *surface* of the water body; on the other hand, the absence of inelastic scattering effects in [FCJ07] can have a visible influence the final appearance of water, as shown in Figure 4 for the Baltic case. The properties of the water have been adjusted according to measurements found in [BSF*03] [Mob94] for our bio-optical model and [BSF*03] in the model by Frisvad et al. In both cases, it is only the changes in the constituents of the waters which yield the different colors. We have additionally performed a numerical analysis of the in-water radiance field, to quantify the influence of each constituent. The results can be seen in Figure 5.

6. Conclusion

We have presented a complete bio-optical model of ocean water based on parameterizing its intrinsic optical properties. Relying on radiative transfer theory, we obtain the resulting in-water light field by extending the rendering algorithm presented in [GMAS05]. The extension can now handle more complex interactions between light and water, including inelastic scattering with anti-Stokes behavior, where the scattered photon absorbs energy from the medium and is re-emitted at *higher* energies. We have additionally studied the influence of the parameters in the apparent optical properties of water in the scene, which are defined by the light field obtained. We have performed an energy-balance analysis, and visual validation of the method has been provided by direct comparison with images by Frisvad et al. [FCJ07], rendering different types of waters based on published constituent data.

We have included Raman scattering by pure water and fluorescence by phytoplankton and CDOM as inelastic scattering events with energy transfers. Even though their combined quantitative contribution to the overall radiance field

is usually less than 2% (see Figure 5), this relatively small percentage does have a clear influence on the apparent optical properties, as Figure 4 shows. We thus argue that these events, usually overlooked in computer graphics literature, are qualitatively important for underwater imagery and should be included in a complete simulation. Other types of inelastic scattering such as *Compton*, *Bragg* or *Brillouin* could also be added, although their influence is more incremental. Other particulate elements could be easily added as well just by including their corresponding absorption and scattering coefficients in the model; however, the three constituents treated here (phytoplankton, minerals and detritus and CDOM) have the most influence in the final radiance field.

The results show how the model developed can easily be used for physically-based simulations of underwater imagery. We believe this work can be of interest not only in the computer graphics community, but in remote sense or oceanographic studies as well.



Figure 4: The influence of inelastic scattering in the apparent optical properties of water (Baltic sea): Left, no inelastic scattering. Center, just chlorophyll inelastic scattering (as in [GMAS05]). Right, all inelastic scattering events included in the simulation.




Figure 5: Radiance distribution of the resulting in-water light field per type of event (Baltic Sea).

Acknowledgments

We thank the anonymous reviewers for their valuable comments. This research was partly done under the sponsorship of the Spanish Ministry of Education and Research through the projects TIN2004-07672-C03-03 and TIN2007-63025.

References

- [BLSS93] BLASI P., LE SAEC B., SCHLICK C.: A rendering algorithm for discrete volume density objects. *Computer Graphics Forum (Eurographics 93)* 12, 3 (1993), 201–210. 4

- [BMP81] BRICAUD A., MOREL A., PRIEUR L.: Absorption by dissolved organic matter of the sea (yellow substance) in the uv and visible domains. *Limnol. Oceanogr.* 26, 1 (1981), 43–53. 3
- [BSF*03] BABIN M., STRAMSKI D., FERRARI G. M., CLAUSTRE H., BRICAUD A., OBOLENSKY G., HOEPEFFNER N.: Variations in the light absorption coefficients of phytoplankton, non-algal particles, and dissolved organic matter in coastal waters around europe. *J. Geophys. Res.* 108(C7), 3211 (2003). 3, 9
- [Cha60] CHANDRASEKHAR S.: *Radiative Transfer*. Dover Publications, Inc., 1960. 2
- [CS04] CEREZO E., SERON F. J.: Rendering natural waters taking fluorescence into account: Research articles. *Comput. Animat. Virtual Worlds* 15, 5 (2004), 471–484. 2
- [FCJ07] FRISVAD J. R., CHRISTENSEN N. J., JENSEN H. W.: Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM Trans. Graph.* 26, 3 (2007), 60. 2, 9
- [Gla95] GLASSNER A. S.: *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995. 2
- [GM83] GORDON H. R., MOREL A.: *Remote Assessment of Ocean Color for Interpretation of Satellite Visible Imagery: A Review*, vol. 4 of *Lecture Notes on Coastal and Estuarine Studies*. Springer-Verlag, New York, 1983. 4
- [GMAS05] GUTIERREZ D., MUNOZ A., ANSON O., SERÓN F. J.: Non-linear volume photon mapping. In *Proc. of the Eurographics Symposium on Rendering Techniques, Konstanz, Germany, June 29 - July 1, 2005* (2005), pp. 291–300. 2, 6, 9
- [GSO03] GREEN R. E., SOSIK H. M., OLSON R. J.: Contributions of phytoplankton and other particles to inherent optical properties in new england continental shelf waters. *Limnol. Oceanogr.* 48, 6 (2003), 2377–2391. 4
- [Haw92] HAWES S.: *Quantum fluorescence efficiencies of marine fulvic and humid acids*. PhD thesis, Dept. of Marine Science, Univ. of South Florida, 1992. 5
- [HG41] HENYEA L., GREENSTEIN J.: Diffuse radiation in the galaxy. *Astrophysics Journal* 93 (1941), 70–83. 4
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *SIGGRAPH 98 Conference Proceedings* (jul 1998), Cohen M., (Ed.), Annual Conference Series, ACM SIGGRAPH, Addison Wesley, pp. 311–320. 7
- [JDZJ08] JAROSZ W., DONNER C., ZWICKER M., JENSEN H. W.: Radiance caching for participating media. *To appear in ACM Transactions on Graphics* (2008). 8
- [Jen01] JENSEN H. W.: *Realistic image synthesis using photon mapping*. A.K. Peters, Natick, Massachusetts, 2001. 2, 6
- [Kir94] KIRK J. T.: *Light and photosynthesis in aquatic ecosystems*. Cambridge University Press, New York, 1994. 3
- [KYNN91] KANEDA K., YUAN G., NAKAMAE E., NISHITA T.: Realistic visual simulation of water surfaces taking into account radiative transfer. In *Proc. of CAD/Graphics 91* (1991), pp. 25–30. 2
- [Maz02] MAZO R. M.: *Brownian Motion: Fluctuations, Dynamics and Applications*, vol. 112 of *International Series of Monographs on Physics*. Oxford University Press Inc., Great Clarendon Street, Oxford, 2002, ch. Einstein-Smoluchowski Theory, pp. 46–62. 4
- [Mob94] MOBLEY C. D.: *Light and Water: Radiative Transfer in Natural Waters*. Academic Press, Inc., San Diego, 1994. 1, 2, 3, 5, 9
- [Mor74] MOREL A.: *Optical Aspects of Oceanography*. Academic Press, New York, 1974. 3, 4
- [Mor88] MOREL A.: Optical modeling of the upper ocean in relation to its biogenous matter content (case i waters). *Journal of Geophysical Research* 93, C9 (1988), 10749–10768. 3
- [NKON90] NAKAMAE E., KANEDA K., OKAMOTO T., NISHITA T.: A lighting model aiming at drive simulators. *Computer Graphics* 24, 4 (Aug. 1990), 395–404. 2
- [NSTN93] NISHITA T., SHIRAI T., TADAMURA K., NAKAMAE E.: Display of the earth taking into account atmosphere scattering. In *Computer Graphics (SIGGRAPH '93 Proceedings)* (1993), vol. 24, pp. 175–182. 2
- [PA01] PREMOZE S., ASHIKHMIN M.: Rendering natural waters. *Comput. Graph. Forum* 20, 4 (2001), 189–199. 2
- [PF97] POPE R. M., FRY E. S.: Absorption spectrum (380–700 nm) of pure water. ii. integrating cavity measurements. *Applied Optics* 36, 33 (1997), 8710–8723. 3
- [Pre76] PREISENDORFER R. W.: *Introduction*, vol. 1 of *Hydrologic Optics*. National Technical Information Service, Springfield, IL, 1976. 1
- [RPC89] ROESLER C. S., PERRY M. J., CARDER K. L.: Modeling in situ phytoplankton absorption from total absorption spectra in productive inland marine waters. *Limnol. Oceanogr.* 34, 8 (1989), 1510–1523. 3
- [RT87] RUSHMEIER H. E., TORRANCE K. E.: The zonal method for calculating light intensities in the presence of a participating medium. *Computer Graphics* 21, 4 (July 1987), 293–302. 2
- [SB81] SMITH R. C., BAKER K. S.: Optical properties of the clearest natural waters (200–800 nm). *Appl. Opt.* 20 (1981), 177–184. 3, 4
- [SBM01] STRAMSKI D., BRICAUD A., MOREL A.: Modeling the inherent optical properties of the ocean based on the detailed composition of planktonic community. *Applied Optics* 40 (2001), 2929–2945. 4, 5
- [SCP94] SPINRAD R. W., CARDER K. L., PERRY M. J. (Eds.): *Ocean Optics*. No. 25 in Oxford Monographs on Geology and Geophysics. Oxford University Press, 1994. 2
- [SGA*07] SUNDSTEDT V., GUTIERREZ D., ANSON O., BANTERLE F., CHALMERS A.: Perceptual rendering of participating media. *ACM Transactions of Applied Perception* 4, 3 (2007). 8
- [SLP87] SATHYENDRANATH S., LAZZARA L., PRIEUR L.: Variations in the spectral values of specific absorption of phytoplankton. *Limnol. Oceanogr.* 32, 2 (1987), 403–415. 3, 4
- [TN95] TADAMURA K., NAKAMAE E.: *Computer Graphics: Developments in Virtual Environments*. Academic Press, 1995, ch. Modeling the colour of Water in Lightning Design, pp. 97–114. 2
- [Wal69] WALRAFEN G. E.: Continuum model of water—an erroneous interpretation. *Journal of Chemical Physics* 50, 1 (January 1969), 567–569. 5

Structured Light in Scattering Media*

Srinivasa G. Narasimhan⁺, Shree K. Nayar⁺⁺, Bo Sun⁺⁺ and Sanjeev J. Koppal⁺

⁺Robotics Institute, Carnegie Mellon University, Pittsburgh, USA

⁺⁺Computer Science Department, Columbia University, New York, USA

Abstract

Virtually all structured light methods assume that the scene and the sources are immersed in pure air and that light is neither scattered nor absorbed. Recently, however, structured lighting has found growing application in underwater and aerial imaging, where scattering effects cannot be ignored. In this paper, we present a comprehensive analysis of two representative methods - light stripe range scanning and photometric stereo - in the presence of scattering. For both methods, we derive physical models for the appearances of a surface immersed in a scattering medium. Based on these models, we present results on (a) the condition for object detectability in light striping and (b) the number of sources required for photometric stereo. In both cases, we demonstrate that while traditional methods fail when scattering is significant, our methods accurately recover the scene (depths, normals, albedos) as well as the properties of the medium. These results are in turn used to restore the appearances of scenes as if they were captured in clear air. Although we have focused on light striping and photometric stereo, our approach can also be extended to other methods such as grid coding, gated and active polarization imaging.

1 Introduction

Structured lighting has received wide attention since the early work of Will and Pennington [28]. Virtually all structured lighting techniques modulate the appearance of a surface (or volume) by projecting a particular pattern of light onto it [15]. Often this modulation makes subsequent processing of acquired images simple; the correspondence problem and issues related to calibration are often alleviated [6; 13; 30] and accurate 3D reconstruction is obtained irrespective of the complex surface texture [20; 21]. As a result, structured lighting has been a key enabling technology for several industrial applications such as assembly, alignment and inspection.

An implicit assumption made in most structured light methods is that light is neither scattered nor absorbed by the medium in which the scene and sources are immersed (as in pure air). This assumption, however, is violated in several media including atmospheric conditions (fog, haze, mist), fluids (murky water, milk) and smoke. An optical image taken in such a medium suffers from significant loss of contrast and attenuation of brightness. Further, this degradation increases exponentially with distance, making it hard to capture meaningful images of scenes that are not near the imaging system. Thus, it is critical to

take into account the effects of scattering while applying structured light methods in such media.

But why use structured light in scattering media at all? Consider, for example, underwater optical imaging. Because of scattering by impurities in water, natural sources such as daylight attenuate completely before reaching significant depths. So, deep underwater (semi-)autonomous vehicles rely on active illumination to explore wreckages¹, find mines, and inspect vessels, docks and rigs, thus playing a key role in maintenance, construction and rescue scenarios. A variety of structured lighting techniques have been developed for these applications, ranging from using thin laser beams [14], to using a number of carefully spaced confocal sources [17], to more sophisticated time-gated [19] and synchronization-gated techniques [7]. While these methods enhance visibility (resulting in better detection of targets), they do not explicitly analyze the appearances of scenes under structured light².

In this paper, we are interested in both enhancing visibility using structured light and also in analyzing the acquired images to recover properties of the scene and the medium. To achieve this, three relevant questions must be addressed. First, what are the scattering effects that result from the interaction of structured light with the medium and the scene? Second, how do we overcome these scattering effects to obtain the results that the structured light methods were traditionally designed for?³ Third, is there additional information that one can extract from these scattering effects that is not possible to obtain using the traditional methods?

We address these questions specifically for two representative techniques - light stripe range scanning and photometric stereo. For each of these, we derive an analytic image formation model that is based on the physics of single scattering. These models describe the interactions of structured light with the medium and the scene. Using the image formation model for light striping, we develop a simple algorithm to reliably detect objects and obtain a 3D reconstruction of the scene in the presence of strong scattering. Based on the image formation model for photometric stereo, we conclude that at least five light source directions (instead of the usual three) are required to reconstruct surface normals and albedos of a lambertian

¹Famously, the 1985 Franco-American expedition discovered and explored the remains of the *Titanic* that sank in 1912 [1].

²Work has also been done on a related but different problem of analyzing the appearances of scenes in scattering media (underwater or the atmosphere) using passive methods [5; 24; 26; 23; 18; 22] that rely on natural illumination external to the medium.

³Note that works that address this question (for instance, laser scanning to reconstruct sea floors [3; 16]) by ignoring scattering effects (first question) are severely limited in their applicability.

*This work is supported by an ONR contract #N00014-05-1-0188. The authors also thank Estuardo Rodas for building the tank used for the experiments in this paper.

object. Interestingly, our method also yields a depth map of the scene, which is not possible using traditional photometric stereo. Further, in both techniques, the interaction of structured light with the medium allows us to estimate the properties of the medium. This result can in turn be used to remove the effects of scattering and compute the appearance of the scene as if seen in clear air.

To verify our methods using real experiments, we have constructed a setup that consists of a glass tank filled with a scattering medium (dilute milk), with a projector (source) and camera placed outside the tank. Note that calibration of this setup requires us to handle light refraction at the medium-glass-air interfaces. We present a calibration procedure that is similar in spirit to [11] and that does not require either explicit geometric calibration of the camera and the projector or the knowledge of refraction locations or refractive indices of media. Although we have focused on light striping and photometric stereo, our results can be used to extend several other techniques such as grid coding [28] and gated [7] and active polarization imaging [25; 10]. We believe that our results can significantly benefit a wide range of underwater [14], aerial and microscopic imaging [9] applications.

2 Single Scattering in Media

In order to keep our techniques tractable, we assume that the scattering medium is homogeneous and not highly dense (for example, murky water, light fog, mist, dilute milk). This allows us to develop simple models based on single scattering. We now define the properties of scattering media [4] and present the single scattering model.

The scattering coefficient β is defined as the fraction of the incident flux scattered by a unit volume of the medium in all directions⁴. The phase function $\mathcal{P}(\alpha)$ defines the angular scattering distribution of the incident flux, where α is the angle between incident and scattered directions. In general, the phase function is smooth and may be represented by a low-order polynomial of $\cos \alpha$ [4]. We use the first-order approximation as given in [4],

$$\mathcal{P}(g, \alpha) = (1/4\pi) (1 + g \cos \alpha), \quad (1)$$

where, $g \in (-1, 1)$ is a parameter that controls the shape of the phase function. Now consider a light ray with radiance L_0 that travels a distance x , gets scattered by a particle at an angle α , before it further travels a distance y to reach the viewer. The intensity of this light ray is attenuated exponentially according to the total distance traveled. Then, the single scattering irradiance at the viewer is given by [4],

$$E_{med} = L_0 \beta \mathcal{P}(g, \alpha) e^{-\beta(x+y)}. \quad (2)$$

For an isotropic point source with radiant intensity I_0 , we may further write $L_0 = I_0/x^2$, while for a collimated beam, L_0 is constant with respect to x . We build upon equations 1 and 2 to derive image formation models for light striping and photometric stereo.

⁴When absorption is present, the scattering coefficient is replaced by the extinction coefficient defined as the fraction of incident flux scattered and absorbed by a unit volume.




Figure 1: **Light striping in media.** A sheet of light swept across a surface produces a brightness discontinuity (black curve). When there is no scattering, the pixel brightness is only due to this discontinuity (red ray). In the presence of scattering, the light plane itself becomes visible (dashed ray) making surface detection hard.

3 Light Striping in Scattering Media

Light stripe range scanning is a technique where a plane (or sheet) of light is swept across an object (figure 1) to obtain its 3D reconstruction. The key observation is that the plane of light intersects the object surface at a curve, producing a large brightness discontinuity. Then, the 3D coordinates of each point on this curve is computed by intersecting the camera ray and the light plane. A critical requirement here is that the intersection curve be detected reliably, which is usually done by thresholding the acquired image. Unfortunately, in the presence of scattering, the entire light plane itself becomes visible and detecting this intersection is not possible by simple thresholding. In this section, we derive the model for image formation when the light plane and the surface are immersed in a scattering medium and develop algorithms for reliable scene detection, and 3D reconstruction and for obtaining a *clear-air* appearance of the scene.

3.1 Image Formation Model

Imagine a light plane sweeping across a surface in a scattering medium. The camera not only receives light reflected by the surface, but also from the medium after scattering (see figure 2). The dashed lines indicate light rays that reach the camera after attenuation and scattering in the medium, but without reaching the surface. Then, the irradiance E_{med} at the camera is exactly given by equation 2. The red line indicates the path traveled by a light ray from the source to the surface and then reflected by the surface toward the camera. The intensity of this ray is exponentially attenuated according to the total distance traveled. Hence, the irradiance E_{surf} at the camera due to this ray is written as⁵,

$$E_{surf} = L_0 e^{-\beta(d_s+d_v)} R, \quad (3)$$

where, R is the radiance (normalized by source intensity) in the absence of scattering. Thus, the image formation model may be compactly written using the Dirac delta function δ as,

$$E = E_{surf} \delta(x = d_s) + E_{med} \delta(x < d_s). \quad (4)$$

⁵Single scattering of the exponentially attenuated surface radiance towards the camera is a minor effect compared to the scattering E_{med} from bright sources, and hence can be safely ignored.




Figure 2: **Image formation in light striping.** The irradiance at the camera is produced by either the light rays that reach the camera after being scattered once by the medium (dashed) or by light rays that are reflected by the surface (solid red). In both cases, the intensities of the rays are attenuated exponentially according to the distance traveled.

3.2 Intersection of Surface and Light Plane

Figure 3 illustrates the profile of the camera irradiance E as a function of the distance x of the source from the surface, according to equation 4. The brightness profile resembles an exponential fall-off followed by a discontinuity at the surface. When there is no scattering ($\beta = 0$), we have $E_{med} = 0$, $E_{surf} = L_0 R$ and hence $E_{surf} \gg E_{med}$. In this case, the brightness profile is a delta function and it is easy to detect the intersection using a threshold, as is done traditionally. For thresholding to work in the presence of scattering, we must have

$$R \gg \beta \mathcal{P}(g, \alpha) e^{\beta(d_s - x + d_v - y)}. \quad (5)$$

However, when scattering is significant (large β), it is mostly the opposite case, $E_{med} \geq E_{surf}$, as shown by the green and blue profiles in figure 3. Thus, the light plane itself becomes brightly visible (see second column in figure 6). In order to detect the intersection of the light plane and the surface, we simply use the brightness profile as a template until a brightness discontinuity is reached at the end. Even for the hard case where the density of the medium is high, this simple scheme performs well.

3.3 Experimental Setup and Calibration

The experimental setup consists of a $20'' \times 20'' \times 10''$ glass tank filled with water (see figure 4(a)). Different quantities of milk are mixed to emulate scattering media with different densities (β). The glass faces are anti-reflection coated to avoid reflections. We used an 8-bit Canon XLIS 3-CCD video camera and an Infocus LP120 1000 ANSI Lumens DLP projector in our experiments. To keep the size of the tank small, the camera and the projector are placed outside the tank. Hence, we need to handle light refractions at the *air-glass-medium* interfaces. Our calibration method is similar in spirit to techniques in [11]. Figure 4(b) illustrates a light plane from the projector shining into the glass tank after refraction. Calibration involves sweeping the light plane across two vertical planar surfaces - the (u, v) - and the (s, t) -planes - placed in the medium. The 3D world coordinates of a few points on these planes are measured *a priori* (the remaining points are interpolated). Then, the equation of each light plane




Figure 3: **Brightness profile for detecting the surface and light plane intersection.** When there is no scattering (red), the profile is a delta function which can be thresholded to detect the intersection. As the density of the medium (β) increases (green and blue), the brightness of the discontinuity (E_{surf}) decreases and the light plane becomes brighter (E_{med}).

is obtained using its line intersections with the (u, v) - and (s, t) -planes. Let this be represented by,

$$Ax + By + Cz + D = 0. \quad (6)$$

Next, we associate with each incoming camera ray (pixel (i, j)), its intersections $P(u, v, r)$ and $Q(s, t, 0)$ with the (u, v) - and the (s, t) -planes respectively (blue line in figure 4(c)). This yields a parametric equation for each camera ray, which is represented by:

$$[x, y, z] = [s, t, 0] + k[u - s, v - t, r - 0], \quad (7)$$

where, k is a scalar parameter. We calibrated our setup with the two planes placed at $z = 0$ inches and $z = 6.0$ inches. To verify calibration accuracy, we reconstructed (as described in Section 3.4) a plane placed at $z = 4.18$ inches with a low RMS error of 0.21 inch (figure 5). *In summary, our method does not require explicit geometric calibration of either the camera or the projector and does not require the position/orientation of the glass face or the refractive indices of media.*

3.4 Scene and Medium Recovery

Once calibrated the setup may be used to recover the 3D structure and *clear-air* appearance of any object in the medium as well as the properties of the medium itself.

3D surface reconstruction: Figure 4(c) shows a top-view (2D) illustration of the light striping setup and the profile of an object's surface. Since a point on the surface lies at the intersection of the reflected ray (blue) and the light plane (red), we may substitute (x, y, z) from equation 7 into equation 6, to solve for the parameter k :

$$k = \frac{As + Bt + D}{A(s - u) + B(t - v) - Cr}. \quad (8)$$

The value of k is then substituted back into equation 7 to obtain the 3D coordinates (x, y, z) of the surface point.

Medium properties: The properties of the medium can be obtained by observing the brightness decay of the light plane without the surface (see profile of E_{med} in figure 3). The distances x and y can be computed using the 3D coordinates of points on the light plane and the dimensions




Figure 4: **Light striping experimental setup and calibration.** (a) The setup consists of a glass tank filled with a scattering medium (dilute milk). The scene of interest is immersed in the medium. A projector illuminates the medium and the scene with planes of light and a video camera views the scene with the effects of scattering. (b) The light plane sweeps (one at a time) two planar surfaces placed vertically in the tank at known distances ($z = 0$ and $z = r$), called the (u, v) - and the (s, t) -planes. The discrete mappings between the light plane and the (u, v) - and (s, t) -planes constitute calibration. Note that no knowledge of the refraction locations or indices is required. (c) The top view of the setup illustrating the intersection of the light plane and the camera ray to yield the 3D coordinates of a surface point.




Figure 5: **Verification of light striping calibration.** (a) Two planes at $z = 0$ and $z = 6.0$ inches are used for calibration. (b) The computed equations of light planes and camera rays are then used to reconstruct a third plane at $z = 4.18$ inches (with RMS error 0.21 inch). The 3D view shows the three vertical planes and a light plane (red) for illustration.

of the tank. Then, equation 2 is nonlinear in the two unknown medium parameters, β and g . Thus, by observing the irradiances E_{med} along a profile on the light plane, we can estimate the two parameters β and g using a nonlinear optimization method ("fminsearch" in MatlabTM).

Scene appearance without scattering: Once the scattering coefficient β is estimated and the 3D surface is reconstructed, the scene appearance without scattering can be computed for each object intersection strip, from equation 3 as,

$$L_0 R = E_{surf} e^{+\beta(d_s + d_v)}, \quad (9)$$

where, E_{surf} is the observed brightness of the object in the presence of scattering. Then, all the intersection strips

are mosaiced to create the appearance of the entire scene as if captured in clear air.

The results of applying the scene and medium recovery algorithms are shown using real experiments in figure 6. The detection of the object intersections and hence the 3D reconstruction obtained under different densities of scattering compare well with the ground truth. Despite the strong effects of scattering, we are able to remove them completely to restore the original scene contrast. Also a comparison to the floodlit images demonstrates that simply using bright sources does not enhance visibility in scattering media, and that structured lighting methods that are designed to focus light on the scene to alleviate blurring and backscattering must be used.

4 Photometric Stereo in Scattering Media

In situations where light stripe scanning takes too long to be practical (for example, dynamic scenes), photometric stereo [29] provides an attractive alternative. Traditionally, photometric stereo is a technique for scene reconstruction (surface normal and albedo) from a small number of images of the scene acquired under different lighting directions. Many variants of this problem exist in vision literature [12; 2], but none of the proposed solutions are effective in scattering media.

In this section, we show how photometric stereo can be extended to scattering media. We choose the simplest version of the problem that assumes the surfaces to be lambertian, the sources distant, interreflections negligible and the camera to be orthographic. In the absence of scattering, it is known that three images of a scene illuminated from different but known directions are sufficient to uniquely determine the surface normals and albedos. We will first determine how many sources are needed in the presence of scattering and then show how scene properties can be recovered from the corresponding images.

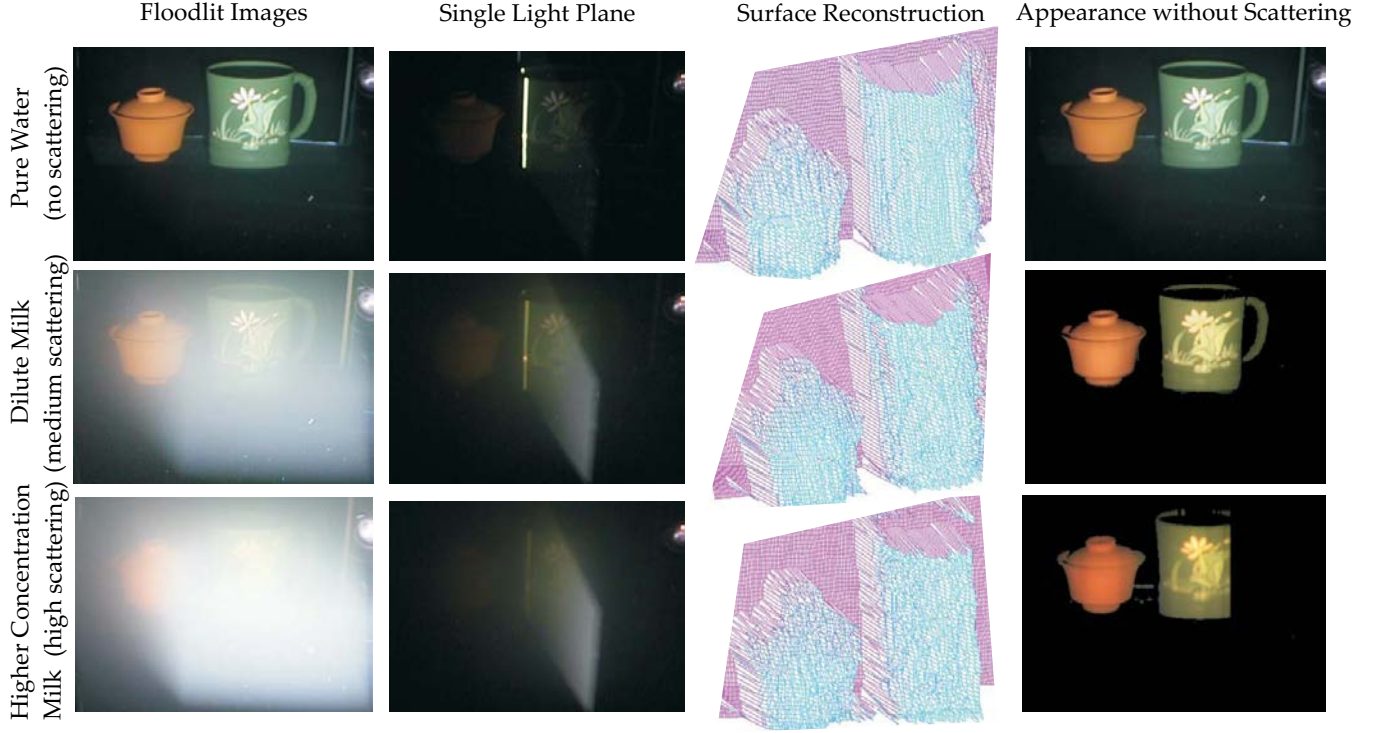


Figure 6: **Experimental results for light striping in scattering media.** The scene consists of two objects immersed in pure water (no scattering, ground truth), water mixed with 6 ml milk (medium scattering) and 15 ml milk (high scattering). The floodlit images (obtained by turning on all pixels in the projector) illustrate the adverse effects due to scattering by the medium. The brightness profile of a single light plane focused on the object confirms the template of the profile model in figure 3. For the two concentrations, our methods estimated $\beta = 0.07 \text{ in}^{-1}$ and 0.16 in^{-1} and $g = 0.9$. In the medium scattering case, our results (3D reconstruction and scene appearance without scattering) are nearly identical to the ground truth (percentage RMS error = 2.1%). In the 15 ml milk case, the green cup is barely visible (especially since its albedo is low) and yet the result is close to the ground truth. The handle of the cup is completely invisible and is hence missed (else, percentage RMS error = 5.5%). Note that our algorithm must be applied to separately to individual color channels. The color difference between the pure water and the other cases is due to white balancing differences between different experiments. **(Please see a video on our website [27] for better visualization and for other examples.)**

4.1 Image Formation Model

Consider the illumination and observation geometry in figure 7. A distant source (direction \mathbf{s}) illuminates a surface point \mathbf{P} with unit normal \mathbf{n} and albedo ρ . A camera observing the surface receives irradiance E_{surf} due to the light reflected by the surface (solid red lines) and irradiance E_{med} due to light scattered by the medium (dashed lines) in the viewing direction. The irradiance E_{surf} is the same as for light striping (see equation 3),

$$E_{surf} = L_0 e^{-\beta d_s} (\rho \mathbf{n} \cdot \mathbf{s}) e^{-\beta d_v}. \quad (10)$$

Here, we have replaced the normalized radiance R by $(\rho \mathbf{n} \cdot \mathbf{s})$ for a lambertian surface. The irradiance E_{med} at the camera due to single scattering by the medium is obtained by integrating the brightness along the viewing direction (see equation 2),

$$E_{med} = \int_0^{d_v} L_0 e^{-\beta x} \beta \mathcal{P}(g, \alpha) e^{-\beta y} dy. \quad (11)$$

Note that α , $\mathcal{P}(g, \alpha)$, β and L_0 are all independent of the integration variable y . Further, we shall also assume




Figure 7: **Image formation for photometric stereo in scattering media.** The sources, viewer and the surface of interest are immersed in the scattering medium. The sources are distant and thus illuminate the surface and the viewing ray in a collimated fashion. The brightness at a pixel is the sum of the contributions from the solid red and the dashed rays.

the source uniformly illuminates the viewing distance d_v . In other words, $x = d_s$ is constant with respect to y (this assumption will be relaxed when we discuss our specific setup). This allows us to simplify equation 11 as,

$$E_{med} = L_0 \mathcal{P}(\alpha) e^{-\beta d_s} (1 - e^{-\beta d_v}). \quad (12)$$




Figure 8: **Refraction of rays in the photometric stereo model.** The sources and camera are outside the scattering medium. The viewing direction of the orthographic camera is normal to the air-medium interface to avoid refractions of incoming camera rays. However, refraction of light rays from the source must be modeled.

Then, the total irradiance E at the camera can be written as the sum of the irradiances E_{med} and E_{surf} :

$$E = L_0 [e^{-\beta(d_s+d_v)} \rho \mathbf{n} \cdot \mathbf{s} + \mathcal{P}(g, \alpha) e^{-\beta d_s} (1 - e^{-\beta d_v})]. \quad (13)$$

For an isotropic point source, $L_0 = I_0/d_s^2$. Equation 13 represents the image formation model for one distant source. Similar equations can be written for each distant source that illuminates the scene.

4.2 Experimental Setup and Calibration

The glass tank described in Section 3 is again used in these experiments and, as before, we place the camera and the sources outside the tank. A 12-bit per channel Canon EOS-20D camera with a 70 – 300 mm zoom lens is placed 20 feet away from the tank and observes the front face of the tank normally (perpendicularly). The field of view occupied by the tank in the image is 2.0 degrees and is hence approximately orthographic.

During calibration, refraction of the light rays from sources at the *air-medium* boundary must be accounted for. Figure 8 shows a schematic of the side view of the setup. The distances d_s and d_v are related using trigonometry,

$$d_v = d_s \cos \alpha. \quad (14)$$

Notice that the light rays that illuminate the viewing ray and the surface travel different distances in the medium (compare the lengths of the dashed parallel rays in figures 7 and 8). Hence, the assumption in simplifying equation 11 that x is constant with respect to y , becomes invalid for our experimental setup. So, an appropriate correction is derived for E_{med} using equation 14 to obtain the irradiance (see appendix A):

$$E = L_0 e^{-\beta d_v(1+1/\cos \alpha)} \rho \mathbf{n} \cdot \mathbf{s} + \frac{L_0 \mathcal{P}(g, \alpha) \cos \alpha}{1 + \cos \alpha} (1 - e^{-\beta d_v(1+1/\cos \alpha)}). \quad (15)$$

We will henceforth call equation 15 as the image formation model. We calibrate our setup using images of a white lambertian sphere in pure water (scattering is minimal). The brightest point on the sphere yields the refracted direction \mathbf{s} (and α) and intensity L_0 of the source.

4.3 Scene and Medium Recovery

Consider a set of images taken of an object under different source directions. In order to find out how many source directions are required to recover the scene and the medium, let us count the number of knowns and unknowns in equation 15. Recall that as part of calibration, the angle α , the source direction \mathbf{s} and intensity L_0 are all estimated *a priori*. Then, the unknowns for each scene point are the surface albedo ρ , unit normal \mathbf{n} , and optical thickness $T_v = \beta d_v$. The medium parameter g in the expression for $\mathcal{P}(g, \alpha)$ (see equation 1) is constant and hence is a global unknown. Thus, there are four unknowns for each scene point and one global unknown. If there are P scene points and L light source directions, the number of unknowns $4P + 1$ must be less than the number of equations PL . So, simple variable counting suggests that a minimum of $L = 5$ is required⁶.

To empirically verify that indeed $L = 5$ suffices (assuming the sources are not in degenerate positions), we performed numerical simulations on 4000 randomly generated combinations of source directions \mathbf{s}_i , surface normals \mathbf{n} , albedos $\rho \in (0, 1)$, optical thicknesses $T_v \in (0, 2)$ and forward scattering parameters $g \in (-1, 1)$, for a single scene point. The MatlabTM function “fminsearch” was used to recover the unknowns by minimizing the sum of squared differences between the simulated values and the model in equation 15. In all trials, the search was initialized with random values for the unknowns. In all cases, the search algorithm converged to the global optimum solution within few seconds. This suggests the presence of a single global minimum of the error function⁷. As a test of robustness, we added uniform random noise (up to 5% of the simulated values) and found that the errors in recovered unknowns were low, as evidenced by the error histograms in figure 10. We also ran the above simulations using only 4 sources, but the global error minimum corresponded to several parameter sets, suggesting that 4 sources are insufficient for unique estimation. *Thus, we conclude that five non-degenerate light source directions are required and sufficient to uniquely estimate the properties of the scene and the medium. In practice, however, more source directions may be used for robustness.*


The experiments performed with our setup are shown in figure 9. Images of a teapot captured in the presence of scattering (by dilute milk) have poor contrast and colors. As expected, applying traditional photometric stereo results in poor results. On the other hand, the surface normals and the albedos obtained using our method⁸ are

⁶In appendix B, we present an interesting but practically limited case where a unique linear solution with four sources is possible.


⁷However, the error function does contain local minima and the search was conducted starting from several (typically 100) initial guesses and the minimum of all the solutions was used.

⁸The non-linear optimization can be executed independently for each pixel. But to speedup execution time, we masked the region where the object is not present ($E_{surf} = 0$) to first estimate the global parameter g , before estimating the 4 parameters for each pixel. As a side note, from our experience, it is critical to use high quality (12 bits per channel HDR) radiometrically linear input images in order to obtain good reconstructions shown.

Pure Water (No scattering) - Ground Truth



(a) Images (2 out of 8) captured in pure water.




(b) Albedo and shape computed using traditional method


Dilute Milk (medium scattering)



(c) Input Images (2 out of 8) captured in dilute milk



(d) Albedo and shape computed using traditional method



(e) Albedo and shape computed using our method

Figure 9: Experimental results of Photometric Stereo in Scattering Media. (a) Two (out of eight) images of a teapot acquired under different lighting directions (depicted in the insets). (b) Results on applying traditional photometric stereo to images in (a) serve as ground truth. (c) The images acquired in dilute milk. Notice the significant loss of contrast. (d) If traditional photometric stereo applied to images in (c), the 3D shape obtained is very flat and the scattering effects are absorbed by the albedos. (e) The results obtained using our algorithm. The percentage RMS error in reconstructing the shape was 6.3%. In addition to surface normals and albedos, our method also yields a depth map, which is not possible using traditional photometric stereo (see figure 10.) The 3D shapes were computed from the surface normals using [8].

very accurate (with only a percentage RMS error of 6.3% in computed shape). *In addition, our method also yields a depth map of the scene $T_v = \beta d_v$, which is not possible using traditional photometric stereo.*

5 Discussion: Real Underwater Scenarios

Since all the experiments in this paper were done using milk as the scattering medium, it is important to identify issues that may arise in real underwater scenarios (lakes, oceans, seas). In initial experiments, we collected water samples from 4 different locations near a Pacific Ocean beach in San Diego, USA. These samples were collected close to the water surface where the impurity levels are generally high. By matching images of these water samples with those of milk, we found that the low to moderate concentration range of milk used in this paper corresponds to the concentrations observed in the ocean waters. Hence, our algorithms can be applied to many underwater scenarios as well. In the future, to make our techniques broadly applicable, we wish to develop underwater deployable systems and improve our algorithms to handle non-homogeneous and dynamic underwater media.

References

- [1] R. D. Ballard. *The Discovery of the Titanic*. Warner Books, 1988.
- [2] R. Basri and D.W. Jacobs. Photometric stereo with general, unknown lighting. In *CVPR*, 2001.
- [3] F. M. Caimi, D. M. Kocak, and V. L. Asper. Developments in laser-line scanned undersea surface mapping and image analysis systems for scientific applications. In *Proc. MTS/IEEE Oceans*, 1996.
- [4] S. Chandrasekhar. *Radiative Transfer*. Dover Publications, Inc., 1960.
- [5] P. C. Y. Chang, J. C. Flitton, K. I. Hopcraft, E. Jakeman, D. L. Jordan, and J. G. Walker. Improving visibility depth in passive underwater imaging by use of polarization. *App. Opt.*, 42 (15), 2003.
- [6] S. Y. Chen and Y. F. Li. Self-recalibration of a colour-encoded light system for automated 3-d measurements. *MeasureSciTech*, 14(1), 2003.
- [7] G.R. Fournier, D. Bonnier, J.L. Forand, and P.W. Pace. Range-gated underwater laser imaging system. *Opt. Eng.*, 32 (9), 1993.
- [8] R.T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *PAMI*, 10(4), 1988.
- [9] E. Fuchs and J. S. Jaffe. Thin laser light sheet microscope for microbial oceanography. *OPTICS EXPRESS*, 10 (2), 2002.
- [10] G. D. Gilbert and J. C. Pernicka. Improvement of underwater visibility by reduction of backscatter with a circular polarization technique. *Applied Optics*, 6 (4):741-746, 1967.
- [11] M.D. Grossberg and S.K. Nayar. The raxel imaging model and ray-based calibration. *IJCV*, 61(2), 2005.
- [12] A. Hertzmann and S.M. Seitz. Shape and materials by example: a photometric stereo approach. In *CVPR*, 2003.
- [13] D.Q. Huynh, R.A. Owens, and P.E. Hartmann. Calibrating a structured light stripe system: A novel approach. *IJCV*, 33(1), 1999.
- [14] J. S. Jaffe, J. McLean, M. P. Strand, and K. D. Moore. Underwater optical imaging: Status and prospects. *Tech. Report, Scripps Institution of Oceanography, La Jolla*, 2002.
- [15] C. Je, S.W. Lee, and R.H. Park. High-contrast color-stripe pattern for rapid structured-light range imaging. In *ECCV*, 2004.
- [16] D. M. Kocak, F. M. Caimi, T. H. Jagielo, and J. Kloske. Laser projection photogrammetry and video system for quantification and mensuration. *MTS/IEEE Oceans*, 2002.




Figure 10: [Left] Simulations show performance of our algorithm for photometric stereo, in the presence of small amounts of noise (uniform random noise up to 5% of the simulated values). The peaks near zero values in the error histograms corresponding to the recovered parameters indicate robustness. [Right] The presence of scattering allows us to compute the scaled depth T_v of every scene point, which is not possible in traditional photometric stereo. The depth map was filtered to remove some noise.

- [17] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas. Synthetic aperture confocal imaging. In *SIG-GRAPH*, 2004.
- [18] S.G. Narasimhan and S.K. Nayar. Vision and the atmosphere. *IJCV*, 48(3).
- [19] P. Naulleau and D. Dilworth. Motion-resolved imaging of moving objects embedded within scattering media by the use of time-gated speckle analysis. *App. Opt.*, 35 (26), 1996.
- [20] Y. Sato, H. Kitagawa, and H. Fujita. Shape measurement of curved objects using multiple slit-ray projections. *PAMI*, 4(6), 1982.
- [21] D. Scharstein and R. Szeliski. High-accuracy stereo depth maps using structured light. In *CVPR*, 2003.
- [22] Y. Y. Schechner and N. Karpel. Attenuating natural flicker patterns. In *Proc. MTS/IEEE Oceans*, 2004.
- [23] Y. Y. Schechner and N. Karpel. Clear underwater vision. In *Proc. CVPR*, 2004.
- [24] J. S. Tyo, M. P. Rowe, Jr. E. N. Pugh, and N. Engheta. Target detection in optically scattering media by polarization-difference imaging. *App. Opt.*, 35 (11), 1996.
- [25] J. G. Walker, P. C. Y. Chang, and K. I. Hopcraft. Visibility depth improvement in active polarization imaging in scattering media. *App. Opt.*, 39 (27), 2000.
- [26] D. Walther, D. R. Edgington, and C. Koch. Detection and tracking of objects in underwater video. In *Proc. CVPR*, 2004.
- [27] Authors Website. <http://www.cs.cmu.edu/~srinivas/>.
- [28] P.M. Will and K.S. Pennington. Grid coding: A preprocessing technique for robot and machine vision. *AI*, 2, 1971.
- [29] R.J. Woodham. Photometric method for determining surface orientation from multiple images. *OptEng*, 19(1), 1980.
- [30] L. Zhang, B. Curless, and S. M. Seitz. Rapid shape acquisition using color structured light and multi-pass dynamic programming. In *The 1st IEEE International Symposium on 3D Data Processing, Visualization, and Transmission*, 2002.

A: Refractions in Photometric Stereo

The expression for E_{surf} is the same as equation 10 with d_s being substituted from equation 14. The irradiance E_{med} is derived by observing that $x = y \cos \alpha$ and $d_v = d_s \cos \alpha$:

$$E_{med} = \int_0^{d_v} L_0 e^{-\beta x} \beta \mathcal{P}(g, \alpha) e^{-\beta y} dy = L_0 \beta \mathcal{P}(g, \alpha) \int_0^{d_v} e^{-\beta(1+1/\cos \alpha)y} dy$$

$$= \frac{L_0 \mathcal{P}(g, \alpha) \cos \alpha}{1 + \cos \alpha} (1 - e^{-\beta(1+1/\cos \alpha)d_v})$$

which is rewritten in equation 15. Note that we have assumed here that the actual source outside the medium

is still distant and does not show an inverse square fall-off within the medium since the distance light traveled within the medium is much less compared to the distance of the source from the medium boundary. Thus, L_0 is constant with respect to d_s .

B: Linear Solution with Four Sources

We assume that (a) the sources and the camera are both immersed in an isotropic medium ($\mathcal{P}(\alpha) = \frac{1}{4\pi}$), and (b) the sources are equidistant from any given surface point (d_s same for all sources, however, d_s may vary for different scene points). Although these assumptions limit practicality, this case serves as an illustration in comparison to the technique without any scattering. Let the intensities observed at a pixel under 4 different sources be E_1 , E_2 , E_3 , and E_4 . From equation 13:

$$E_i = L_0 e^{-\beta(d_{s_i}+d_v)} \rho \mathbf{n} \cdot \mathbf{s}_i + \frac{L_0}{4\pi} e^{-\beta d_{s_i}} (1 - e^{-\beta d_v}), \quad (16)$$

where, the subscript i denotes the source. For this special case, $d_{s_1} = d_{s_2} = d_{s_3} = d_{s_4} = d_s$. The second term in equation 16 can be eliminated by considering pair-wise differences between intensities to obtain a set of equations that are linear in the surface normal \mathbf{n} :

$$\begin{pmatrix} E_1 - E_2 \\ E_1 - E_3 \\ E_1 - E_4 \end{pmatrix} = L_0 e^{-\beta(d_s+d_v)} \rho \begin{pmatrix} \mathbf{s}_1 - \mathbf{s}_2 \\ \mathbf{s}_1 - \mathbf{s}_3 \\ \mathbf{s}_1 - \mathbf{s}_4 \end{pmatrix} \mathbf{n}. \quad (17)$$

In matrix and vector notation, we rewrite equation 17,

$$\Delta \mathbf{E} = L_0 e^{-\beta(d_s+d_v)} \rho \Delta \mathbf{S} \mathbf{n}, \quad (18)$$

where, $\Delta \mathbf{E}$ and \mathbf{n} are 3×1 column vectors and $\Delta \mathbf{S}$ is a 3×3 matrix. By applying $\Delta \mathbf{S}^{-1}$, we obtain the surface normal \mathbf{n} and the attenuated appearance $L_0 e^{-\beta(d_s+d_v)} \rho$. Thus, all the terms except for βd_s and βd_v are estimated or known in equation 16. If a relationship between d_s and d_v is known (such as equation 14), then βd_s and βd_v can be estimated using equation 16 and in turn the albedo ρ can be estimated from the product $L_0 e^{-\beta(d_s+d_v)} \rho$. Thus, in the case of isotropic scattering and equidistant sources, four light source directions are required to compute surface normals, albedos and optical thickness βd_v (or scaled distance).

On Controlling Light Transport in Poor Visibility Environments

Mohit Gupta, Srinivasa G. Narasimhan
Carnegie Mellon University
School of Computer Sc., Pittsburgh 15232, USA
{mohitg, srinivas}@cs.cmu.edu

Yoav Y. Schechner
Technion - Israel Institute of Technology
Dep. of Electrical Eng., Haifa 32000, Israel
yoav@ee.technion.ac.il

Abstract

Poor visibility conditions due to murky water, bad weather, dust and smoke severely impede the performance of vision systems. Passive methods have been used to restore scene contrast under moderate visibility by digital post-processing. However, these methods are ineffective when the quality of acquired images is poor to begin with. In this work, we design active lighting and sensing systems for controlling light transport before image formation, and hence obtain higher quality data. First, we present a technique of polarized light striping based on combining polarization imaging and structured light striping. We show that this technique out-performs different existing illumination and sensing methodologies. Second, we present a numerical approach for computing the optimal relative sensor-source position, which results in the best quality image. Our analysis accounts for the limits imposed by sensor noise.

1. Introduction

Computer vision systems are increasingly being deployed in domains such as surveillance and transportation (terrestrial, underwater or aerial). To be successful, these systems must perform satisfactorily in common poor visibility conditions including murky water, bad weather, dust and smoke. Unfortunately, images captured in these conditions show severe contrast degradation and blurring, making it hard to perform meaningful scene analysis.

Passive methods for restoring scene contrast [14, 18, 22] and estimating 3D scene structure [4, 12, 28] rely on post-processing based on the models of light transport in natural lighting. Such methods do not require special equipment and are effective under moderate visibility [12], but are of limited use in poor visibility environments. Very often, there is simply not enough useful scene information in images. For example, in an 8-bit camera, the intensity due to dense fog might take up 7 bits, leaving only 1 bit for scene radiance. Active systems, on the other hand, give us flexibility in lighting and/or camera design, allowing us to control the light transport in the environment for better image quality. Figure 1 illustrates the significant increase in image quality using our technique. In this experiment, the




Figure 1. Polarized light striping versus flood-lighting. In this experiment, the scene is comprised of objects immersed in murky water. Using the polarized light striping approach, we can control the light transport *before* image formation for capturing the same scene with better color and contrast. High-resolution images can be downloaded from the project web-page [7].

scene comprised of objects immersed in murky water.

While propagating within a medium such as murky water or fog, light gets absorbed and scattered. Broadly speaking, light transport [2] can be classified based on three specific pathways: (a) from the light source to the object, (b) from the object to the sensor and (c) from the light source to the sensor without reaching the object (see Figure 2). Of these, the third pathway causes loss of contrast and effective dynamic range (for example, the backscatter of car headlights in fog), and is thus undesirable.

We wish to build active illumination and sensing systems that maximize light transport along the first two pathways while simultaneously minimizing transport along the third. To this end, we exploit some real world observations. For example, while driving in foggy conditions, flood-lighting the road ahead with a *high-beam* may reduce visibility due to backscatter. On the other hand, underwater divers realize that maintaining a good separation between the source and the camera reduces backscatter, and improves visibil-




Figure 2. Light transport in scattering media for different source and sensor configurations. (a) Illustration of the three light transport components. (b) The backscatter B reduces the image contrast. The amount of backscatter increases with the common backscatter volume. (c) By changing the relative placement of the sensor and source, we can modulate the light transport components for increasing the image contrast. (d) The common backscatter volume can be reduced by using light stripe scanning as well.

ity [21, 9]. Polarization filters have also been used to reduce contrast loss due to haze and murky water [20, 24, 19, 6]. Based on these observations, we attempt to address two key questions. First, which illumination and sensing modality allows us to modulate the three light transport pathways most effectively? Second, what is the “optimal” placement of the source and the sensor? This paper has two main contributions:

(1) We present an active imaging technique called *polarized light striping* and show that it performs better than previous techniques such as flood-lighting, unpolarized light striping [10, 15, 9], and high frequency illumination based separation of light transport components [16].

(2) We derive a numerical approach for computing the *optimal relative sensor-source position* in poor visibility conditions. We consider a variety of illumination and sensing techniques, while accounting for the limits imposed by sensor noise. Our model can be used for improving visibility in different outdoor applications. It is useful for tasks such as designing headlights for vehicles (terrestrial and underwater). We validate our approach in real experiments.

2. How to Illuminate and Capture the Scene?

In this section, we present an active imaging technique: polarized light striping. We also analyze the relative merits of different existing techniques, and show that polarized light striping outperforms them.

While propagating through a medium, light gets absorbed and scattered (Figure 2). The image irradiance at a particular pixel is given as a sum of the three components, the direct signal (D), the indirect signal (A) and the backscatter (B):

$$E(x, y) = \underbrace{D(x, y) + A(x, y)}_{\text{Signal}} + \underbrace{B(x, y)}_{\text{Backscatter}}. \quad (1)$$

The total signal S is

$$S(x, y) = D(x, y) + A(x, y). \quad (2)$$




Figure 3. Our experimental setup consisting of a glass tank, filled with moderate to high concentrations of milk (four times as those in [15]). An LCD projector illuminates the medium with polarized light. The camera (with a polarizer attached) observes a contrast chart through the medium.

The backscatter B degrades visibility and depends on the optical properties of the medium such as the extinction coefficient and the phase function. The direct and the indirect components (D and A) depend on both the object reflectance and the medium. Our goal is to design an active illumination and sensing system that modulates the components of light transport effectively. Specifically, we want to maximize the signal S , while minimizing the backscatter B .

We demonstrate the effectiveness of different imaging techniques in laboratory experiments. Our experimental setup consists of a $60 \times 60 \times 38 \text{ cm}^3$ glass tank filled with dilute milk (see Figure 3). The glass facades are anti-reflection coated to avoid stray reflections.¹ The scene consists of objects immersed in murky water or placed behind the glass tank. A projector illuminates the scene and a camera fitted with a polarizer observes the scene. We use a Sony VPL-HS51A, Cineza 3-LCD video projector. The red and the green light emitted from the projector are inherently polarized channels. If we want to illuminate the scene with blue light, we place a polarizer in front of the projector. We use a 12-bit Canon EOS1D Mark-II camera, and a Kodak contrast chart as the object of interest to demonstrate the contrast loss or enhancement for different techniques.

¹Imaging into a medium through a flat interface creates a non-single viewpoint system. The associated distortions are analyzed in [25].




Figure 4. Limitations of the high frequency illumination based method. A shifting checkerboard illumination pattern was used with the checker size of 10×10 pixels. (a) Maximum image (b) Minimum image (global component) (c) Direct component (d) Direct component obtained using lower frequency illumination (checker size of 20×20 pixels). The direct component images have low SNR in the presence of moderate to heavy volumetric scattering. The global image is approximately the same as a flood-lit image, and hence, suffers from low contrast. This experiment was conducted in moderate scattering conditions, same as the second row of Figure 6.




Figure 5. The relative direct component of the signal reduces with increasing optical thickness of the medium. This plot was calculated using simulations, with a two-term Henyey-Greenstein scattering phase function [8] for a parameter value of 0.8.

High-frequency illumination: Ref. [16] presented a technique to separate direct and global components of light transport using high frequency illumination, with good separation results for inter-reflections and sub-surface scattering. What happens in the case of light transport in volumetric media? Separation results in the presence of moderate volumetric scattering are illustrated in Figure 4. The direct component is the direct signal (D), whereas the global component is the sum of indirect signal (A) and the backscatter (B), as shown in Figure 2. Thus, this method seeks the following separation:

$$E(x, y) = \underbrace{D(x, y)}_{\text{Direct}} + \underbrace{A(x, y) + B(x, y)}_{\text{Global}}. \quad (3)$$

However, to achieve the best contrast, we wish to separate the signal $D + A$ from the backscatter B . As the medium becomes more strongly scattering, the ratio $\frac{D}{S}$ falls rapidly due to heavy attenuation and scattering, as illustrated in Figure 5. This plot was estimated using numerical simulations using the single scattering model of light transport.² Consequently, for moderate to high densities of the

²With multiple scattering, the ratio falls even more sharply.

medium, the direct image suffers from low signal-to-noise-ratio (SNR), as shown in Figure 4. Further, the indirect signal (A) remains unseparated from the backscatter B , in the global component. Thus, the global image is similar to a flood-lit image, and suffers from low contrast.

Polarized flood-lighting: Polarization imaging has been used to improve image contrast [19, 23, 6] in poor visibility environments. It is based on the principle that the backscatter component is partially polarized, whereas the scene radiance is assumed to be unpolarized. Using a sensor mounted with a polarizer, two images can be taken with two orthogonal orientations of the polarizer:

$$E_b = \frac{D + A}{2} + \frac{B(1 - p)}{2} \quad (4)$$

$$E_w = \frac{D + A}{2} + \frac{B(1 + p)}{2}, \quad (5)$$

where p is the degree of polarization (DOP) of the backscatter. Here, E_b and E_w are the ‘best-polarized image’ and the ‘worst-polarized image’, respectively. Thus, using optical filtering alone, backscatter can be removed partially, depending on the value of p . Further, it is possible to recover an estimate of the signal \hat{S} in a post-processing step [19]:

$$\hat{S} = E_b \left(1 + \frac{1}{p}\right) + E_w \left(1 - \frac{1}{p}\right). \quad (6)$$

However, in optically dense media, heavy backscatter due to flood-lighting can dominate the signal, making it impossible for the signal to be recovered. This is illustrated in Figure 6, where in the case of flood-lighting under heavy scattering, polarization imaging does not improve visibility.

Light stripe scanning: Here, a thin sheet of light is scanned across the scene. In comparison to the above approaches, the common backscatter volume is considerably




Figure 6. Comparison of various illumination and sensing techniques (zoom into the marked areas to better assess the image quality). Flood-lit images suffer from a severe loss of contrast, specially in the presence of heavy scattering (a,b). *Polarized light striping* achieves a significant increase in image contrast, even in the presence of heavy scattering (a-d). In moderate scattering, fine details (text) are recovered more reliably in (g) and (h), as compared to (e). See (i), (j), (k) and (l) for close-ups of the marked areas in (e), (f), (g) and (h) respectively. The moderate scattering experiment was conducted under the same conditions as the experiment in Figure 4.

reduced (see Figure 2d). The sheet of light intersects the object to create a stripe that is detected using a gradient operator.³ All stripes are then mosaiced to create a composite image CI [10, 15, 9]. Alternatively, the composite image can be obtained by simply selecting the maximum value at each pixel over all the individual light stripe images SI_k :

$$CI(x, y) = \max_k \{SI_k(x, y)\}. \quad (7)$$

³In our particular implementation, the projector illuminates a single plane and has low power. We compensate for this by increasing the exposure time of the camera.

Polarized light striping: We propose polarized light striping as a technique that combines the advantages of polarization imaging and light striping, and thus, is applicable for an extended range of medium densities. Earlier, we demonstrated that light striping reduces the amount of backscatter. However, reliable localization of the object stripes (by using gradient operator or by selecting the maximum pixel value, as in Eq. 7) is severely impeded due to strong backscatter. This is illustrated in Figure 7.

To enable reliable detection of the object stripes even in the presence of strong scattering, we use polarization imaging in conjunction with light striping. A high DOP of the




Figure 7. Unpolarized versus polarized light stripe scanning. (a) Ray diagram illustrating light stripe scanning, adapted from [15]. (b) The camera observes a light stripe (1 out of 30) without a polarizer. The visible light plane is the backscatter and impedes reliable detection of the object stripe. (c) Through a polarizer, there is a considerable reduction in backscatter. The light plane-object intersection becomes more distinct, thus enabling its reliable delineation. (d) The removed backscatter (difference of (b) and (c)). Video of a complete scan can be downloaded from the project web-page [7].

backscatter is essential for removing the backscatter using polarization filtering (Eq. 4), or to recover a reliable estimate of the signal using post-processing (Eq. 6). In our experiments, the camera observes the scene through a polarization filter and the light sheets irradiating the scene are polarized. Since the incident illumination is completely polarized, the DOP of the backscatter is high (see appendix). This results in a significant reduction in the amount of backscatter, and thus, enables reliable detection of the stripes.⁴ This is shown in Figure 7. We compare the results of polarized light striping versus previous illumination and sensing techniques in Figure 6. Notice especially the differences in the contrast under strong scattering. Another result is shown in Figure 1.

3. Optimal Camera-Source Placement

Conventional wisdom from the underwater imaging literature suggests maximizing the sensor-source separation to reduce the backscatter, and hence, increase the image contrast [9, 21] (see Figure 2). However, this does not take into account the limitations posed by measurement noise. Indeed, placing the source and the sensor far from each other or the scene results in strong attenuation of light, and a low SNR. In this section, we investigate this trade-off between image contrast and SNR to compute the optimal relative sensor-source positions.

3.1. Quality Measures

In order to formalize the notion of “optimal”, we define various image quality measures for different imaging and illumination techniques. These quality measures serve as objective functions which can be maximized to find the optimal placement of the source and the camera.

Contrast Quality Measure: A major goal of an imaging system is to maximize the image contrast. Analogous to [1,

5], we define the contrast quality measure, $CQM(x, y)$ as the ratio of the signal $S(x, y)$ to the total intensity $E(x, y)$:

$$CQM(x, y, p) = \frac{S}{S + B(1 - p)}. \quad (8)$$

This measure takes polarization imaging into account by defining the total intensity as that of the *best polarized image*, as in Eq. (4). In the absence of a polarizer, $p = 0$.

Delineation of light plane-scene intersection: Success of light striping in scattering media relies on reliable delineation of the object stripe. One scheme is to detect a brightness discontinuity in the intensity profile across the stripe edge. Thus, for a light stripe scanning system, we define a gradient quality measure (GQM) along the edge of the stripe in terms of the strength of gradient across the stripe edge. Consider Figure 8c; since the scene point O' does not have the direct component D or the backscatter component $B(1 - p)$, the normalized difference in intensity of O and O' is given as:

$$GQM(x, y, p) = \frac{D + B(1 - p)}{D + A + B(1 - p)}. \quad (9)$$

SNR dependent weighting: An image with high contrast but low overall intensity may result in a low SNR, and hence be of limited use. Thus, we define an SNR dependent weight W as a monotonically increasing function of the total signal value S . The quality measures (CQM and GQM) are weighted by W so that signal values in the low SNR range are penalized. For example, W can be a linear function of S . For more flexibility, we use a sigmoid function of S :

$$W(x, y) = \frac{1}{1 + e^{-\left(\frac{S - \mu}{z}\right)}}, \quad (10)$$

where μ is the shift and z is the steepness of the sigmoid. For example, μ can be the dark current offset. Similarly, if the noise is derived from a Gaussian distribution, z can be the standard deviation. In addition, we should account for the effect of post-processing on image noise [24, 17].

⁴Polarization imaging was previously used with phase-shifted structured illumination for improved reconstruction of translucent objects [3].




Figure 8. Simulating image formation for finding the optimal sensor-source configuration. (a) A schematic view of the volume. We use a point light source (L) and a pinhole camera (C). The object is Lambertian, with reflectance R . (b) We calculate D , A and B according to Eqs. (11-13). (c) In the case of light striping, the point O' is not getting directly irradiated by the source. Also, the viewing ray from O' does not intersect the common backscatter volume. Thus, the direct component and the backscatter component at O' are null. This results in a brightness gradient across the stripe edge. The strength of the gradient is given by Eq. 9.

3.2. Simulations

Consider an underwater scenario where a remote operated vehicle (ROV) wants to capture images at a given distance. Given an approximate estimate of the object albedo, medium scattering parameters [13] and sensor noise, we can simulate the image formation process. To illustrate the concept, we simulate the image formation process for our experimental setup. The Lambertian object reflectance was assumed to be 0.6. For different source-camera configurations, we compute the appropriate quality measure described above. Then, the optimal configuration is the one that maximizes the quality measure.

Figure 8 illustrates the image formation geometry. In our experiments and simulations, the scene and camera remain fixed, while the source is moved to vary the sensor-source separation d_{LC} . Point O on the object is being observed by the camera. Points X and Y are in the medium. The distances d_{LO} , d_{CO} , d_{LX} , d_{XO} , d_{CO} , d_{LY} and d_{YC} , and the angles ϕ , α , γ , θ are as illustrated in Figure 8. To keep our simulations simple, we assume a single scattering model of light transport and a homogeneous medium. The individual components of light transport are then given by:

$$D = \frac{I_0}{d_{LO}^2} e^{-\sigma(d_{LO}+d_{CO})} R(\phi) \quad (11)$$

$$A = \int_V \frac{I_0}{d_{LX}^2} e^{-\sigma(d_{LX}+d_{XO}+d_{CO})} F(\alpha) R(\gamma) dV \quad (12)$$

$$B = \int_O^C \frac{I_0}{d_{LY}^2} e^{-\sigma(d_{LY}+d_{YC})} F(\theta) dY, \quad (13)$$



where I_0 is the source radiance, σ is the extinction coefficient, R is the Lambertian object reflectance, F is the scattering phase function (we use the two-term Henyey-Greenstein function [8]) and V is the illuminated volume.

Polarized images, E_b and E_w are simulated according to Eqs. (4-5). This requires knowledge of the DOP of the

backscatter p . Using our experimental setup, we estimated p to be approximately 0.8, from the regions of the image without any object. We can also compute p analytically, given the dependence of the DOP of scattered light on the scattering angle, such as given in the Appendix.

Optimal configuration for flood-lighting: Let us find the configuration that is optimal in terms of both image contrast and noise. We plot the product of the CQM and W versus the sensor-source separation d_{LC} (Figure 9a). The tradeoff between contrast and SNR results in a local maximum. Notice that polarization improves image quality as compared to unpolarized imaging. However, since the DOP (and hence, the amount of contrast enhancement) is similar for all sensor-source positions, the location of the peak remains the same. The curve for the ideal case of zero noise increases monotonically. However, for real world scenarios, where measurement noise places limits on the sensor's abilities, our approach can yield an optimal placement. This is illustrated in Figure 9 (b-c). The image taken using the optimal separation (40 cms) has high contrast and low noise. On the other hand, notice the significant noise in the image taken using a large separation (60 cms).

Optimizing the light stripe scan: The case of light stripe scanning is more interesting. Instead of illuminating the whole scene at once, we illuminate it using one sheet of light at a time. We want to find the optimal light stripe scan. Should we scan the scene (a) by rotating the source, (b) by translating it, or (c) a combination thereof? To answer this, we plot the product of the GQM and the W for our setup (Figure 10). We observe different optimal separations for different (3 out of 30) stripe locations. Figure 10 (e) shows the high-contrast image acquired using the results of the simulations. The camera and the projector were placed



(b) Large Separation

(c) Optimal Separation

Figure 9. Optimal sensor-source configuration for flood-lighting. (a) Plot of $CQM \times W$ versus d_{LC} for our experimental setup. The tradeoff between contrast and SNR results in a maximum. (b) Large separation (60 cms) results in heavy image noise (c) Optimal separation (40 cms) results in a high contrast, low noise image (zoom into the marked area). Both the frames were captured with the same exposure time.

at a small distance from the facade of the glass tank in real experiments. By carefully choosing the light rays, we can simulate a light source and a sensor placed on the glass facade, as assumed in the simulations. The optimal scan for polarized light striping is the same as unpolarized light striping, but results in better image quality.

4. Discussion

We study new ways to control light transport for the purpose of capturing better quality data in poor visibility environments. With existing techniques for measurement of medium scattering [13] and polarization properties [27], our simulation-based approach can be used to adapt the illumination and sensing system in-situ. Post-processing approaches are expected to recover the scene when applied to the images acquired using our system. Our analysis focused

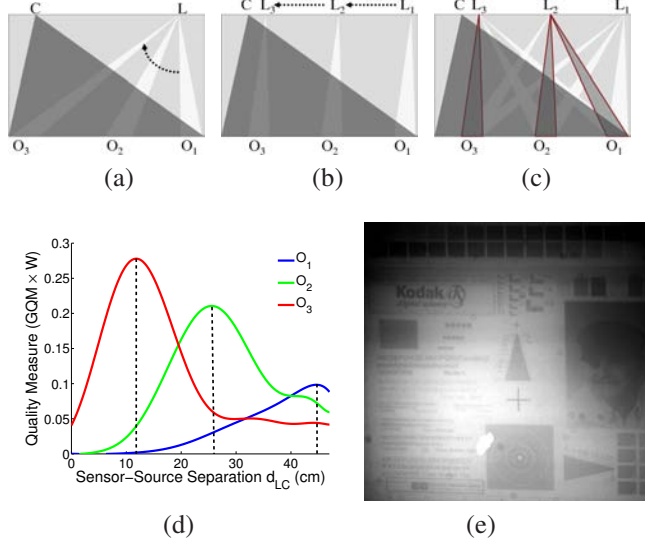


Figure 10. We can scan the scene (a) by rotating the source, (b) by translating it, or (c) a combination thereof. (d) Plot of $GQM \times W$ versus d_{LC} for different stripe locations O_1 , O_2 and O_3 , for our setup. We can notice different optimal separations for these stripe locations. (e) A high contrast image resulting from the optimal light stripe scan designed using simulations.

on a single divergent source and a single camera. It is worth extending our analysis to multiple cameras and sources [11]. More broadly, we believe that better control of the light transport can be achieved with greater flexibility in choosing illumination and viewing rays.

Acknowledgments

The authors thank Shahriar Negahdaripour for helpful discussions. This research was supported in parts by Grants # ONR N00014-08-1-0330, NSF CAREER IIS-0643628, NSF CCF-0541307 and the US-Israel Binational Science Foundation (BSF) Grant # 2006384. Yoav Schechner is a Landau Fellow - supported by the Taub Foundation. Yoav's work was conducted in the Ollendorff Minerva Center. Minerva is funded through the BMBF.

References

- [1] F. M. Caimi, F. R. Dalgleish, T. E. Giddings, J. J. Shirron, C. Mazel, and K. Chiang. Pulse versus CW laser line scan imaging detection methods: Simulation results. *In Proc. IEEE OCEANS*, pages 1–4, 2007. 5
- [2] S. Chandrasekhar. *Radiative Transfer*. Dover Publications, Inc., 1960. 1
- [3] T. Chen, H. P. A. Lensch, C. Fuchs, and H.-P. Seidel. Polarization and phase-shifting for 3D scanning of translucent objects. *In Proc. IEEE CVPR*, pages 1–8, 2007. 5
- [4] F. Cozman and E. Krotkov. Depth from scattering. *In Proc. IEEE CVPR*, pages 801–806, 1997. 1

- [5] T. E. Giddings, J. J. Shirron, and A. Tirat-Gefen. EODES-3: An electro-optic imaging and performance prediction model. *In Proc. IEEE OCEANS*, 2:1380–1387, 2005. 5
- [6] G. D. Gilbert and J. C. Pernicka. Improvement of underwater visibility by reduction of backscatter with a circular polarization technique. *Applied Optics*, 6(4):741–746, 1967. 2, 3
- [7] M. Gupta and S. G. Narasimhan. Light transport web-page. <http://graphics.cs.cmu.edu/projects/LightTransport/>. 1, 5
- [8] V. I. Haltrin. One-parameter two-term henyeey-greenstein phase function for light scattering in seawater. *Applied Optics*, 41(6):1022–1028, 2002. 3, 6
- [9] J. Jaffe. Computer modeling and the design of optimal underwater imaging systems. *IEEE Journal of Oceanic Engineering*, 15(2):101–111, 1990. 2, 4, 5
- [10] D. M. Kocak and F. M. Caimi. The current art of underwater imaging with a glimpse of the past. *MTS Journal*, 39:5–26, 2005. 2, 4
- [11] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas. Synthetic aperture confocal imaging. *ACM Trans. Graph.*, 23(3):825–834, 2004. 7
- [12] S. G. Narasimhan. Models and algorithms for vision through the atmosphere. In *Columbia Univ. Dissertation*, 2004. 1
- [13] S. G. Narasimhan, M. Gupta, C. Donner, R. Ramamoorthi, S. K. Nayar, and H. W. Jensen. Acquiring scattering properties of participating media by dilution. *ACM Trans. Graph.*, 25(3):1003–1012, 2006. 6, 7
- [14] S. G. Narasimhan and S. K. Nayar. Contrast restoration of weather degraded images. 25(6):713–724, 2003. 1
- [15] S. G. Narasimhan, S. K. Nayar, B. Sun, and S. J. Koppal. Structured light in scattering media. In *In Proc. IEEE ICCV*, pages 420–427, 2005. 2, 4, 5
- [16] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.*, 25(3):935–944, 2006. 2, 3
- [17] Y. Y. Schechner and Y. Averbuch. Regularized image recovery in scattering media. *IEEE Trans. PAMI*, 29(9):1655–1660, 2007. 5
- [18] Y. Y. Schechner and N. Karpel. Recovery of underwater visibility and structure by polarization analysis. *IEEE Journal of Oceanic Engineering*, 30(3):570–587, 2005. 1
- [19] Y. Y. Schechner, S. G. Narasimhan, and S. K. Nayar. Polarization-based vision through haze. *Applied Optics*, 42(3):511–525, 2003. 2, 3
- [20] W. A. Shurcliff and S. S. Ballard. *Polarized Light*, pages 98–103. Van Nostrand, Princeton, N.J., 1964. 2
- [21] B. Skerry and H. Hall. *Successful Underwater Photography*. New York: Amphoto books, 2002. 2, 5
- [22] K. Tan and J. P. Oakley. Physics-based approach to color image enhancement in poor visibility conditions. *JOSA A*, 18(10):2460–2467, 2001. 1
- [23] T. Treibitz and Y. Y. Schechner. Instant 3Descatter. In *Proc. IEEE CVPR*, volume 2, pages 1861–1868, 2006. 3, 8
- [24] T. Treibitz and Y. Y. Schechner. Active polarization descattering. *IEEE Trans. PAMI*, To appear, 2008. 2, 5
- [25] T. Treibitz, Y. Y. Schechner, and H. Singh. Flat refractive geometry. In *Proc. IEEE CVPR*, 2008. 2




Figure 11. Variation of the DOP of the scattered light, DOP_B , on the scattering angle. For low values of DOP_L , the curve qualitatively resembles that of Rayleigh scattering. For a completely polarized source, the curve is flatter, with an average value of ≈ 0.8 for backscattering angles ($> 90^\circ$).

- [26] H. van de Hulst. *Light Scattering by Small Particles*. Chapter 5. Wiley, New York, 1957. 8
- [27] K. J. Voss and E. S. Fry. Measurement of the mueller matrix for ocean water. *Applied Optics*, 23:4427–4439, 1984. 7, 8
- [28] S. Zhang and S. Negahdaripour. 3D shape recovery of planar and curved surfaces from shading cues in underwater images. *IEEE Journal of Oceanic Engineering*, 27:100–116, 2002. 1

A. Degree of Polarization of Scattering

In this appendix, we study the dependence of the DOP of the scattered light, DOP_B , on the scattering angle and the DOP of the incident light, DOP_L . We consider only the vertical and horizontal polarized components of linearly polarized light. Hence, we consider the first 2×2 sub-matrix of the full 4×4 Mueller matrix. Polarization properties of scattered light can be characterized by the Mueller matrix [26]:

$$\begin{bmatrix} I_B \\ Q_B \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix} \begin{bmatrix} I_L \\ Q_L \end{bmatrix}, \quad (14)$$

where I_L is the sum, and Q_L is the difference of the horizontal and vertically polarized components of the incident light. Similarly, I_B and Q_B are the sum and difference respectively of the scattered light. Note that $DOP = \frac{Q}{I}$. Consequently, based on Eq. (14):

$$DOP_B = \frac{m_{21} + m_{22} DOP_L}{m_{11} + m_{12} DOP_L}. \quad (15)$$

Using the above equation and the measured Mueller matrix data for ocean water [27], we plot DOP_B versus the scattering angle in Figure 11. For comparison, we also plot the behavior for Rayleigh scattering. For low values of DOP_L (natural light), the curve qualitatively resembles that of Rayleigh scattering. On the other hand, for a completely polarized source (for example, an LCD projector), the curve is flatter, with an average value of ≈ 0.8 for backscattering angles. Interestingly, this agrees with the observation made in [23] as well.



Vision and the Atmosphere

SRINIVASA G. NARASIMHAN AND SHREE K. NAYAR
Computer Science Department, Columbia University, New York, NY, USA

srinivas@cs.columbia.edu

nayar@cs.columbia.edu

Received January 31, 2001; Revised October 16, 2001; Accepted December 4, 2001

Abstract. Current vision systems are designed to perform in clear weather. Needless to say, in any outdoor application, there is no escape from “bad” weather. Ultimately, computer vision systems must include mechanisms that enable them to function (even if somewhat less reliably) in the presence of haze, fog, rain, hail and snow.

We begin by studying the visual manifestations of different weather conditions. For this, we draw on what is already known about atmospheric optics, and identify effects caused by bad weather that can be turned to our advantage. Since the atmosphere modulates the information carried from a scene point to the observer, it can be viewed as a mechanism of visual information coding. We exploit two fundamental scattering models and develop methods for recovering pertinent scene properties, such as three-dimensional structure, from one or two images taken under poor weather conditions.

Next, we model the chromatic effects of the atmospheric scattering and verify it for fog and haze. Based on this chromatic model we derive several geometric constraints on scene color changes caused by varying atmospheric conditions. Finally, using these constraints we develop algorithms for computing fog or haze color, depth segmentation, extracting three-dimensional structure, and recovering “clear day” scene colors, from two or more images taken under different but unknown weather conditions.

Keywords: physics based vision, atmosphere, bad weather, fog, haze, visibility, scattering, attenuation, airlight, overcast sky, scene structure, defog, dehaze

1. Computer Vision and the Weather

Virtually all work in computer vision is based on the premise that the observer is immersed in a transparent medium (air). It is assumed that light rays reflected by scene objects travel to the observer without attenuation or alteration. Under this assumption, the brightness of an image point depends solely on the brightness of a single point in the scene. Quite simply, existing vision sensors and algorithms have been created only to function on “clear” days. A dependable vision system however must reckon with the entire spectrum of weather conditions, including, haze, fog, rain, hail and snow.

The study of the interaction of light with the atmosphere (and hence weather) is widely known as atmospheric optics. Atmospheric optics lies at the heart of

the most magnificent visual experiences known to man, including, the colors of sunrise and sunset, the blueness of the clear sky, and the rainbow (see Minnaert (1954) and Henderson (1977)). The literature on this topic has been written over the past two centuries. A summary of where the subject as a whole stands would be too ambitious a pursuit. Instead, our objective will be to sieve out of this vast body of work, models of atmospheric optics that are of direct relevance to computational vision. Our most prominent sources of background material are the works of McCartney (1975) and Middleton (1952) whose books, though dated, serve as excellent reviews of prior work.

The key characteristics of light, such as its intensity and color, are altered by its interactions with the atmosphere. These interactions can be broadly classified

into three categories, namely, *scattering*, *absorption* and *emission*. Of these, scattering due to suspended particles is the most pertinent to us. As can be expected, this phenomenon leads to complex visual effects. So, at first glance, atmospheric scattering may be viewed as no more than a hindrance to an observer. However, it turns out that bad weather can be put to good use. The farther light has to travel from its source (say, a surface) to its destination (say, a camera), the greater it will be effected by the weather. Hence, bad weather could serve as a powerful means for coding and conveying scene structure. This observation lies at the core of our investigation; we wish to understand not only what bad weather does to vision but also what it can do for vision.

Surprisingly little work has been done in computer vision on weather related issues. An exception is the work of Cozman and Krotkov (1997) which uses the scattering models in McCartney (1975) to compute depth cues. Their algorithm assumes that all scene points used for depth estimation have the same intensity on a clear day. Since scene points can have their own reflectances and illuminations, this assumption is hard to satisfy in practice.

Research in image processing has been geared towards restoring contrast of images degraded by bad weather. Note that bad weather effects depend strongly on the depths of scene points. Hence, simple contrast enhancement techniques such as histogram equalization and contrast stretching do not suffice here. Oakley and Satherley (1998) use separately measured range data and describe an algorithm to restore contrast of atmospherically degraded images based on the principles of scattering. However, they approximate the distribution of radiances in the scene by a single gaussian with known variance. Kopeika (1998) and Yitzhaky et al. (1998) restore image contrast using weather predicted atmospheric modulation transfer function and an a priori estimate of the distance from which the scene was imaged.

The goal of our work is to lay the foundation for interpreting scenes from one or more images taken under bad weather conditions. We discuss various types of weather conditions and their formation processes. We summarize two models of atmospheric scattering—attenuation and airlight—that are most pertinent to us. Using these models, we develop algorithms that recover complete depth maps of scenes without requiring any prior information about the properties of the scene points or atmospheric conditions.

Next, we study the color effects of atmospheric scattering. A new model that describes the appearance of scene colors under bad weather is presented and verified for fog and haze. Based on this color model, we develop several geometric constraints on scene-color changes, caused by varying atmospheric conditions. Using these constraints, we present methods to recover structure as well as “clear day” scene colors from images taken under poor weather conditions. All of these methods only require changes in weather conditions and accurate measurement of scene irradiance, and not any prior information about the scene points or weather conditions.

2. Bad Weather: Particles in Space

Weather conditions differ mainly in the types and sizes of the particles involved and their concentrations in space. A great deal of effort has gone into measuring particle sizes and concentrations for a variety of conditions (see Table 1). Given the small size of air molecules, relative to the wavelength of visible light, scattering due to air is rather minimal. We will refer to the event of pure air scattering as a *clear* day (or night). Larger particles produce a variety of weather conditions which we will briefly describe below.

Haze. Haze is constituted of *aerosol* which is a dispersed system of small particles suspended in gas. Haze has a diverse set of sources including volcanic ashes, foliage exudation, combustion products, and sea salt (see Hidy (1972)). The particles produced by these sources respond quickly to changes in relative humidity and act as nuclei (centers) of small water droplets when the humidity is high. Haze particles are larger than air molecules but smaller than fog droplets. Haze tends to produce a distinctive gray or bluish hue and is certain to effect visibility.

Table 1. Weather conditions and associated particle types, sizes and concentrations (adapted from McCartney (1975)).

Condition	Particle type	Radius (μm)	Concentration (cm^{-3})
Air	Molecule	10^{-4}	10^{19}
Haze	Aerosol	10^{-2} –1	10^3 – 10
Fog	Water droplet	1–10	100–10
Cloud	Water droplet	1–10	300–10
Rain	Water drop	10^2 – 10^4	10^{-2} – 10^{-5}

Fog. Fog evolves when the relative humidity of an air parcel reaches saturation. Then, some of the nuclei grow by condensation into water droplets. Hence, fog and certain types of haze have similar origins and an increase in humidity is sufficient to turn haze into fog. This transition is quite gradual and an intermediate state is referred to as *mist*. While perceptible haze extends to an altitude of several kilometers, fog is typically just a few hundred feet thick. A practical distinction between fog and haze lies in the greatly reduced visibility induced by the former. There are many types of fog (ex., radiation fog, advection fog, etc.) which differ from each other in their formation processes (Myers, 1968).

Cloud. A cloud differs from fog only in existing at higher altitudes rather than sitting at ground level. While most clouds are made of water droplets like fog, some are composed of long ice crystals and ice-coated dust grains. Details on the physics of clouds and precipitation can be found in Mason (1975). For now, clouds are of less relevance to us as we restrict ourselves to vision at ground level rather than high altitudes.

Rain and snow. The process by which cloud droplets turn to rain is a complex one (Mason, 1975). When viewed up close, rain causes random spatial and temporal variations in images and hence must be dealt with differently from the more static weather conditions mentioned above. Similar arguments apply to snow, where the flakes are rough and have more complex shapes and optical properties (Koenderink and Richards, 1992; Ohtake, 1970). Snow too, we will set aside for now.

3. Mechanisms of Atmospheric Scattering

The manner in which a particle scatters incident light depends on its material properties, shape and size. The exact form and intensity of the scattering pattern varies dramatically with particle size (Minnaert, 1954). As




Figure 1. A particle in the path of an incident light wave abstracts and reradiates incident energy. It therefore behaves like a point source of light. The exact scattering function is closely related to the ratio of particle size to wavelength of incident light. (Adapted from Minnaert (1954)).

seen in Fig. 1, a small particle (about $1/10 \lambda$, where λ is the wavelength of light) scatters almost equally in the forward (incidence) and backward directions, a medium size particle (about $1/4 \lambda$) scatters more in the forward direction, and a large particle (larger than λ) scatters almost entirely in the forward direction. Substantial theory has been developed to derive scattering functions and their relations to particle size distributions (Mie, 1908; Hulst, 1957; Chandrasekhar, 1960; Chu and Hogg, 1968; Rensch and Long, 1970; Nieto-Vesperinas and Dainty, 1990).

Figure 1 illustrates scattering by a single particle. Clearly, particles are accompanied in close proximity by numerous other particles. However, the average separation between atmospheric particles is several times the particle size. Hence, the particles can be viewed as *independent* scatterers whose scattered intensities do not interfere with each other. This does not imply that the incident light is scattered only by a single particle. *Multiple* scatterings take place and any given particle is exposed not only to the incident light but also light scattered by other particles. A simple analogy is the inter-reflections between scene points. In effect, multiple scattering causes the single scattering functions in Fig. 1 to get smoother and less directional.

Now, consider the simple illumination and detection geometry shown in Fig. 2. A unit volume of scattering medium with suspended particles is illuminated with spectral irradiance $E(\lambda)$ per cross section area. The radiant intensity $I(\theta, \lambda)$ of the unit volume in the direction θ of the observer is (see McCartney (1975)):

$$I(\theta, \lambda) = \beta(\theta, \lambda)E(\lambda), \quad (1)$$

where, $\beta(\theta, \lambda)$ is the *angular scattering coefficient*. The radiant intensity $I(\theta, \lambda)$ is the flux radiated per unit solid angle, per unit volume of the medium. The irradiance $E(\lambda)$ is, as always, the flux incident on the volume




Figure 2. A unit volume of randomly oriented suspended particles illuminated and observed.

per unit cross-section area. The total flux scattered (in all directions) by this volume is obtained by integrating over the entire sphere:

$$\phi(\lambda) = \beta(\lambda)E(\lambda), \quad (2)$$

where, $\beta(\lambda)$ is the *total scattering coefficient*. It represents the ability of the volume to scatter flux of a given wavelength in all directions. It is generally assumed that the coefficient $\beta(\lambda)$ is constant (homogeneous medium) over horizontal paths. To satisfy this constraint, we will restrict ourselves to the case where the observer is at (or close to) ground level and is interested not in the sky but other objects on (or close to) ground level. Also, we will assume that the atmosphere is more or less homogeneous in the scene of interest.

3.1. Attenuation

The first mechanism that is relevant to us is the attenuation of a beam of light as it travels through the atmosphere. This causes the radiance of a scene point to fall as its depth from the observer increases. Here, we will summarize the derivation of the attenuation model given in McCartney(1975). Consider a collimated beam of light incident on the atmospheric medium, as shown in Fig. 3. The beam is assumed to have unit cross-sectional area. Consider the beam passing through an infinitesimally small sheet (lamina) of thickness dx . The fractional change in irradiance at location x can be written as:

$$\frac{dE(x, \lambda)}{E(x, \lambda)} = -\beta(\lambda) dx. \quad (3)$$




Figure 3. Attenuation of a collimated beam of light by suspended particles. The attenuation can be derived by viewing the medium as a continuum of thin sheets.

By integrating both sides between the limits $x = 0$ and $x = d$ we get:

$$E(d, \lambda) = E_o(\lambda)e^{-\beta(\lambda)d}, \quad (4)$$

where, $E_o(\lambda)$ is the irradiance at the source ($x = 0$). This is Bouguer's exponential law of attenuation (Bouguer, 1729). At times, attenuation due to scattering is expressed in terms of *optical thickness*, $T = \beta(\lambda)d$. The utility of Bouguer's law is somewhat limited as it assumes a collimated source of incident energy. This is easily remedied by incorporating the inverse-square law for diverging beams from point sources:

$$E(d, \lambda) = \frac{I_o(\lambda)e^{-\beta(\lambda)d}}{d^2}, \quad (5)$$

where, $I_o(\lambda)$ is the radiant intensity of the point source. This is Allard's law (Allard, 1876). (See Hardy (1967) for an analysis of the applicability of the inverse square criterion for sources of various sizes.)

In deriving Allard's law, we have assumed that all scattered flux is removed from the incident energy. The fraction of energy that remains is called *direct transmission* and is given by expression (5). We have ignored the flux scattered in the forward direction (towards the observer) by each particle. Fortunately, this component is small in vision applications since the solid angles subtended by the source and the observer with respect to each other are small (see Middleton (1949)). In the remainder of the paper, we refer to the terms direct transmission model and attenuation model interchangeably.

Finally, in some situations such as heavy fog, the exponential law may not hold due to significant multiple scatterings of light by atmospheric particles. We will assume here that once light flux is scattered out of a column of atmosphere (seen by a pixel, say), it does not re-enter the *same* column (or only an insignificant amount does). Multiple scattering can also cause blurring in the image of a scene. In other words, the flux scattered out of an atmospheric column (visible to a pixel) enters another column (seen by a *neighboring* pixel). In this work, we do not model the blurring effects of multiple scattering.

3.2. Airlight

A second mechanism causes the atmosphere to behave like a source of light. This phenomenon is called airlight (Koschmieder, 1924) and it is caused by the scattering of environmental illumination by particles in the

atmosphere. The environmental illumination can have several sources, including, direct sunlight, diffuse skylight and light reflected by the ground. While attenuation causes scene radiance to decrease with pathlength, airlight increases with pathlength. It therefore causes the apparent brightness of a scene point to increase with depth. We now build upon McCartney's (1975) derivation of airlight as a function of pathlength.

Consider the illumination and observation geometry shown in Fig. 4. The environmental illumination along the observer's line of sight is assumed to be constant but unknown in direction, intensity and spectrum. In effect, the cone of solid angle $d\omega$ subtended by a single receptor at the observer's end, and truncated by a physical object at distance d , can be viewed as a source of airlight. The infinitesimal volume dV at distance x from the observer may be written as the product of the cross section area, $d\omega x^2$, and thickness dx :

$$dV = d\omega x^2 dx. \quad (6)$$

Irrespective of the exact type of environmental illumination incident upon dV , its intensity due to scattering in the direction of the observer is:

$$dI(x, \lambda) = dV k\beta(\lambda) = d\omega x^2 dx k\beta(\lambda), \quad (7)$$

where, $\beta(\lambda)$ is the total scattering coefficient and the proportionality constant k accounts for the exact nature of the illumination and the form of the scattering function.




Figure 4. The cone of atmosphere between an observer and an object scatters environmental illumination in the direction of the observer. It therefore acts like a source of light, called airlight, whose brightness increases with pathlength.

If we view element dV as a source with intensity $dI(x, \lambda)$, the irradiance it produces at the observer's end, after attenuation due to the medium, is given by (5):

$$dE(x, \lambda) = \frac{dI(x, \lambda)e^{-\beta(\lambda)x}}{x^2}. \quad (8)$$

We can find the radiance of dV from its irradiance as:

$$dL(x, \lambda) = \frac{dE(x, \lambda)}{d\omega} = \frac{dI(x, \lambda)e^{-\beta(\lambda)x}}{d\omega x^2}. \quad (9)$$

By substituting (7) we get $dL(x, \lambda) = k\beta(\lambda)e^{-\beta(\lambda)x} dx$. Now, the total radiance of the pathlength d from the observer to the object is found by integrating this expression between $x = 0$ and $x = d$:

$$L(d, \lambda) = k(1 - e^{-\beta(\lambda)d}). \quad (10)$$

If the object is at an infinite distance (at the *horizon*), the radiance of airlight is maximum and is found by setting $d = \infty$ to get $L_\infty(\lambda) = k$. Therefore, the radiance of airlight for any given pathlength d is:

$$L(d, \lambda) = L_\infty(\lambda)(1 - e^{-\beta(\lambda)d}). \quad (11)$$

As expected, the radiance of airlight for an object right in front of the observer ($d = 0$) equals zero. Of great significance to us is the fact that the above expression no longer includes the unknown factor k . Instead, we have the airlight radiance $L_\infty(\lambda)$ at the horizon, which is an observable.

3.3. Overcast Sky Illumination

Allard's attenuation model in (5) is in terms of the radiant intensity of a *point source*. This formulation does not take into account the sky illumination and its reflection by scene points. We make two simplifying assumptions regarding the illumination received by a scene point. Then, we reformulate the attenuation model in terms of sky illumination and the *BRDF* of scene points.

Usually, the sky is overcast under foggy conditions. So, we use the overcast sky model for environmental illumination (Gordon and Church, 1966; Moon and Spencer, 1942). We also assume that the irradiance at each scene point is dominated by the radiance of the sky, and that the irradiance due to other scene points is not significant. In Appendix A, we show that the

attenuated irradiance at the observer is given by,

$$E(d, \lambda) = g \frac{L_\infty(\lambda) \rho(\lambda) e^{-\beta(\lambda)d}}{d^2}. \quad (12)$$

where $L_\infty(\lambda)$ is the horizon radiance. $\rho(\lambda)$ represents the sky aperture (the cone of sky visible to a scene point), and the reflectance of the scene point in the direction of the viewer. g represents the optical settings of the camera (aperture, for instance). Note that we refer to (5) as the direct transmission model while dealing with images of light sources taken at night. However, while dealing with images of scenes taken during daylight, we refer to (12) as the direct transmission model.

4. Depths of Light Sources from Attenuation

Consider the image of an urban setting taken at *night* (see Fig. 5). Environmental illumination of the scene due to sunlight, skylight and reflected ground light are minimal and hence airlight can be safely ignored. The bright points in the image are mainly sources of light such as street lamps and windows of lit rooms. On a clear night, these sources are visible to a distant observer in their brightest and clearest forms. As haze or fog sets in, the radiant intensities of the sources diminish due to attenuation. Our goal here is to recover the relative depths of the sources in the scene from two images taken under different (unknown) atmospheric conditions.

Since environmental illumination is negligible at night, the image irradiance of a light source in the scene can be expressed using the attenuation model (5) as:

$$E(d, \lambda) = g \frac{I_o(\lambda) e^{-\beta(\lambda)d}}{d^2}, \quad (13)$$




Figure 5. The relative depths of sources of unknown intensities can be recovered from two images taken under different but unknown atmospheric conditions.

where, $I_o(\lambda)$ is the radiant intensity of the source, d is the distance between the source and the camera and the constant gain g accounts for the optical parameters (aperture, for instance) of the camera. It is important to note that $\beta(\lambda)$ is the total scattering coefficient and not the angular one. We are assuming here that the lines of sight are not too inclined and hence all lines of sight pass through the same atmospheric conditions. This removes all dependence on the exact form of the scattering function; the attenuation is determined by a single coefficient $\beta(\lambda)$ that is independent of viewing direction.

If the detector of the camera has spectral response $s(\lambda)$, the final image brightness recorded is determined as:

$$E' = \int s(\lambda) E(d, \lambda) d\lambda = \int g s(\lambda) \frac{I_o(\lambda) e^{-\beta(\lambda)d}}{d^2} d\lambda. \quad (14)$$

For the visible light spectrum, the relationship between the scattering coefficient β , and the wavelength λ , is given by the inverse power law (analogous to Rayleigh's law for small air particles):

$$\beta(\lambda) = \frac{\text{Constant}}{\lambda^\gamma}, \quad (15)$$

where $\gamma \in [0, 4]$. Fortunately, for fog and dense haze, $\gamma \approx 0$ (see Middleton (1952) and McCartney (1975)). In these cases, β does *not* change appreciably with wavelength. Furthermore, since the spectral bandwidth of the camera is rather limited (visible light range for a gray-scale camera, and even narrower spectral bands when the camera is color), we will assume the scattering coefficient $\beta(\lambda)$ to be constant over this bandwidth. Then, we have:

$$E' = g \frac{e^{-\beta d}}{d^2} \int s(\lambda) I(\lambda) d\lambda = g \frac{e^{-\beta d}}{d^2} I'. \quad (16)$$

Now consider two different weather conditions, say, mild and dense fog. Or, one of the conditions could be clear with $\beta = 0$. In either case we have two different attenuation coefficients, β_1 and β_2 . If we take the ratio of the two resulting image brightness values, we get:

$$R = \frac{E'_1}{E'_2} = e^{-(\beta_1 - \beta_2)d}. \quad (17)$$

Using the natural log, we obtain: $R' = \ln R = -(\beta_1 - \beta_2)d$. This quantity is independent of the sensor gain and the radiant intensity of the source. In fact, it is

nothing but the *difference in optical thicknesses* (DOT) of the source for two weather conditions. In the atmospheric optics literature, the term DOT is used as a quantitative measure of the “change” in weather conditions. Now, if we compute the DOTs of two different light sources in the scene (see Fig. 5) and take their ratio, we determine the relative depths of the two source locations:

$$\frac{R'_i}{R'_j} = \frac{d_i}{d_j} \quad (18)$$

Hence, the relative depths of all sources (with unknown radiant intensities) in the scene can be computed from two images taken under unknown but different haze or fog conditions. Since we may not entirely trust the DOT computed for any single source, the above calculation may be made more robust by using:

$$\frac{R'_i}{\sum_{j=0}^{j=N} R'_j} = \frac{d_i}{\sum_{j=0}^{j=N} d_j} \quad (19)$$

By setting the denominator on the right hand side to an arbitrary constant we have computed the depths of all sources in the scene up to a scale factor.

Figure 6 shows experimental results on the recovery of light sources from night images. This experiment and all subsequent ones are based on images acquired using a Nikon N90s SLR camera and a Nikon LS-2000 slide scanner. All images are linearized using the radiometric response curve of the imaging system that is computed off-line using a color chart. Figure 6(a) shows a clear day image of a scene with five lamps. This image is provided only to give the reader an idea of where the lamps are located in the scene. Figures 6(b) and (c) are clear night and foggy night images of the same scene. The above algorithm for depth estimation was used to recover the locations of all five light sources up to a scale factor. Figure 6(d) shows different perspectives of the recovered coordinates of the lamps in three-dimensional space. The poles and the ground plane are added only to aid visualization of the results.

5. Structure from Airlight

Under dense fog and close by objects or mild fog and distant objects, attenuation of object brightness is severe and airlight is the main cause of image irradiance. Also, in the case of dense haze around noon, airlight dominates. In such cases, airlight causes object brightness to increase with distance from the observer. Here,

we present a simple method for computing scene structure from a single airlight image. A different but related method for computing depth cues was proposed by Cozman and Krotkov (1997).

Let a scene point at depth d produce airlight radiance $L(d, \lambda)$. If our camera has a spectral response $s(\lambda)$, the final brightness value recorded for the scene point is:

$$E'(d) = \int g s(\lambda) L(d, \lambda) d\lambda, \quad (20)$$

where, g accounts for the constant of proportionality between scene radiance and image irradiance (Horn, 1986). Substituting the model for airlight given by (11) we get:

$$E'(d) = \int g s(\lambda) L_\infty(\lambda) (1 - e^{-\beta(\lambda)d}) d\lambda \quad (21)$$

where, $L_\infty(\lambda)$ is again the radiance of airlight at the horizon. As before, we will assume that the scattering coefficient $\beta(\lambda)$ is more or less constant over the spectral band of the camera. This allows us to write:

$$E'(d) = E_\infty (1 - e^{-\beta d}). \quad (22)$$

Let us define:

$$S = \frac{E_\infty - E'(d)}{E_\infty}. \quad (23)$$

By substituting (22) in the above expression and taking the natural logarithm, we get:

$$S' = \ln S = -\beta d. \quad (24)$$

Hence, the three-dimensional structure of the scene can be recovered up to a scale factor (the scattering coefficient β) from a single image. Clearly, at least a small part of the horizon must be visible to obtain E_∞ . If so, this part is easily identified as the brightest region of the image. If there is a strong (directional) sunlight component to the illumination, scattering would be greater in some directions and airlight could be dependent on viewing direction. This problem can be alleviated by using the horizon brightness E_∞ that lies closest to the scene point under consideration. Figure 7 shows the structure of an urban setting computed from a hazy image taken around noon, and the structure of a mountain range computed using a foggy image. Given that some of the objects are miles away




Figure 6. Relative depths from brightness decay of point sources at night. (a) A scene with five light sources (street lamps). This image is shown only to convey the relative locations of the sources to the reader. (b) An image of the scene taken on a clear night. (c) An image of the scene taken on a foggy night. The three-dimensional coordinates of the five sources were computed from images (b) and (c). (d) Rotated graphical illustrations used to demonstrate the computed lamp coordinates (small bright spheres). The lamp poles and the ground plane are added only to aid visualization.

from the camera, such scene structures are hard to compute using stereo or structure from motion. An interesting study of the visibility of distant mountains taking into account earth's curvature can be found in Porch (1975).

6. Dichromatic Atmospheric Scattering

Thus far, we have not taken into account the chromatic effects of atmospheric scattering. Furthermore, we have described attenuation and airlight separately.

However, in most situations the effects of both attenuation and airlight coexist. In the remainder of the paper, we discuss the chromatic effects of atmospheric scattering that include both attenuation and airlight, and hence develop a general framework for analyzing color images taken in bad weather. For this, we first present a new model that describes the appearance of scene colors in poor visibility conditions.

As we know, attenuation causes the radiance of the surface to decay as it travels to the observer. In addition, if the particle sizes are comparable to the wavelengths




Figure 7. Structure from one image taken under dense fog/haze. (Left column) (a) Image of an urban scene taken under noon haze. (b) Depth map of the scene computed using the image in (a). (c) A three-dimensional rendering of the scene. (Right column) (d) Image of a mountain range taken under foggy conditions. (e) Depth map computed from the image in (d). (f) A three-dimensional rendering of the scene. Some of the objects in these scenes are several kilometers away from the camera.

of the reflected light, the spectral composition of the reflected light can be expected to vary as it passes through the medium. For fog and dense haze, these shifts in the spectral composition are minimal (see Middleton (1952) for details), and hence we may assume the hue of direct transmission to be *independent of the depth* of the reflecting surface. The hue of airlight depends on the particle size distribution and tends to be gray or light blue in the case of haze and fog. Therefore, the final spectral distribution $E(d, \lambda)$ received by the observer is a sum of the distributions $D(d, \lambda)$ of directly transmitted light and $A(d, \lambda)$ of airlight, which are de-

termined by the attenuation model (12) and the airlight model (11) respectively:

$$\begin{aligned} E(d, \lambda) &= D(d, \lambda) + A(d, \lambda), \\ D(d, \lambda) &= g \frac{e^{-\beta(\lambda)d}}{d^2} L_{\infty}(\lambda) \rho(\lambda), \\ A(d, \lambda) &= g (1 - e^{-\beta(\lambda)d}) L_{\infty}(\lambda). \end{aligned} \quad (25)$$

Here, $L_{\infty}(\lambda)$ is the radiance of the horizon ($d = \infty$), and g is a constant that accounts for the optical settings of the imaging system. $\rho(\lambda)$ represents the reflectance

properties and sky aperture of the scene point. We refer to the above expression as the *dichromatic atmospheric scattering model*. It is similar in its spirit to the dichromatic reflectance model (Shafer, 1985) that describes the spectral effects of diffuse and specular surface reflections. A fundamental difference here is that one of our chromatic components is due to surface and volume scattering (transmission of reflected light) while the other is due to pure volume scattering (airlight). If a chromatic filter with a spectral response $f(\lambda)$ is incorporated into the imaging system, image irradiance is obtained by multiplying (25) by $f(\lambda)$ and integrating over λ :

$$E^{(f)}(d) = D^{(f)}(d) + A^{(f)}(d). \quad (26)$$

In the case of a color image detector several such filters (say, red, green and blue) with different sensitivities are used to obtain a color measurement vector. The dichromatic model can then be written as:

$$\mathbf{E}(d) = \mathbf{D}(d) + \mathbf{A}(d) \quad (27)$$

where, $\mathbf{E} = [E^{(f_1)}, E^{(f_2)}, \dots, E^{(f_n)}]^T$. As we mentioned earlier (see (15)), for fog and haze, the dependence of the scattering coefficient $\beta(\lambda)$ on the wavelength (within the small bandwidth of the camera) of light tends to be rather small. Therefore, except in the case of certain types of metropolitan haze, we may assume the scattering coefficient to be constant with respect to wavelength ($\beta(\lambda) = \beta$). Then, expression (26) may be simplified as:

$$E^{(f)}(d) = p'(d)D^{(f)} + q'(d)A^{(f)}, \quad (28)$$

where:

$$\begin{aligned} D^{(f)} &= \int g f(\lambda) L_\infty(\lambda) \rho(\lambda) d\lambda, \\ A^{(f)} &= \int g f(\lambda) L_\infty(\lambda) d\lambda, \\ p'(d) &= \frac{e^{-\beta d}}{d^2}, \quad q'(d) = (1 - e^{-\beta d}). \end{aligned} \quad (29)$$

Here, $D^{(f)}$ is the image irradiance due to the scene point *without* atmospheric attenuation and $A^{(f)}$ is the image irradiance at the horizon in the presence of bad weather. We are assuming here that the clear and bad weather have illuminations with similar spectral distributions. Hence, the color measurement given by (27) can be rewritten as: $\mathbf{E}(d) = p'(d)\mathbf{D} + q'(d)\mathbf{A}$. Since the intensity of illumination (or magnitude of the illumination




Figure 8. Dichromatic atmospheric scattering model. The color \mathbf{E} of a scene point on a foggy or hazy day, is a linear combination of the direction $\hat{\mathbf{D}}$ of direct transmission color, and the direction $\hat{\mathbf{A}}$ of airlight color.

spectrum) at a scene point is expected to vary between clear and bad weather, it is more convenient to write:

$$\mathbf{E}(d) = m|L_\infty|p'(d)\hat{\mathbf{D}} + n|L_\infty|q'(d)\hat{\mathbf{A}} \quad (30)$$

where $\hat{\mathbf{D}}$ and $\hat{\mathbf{A}}$ are unit vectors and m and n are scalars. $|L_\infty|$ is the magnitude of the illumination spectrum. For convenience, the dichromatic model is re-written as:

$$\mathbf{E} = p\hat{\mathbf{D}} + q\hat{\mathbf{A}}, \quad (31)$$

where p is the magnitude of direct transmission, and q is the magnitude of airlight (see Fig. 8). From (30) we have,

$$p = \frac{E_\infty r e^{-\beta d}}{d^2}, \quad q = E_\infty (1 - e^{-\beta d}). \quad (32)$$

where $E_\infty = n|L_\infty|$, is termed as the sky intensity and $r = m/n$ is a function that depends on the properties of the scene point (reflectance and sky aperture). For our analysis, the exact nature of r is not important; it suffices to note that r does not depend on the weather condition β .¹ This simplified dichromatic scattering model will prove useful in the coming sections when we attempt to recover scene structure and remove weather effects from images.

It is easy to see that the simplified dichromatic model (31) is *linear* in color space. In other words, $\hat{\mathbf{D}}$, $\hat{\mathbf{A}}$ and \mathbf{E} lie on the same *dichromatic plane* in color space. As




Figure 9. For fog and haze, the transmittance ($e^{-\beta(\lambda)d}$) does not vary appreciably with wavelength within the visible spectrum. The plots were generated using the atmospheric transmission software MODTRAN 4.0, with a fixed viewing geometry (distance, d and viewing directions are fixed).

stated earlier, we impose the restriction that the hue of illumination under various weather conditions remains the same although its intensity can vary. It follows that the unit vectors $\hat{\mathbf{D}}$ and $\hat{\mathbf{A}}$ do not change due to different atmospheric conditions (say, mild fog and dense fog). Therefore, the colors of any scene point, observed under different atmospheric conditions, lie on a single dichromatic plane (see Fig. 10(a)).

We performed simulations using the atmospheric transmission software MODTRAN 4.0 (Acharya et al., 1999) to verify that the scattering coefficient does not vary with wavelength within the visible spectrum ($0.4\mu\text{--}0.7\mu$). Figure 9 shows plots of transmittance ($e^{-\beta(\lambda)d}$) for a particular viewing geometry in fog and haze respectively. The distance from the observer to the scene was fixed at $d = 0.2$ km and the viewing direction was fixed at 5 degrees off the ground plane. The plots show that the variation in β is very small within the visible spectrum.

Experiments with real scenes (shown in Fig. 17) were performed to verify this model under three different fog and haze conditions. The sky was overcast in all these conditions. The images used contained around half a million pixels. The dichromatic plane for each pixel was computed by fitting a plane to the colors of that pixel, observed under the three atmospheric conditions. The error of the plane-fit was computed in terms of the angle between the observed color vectors and the estimated plane. The average absolute error (in de-




Figure 10. (a) Dichromatic plane geometry and its verification. The observed color vectors E_i of a scene point under different (two in this case) foggy or hazy conditions lie on a plane called the dichromatic plane. (b) Experimental verification of the dichromatic model with two scenes imaged under three different foggy and hazy conditions, respectively. The error was computed as the mean angular deviation (in degrees) of the observed scene color vectors from the estimated dichromatic planes, over half a million pixels in the images.

grees) for all the pixels in each of the two scenes is shown in Fig. 10(b). The small error values indicate that the dichromatic model indeed works well for fog and haze.

7. Weather Removal and Structure using Chromatic Decomposition

Consider color images of a scene taken under clear weather and foggy or hazy weather. Assume that the clear day image is taken under environmental illumination with similar spectral characteristics as the bad weather image. If not, a white patch in the scene may be used to apply the needed color corrections. The sky in the bad weather image reveals the *direction* of the airlight color $\hat{\mathbf{A}}$. The *direction* of the color $\hat{\mathbf{D}}$ of each scene point is revealed by the clear weather image. Therefore, Eq. (31) can be used to decompose the bad weather color \mathbf{E} at each pixel into its two components and determine the airlight magnitude $q(d)$. The resulting airlight image is then used to compute a depth map as described in Section 5. Figure 11 shows experimental results obtained using the above decomposition method. Figure 12 demonstrates a simple form of weather removal by defogging windows of buildings.

In computing depth from the airlight component, we have assumed that the atmosphere itself is uniformly illuminated. Consider a pathlength that extends from a point on a building to an observer. Clearly, atmospheric points closer to the building see less of the sky due to occlusion by the building. This effect increases towards the foot of the building. Some of the errors in the computed structure can be attributed to this *illumination*

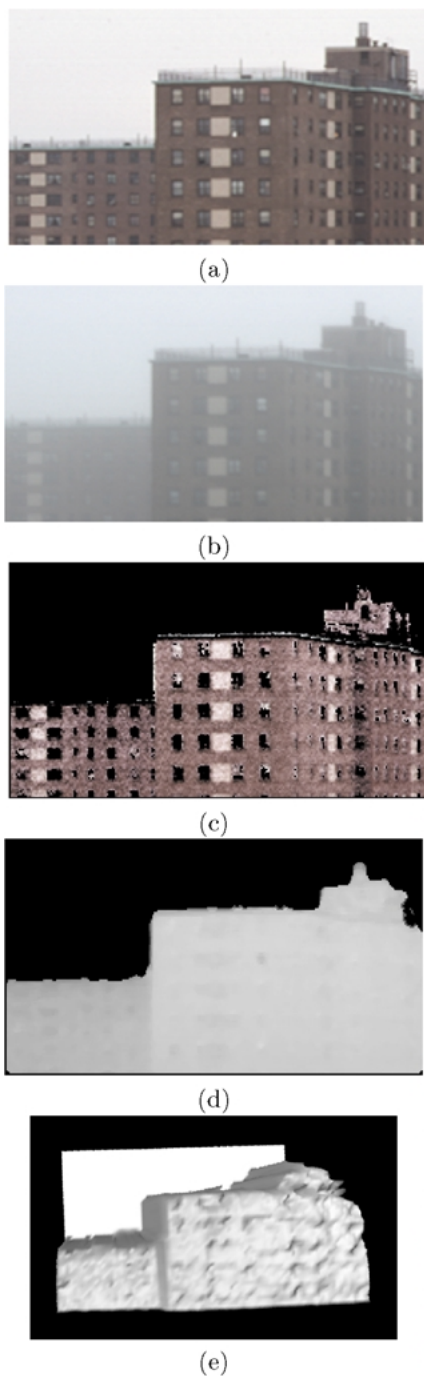


Figure 11. Structure from chromatic decomposition. (a) Clear day image of buildings. (b) Foggy day image of the same scene. (c) The direct transmission component (brightened) estimated by the chromatic decomposition algorithm. Black and gray points (windows) are discarded due to lack of color. (d) Depth map of the scene computed from the airlight component (depths of window areas are interpolated). (e) A three-dimensional rendering of the computed depth map.



Figure 12. Demonstration of fog removal. (a) A clear day image of a building taken under an overcast sky. The color directions (and not magnitudes) of scene points (non-window regions) are recorded as true colors or clear day colors. (b) A foggy day image of the same scene again captured under an overcast sky. Note: Even though both images in (a) and (b) were taken on overcast days (i.e., spectral composition of the daylight on both days are more or less identical), the horizon brightnesses (and/or camera exposure parameters) can vary. (c) The true colors recorded were used to decompose the foggy image into direct transmission and airlight components. The airlight component was subtracted from the window regions to demonstrate a simple form of weather removal.

occlusion effect (see Appendix B for a more detailed treatment).

Finally, there are certain limitations to this type of decomposition. First, we cannot decompose (31) if both the airlight and scene points have the same color. Also, this algorithm for chromatic decomposition is restrictive since it requires a clear day image of the scene. In the remainder of the paper, we develop more general constraints and algorithms to compute structure as well as recover “clear day” colors, without requiring a clear day image of the scene.

8. Computing the Direction of Airlight Color

The direction of airlight (fog or haze) color can be simply computed by averaging a patch of the sky on a foggy or hazy day (as was done in Section 7), or from scene points whose direct transmission color is *black*.² However, these methods necessitate either (a) the inclusion of a part of the sky (which is more prone to color saturation or clipping) in the image or (b) a clear day image of the scene with *sufficient* black points to yield a robust estimate of the direction of airlight color. Here, we present a method that does not require either the sky or a clear day image, to compute the direction of airlight color.

Figure 13 illustrates the dichromatic planes for two scene points P_i and P_j , with different direct transmission colors $\hat{\mathbf{D}}^{(i)}$ and $\hat{\mathbf{D}}^{(j)}$. The dichromatic planes Q_i




Figure 13. Intersection of two different dichromatic planes yields the direction $\hat{\mathbf{A}}$ of airlight color.

and Q_j are given by their normals,

$$\begin{aligned}\mathbf{N}_i &= \mathbf{E}_1^{(i)} \times \mathbf{E}_2^{(i)}, \\ \mathbf{N}_j &= \mathbf{E}_1^{(j)} \times \mathbf{E}_2^{(j)}.\end{aligned}\quad (33)$$

Since the direction $\hat{\mathbf{A}}$ of the airlight color is the same for the entire scene, it must lie on the dichromatic planes of all scene points. Hence, $\hat{\mathbf{A}}$ is given by the intersection of the two planes Q_i and Q_j ,

$$\hat{\mathbf{A}} = \frac{\mathbf{N}_i \times \mathbf{N}_j}{\|\mathbf{N}_i \times \mathbf{N}_j\|}. \quad (34)$$

In practice, scenes have several points with different colors. Therefore, we can compute a robust intersection of several dichromatic planes by minimizing the objective function

$$\epsilon = \sum_i (\mathbf{N}_i \cdot \hat{\mathbf{A}})^2. \quad (35)$$

Thus, we are able to compute the color of fog or haze using only the observed colors of the scene points under two atmospheric conditions, and not relying on a patch of the sky being visible in the image.

We verified the above method for the two scenes shown in Fig. 17. First, the direction of airlight color was computed using (35). Then, we compared it with the direction of the airlight color obtained by averaging an unsaturated patch of the sky. For the two scenes, the angular deviations were found to be 1.2° and 1.6° respectively. These small errors in the computed directions of airlight color indicate the robustness of the method.

9. Dichromatic Constraints for Iso-depth Scene Points

In this section, we derive a simple constraint for scene points that are at the same depth from the observer. This constraint can then be used to segment the scene based on depth, *without* knowing the actual reflectances of the scene points and their sky apertures. For this, we first prove the following lemma.

Lemma. *Ratios of the direct transmission magnitudes for points under two different weather conditions are equal, if and only if the scene points are at equal depths from the observer.*

Proof: Let β_1 and β_2 be two unknown weather conditions with horizon brightness values $E_{\infty 1}$ and $E_{\infty 2}$.

Let P_i and P_j be two scene points at depths d_i and d_j , from the observer. Also, let $r^{(i)}$ and $r^{(j)}$ represent sky apertures and reflectances of these points.

From (32), the direct transmission magnitudes of P_i under β_1 and β_2 , can be written as

$$p_1^{(i)} = \frac{E_{\infty_1} r^{(i)} e^{-\beta_1 d_i}}{d_i^2}, \quad p_2^{(i)} = \frac{E_{\infty_2} r^{(i)} e^{-\beta_2 d_i}}{d_i^2}.$$

Similarly, the direct transmission magnitudes of P_j under β_1 and β_2 , are

$$p_1^{(j)} = \frac{E_{\infty_1} r^{(j)} e^{-\beta_1 d_j}}{d_j^2}, \quad p_2^{(j)} = \frac{E_{\infty_2} r^{(j)} e^{-\beta_2 d_j}}{d_j^2}.$$

Then, we immediately see that the relation:

$$\frac{p_2^{(i)}}{p_1^{(i)}} = \frac{p_2^{(j)}}{p_1^{(j)}} = \left(\frac{E_{\infty_2}}{E_{\infty_1}} \right) e^{-(\beta_2 - \beta_1)d}, \quad (36)$$

holds if and only if $d_i = d_j = d$. \square

So, if we have the ratio of direct transmissions for each pixel in the image, we can group the scene points according to their depths from the observer. But how do we compute this ratio for any scene point without knowing the actual direct transmission magnitudes?

Consider the dichromatic plane geometry for a scene point P , as shown in Fig. 14. Here, we denote a vector by the line segment between its end points. Let p_1 and p_2 be the unknown direct transmission magnitudes of P




Figure 14. Geometric constraint for iso-depth scene points. The ratio p_2/p_1 of the direct transmissions for a scene point under two different atmospheric conditions is equal to the ratio $|E_2 A_1|/|E_1 O|$ of the parallel sides. Shaded triangles are similar.

under β_1 and β_2 , respectively. Similarly let q_1 and q_2 be the unknown airlight magnitudes for P under β_1 and β_2 .

We define a magnitude $|OA_t|$ on the airlight vector such that $\overline{E_2 A_t} \parallel \overline{E_1 O}$. Also, since the *direction* of direct transmission color for a scene point does not vary due to different atmospheric conditions, $\overline{E_1 A_1} \parallel \overline{E_2 A_2}$. Here A_1 and A_2 correspond to the end points of the airlight magnitudes of P under β_1 and β_2 , as shown in Fig. 14. Thus, $\triangle E_1 O A_1 \sim \triangle E_2 A_t A_2$. This implies,

$$\frac{p_2}{p_1} = \frac{q_2 - |OA_t|}{q_1} = \frac{|E_2 A_t|}{|E_1 O|}. \quad (37)$$

Since the right hand side of (37) can be computed using the observed color vectors of the scene point P , we can compute the ratio (p_2/p_1) of direct transmission magnitudes for P under two atmospheric conditions. Therefore, from (36), we have a simple method to find points at the same depth, *without* having to know their reflectances and sky apertures.

Let us now consider the numerical stability of the direct transmission ratio (37). Under heavy fog/haze (or when the dynamic range of the sensor is low), the direct transmission magnitudes are low and their ratio could be unstable. In such cases, the ratio constraint can be supported by another constraint for depth segmentation we describe briefly. Consider the dichromatic planes of two different scene points as illustrated in Fig. 15. It




Figure 15. Another geometric constraint for two scene points to be equidistant from the observer. The dichromatic planes for the two points P_1 and P_2 are shown. Note that superscripts denote scene points while subscripts denote weather conditions. Shaded triangles are similar if and only if P_1 and P_2 are equidistant from the observer. Hence the iso-depth constraint is $\overline{E_1^{(1)} E_1^{(2)}} \parallel \overline{E_2^{(1)} E_2^{(2)}}$.

can be shown (using the geometric analysis in Fig. 14) that the shaded triangles are similar if and only if the two scene points are at equal depths from the observer. Therefore, the constraint for two scene points to be iso-depth is given in terms of observables,

$$\overline{E_1^{(1)} E_1^{(2)}} \parallel \overline{E_2^{(1)} E_2^{(2)}}. \quad (38)$$

Using the constraints in (36) and (38) a sequential labeling like algorithm can be used to efficiently segment scenes into regions of equal depth.

10. Scene Structure from Two Bad Weather Images

We extend the direct transmission ratio constraint given in (36) one step further and present a method to construct the complete structure of an arbitrary scene, from two images taken under poor weather conditions.

From (36), the ratio of direct transmissions of a scene point P under two atmospheric conditions, is given by

$$\frac{p_2}{p_1} = \frac{E_{\infty_2}}{E_{\infty_1}} e^{-(\beta_2 - \beta_1)d}. \quad (39)$$

Note that we have already computed the left hand side of the above equation using (37). Taking natural logarithms on both sides, we get

$$(\beta_2 - \beta_1)d = \ln \left(\frac{E_{\infty_2}}{E_{\infty_1}} \right) - \ln \left(\frac{p_2}{p_1} \right). \quad (40)$$

So, if we know the horizon brightness values, E_{∞_1} and E_{∞_2} , then we can compute the scaled depth $(\beta_2 - \beta_1)d$ at P . As before, $(\beta_2 - \beta_1)d$ is just the difference in optical thicknesses (DOT) for the pathlength d , under the two weather conditions.

10.1. Estimation of E_{∞_1} and E_{∞_2}

The expression for scaled depth give in (40), includes the horizon brightness values, E_{∞_1} and E_{∞_2} . These two terms are *observables* only if some part of the sky is visible in the image. However, the brightness values within the region of the image corresponding to the sky, cannot be trusted since they are prone to intensity saturation and color clipping. Here, we estimate E_{∞_1}

and E_{∞_2} using only points in the “non-sky” region of the scene.

Let q_1 and q_2 denote the magnitudes of airlight for a scene point P under atmospheric conditions β_1 and β_2 . Using (32), we have

$$q_1 = E_{\infty_1}(1 - e^{-\beta_1 d}), \quad q_2 = E_{\infty_2}(1 - e^{-\beta_2 d}). \quad (41)$$

Therefore,

$$\frac{E_{\infty_2} - q_2}{E_{\infty_1} - q_1} = \frac{E_{\infty_2}}{E_{\infty_1}} e^{-(\beta_2 - \beta_1)d}. \quad (42)$$

Substituting (39), we can rewrite the previous equation as

$$\left(\frac{p_2}{p_1} \right) = \frac{q_2 - c}{q_1}, \quad \text{where, } c = E_{\infty_2} - \left(\frac{p_2}{p_1} \right) E_{\infty_1}. \quad (43)$$

Comparing (43) and (37), we get $c = |OA_i|$ (see Fig. 14). hence, the expression for c in (43) represents a straight line equation in the unknown parameters, E_{∞_1} and E_{∞_2} . Now consider several pairs of $\{c^{(i)}, (p_2^{(i)}/p_1^{(i)})\}$ corresponding to scene points P_i , at different depths. Then, the estimation of E_{∞_1} and E_{∞_2} is reduced to a *line fitting* problem. Quite simply, we have shown that the horizon brightnesses under different weather conditions can be computed using *only* non-sky scene points.

Since both the terms on the right hand side of (40) can be computed for every scene point, we have a simple algorithm for computing the scaled depth at each scene point, and hence the complete scene structure, from two images taken under different atmospheric conditions.

10.2. Experimental Results

We now present results showing scene structure recovered from both synthetic and real images. The synthetic scene we used is shown on the left side of Fig. 16(a) as a 200×200 pixel image with 16 color patches. The colors in this image represent the direct transmission or “clear day” colors of the scene. We assigned a random depth value to each color patch. The rotated 3D structure of the scene is shown on the right side of Fig. 16(a). Then, two different levels of fog ($\beta_1/\beta_2 = 0.67$) were added to the synthetic scene according to the dichromatic model. To test robustness, we added noise to the foggy images. The noise was randomly selected from a uniformly distributed color cube of dimension 10. The




Figure 16. Experiments with a synthetic scene. (a) On the left, a 200×200 pixel image representing a synthetic scene with 16 color patches, and on the right, its rotated 3D structure. (b) Two levels of fog ($\beta_1/\beta_2 = 0.67$) are added to the synthetic image according to the dichromatic model. To test robustness, noise is added by random selection from a uniformly distributed color cube of dimension 10. (c) The recovered structure (3×3 median filtered).

resulting two foggy (and noisy) images are shown in Fig. 16(b). The structure shown in 16(c) is recovered from the two foggy images using the technique we described above.

Simulations were repeated for the scene in Fig. 16(a) for two relative scattering coefficient values (β_1/β_2), and three different noise levels. Once again, the noise was randomly selected from a uniformly distributed color cube of dimension η . Table 2 shows results of simulations for two parameter sets $\{\beta_1/\beta_2, E_{\infty_1}, E_{\infty_2}\} = \{0.5, 100, 255\}$ and $\{0.67, 200, 400\}$. The computed




Figure 17. Structure from two bad weather images. (a) A scene imaged under two different foggy conditions. (b) Depth map computed from images in (a). (c) Another scene imaged under two different hazy conditions. (d) Depth map computed from images in (c). All these images were captured under overcast sky conditions.

values for E_{∞_1} , E_{∞_2} , and the percentage RMS error in the recovered scaled depths, computed over all 200×200 pixels are given. These results show that our method for recovering structure is robust for reasonable amounts of noise.

Experiments with two real scenes under foggy and hazy conditions are shown in Fig. 17. The first of the two scenes was imaged under two foggy conditions, and is shown in 17(a). The second scene was imaged

Table 2. Simulations were repeated for the scene in Fig. 16(a), for two sets of parameter values, and three different noise levels. Noise was randomly selected from a uniformly distributed color cube of dimension η .

Noise (η)	0	5	10	15
Actual values $\{\beta_1/\beta_2, E_{\infty_1}, E_{\infty_2}\} = \{0.5, 100, 255\}$				
Estimated E_{∞_1}	100	108.7	109.2	119.0
Estimated E_{∞_2}	255	262.7	263.6	274.0
Depth error (%)	0.0	7.14	11.7	15.3
Actual values $\{\beta_1/\beta_2, E_{\infty_1}, E_{\infty_2}\} = \{0.67, 200, 400\}$				
Estimated E_{∞_1}	200	204.3	223.7	249.5
Estimated E_{∞_2}	400	403.8	417.5	444.2
Depth error (%)	0.0	12.3	15.3	17.8

under two hazy conditions as shown in 17(c). Figure 17(b) and (d) shows the corresponding recovered depth maps.

11. Clear Day Scene Colors

As we stated in the beginning of the paper, most outdoor vision applications perform well only under clear weather. Any discernible amount of scattering due to fog or haze in the atmosphere, hinders a clear view of the scene. Earlier we presented a simple form of weather removal that requires a clear day image of the scene (see Figs. 11 and 12 in Section 7). In this section, we compute the scene colors as they would appear on a clear but overcast day from two bad weather images. More precisely, we compute the direct transmission colors of the entire scene using minimal a priori scene information. For this, we first show that, given additional scene information (airlight or direct transmission vector) at a *single* point in the scene, we can compute the clear day colors of the *entire* scene from two bad weather images.

Consider the dichromatic model given in (31). The color of a scene point P_i under weather condition β is,

$$\mathbf{E}^{(i)} = p^{(i)}\hat{\mathbf{D}}^{(i)} + q^{(i)}\hat{\mathbf{A}}, \quad (44)$$

where $p^{(i)}$ is the direct transmission magnitude, and $q^{(i)}$ is the airlight magnitude of P_i . Suppose that the direction $\hat{\mathbf{D}}^{(i)}$ of direct transmission color for a *single* point P_i is given. Besides, the direction $\hat{\mathbf{A}}$ of airlight color for the entire scene can be estimated using (35). Therefore, the coefficients $p^{(i)}$ and $q^{(i)}$ can be computed using (44). Furthermore, the optical thickness βd_i of P_i can be computed from (32).

Since we have already shown how to compute the scaled depth of every scene point (see (40)), the relative depth d_j/d_i of any other scene point P_j with respect to P_i can be computed using the ratio of scaled depths. Hence, the optical thickness and airlight for the scene point P_j , under the same atmospheric condition are given by

$$\begin{aligned} \beta d_j &= \beta d_i (d_j/d_i), \\ q^{(j)} &= E_{\infty}(1 - e^{-\beta d_j}). \end{aligned} \quad (45)$$

Finally, the direct transmission color vector of P_j can be computed as

$$p^{(j)}\hat{\mathbf{D}}^{(j)} = \mathbf{E}^{(j)} - q^{(j)}\hat{\mathbf{A}}. \quad (46)$$

Thus, given a *single* measurement (in this case, the direction of direct transmission color of a single scene point), we have shown that the direct transmission and airlight color vectors of any other point, and hence the entire scene can be computed. But how do we specify the clear day color of any scene point without actually capturing the clear day image?

For this, we assume that there exists at least one scene point whose direct transmission color \mathbf{D} lies on




Figure 18. The observed color \mathbf{E} of a scene point, its airlight direction $\hat{\mathbf{A}}$ and clear day color direction $\hat{\mathbf{D}}$ are shown in the R-G-B color cube. \tilde{q} is the distance from \mathbf{E} to a surface of the cube along negative $\hat{\mathbf{A}}$. For scene points whose clear day colors do not lie on the cube surface, \tilde{q} is greater than the true airlight magnitude q .




Figure 19. [(a) and (c)] Clear day scene colors recovered from the two foggy and hazy images shown in Fig. 17(a) and (c) respectively. The colors in some of the dark window interiors are dominated by airlight and thus their clear day colors are computed to be black. The images are median filtered to reduce noise and brightened for display purposes. [(b) and (d)] Actual clear day images of the scenes are shown for qualitative comparison. Note: The clear day images on the right and the bad weather images (Fig. 17) were captured on different days. Some differences between actual and recovered clear day colors are due to the different spectral distributions of illumination in the scene, during image acquisition.

the surface of the color cube (including origin or black) and we wish to identify such point(s) in the scene automatically. Consider the R-G-B color cube in Fig. 18. If the clear day color of a scene point lies *on* the surface





Figure 20. (a) and (b) Foggy images of a scene under an overcast sky. (c) Defogged image. (d) Actual clear day image taken under a partly cloudy sky.



Relative Depth	Ground Truth	Computed Value
d_3 / d_1	9.55	10.457
d_3 / d_2	1.3687	1.354
d_3 / d_4	1.489	1.61
d_3 / d_5	1.083	1.124

(b)

Figure 21. (a) Depth map computed using images in Fig. 20(a) and (b). Depth map is brightened for display purposes. (b) Table comparing the computed relative depths with ground truth relative depths of 5 different regions, $d_1 - d_4$, in the scene. The relative depths are averaged over small neighborhoods. Note that scaled depth can be computed only approximately due to the illumination occlusion problem (see Appendix B for more details). The depths in some window interiors are not reliable since they changed during acquisition of images over time.

of the color cube, then the computed \tilde{q} is equal to the airlight magnitude q of that point. However, if it lies *within* the color cube, then clearly $\tilde{q} > q$. For each point P_i , we compute $\tilde{q}^{(i)}$ and optical thickness $\tilde{\beta}_1 d_i$. Note that $\tilde{\beta}_1 d_i$ may or may not be the correct optical thickness. We normalize the optical thicknesses of the scene points by their scaled depths (DOTs) to get

$$\tilde{\alpha}_i = \frac{\tilde{\beta}_1 d_i}{(\beta_2 - \beta_1) d_i}. \quad (47)$$

For scene points that do not lie on the color cube surface, $\tilde{\alpha}_i$ is greater than what it should be. Since we have assumed that there exists at least one scene point whose clear day color is on the surface of the cube, it must be the point that has the minimum $\tilde{\alpha}_i$. So, $\tilde{q}^{(i)}$ of that point is its true airlight. Hence, from (45), the airlights and direct transmission colors of the entire scene can be computed without using a clear day image. For robustness, we use k least $\tilde{\alpha}_i$'s. We call this the *Color Cube Boundary Algorithm*.

Figure 19 illustrates experiments with real scenes. Usually in urban scenes, window interiors have very little color of their own. Their intensities are solely due to airlight and not due to direct transmission. In other words, their direct transmission color is black (the origin of the color cube). We detected such points in the scene using the above technique and recovered the clear day colors of foggy and hazy scenes. A second result is shown in Figs. 20 and 21.

12. Summary

Research in atmospheric optics has been around for over two centuries. The physical processes that govern the effects of atmospheric scattering on scene appearance are well established. This article is just an initial attempt at understanding and exploiting the manifestations of weather in order to interpret, recover and render scenes under various atmospheric conditions. We summarized existing models in atmospheric optics and proposed new ones, keeping in mind the constraints faced by most vision applications. We presented several simple algorithms for recovering scene structure from one or two bad weather images and demonstrated that bad weather can be put to good use. Using scene structure, algorithms to remove weather effects were developed. We intend to use these results as building blocks for developing more advanced weather-tolerant vision techniques. Potential applications of this work are in outdoor surveillance, navigation, underwater explorations and image based rendering.

Appendix A: Direct Transmission Under Overcast Skies

We present an analysis of the effect of sky illumination and its reflection by a scene point, on the direct transmission from the scene point. For this, we make two simplifying assumptions on the illumination received by scene points. Usually, the sky is overcast under foggy conditions. So we use the overcast sky model




Figure 22. The illumination geometry of a scene point P with surface normal $\hat{\mathbf{n}}$. The irradiance of P is due to the airlight radiance of its sky aperture Ω .

(Gordon and Church, 1966; IRIA, 1978) for environmental illumination. We also assume that the irradiance of each scene point is dominated by the radiance of the sky, and that the irradiance due to other scene points is not significant. See Langer and Zucker's work (1994) for a related analysis.

Consider the illumination geometry shown in Fig. 22. Let P be a point on a surface and $\hat{\mathbf{n}}$ be its normal. We define the *sky aperture* Ω of point P , as the cone of sky visible from P . Consider an infinitesimal patch of the sky, of size $\delta\theta$ in polar angle and $\delta\phi$ in azimuth as shown in Fig. 22. Let this patch subtend a solid angle $\delta\omega$ at P . For overcast skies, Moon (Moon and Spencer, 1942) and Gordon (Gordon and Church, 1966) have shown that the radiance of the infinitesimal cone $\delta\Omega$, in the direction (θ, ϕ) is given by $L(\theta, \phi) = L_\infty(\lambda)(1 + 2 \cos \theta)\delta\omega$, where $\delta\omega = \sin \theta \delta\theta \delta\phi$. Hence, the irradiance at P due to the entire aperture Ω , is given by

$$E(\lambda) = \iint_{\Omega} L_\infty(\lambda)(1 + 2 \cos \theta) \cos \theta \sin \theta d\theta d\phi, \quad (48)$$

where $\cos \theta$ accounts for foreshortening (Horn, 1986). If R is the *BRDF* of P , then the radiance from P toward the observer can be written as

$$L_o(\lambda) = \iint_{\Omega} L_\infty(\lambda) f(\theta) R(\theta, \phi, \lambda) d\theta d\phi, \quad (49)$$

where $f(\theta) = (1 + 2 \cos \theta) \cos \theta \sin \theta$. Let σ be the projection of a unit patch around P , on a plane perpendicular to the viewing direction. Then, the radiant intensity of P is given by $I_o(\lambda) = \sigma L_o(\lambda)$. Since $L_\infty(\lambda)$ is a constant with respect to θ and ϕ , we can factor it

out of the integral and write concisely as

$$I_o(\lambda) = L_\infty(\lambda) \rho(\lambda), \quad (50)$$

where

$$\rho(\lambda) = \sigma \iint_{\Omega} f(\theta) R(\theta, \phi, \lambda) d\theta d\phi. \quad (51)$$

The term $\rho(\lambda)$ represents the sky aperture and the reflectance in the direction of the viewer. Substituting for $I_o(\lambda)$ in the direct transmission model in (5), we obtain

$$E(d, \lambda) = g \frac{L_\infty(\lambda) \rho(\lambda) e^{-\beta(\lambda)d}}{d^2}, \quad (52)$$

where g represents the optical setting of the camera (exposure, for instance). We have thus formulated the direct transmission model in terms of overcast sky illumination and the reflectance of the scene points.

Appendix B: Illumination Occlusion Problem

In deriving the expression for the radiance due to airlight in Section 3.2, we assumed that the atmosphere is illuminated uniformly regardless of the type of illumination. This is not always true since not all points in the atmosphere “see” the same solid angle of the sky. In fact, the scene itself occludes part of the sky hemisphere visible to a point in the atmosphere. For explanation purposes, consider a scene with a single building. The solid angle subtended at any point in the atmosphere by the sky is called its sky aperture. As seen in Fig. 23, this solid angle decreases as the distance increases from the observer for any given pathlength. Similarly, the solid angle is smaller for points near the bottom of the building.

We now present a simplified analysis of this effect. We assume that the atmosphere is illuminated mainly by overcast skylight (ground light is ignored here). Then, the irradiance received by any point in the atmosphere is given by (see Eq. (48)),

$$\begin{aligned} E &= E^{(\text{hemisphere})} - E^{(\text{occluded})}, \\ E^{(\text{occluded})} &= \int_{-\phi}^{\phi} \int_0^{\theta} L_\infty(1 + 2 \cos \theta) \\ &\quad \times \cos \theta \sin \theta d\theta d\phi, \\ E^{(\text{hemisphere})} &= \int_{-\pi}^{\pi} \int_0^{\pi/2} L_\infty(1 + 2 \cos \theta) \\ &\quad \times \cos \theta \sin \theta d\theta d\phi, \end{aligned} \quad (53)$$




Figure 23. The scene occludes the sky aperture of points in the atmosphere. As a result points in the atmosphere are not uniformly illuminated by the sky.

where $E^{(hemisphere)}$ is the irradiance the point would receive from the entire sky hemisphere (as if there were no occlusions). $E^{occluded}$ is the irradiance the point would have received from the occluded part. θ and ϕ denote the polar and azimuth of the occluded region. The above equation simplifies to

$$E = L_{\infty} \frac{7\pi - 7\phi \cos^2 \theta (3 + 4 \cos \theta)}{3}. \quad (54)$$

To correct for the radiance of airlight in Section 3.2, we multiply by the fraction of irradiance received by each point and rewrite the airlight radiance (10) of a pathlength d as

$$\begin{aligned} L(d, \lambda) &= k(1 - e^{-\beta(\lambda)d}) - \int_0^d k \left(\frac{\phi \cos^2 \theta (3 + 4 \cos \theta)}{\pi} \right) \\ &\quad \times \beta(\lambda) e^{-\beta(\lambda)x} dx. \end{aligned} \quad (55)$$

Note here that both θ and ϕ depend on the depth from the observer x (see Fig. 23). In other words, the integral in the previous equation depends on the exact extent of occlusion by the scene. In our experiments, we have assumed uniform illumination of the atmosphere and thus some of the errors in the depth maps can be attributed to this effect.

Acknowledgments

This work was supported in parts by a DARPA/ONR HumanID Contract (N00014-00-1-0916), an NSF Award (IIS-99-87979), and a DARPA/ONR MURI Grant (N00014-95-1-0601). The authors thank Jan Koenderink of Utrecht University for pointers to early work on atmospheric optics. The authors also thank Yoav Schechner for the discussions on this topic that helped improve the paper. Some of the results presented in this paper have appeared in the proceedings of the IEEE International Conference on Computer Vision'1999 (Nayar and Narasimhan, 1999), IEEE Conference on Computer Vision and Pattern Recognition'2000 (Narasimhan and Nayar, 2000) and SPIE Conference on Human Vision and Electronic Imaging'2001 (Narasimhan and Nayar, 2001).

Notes

1. We do not handle situations where wet materials may appear darker than dry materials.
2. Sky and black points take on the color of airlight on a bad weather day.

References

- Acharya, P.K., Berk, A., Anderson, G.P., Larsen, N.F., Tsay, S.C., and Stamnes, K.H. 1999. Modtran4: Multiple scattering and BRDF upgrades to modtran. In *SPIE Proc. Optical Spectroscopic Techniques and Instrumentation for Atmospheric and Space Research III*, p. 3756.
- Allard, E. 1876. *Memoire sur l'intensite' et la portee des phares*. Dunod: Paris.
- Bouguer, P. 1729. *Traite' d'optique sur la gradation de la lumiere*.
- Chandrasekhar, S. 1960. *Radiative Transfer*. Dover Publications: New York.
- Chu, T.S. and Hogg, D.C. 1968. Effects of precipitation on propagation at 0.63, 3.5 and 10.6 microns. *The Bell System Technical Journal*.
- Cozman, F. and Krotkov, E. 1997. Depth from scattering. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition*, vol. 31, pp. 801–806.
- Gordon, J. and Church, P. 1966. Overcast sky luminances and directional luminous reflectances of objects and backgrounds under overcast skies. *Applied Optics*, 5:919.
- Hardy, A.C. 1967. How large is a point source? *Journal of Optical Society of America*, 57(1).
- Henderson, S.T. 1977. *Daylight and its Spectrum*. Wiley: New York.
- Hidy, G.M. 1972. *Aerosols and Atmospheric Chemistry*. Academic Press: New York.
- Horn, B.K.P. 1986. *Robot Vision*. The MIT Press: Cambridge, MA.
- IRIA. 1978. *The Infrared Handbook*. Infrared Information and Analysis Center, Environmental Research Institute of Michigan.

- Koenderink, J.J. and Richards, W.A. 1992. Why is snow so bright? *Journal of Optical Society of America*, 9(5):643–648.
- Kopeika, N.S. 1998. *A System Engineering Approach to Imaging*. SPIE Press.
- Koschmieder, H. 1924. Theorie der horizontalen sichtweite. *Beitr. Phys. Freien Atm.*, 12:33–53, 171–181.
- Langer, M.S. and Zucker, S.W. 1994. Shape from shading on a cloudy day. *JOSA-A*, 11(2):467–478.
- Mason, B.J. 1975. *Clouds, Rain, and Rainmaking*. Cambridge University Press: Cambridge.
- McCartney, E.J. 1975. *Optics of the Atmosphere: Scattering by Molecules and Particles*. John Wiley and Sons: New York.
- Middleton, W.E.K. 1949. The effect of the angular aperture of a telephotometer on the telephotometry of collimated and non-collimated beams. *Journal of Optical Society of America*, 39:576–581.
- Middleton, W.E.K. 1952. *Vision Through the Atmosphere*. University of Toronto Press.
- Mie, G. 1908. A contribution to the optics of turbid media, especially colloidal metallic suspensions. *Ann. of Physics*, 25(4):377–445.
- Minnaert, M. 1954. *The Nature of Light and Color in the Open Air*. Dover: New York.
- Moon, P. and Spencer, D.E. 1942. Illumination from a non-uniform sky. *Illum Eng.*, 37:707–726.
- Myers, J.N. 1968. Fog. *Scientific American*, pp. 75–82.
- Narasimhan, S.G. and Nayar, S.K. 2000. Chromatic framework for vision in bad weather. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Narasimhan, S.G. and Nayar, S.K. 2001. Vision and the weather. In *Proceedings of SPIE Conference on Human Vision and Electronic Imaging VI*, p. 4299.
- Nayar, S.K. and Narasimhan, S.G. 1999. Vision in bad weather. In *Proceedings of the 7th International Conference on Computer Vision*.
- Nieto-Vesperinas, M. and Dainty, J.C. 1990. *Scattering in Volumes and Surfaces*. North-Holland: New York.
- Oakley, J.P. and Satherley, B.L. 1998. Improving image quality in poor visibility conditions using a physical model for degradation. *IEEE Trans. on Image Processing*, 7.
- Ohtake, T. 1970. Factors affecting the size distribution of raindrops and snowflakes. *Journal of Atmospheric Science*, 27:804–813.
- Porch, W.M. 1975. Visibility of distant mountains as a measure of background aerosol pollution. *Applied Optics*, 14.
- Rensch, D.B. and Long, R.K. 1970. Comparative studies of extinction and backscattering by aerosols, fog, and rain at 10.6 and 0.63 microns. *Applied Optics*, 9(7).
- Shafer, S. 1985. Using color to separate reflection components. *Color Research and Applications*, pp. 210–218.
- Van De Hulst, 1957. *Light Scattering by Small Particles*. John Wiley and Sons: New York.
- Yitzhaky, Y., Dror, I., and Kopeika, N.S. 1998. Restoration of atmospherically blurred images according to weather-predicted atmospheric modulation transfer function. *Optical Engineering*, 36.

Contrast Restoration of Weather Degraded Images

Srinivasa G. Narasimhan and Shree K. Nayar

Abstract—Images of outdoor scenes captured in bad weather suffer from poor contrast. Under bad weather conditions, the light reaching a camera is severely scattered by the atmosphere. The resulting decay in contrast varies across the scene and is exponential in the depths of scene points. Therefore, traditional space invariant image processing techniques are not sufficient to remove weather effects from images. In this paper, we present a physics-based model that describes the appearances of scenes in uniform bad weather conditions. Changes in intensities of scene points under different weather conditions provide simple constraints to detect depth discontinuities in the scene and also to compute scene structure. Then, a fast algorithm to restore scene contrast is presented. In contrast to previous techniques, our weather removal algorithm does not require any a priori scene structure, distributions of scene reflectances, or detailed knowledge about the particular weather condition. All the methods described in this paper are effective under a wide range of weather conditions including haze, mist, fog, and conditions arising due to other aerosols. Further, our methods can be applied to gray scale, RGB color, multispectral and even IR images. We also extend our techniques to restore contrast of scenes with moving objects, captured using a video camera.

Index Terms—Physics-based vision, atmosphere, bad weather, fog, haze, visibility, scattering, attenuation, airlight, overcast sky, scene structure, defog, dehaze, contrast restoration, shape from X, shape from weather, scene reconstruction.

1 TOWARD WEATHER-FREE VISION

HUMAN perception of scene color and contrast through the atmosphere has been extensively studied [14], [15], [17], [18]. For centuries, artists have rendered their paintings with an “atmospheric or aerial perspective” [7]. They illustrate, in their paintings, optical phenomena such as the bluish haze of distant mountains and reduced visibility under adverse weather conditions such as mist, fog, rain, and snow. Leonardo da Vinci’s paintings often contain an atmospheric perspective of the background scene [26], where farther scene points were painted brighter and bluer. While these optical phenomena can be argued to be aesthetically pleasing to humans, they are often hindrances to the satisfactory working of a computer vision system.

Most outdoor vision applications such as surveillance, terrain classification, and autonomous navigation require robust detection of image features. Under bad weather conditions, however, the contrast and color of images are drastically altered or degraded. Hence, it is imperative to remove weather effects from images in order to make vision systems more reliable. Unfortunately, the effects of bad weather increase exponentially with the distances of scene points from the sensor. As a result, conventional space invariant filtering techniques fail to adequately remove weather effects from images.

Recently, there has been an increased interest in the image processing and vision communities on issues related to imaging under bad weather. Kopeika [13] and Yitzhaky et al. [38] deblur atmospherically degraded

images using a weather-predicted atmospheric modulation transfer function, and an a priori estimate of the distance from which the scene is imaged. Oakley and Satherley [25] and Tan and Oakley [33], [34] describe a physics-based method to restore scene contrast without using predicted weather information. However, they assume that scene depths are known beforehand, and they approximate the distribution of radiances in the scene by a single Gaussian with known variance. Another work by Grewe and Brooks [9] uses wavelet-based fusion of multiple bad weather images to get a less blurred image.

Narasimhan and Nayar [21] analyze color variations in the scene under different weather conditions based on the dichromatic atmospheric scattering model proposed in [24]. Using constraints on scene color changes, they compute complete 3D structure and recover clear day scene colors from two or more bad weather images [23]. However, they assume that the atmospheric scattering properties do *not* change with the wavelength of light. This property holds over the visible spectrum only for certain weather conditions such as fog and dense haze. For several aerosols, however, scattering strongly depends on the wavelength of incident light. Furthermore, scene recovery using the dichromatic model is ambiguous for scene points whose colors match the color of fog or haze.

Polarization has been used as a cue to reduce haze in images based on the effects of scattering on light polarization [2], [5], [27]. In many works [4], [28], the radiation from the object of interest is assumed to be polarized, whereas the natural illumination scattered toward the observer (airlight) is assumed to be unpolarized. In other works [6], [29], [36], the radiation from the scene of interest is assumed to be unpolarized, whereas airlight is assumed to be partially polarized. Polarizing filters are, therefore, used widely by photographers to reduce haziness in landscape images, where the radiance from the landscapes is generally unpolarized.

• The authors are with the Computer Science Department, Columbia University, 500 West 120th Street, Room 450, New York, NY 10027.
E-mail: {srinivas, nayar}@cs.columbia.edu.

Manuscript received 6 Feb. 2002; revised 25 Sept. 2002; accepted 4 Oct. 2002.
Recommended for acceptance by R. Beveridge.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 115843.

However, polarization filtering alone does not ensure complete removal of haze. Schechner et al. [29], [30] further analyzed two or more polarization filtered images to compute scene structure and dehaze images. The effectiveness of polarization as a cue to remove weather effects is limited under dense fog and mist with overcast sky illumination since scattered light is mostly depolarized.

In this paper, we present a physics-based method to restore contrast of a scene from two or more images taken in uniform bad weather conditions. A monochrome atmospheric scattering model that describes how scene intensities are affected by homogeneous weather conditions is presented. This model is valid in both the visible and near-IR spectra, and for a wide range of weather conditions such as mist, haze, fog, and other aerosols. The model does not require the scattering properties of the atmosphere to be constant with respect to wavelength of light over a large spectral range (for example, the range of visible spectrum, as in [23]). Since we are interested in a short range of distances (of the order of a few kilometers), we assume that the weather condition does not change spatially in the field of view.¹

Using the monochrome weather model, we show how contrast of a scene degrades with distance. We conclude that standard contrast enhancement techniques can only handle a scene or a region within a scene at a fixed distance from the sensor. A simple contrast restoration technique similar to contrast stretching is derived for scenes where depth segmentation is known a priori. Changes in scene intensities, observed under different weather conditions, present strong physical constraints regarding scene structure. These constraints are exploited to automatically detect depth discontinuities in the scene and also to recover complete scene structure from two images taken under different weather conditions during daytime. Using the computed structure, contrast is restored from a single weather-degraded image of the scene. Unlike previous methods for contrast restoration, we do not need accurately predicted weather information or prior distributions on scene radiances. We extend our algorithms to handle video and describe a simple heuristic to restore contrasts of moving objects in the scene whose depths are unknown.

The entire analysis in this paper is done for monochrome (single narrow spectral band) images. However, the same methods can be applied independently to images with multiple spectral bands. We show that our methods can be applied to images taken using gray scale, wide-band RGB, multispectral, and also narrow-band IR cameras.

2 ATMOSPHERIC SCATTERING MODELS

Scattering of light by physical media has been one of the main topics of research in the atmospheric optics and astronomy communities. In general, the exact nature of scattering is highly complex and depends on the types, orientations, sizes, and distributions of particles constituting the media, as well as wavelengths, polarization states, and directions of the incident light [3], [10]. Here, we focus on two models—attenuation and airlight, which form the basis of our work. Since we are interested in a short range of

distances (of the order of a few kilometers), we assume that the properties of the weather condition (say, type of particles and their density) does not change spatially. In other words, we only consider homogeneous atmospheres in this paper. Also, given the limited dynamic range of the sensors (say, 8 bits per pixel), we do not explicitly model multiple scattering or blurring effects of bad weather [23].

2.1 Attenuation and Airlight

The attenuation model describes the way light gets attenuated as it traverses from a scene point to the observer. Due to atmospheric scattering, a fraction of light flux is removed from the incident beam. The unscattered flux, called direct transmission, is transmitted to the observer. The attenuated irradiance at the observer is given by (see [17], [21]),

$$E_{dt}(d, \lambda) = \frac{E_{\infty}(\lambda) r(\lambda) e^{-\beta(\lambda)d}}{d^2}, \quad (1)$$

where, d is the depth of the scene point from the observer and λ is the wavelength. $\beta(\lambda)$ is called the *scattering coefficient* of the atmosphere; it represents the ability of a unit volume of atmosphere to scatter light in all directions. $\beta(\lambda)d$ is called the *optical depth* of the scene point. E_{∞} is the *horizon brightness* and r is a function that describes the reflectance properties and the sky aperture² of the scene point. The attenuated irradiance is illustrated by the solid arrow from the scene to the sensor in Fig. 1. The sky is assumed to be mostly cloudy or overcast and that the radiance from the sky varies smoothly with respect to the polar and azimuth angles of the hemisphere [8], [11], [20]. For more details, we refer the reader to [21].

The second atmospheric scattering model we consider is called the *airlight* model. The airlight model quantifies how a column of atmosphere acts as a light source by reflecting environmental illumination towards an observer. Refer to the dotted arrows in Fig. 1. The light reflected into the line of sight is integrated along the entire path length, d , from the scene to the observer. Then, the irradiance due to airlight is given by (see [16]),

$$E_a(d, \lambda) = E_{\infty}(\lambda) \left(1 - e^{-\beta(\lambda)d}\right). \quad (2)$$

The total irradiance E received by the sensor is the sum of irradiances due to attenuation and airlight, respectively,

$$E(d, \lambda) = E_{dt}(d, \lambda) + E_a(d, \lambda). \quad (3)$$

2.2 Wavelength Dependence of Scattering

Generally, different wavelengths of light are scattered differently by atmospheric particles. Interesting atmospheric phenomena such as the blueness of the sky and the bluish haze of distant mountains are examples of the wavelength selective behavior of atmospheric scattering [12], [18]. In these cases, the blue wavelengths are scattered more compared to other visible wavelengths. On the other hand, fog and dense haze scatter all visible wavelengths more or less the same way.

Over the visible spectrum, Rayleigh's law of atmospheric scattering provides the relationship between the scattering coefficient β and the wavelength λ [16]:

2. Solid angle subtended by the area of sky visible to a scene point.

1. Multiple scattering effects are not taken into account in this model and hence, for highly dense weather conditions, the model will not be effective. Also, this model does not take into account blurring effects of bad weather due to turbulence.




Fig. 1. Scattering of light by atmospheric particles can be described by two models—direct transmission (or attenuation) and airlight. Direct transmission is the attenuated irradiance received by the sensor from the scene point along the line of sight. Airlight is the total amount of environmental illumination (sunlight, skylight, ground light) reflected into the line of sight by atmospheric particles.

$$\beta(\lambda) \propto \frac{1}{\lambda^\gamma}, \quad (4)$$

where $0 \leq \gamma \leq 4$ depending on the exact particle size distribution in the atmosphere. For pure air, the constituent particle (molecules) sizes are very small compared to the wavelength of light and hence, there is a strong wavelength dependence of scattering. In this case, $\gamma = 4$; short (blue) wavelengths dominate and we see the clear blue sky. For fog, the constituent particle (water droplets) sizes are large compared to the wavelength of light and, hence, the scattering coefficient does not depend on wavelength. So, for fog, $\gamma \approx 0$; all wavelengths are scattered equally and we see grayish (or white) fog. A wide gamut of atmospheric conditions arise from aerosols whose particle sizes range between minute air molecules ($10^{-4} \mu m$) and large fog droplets ($1 - 10 \mu m$). Such aerosols (e.g., mild haze) show a significant wavelength selectivity ($0 < \gamma < 4$).

2.3 Weather Conditions and Camera Response

Different cameras measure irradiance over different color bands. Some examples include gray-scale cameras (entire visible spectrum), conventional color cameras (three broad bands R, G, and B), and multispectral cameras (multiple narrow color bands). In the appendix, we derive an expression for the brightness recorded by a monochrome (narrow spectral band) camera, using (3). In this derivation, we assume that the scattering coefficient β remains constant within the spectral bandwidth of the monochrome camera.

Keeping the above assumption in mind, we now discuss under what weather conditions can our methods be applied to various sensors. Recall from Section 2.2 that the scattering coefficient for fog and dense haze remains more or less constant over the visible spectrum. Accordingly, a broadband RGB or gray-scale camera suffices to analyze images taken in fog and dense haze. For other aerosols such as mild haze, multispectral cameras or cameras fitted with narrow-band filters should be used in order to apply our methods satisfactorily. Finally, scattering coefficients of most weather conditions vary significantly in the near-IR spectrum [37] and hence, narrow-band IR cameras have to be used for the analysis beyond the visible wavelengths. In other words, the greater the variation in the scattering coefficient with respect

to wavelength, the narrower the spectral bandwidth needed for effective results.

We would like to clarify that multiple color channels are not required for our algorithms. We can, however, apply the methods we describe in this paper to each color channel of the sensor independently. This is in contrast to previous methods that required at least three color channels (say, R, G, and B) over which the scattering coefficient had to be equal [21].

3 CONTRAST DEGRADATION IN BAD WEATHER

In this section, we show how contrast degrades in poor visibility conditions as a function of both the scattering coefficient of the atmosphere and the distance of the scene from the sensor. Consider an image taken in bad weather. The brightness at any pixel recorded by a monochrome camera is derived in the appendix:

$$E = I_\infty \rho e^{-\beta d} + I_\infty (1 - e^{-\beta d}), \quad (5)$$

where I_∞ is termed as sky intensity. We call ρ the *normalized radiance* of a scene point; it is a function of the scene point reflectance (BRDF), normalized sky illumination spectrum, and the spectral response of the camera, *but not* the weather condition defined by (β, I_∞) (see the appendix).

Using (5), we formulate the image contrast between two adjacent scene points as a function of the amount of scattering and their distance from the observer. Consider two adjacent scene points P_i and P_j at the same depth d from a sensor. Their pixel intensities are given by,

$$\begin{aligned} E^{(i)} &= I_\infty \rho^{(i)} e^{-\beta d} + I_\infty (1 - e^{-\beta d}), \\ E^{(j)} &= I_\infty \rho^{(j)} e^{-\beta d} + I_\infty (1 - e^{-\beta d}). \end{aligned} \quad (6)$$

The observed contrast between P_i and P_j can be defined as,

$$\frac{E^{(i)} - E^{(j)}}{E^{(i)} + E^{(j)}} = \frac{\rho^{(i)} - \rho^{(j)}}{\rho^{(i)} + \rho^{(j)} + 2(e^{\beta d} - 1)}. \quad (7)$$

This shows that the contrast degrades *exponentially* with the scattering coefficient β and the depths of scene points in bad weather. As a result, conventional space-invariant image processing techniques cannot be used to completely remove

weather effects. Note that other formulations for image contrast (e.g., MTF, log intensity) [13] can also be used to illustrate the exponential contrast decay.

4 CONTRAST RESTORATION OF ISO-DEPTH REGIONS

We now describe a simple method to restore scene contrast from one bad weather image using depth segmentation of the scene. We define depth segmentation as the extraction of iso-depth regions in the scene. Note, this does *not* mean that actual scene depths have to be known. In several situations, it may be easy to interactively provide the necessary segmentation. For instance, in urban scenes with frontal views of buildings, a user can easily mark out regions that roughly have the same depths. Later, we will present two automatic depth segmentation techniques using images taken under different weather conditions.

Consider an image taken in bad weather. The brightness at any pixel recorded by a monochrome camera is given by,

$$E = I_{\infty} \rho e^{-\beta d} + I_{\infty} (1 - e^{-\beta d}). \quad (8)$$

Now, consider two scene points P_i and P_j at the same depth d from a sensor. The observed contrast between P_i and P_j is given by (7). Eliminating the unknown $e^{-\beta d}$ from (6), we obtain,

$$\frac{1 - \rho^{(i)}}{1 - \rho^{(j)}} = \frac{I_{\infty} - E^{(i)}}{I_{\infty} - E^{(j)}}. \quad (9)$$

For robustness, we consider all the pixels at the same depth,

$$\frac{1 - \rho^{(i)}}{\sum_j (1 - \rho^{(j)})} = \frac{I_{\infty} - E^{(i)}}{\sum_j (I_{\infty} - E^{(j)})}. \quad (10)$$

Then, the normalized radiance of any scene point is obtained using,

$$\rho^{(i)} = 1 - \left(\sum_j 1 - \sum_j \rho^{(j)} \right) \frac{I_{\infty} - E^{(i)}}{\sum_j (I_{\infty} - E^{(j)})}. \quad (11)$$

This procedure is repeated independently for each depth in the scene. So, if we have a priori depth segmentation of the scene and have measured the sky intensity I_{∞} , then $\rho^{(i)}$ can be computed up to a linear factor $\sum_j \rho^{(j)}$. Since ρ is independent of the weather condition, we have restored the contrast of the scene using just *one* bad weather image.

What can we do if we do not have the sky intensity I_{∞} ? Let us assume that P_j has the minimum brightness within the iso-depth scene points: $E^{(j)} = E^{\min}$. Also, since scene point brightnesses are much lower than sky brightness on an overcast day, i.e., $(\forall i), I_{\infty} > E^{(i)}$, we can set $I_{\infty} = E^{\max}$. Then, the right-hand side of (9) is just contrast stretching the inverted bad weather image. In other words, by setting $\rho^{\min} = 0$ and $\rho^{\max} = 1$, contrast stretching (or histogram stretching) each iso-depth region in the image can restore contrast. Even though contrast is restored at each depth satisfactorily, the image can look unrealistic. In summary, simple image processing techniques such as contrast stretching can be effective for scenes that are at the same depth from the sensor (e.g., a planar scene at a fixed distance from the camera). Clearly, for scenes with

significant depth variations, this simple method will not be effective.

5 DEPTH EDGES FROM TWO WEATHER CONDITIONS

In this section, we present a simple cue to automatically locate the depth edges (discontinuities) present in a scene from two monochrome images taken under different but unknown weather conditions. In other words, we present a method to label image edges as reflectance edges and depth edges. Several researchers have pursued the problem of classifying different types of edges (diffuse, specular, and occluding) based on image intensity/color cues [31], [35] and polarization cues [1]. As we shall show, changes in weather conditions can be exploited as a cue to differentiate depth edges from reflectance edges.

Note that closed contours of depth edges can be used for depth segmentation. In outdoor surveillance applications, video cameras capture the same scene (albeit with moving objects) over long periods of time during which the weather may change. Also, depth edges in the static portion of any scene have to be computed just once and not for every video frame. Hence, we see this as an initialization step that needs to be done before applying the contrast restoration algorithm of Section 4 to all frames.

Consider a small image neighborhood corresponding to scene points that are at the same depth from an observer (i.e., no depth edges present). We call such a neighborhood an *iso-depth neighborhood*. From (5), the average brightness of an iso-depth neighborhood is,

$$\bar{E} = [I_{\infty} e^{-\beta d}] \bar{\rho} + [I_{\infty} (1 - e^{-\beta d})], \quad (12)$$

and the standard deviation of the neighborhood is,

$$\sigma_E = \sqrt{\frac{1}{n} \sum_{i=1}^n (E^{(i)} - \bar{E})^2}. \quad (13)$$

Using (5), we simplify to obtain,

$$\sigma_E = I_{\infty} e^{-\beta d} \sqrt{\frac{1}{n} \sum_{i=1}^n (\rho^{(i)} - \bar{\rho})^2}. \quad (14)$$

Normalizing the pixel values in the neighborhood, we get,

$$\frac{E^{(i)} - \bar{E}}{\sigma_E} = \frac{(\rho^{(i)} - \bar{\rho})}{\sqrt{\frac{1}{n} \sum_{i=1}^n (\rho^{(i)} - \bar{\rho})^2}}. \quad (15)$$

For iso-depth neighborhoods, clearly the above equation is *invariant* to the weather condition (β, I_{∞}) . More importantly, the invariance does *not* hold for a neighborhood that contains depth edges. This is easily explained as follows: The airlight does not remain constant across a neighborhood with depth discontinuities. Hence, subtracting the mean (as in (15)) will not remove the airlight completely.

Now, let us consider two images captured under different weather conditions. We assume that the two images are taken under similar daylight distributions. However, the magnitudes of the distributions (I_{∞}) may vary. In other words, the shadow edges (if any) appear at the same pixel

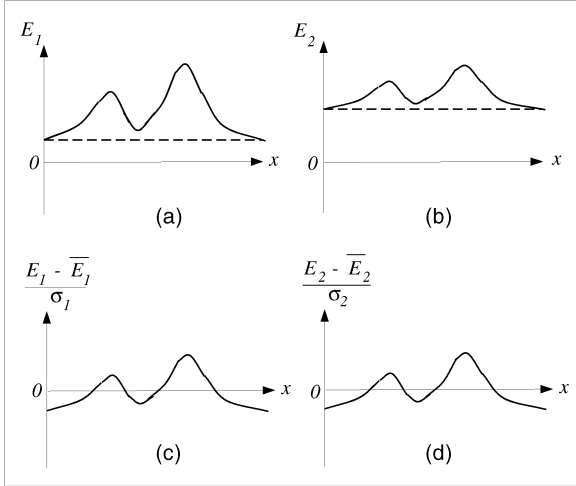


Fig. 2. Invariance of iso-depth neighborhoods to weather conditions. (a) and (b) Signals representing the intensities of a neighborhood of iso-depth scene points in two weather conditions. Airlight (dashed lines) is constant for the entire neighborhood. (c) and (d) Normalized signals in the two weather conditions match exactly.

location in both the images. Figs. 2a and 2b illustrate the brightnesses within an iso-depth neighborhood under two weather conditions. Figs. 2c and 2d show that the normalized signals under the two weather conditions match perfectly. On the other hand, Fig. 3 illustrates that normalized signals of scene neighborhoods that contain depth edges, do not match. Normalized SSD can be used to determine the quality of the match. Note that (15) still holds if we apply a more robust estimate of mean and standard deviation (for e.g., median of absolute deviations from the neighborhood median).

It is interesting to note what happens if we treat the entire image as a single neighborhood. Applying normalized SSD to two images of a scene, a poor match implies that the weather condition *changed between* the two images and a good match

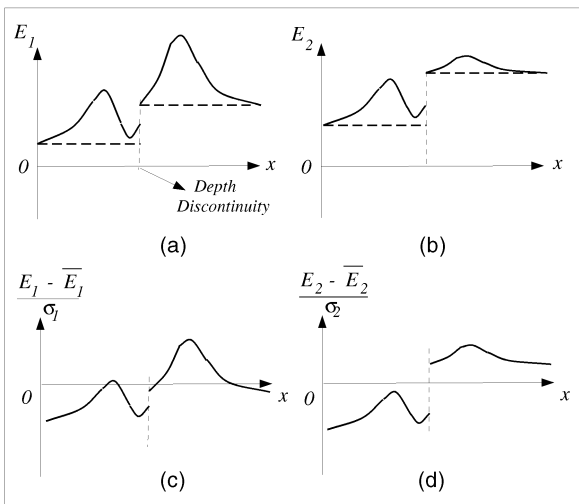


Fig. 3. Illustration of scene intensities of a neighborhood that has a depth edge. (a) and (b) Signals representing the intensities of the neighborhood under two weather conditions. Airlight (dashed lines) varies across the neighborhood. (c) and (d) Normalized signals in the two weather conditions do not match.




Fig. 4. Classification of image edges into reflectance edges and depth edges. (a) and (b) Images of the same scene captured under different fog conditions (half an hour apart). (c) The image in (a) is histogram equalized to aid visualization of depth edges (shown using arrows). (d) White pixels denote depth edges and black pixels denote reflectance edges. Note that the edge detector was applied to the original image in (a) and not the histogram equalized image in (c).

implies otherwise. For this, the scene should have at least two different depths and the images should be linearized using the radiometric response function of the camera. This cue is helpful in deciding which frames can be used to compute depth edges in a video sequence.

Fig. 4 shows the experimental results of classifying image edges into *reflectance edges* and *depth edges* for a real scene captured under two different foggy conditions. The time between the capture of the images was about half an hour. The edge map of one of the images was computed using the Canny edge detector. For each edge pixel, we considered 15×15 neighborhoods around the pixel in the two images. We applied normalized SSD to match these neighborhoods. For the depth edges, the normalized SSD value was high; for the reflectance edges, the value was low. The depth edges are shown in white and reflectance edges are shown in black (Fig. 4d). Note if both reflectance edges and depth edges are within the same neighborhood, this method may misclassify the reflectance edges as depth edges. Also, note that shadow edges (if any) will not be distinguished from

reflectance edges. Finally, this method to classify edges can be sensitive to noise, especially under poor weather conditions. Under poor weather conditions, due to the limited dynamic range of the sensor (typically 8 bits), the direct transmission (signal) to airlight (noise) ratio can be so low that the direct transmission magnitude can be compared to the sensor noise level. In this case, the results produced by the method may not be trusted.

6 SCENE STRUCTURE FROM TWO WEATHER CONDITIONS

In the previous section, we described a method to locate depth discontinuities from two bad weather images. Note, however, that normalized SSD is effective only in textured neighborhoods (reflectance edges and depth discontinuities). In other words, normalized SSD is not reliable for “flat” intensity regions and regions where depth changes are *gradual*. Moreover, due to the blurring seen in images taken under *poor* visibility conditions, the edge maps may not be reliable enough to create closed contours of depth discontinuities (needed for depth segmentation).

In this section, we present a method to compute complete structure of an arbitrary scene, from two images taken under different weather conditions. In contrast to the methods proposed in [21], [24] that *require* color images (three color channels), our algorithm can be applied to both gray scale as well as color images.

Consider the observed pixel values E_1 and E_2 of a scene point under two weather conditions (β_1, I_{∞_1}) and (β_2, I_{∞_2}) . Let us examine how the brightness of this scene point changes from the first weather condition to the second. From (5),

$$\begin{aligned} E_1 &= I_{\infty_1} \rho e^{-\beta_1 d} + I_{\infty_1} (1 - e^{-\beta_1 d}) \\ E_2 &= I_{\infty_2} \rho e^{-\beta_2 d} + I_{\infty_2} (1 - e^{-\beta_2 d}). \end{aligned} \quad (16)$$

Eliminating ρ from (16) we get,

$$E_2 = \left[\frac{I_{\infty_2}}{I_{\infty_1}} e^{-(\beta_2 - \beta_1)d} \right] E_1 + \left[I_{\infty_2} (1 - e^{-(\beta_2 - \beta_1)d}) \right], \quad (17)$$

which is *linear* in E_1 and E_2 . Also, for the two weather conditions, the coefficients of the linear equation depend *only* on scene depth. In other words, for iso-depth scene points, the plot of E_1 versus E_2 is a straight line.

Another significant constraint results from our physical model that suggests a means of estimating sky intensities. Interestingly, if we substitute $E_1 = I_{\infty_1}$ in (17), we get $E_2 = I_{\infty_2}$, irrespective of the depth d . Therefore, the point $(I_{\infty_2}, I_{\infty_1})$ lies on all the straight lines corresponding to different depths in the scene (see Fig. 5). In other words, the intersection of straight lines corresponding to different depths yields the sky intensities I_{∞_1} and I_{∞_2} .

The iso-depth lines in the plot of E_1 versus E_2 can be detected using the Hough transform. Then, the intersection (sky intensities) of the iso-depth lines can be computed using a least squares line-fitting algorithm. A problem arises if the iso-depth lines are not detected correctly in the plot of E_1 versus E_2 . In other words, scene depths can change smoothly and the iso-depth lines could “bunch up.” In order to




Fig. 5. Plot of the pixel values E_1 observed under one weather condition versus the corresponding pixel values E_2 observed under another weather condition. Each line represents all the scene points at the same depth from the sensor. All iso-depth lines intersect at the horizon brightnesses $(I_{\infty_1}, I_{\infty_2})$ of the two weather conditions.

compute sky intensities, we just divide the two images into blocks and within each block we fit lines to the (E_2, E_1) pairs of scene points. If the fit is good, we decide that the scene points in the block are at the same depth. Finally, we use at least two such iso-depth blocks to estimate sky the intensities.

Substituting the values of I_{∞_1} and I_{∞_2} in (17), we obtain the scaled depth of *each* scene point:

$$(\beta_2 - \beta_1)d = -\ln \frac{I_{\infty_2} - E_2}{I_{\infty_1} - E_1} - \ln \frac{I_{\infty_1}}{I_{\infty_2}}. \quad (18)$$

Thus, we have computed the depth map of a scene from two images taken under different weather conditions.

7 CONTRAST RESTORATION USING SCENE STRUCTURE

In Section 4, we described a method to restore scene contrast given a depth segmentation of the scene. This method is simple and effective for scenes where depth changes are abrupt (for example, an urban scene with frontal views of buildings). However, it is hard to define good depth segmentation when scene depths change *gradually* (for instance, a natural scene with mountains or an urban scene with an oblique view of a road). In this section, we present a method to restore contrast of an *arbitrary* scene using *scaled depths* (18) of scene points.

We assume that there exists a patch E_{zero} in the scene whose direct transmission is zero. This can happen in two instances. First, E_{zero} can be a black patch with its scene radiance equal to zero. Note that the black scene patch will not appear black in the image due to the addition of airlight. Second, E_{zero} could be a distant scene patch that is completely invisible due to strong airlight. In other words, this distant scene patch has zero direct transmission and its contrast cannot be restored from a bad weather image.

We can either mark such a patch manually or detect one automatically from the image. To detect zero direct transmission patches automatically in weather degraded images, we use the method described in [21]. Since the




Fig. 6. Generating a synthetic scene—a stack of discs textured with random gray dots. (a) On the left is the 3D structure and on the right is an image of the top view of the scene. The gray levels on the structure are used only to illustrate the disks better. (b) Two different amounts of fog and noise ($\sigma = 3.0$ gray levels) are added to the image in (a).

apparent brightness of the patch E_{zero} is solely due to airlight, its optical depth can be computed as,

$$\beta d_{zero} = -\ln(1 - E_{zero}/I_{\infty}). \quad (19)$$

Then, the optical depth of *any other* scene point P_i is obtained using,

$$\beta d_i = (\beta d_{zero}) \left(\frac{d_i}{d_{zero}} \right), \quad (20)$$

where the second term can be computed using the ratio of scaled depths (see (18)). Then, the normalized radiance ρ_i of the scene point P_i is estimated using (5). Recall that ρ does not depend on the weather condition (β, I_{∞}). Thus, by computing ρ for each scene point, we restore contrast of the entire scene.

Note that structure computation requires two images to be taken under different weather conditions, but under similar daylight spectra. However, once scene structure is computed, contrast can be restored from a single image of the scene taken under arbitrary weather and illumination conditions.

8 EXPERIMENTAL RESULTS

We performed experiments with both synthetic and real scenes. Fig. 6a shows a synthetic scene consisting of a stack of cylinders with random brightness values. To this image, two different amounts of fog are added according to the model described in (5). To this image, Gaussian random noise of $\sigma = 3.0$ gray levels was added. Figs. 7a and 7b shows the results of applying the structure computation and contrast restoration algorithms to images of the synthetic scene. The error in the recovered depth map was less than 1 percent.

Results of experiments performed on images of real scenes are shown in Figs. 8 and 9. The images were captured

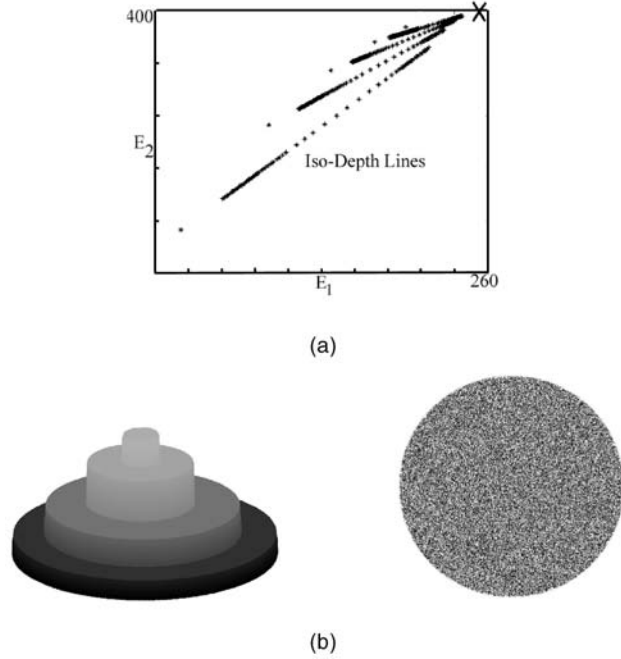


Fig. 7. Experiments on a synthetic scene—a stack of discs textured with random gray dots. Two images of a synthetic scene with different amounts of fog are shown in Fig. 6b. (a) Iso-depth lines shown in the plot of pixel values under the first weather condition versus the corresponding pixel values under the second weather condition. X mark shows the intersection $(I_{\infty_2}, I_{\infty_1})$ of all the iso-depth lines. (b) The recovered structure and contrast restored image. Compare (b) with the original synthetic scene in Fig. 6a.

using a Professional KODAK DCS 315 digital camera. Multiple exposures of the same scene were acquired and the radiometric response function of the camera was computed using the method proposed in [19]. Then, the multiple exposed images were linearized and combined using simple weighted averaging to obtain a high dynamic range image of the scene. Fig. 8a shows two high dynamic range images of the same scene captured under different conditions of mist (light and moderate). The depth map computed using the algorithm mentioned in Section 6 is shown in Fig. 8b. The mist was removed using the contrast restoration algorithm mentioned in Section 7. Notice the windows of the farther buildings that are clearly visible in Fig. 8d as compared to the images in Fig. 8a.

We performed experiments under rainy conditions also. Since we are interested in a far away scene, the captured image does seem like a foggy or a misty one due to spatio-temporal averaging in the sensor. The results proved that we can apply our algorithm to rainy images of faraway scenes as well. In this case, we just captured one image under mild rain conditions shown in Fig. 9a. The depth map, precomputed from the misty images shown in Fig. 8a, was used to restore the contrast of the rainy day image. Thus, changes in weather conditions are required only to compute scene structure whereas contrast restoration can be applied to a single image of that scene taken under arbitrary weather conditions. Compare the results of our algorithm (Fig. 9d) with conventional histogram equalization (Fig. 9b).

In general, removing the spatio-temporal effects of rain is a much harder problem compared to more stable weather




Fig. 8. Structure computation and restoration of image contrast from two images taken under poor visibility conditions. The depth map is median filtered and averaged to reduce noise. Notice the significant increase in contrast in the farther buildings. Contrast stretching is applied to all the images for display purposes. (a) Images taken at 3 P.M. and 4 P.M. under poor visibility conditions (mist). (b) Computed depth map. (c) Contrast restore image.

conditions such as fog, mist, and haze. The brightnesses due to raindrops in the scene cannot be modeled using the simple direct transmission and airlight models used in this paper.

8.1 Experiments with Video: Moving Objects

Consider an outdoor surveillance video camera capturing a scene (with moving objects) over an extended period of time. We would like to process this video in real-time to obtain a weather-free video. Note that our algorithms cannot remove temporal effects of rain from a video of a rainy scene. For the purposes of discussion, we define the static part of the scene as the *background* and the moving objects in the scene as the *foreground*. The foreground objects can be separated from the background using any background subtraction method (for instance, [32]). Then, weather-free video is obtained using an algorithm that has the following two stages:

- **Initialization stage.** We first detect any change in weather condition using normalized SSD (Section 5). Then, the two frames that correspond to the different weather conditions are used to compute scaled depths of the background scene (Section 6).
- **Contrast Restoration.** Note that the methods we described hitherto cannot be used to restore contrast of moving objects since their depths are unknown. Therefore, heuristics are needed to assign depths to foreground objects. One conservative heuristic is to examine the depths in a neighborhood around each moving object and assign the minimum depth to it.

The algorithm presented in Section 7 can then be applied to the entire frame to restore scene contrast.

Experimental results with a video of a traffic scene taken under foggy conditions are shown in Fig. 10. We used an off-the-shelf 8-bit digital video camera and captured two short video clips half an hour apart. As described in previous experiments, we linearized the frames with the radiometric response function of the video camera. We averaged 100 frames in each video clip to reduce noise and used the resulting images to compute structure of the background scene (buildings). The scaled depths in the road region were linearly interpolated using scaled depth values at pixels on the left and right corners of the road. Then, contrasts of buildings, the road, and moving vehicles were restored for each frame of the video. Notice the significant increase in contrast at various depths in the scene (Figs. 10d and 11). Compare our method to histogram equalization in Fig. 12. In our current implementation, contrast restoration was applied to the video offline.

9 SUMMARY

In this paper, we addressed the problem of restoring the contrast of atmospherically degraded images and video. We presented methods to locate depth discontinuities and to compute structure of a scene, from two images captured under different weather conditions. Using either depth segmentation (regions within closed contours of depth edges)




Fig. 9. Contrast restoration from one bad weather (in this case, rain) image and precomputed scene structure. Scene structure can be computed using two images taken under possibly different weather conditions (say, mist). The depth map computed from two misty images (Fig. 8a) was used to restore contrast from just one image of the same scene under rain. The rainy image shown in (a) and the misty images shown in Fig. 8a were captured on different days. (b) Applying histogram equalization to the entire rainy image does not enhance contrast in all depths. (c) Contrast restoration using the algorithm proposed in Section 7.

or scene structure (scaled depths), we then showed how to restore contrast from any image of the scene taken in bad weather. Note, although structure computation requires changes in weather, the contrast restoration algorithms do not. The entire analysis is presented for monochrome images. However, our methods can be applied to images captured using multispectral cameras, IR cameras, and the usual broadband RGB and gray-scale cameras.

APPENDIX

MONOCHROME CAMERA SENSING IN BAD WEATHER

In this section, we derive an expression for the intensity E , of a scene point under bad weather, recorded by a camera within a narrow wavelength band $(\lambda, \lambda + \delta\lambda)$. From (3) we write,

$$E = \int_{\lambda}^{\lambda+\delta\lambda} s(\lambda)(E_{dt}(d, \lambda) + E_a(d, \lambda)) d\lambda, \quad (21)$$

where $s(\lambda)$ is the spectral response of the camera. We assume that the scattering coefficient β does not change appreciably over the narrow spectral band and write,

$$E = \frac{e^{-\beta d}}{d^2} \int_{\lambda}^{\lambda+\delta\lambda} E_{\infty}(\lambda)s(\lambda)r(\lambda)d\lambda \dots + (1 - e^{-\beta d}) \int_{\lambda}^{\lambda+\delta\lambda} E_{\infty}(\lambda)s(\lambda)d\lambda. \quad (22)$$

The sky illumination spectrum can be written as,

$$E_{\infty}(\lambda) = I'_{\infty} \hat{E}_{\infty}(\lambda), \quad (23)$$

where I'_{∞} is the magnitude of the sky illumination spectrum and $\hat{E}_{\infty}(\lambda)$ is the normalized sky illumination spectrum. Letting

$$\begin{aligned} g &= \int_{\lambda}^{\lambda+\delta\lambda} \hat{E}_{\infty}(\lambda)s(\lambda)d\lambda, \\ \rho &= \frac{1}{gd^2} \int_{\lambda}^{\lambda+\delta\lambda} \hat{E}_{\infty}(\lambda)s(\lambda)r(\lambda)d\lambda, \\ I_{\infty} &= I'_{\infty}g, \end{aligned} \quad (24)$$

we rewrite the final brightness at any pixel as,

$$E = I_{\infty} \rho e^{-\beta d} + I_{\infty} (1 - e^{-\beta d}), \quad (25)$$

where I_{∞} is termed as sky intensity. Note that ρ is a function of normalized sky illumination spectrum, scene point reflectance, and the spectral response of the camera, *but not* the weather condition β . The algorithm we present in the paper recovers ρ for each pixel to restore scene contrast.

Let us now examine the wavelength range in which this model can be applied. By changing the limits of integration to $[\lambda_1, \lambda_2]$, and assuming the scattering coefficient to be constant over this wavelength band, we can use the same model for a




Fig. 10. (a) Scene imaged at 5:00 P.M. (b) Scene imaged at 5:30 P.M. (c) Depth map computed from images (a) and (b). (d) Contrast restored using image (b). Experiments with videos of a traffic scene on a foggy day. (a) and (b) Two short video clips were captured half an hour apart using an 8-bit video camera. 100 frames were averaged to reduce noise. Note that the vehicles on the road in the two images are different. (c) The depth map was computed for the background image using the algorithm presented in Section 5. The scaled depths of the region corresponding to the road were linearly interpolated using scaled depth values at pixels on the left and right corners of the road. (d) The defogged (contrast restored) image obtained from the video frame in (b). Compare the contrast restored image with the histogram equalized image in Fig. 12.




Fig. 11. Zoomed in regions of the frame (see the marked rectangles in Fig. 10b) demonstrate the significant increase in contrast at various depths of the scene. Note that different amounts of fog were removed at different depths. Also, notice the better contrast of moving objects (vehicles).

black and white camera (entire visible range), or smaller color bands (R, G, B) for a color camera, or narrow band multi-spectral cameras. Thus, for removal of fog and dense haze, we can use RGB color or gray-scale cameras whereas we must use narrow spectral band cameras for the removal of many aerosols.

ACKNOWLEDGMENTS

This work was supported in part by a DARPA/ONR HumanID Contract (N00014-00-1-0916) and an US National Science Foundation Award (IIS-99-87979). The authors thank Yoav Schechner for valuable discussions on this topic. Some of the results presented in this paper have appeared in the



Fig. 12. Artificial contrast enhancement (not contrast restoration) of a foggy image by histogram equalization does not remove fog completely at all depths in the scene. In this example, farther depths have poorer contrast when compared to the nearer depths.

Proceedings of the 2001 IEEE Conference on Computer Vision and Pattern Recognition [22].

REFERENCES

- [1] T.E. Boult and L.B. Wolff, "Physically-Based Edge Labelling," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1991.
- [2] B. Cairns, B.E. Carlson, A.A. Lacis, and E.E. Russell, "An Analysis Of Ground-Based Polarimetric Sky Radiance Measurements," *Proc. SPIE*, vol. 3121, 1997.
- [3] S. Chandrasekhar, *Radiative Transfer*. Dover Publications, Inc., 1960.
- [4] D.B. Chenault and J.L. Pezzaniti, "Polarization Imaging through Scattering Media," *Proc. SPIE*, vol. 4133, 2000.
- [5] K.L. Coulson, "Polarization of Light in the Natural Environment," *Proc. SPIE*, vol. 1166, 1989.
- [6] L.J. Denes, M. Gottlieb, B. Kaminsky, and P. Metes, "Aotf Polarization Difference Imaging," *Proc. SPIE*, vol. 3584, 1998.
- [7] S.D. Gedzelman, "Atmospheric Optics in Art," *Applied Optics*, vol. 30, 1991.
- [8] J. Gordon and P. Church, "Overcast Sky Luminances and Directional Luminous Reflectances of Objects and Backgrounds under Overcast Skies," *Applied Optics*, vol. 5, p. 919, 1966.
- [9] L.L. Grewe and R.R. Brooks, "Atmospheric Attenuation Reduction through Multisensor Fusion," *Sensor Fusion: Architectures, Algorithms, and Applications II*, *Proc. SPIE*, vol. 3376, Apr. 1998.
- [10] Van De Hulst, *Light Scattering by Small Particles*. John Wiley and Sons, 1957.
- [11] R.L. Lee Jr., "Horizon Brightness Revisited: Measurements and a Model of Clear-Sky Radiances," *Applied Optics*, vol. 20, pp. 4620-4628, 1994.
- [12] N.S. Kopeika, "General Wavelength Dependence of Imaging through the Atmosphere," *Applied Optics*, vol. 20, no. 9, May 1981.
- [13] N.S. Kopeika, *A System Engineering Approach to Imaging*. SPIE Press, 1998.
- [14] D.K. Lynch, "Step Brightness Changes of Distant Mountain Ridges and Their Perception," *Applied Optics*, vol. 30, 1991.
- [15] S. Mahadev and R.C. Henry, "Color Perception through Atmospheric Haze," *J. Optical Soc. Am. A*, vol. 17, no. 5, May 2000.
- [16] E.J. McCartney, *Optics of the Atmosphere: Scattering by Molecules and Particles*. John Wiley and Sons, 1975.
- [17] W.E.K. Middleton, *Vision through the Atmosphere*. Univ. of Toronto Press, 1952.
- [18] M. Minnaert, *The Nature of Light and Color in the Open Air*. Dover Publications, Inc., 1954.
- [19] T. Mitsunaga and S.K. Nayar, "Radiometric Self Calibration," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1999.
- [20] P. Moon and D.E. Spencer, "Illumination from a Non-Uniform Sky," *Illuminating Eng.*, vol. 37, pp. 707-726, 1942.
- [21] S.G. Narasimhan and S.K. Nayar, "Chromatic Framework for Vision in Bad Weather," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2000.
- [22] S.G. Narasimhan and S.K. Nayar, "Removing Weather Effects from Monochrome Images," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2001.
- [23] S.G. Narasimhan and S.K. Nayar, "Vision and the Atmosphere," *Int'l J. Computer Vision*, vol. 48, no. 3, pp. 233-254, Aug. 2002.
- [24] S.K. Nayar and S.G. Narasimhan, "Vision in Bad Weather," *Proc. Seventh Int'l Conf. Computer Vision*, 1999.
- [25] J.P. Oakley and B.L. Satherley, "Improving Image Quality in Poor Visibility Conditions Using a Physical Model for Degradation," *IEEE Trans. Image Processing*, vol. 7, Feb. 1998.
- [26] "Museum of Science," Leonardo's Perspective. <http://www.mos.org/sln/Leonardo/InvestigatingAerialP.html>, 1997.
- [27] M.J. Rakovic, G.W. Kattawar, M. Mehrubeoglu, B.D. Cameron, L.V. Wang, S. Rastegar, and G.L. Cote, "Light Backscattering Polarization Patterns from Turbid Media: Theory and Experiment," *Applied Optics*, vol. 38, 1999.
- [28] M.P. Rowe, E.N. Pugh Jr., J.S. Tyo, and N. Engheta, "Polarization-Difference Imaging: A Biologically Inspired Technique for Observation through Scattering Media," *Optical Letters*, vol. 20, 1995.
- [29] Y.Y. Schechner, S.G. Narasimhan, and S.K. Nayar, "Instant Dehazing of Images Using Polarization," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2001.
- [30] Y.Y. Schechner, S.G. Narasimhan, and S.K. Nayar, "Polarization Based Vision through Haze," *Applied Optics*, special issue: light and color in the open air, vol. 42, no. 3, Jan. 2003.
- [31] S. Shafer, "Using Color to Separate Reflection Components," *Color Research and Applications*, pp. 210-218, 1985.
- [32] C. Stauffer and W.E.L. Grimson, "Adaptive Background Mixture Models for Real-Time Tracking," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1999.
- [33] K. Tan and J.P. Oakley, "Enhancement of Color Images in Poor Visibility Conditions," *Proc. Int'l Conf. Image Processing*, vol. 2, Sept. 2000.
- [34] K. Tan and J.P. Oakley, "Physics Based Approach to Color Image Enhancement in Poor Visibility Conditions," *J. Optical Soc. Am. A*, vol. 18, no. 10, pp. 2460-2467, Oct. 2001.
- [35] S. Ullman, "On the Visual Detection of Light Sources," *Biological Cybernetics*, pp. 205-212, 1976.
- [36] J.G. Walker, P.C.Y. Chang, and K.I. Hopcraft, "Visibility Depth Improvement in Active Polarization Imaging in Scattering Media," *Applied Optics*, vol. 39, 1995.
- [37] W.L. Wolfe and G.J. Zissis, *The Infrared Handbook*. Prepared for Office of Naval Research, Dept. of the Navy, 1978.
- [38] Y. Yitzhaky, I. Dror, and N.S. Kopeika, "Restoration of Atmospherically Blurred Images According to Weather-Predicted Atmospheric Modulation Transfer Function," *Optical Eng.*, vol. 36, Nov. 1998.



Srinivasa G. Narasimhan received the MS degree in computer science from Columbia University in 1999. He is currently pursuing the PhD degree in computer science at the Columbia University Vision Laboratory. His research is focused on the development of physics-based models and algorithms for scene understanding in bad weather. His paper received the Best Paper Honorable Mention Award at the 2000 IEEE Computer Vision and

Pattern Recognition Conference.



Shree K. Nayar received the PhD degree in electrical and computer engineering from the Robotics Institute at Carnegie Mellon University in 1990. He is the T.C. Chang Professor of Computer Science at Columbia University. He currently heads the Columbia Automated Vision Environment (CAVE), which is dedicated to the development of advanced computer vision systems. His research is focused on three areas, namely, novel vision sensors, physics-based

models for vision, and algorithms for scene understanding. His work is motivated by applications in the fields of computer vision, computer graphics, human-machine interfaces, and robotics. Dr. Nayar has authored and coauthored papers that have received the Best Paper Honorable Mention Award at the 2000 IEEE Computer Vision and Pattern Recognition Conference (CVPR), the David Marr Prize at the 1995 International Conference on Computer Vision (ICCV) held in Boston, Siemens Outstanding Paper Award at the 1994 IEEE CVPR Conference held in Seattle, the 1994 Annual Pattern Recognition Award from the Pattern Recognition Society, Best Industry Related Paper Award at the 1994 International Conference on Pattern Recognition held in Jerusalem, and the David Marr Prize at the 1990 ICCV held in Osaka. He holds several US and international patents for inventions related to computer vision and robotics. He was the recipient of the David and Lucile Packard Fellowship for Science and Engineering in 1992, the National Young Investigator Award from the US National Science Foundation in 1993, and the Excellence in Engineering Teaching Award from the Keck Foundation in 1995.

► **For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.**

Polarization-based vision through haze

Yoav Y. Schechner, Srinivasa G. Narasimhan, and Shree K. Nayar

We present an approach for easily removing the effects of haze from passively acquired images. Our approach is based on the fact that usually the natural illuminating light scattered by atmospheric particles (airlight) is partially polarized. Optical filtering alone cannot remove the haze effects, except in restricted situations. Our method, however, stems from physics-based analysis that works under a wide range of atmospheric and viewing conditions, even if the polarization is low. The approach does not rely on specific scattering models such as Rayleigh scattering and does not rely on the knowledge of illumination directions. It can be used with as few as two images taken through a polarizer at different orientations. As a byproduct, the method yields a range map of the scene, which enables scene rendering as if imaged from different viewpoints. It also yields information about the atmospheric particles. We present experimental results of complete dehazing of outdoor scenes, in far-from-ideal conditions for polarization filtering. We obtain a great improvement of scene contrast and correction of color. © 2003 Optical Society of America

OCIS codes: 290.1310, 330.0330, 260.5430, 100.2000, 100.3020, 150.5670.

1. Introduction

Recently there has been a growing interest in the analysis of images acquired in poor-visibility conditions. The main objective has been to enhance^{1–4} images taken in poor visibility and even restore clear-day visibility of the scene.^{5–7} Some methods are based on specialized radiation sources and detection hardware.^{8,9} For natural light, visibility degradation effects due to haze vary as distances to the objects increase,^{10,11} and are referred to as aerial perspective.¹² For this reason, some image-enhancement methods proposed in the past require prior information about the distances of objects^{3,4} or about the scene colors.⁴

It has been observed that aerial perspective can actually be exploited to obtain an estimated range map^{6,7,13} of the scene. Computer vision methods have restored clear-day visibility of scenes using neither special radiation sources nor exact external knowledge about the scene structure or aerosols.^{5,7} These methods relied only on the acquired images but required weather conditions to change between im-

age acquisitions. This can take too long to make dehazing practical. In this paper we describe an approach that does not need the weather conditions to change and can thus be applied instantly.

Our approach is based on analyzing images taken through a polarizer. Polarization filtering has long been used in photography through haze.¹⁴ Relying only on optical filtering is, however, restrictive: It is sufficient only on clear days, with weak light scattering (mainly due to air molecules), when the Sun is $\approx 90^\circ$ to the viewing direction.^{14,15} In these situations photographers have set the polarization filter at an orientation that best improves image contrast. In general, however, polarization filtering alone *cannot* remove the haze from images. Our method further analyzes optically filtered images to obtain significantly better results.

The effects of scattering on light polarization have been extensively studied.^{16–21} Polarization has mainly been considered in the context of artificial illumination,^{22–26} where the signal to be recovered is associated with polarized light while the light scattered by the medium is associated with depolarized light.^{22–27} It has also been observed that an object masked by scattered light is enhanced by a linear superposition of polarization-filtered images.²⁸ In contrast we show that in hazy conditions, polarization can be associated with scattering of ambient illumination, rather than the signal (object radiance). Solving inverse problems associated with polarization-filtered images has proved to be useful in other regimes of scene analysis. For example, it

The authors are with Columbia Automated Vision Environment, Department of Computer Science, Columbia University, New York, New York 10027. Y. Schechner's e-mail address is yoav@cs.columbia.edu.

Received 18 January 2002; revised manuscript received 8 April 2002.

0003-6935/03/030511-15\$15.00/0

© 2003 Optical Society of America




Fig. 1. (Dashed rays) Light coming from the illuminant (e.g., Sun) and scattered toward the camera by atmospheric particles is the airlight (path radiance) A . The airlight increases with the distance z of the object. (Solid ray) The light emanating from the object is attenuated along the line of sight as z increases, leading to the direct transmission D . Without scattering, object radiance would have been L^{object} . The scene is imaged through a polarizing filter at angle α . The polarization component parallel to the plane of incidence is best transmitted through the filter at $\alpha = \theta_{\parallel}$.

was used to separate transparent and semireflected scenes,^{29,30} analyze specularities,^{31–33} and classify materials.³⁴

In this paper we describe an image-formation model that accounts for natural polarization effects in imaging through haze. We then invert the image-formation model to recover the dehazed scene and also to obtain information about scene structure and atmospheric properties. Our approach *does not* require modeling of the scattering particles' size or their precise scattering mechanisms. The principle is simple: The image is composed of two unknown components—the scene radiance in the absence of haze, and airlight (the natural illumination scattered toward the viewer). To recover these two unknowns, we need two independent images. We easily obtain these images because airlight is usually partially polarized. The method requires only that the airlight induce *detectable* partial polarization. We demonstrate removal of haze effects from real scenes in situations in which pure optical filtering (without applying our algorithm) does *not* suffice at all.

2. Theoretical Background

As depicted in Fig. 1, when imaging through the atmosphere we sense two sources. The first source is the scene object whose radiance is attenuated by scattering. The corresponding signal reaching the camera is called the *direct transmission*. The second source is the ambient illumination (Sun, sky, and so on). The part of the illumination scattered toward the camera by aerosols is called the *airlight*.^{7,10,11,35,36} It is also referred to in the literature as path radiance¹⁰ and veiling light.³⁷ In this section we describe each of these signals and the polarization effects associated with them.

A. Airlight Polarization

One of the causes of image degradation associated with atmospheric scattering is airlight. The atmosphere scatters light coming from the illumination sources (e.g., the Sun) toward the viewer⁷ (see Fig. 1). The airlight increases with the distance z from the object. As discussed in Subappendix A.2,

$$A = A_{\infty}[1 - t(z)], \quad (1)$$

where A_{∞} is the airlight radiance corresponding to an object at an infinite distance, e.g., the horizon. $t(z)$ is the transmittance of incoherent light,³⁵ given by

$$t(z) = \exp \left[- \int_0^z \beta(z') dz' \right], \quad (2)$$

where β is the coefficient of extinction due to scattering and absorption (see Subappendix A.2). When the extinction coefficient is distance invariant, $\beta(z') = \beta$, then

$$t(z) = \exp(-\beta z). \quad (3)$$

Assume for a moment that the illumination of any scattering particle comes from one direction (one illumination source). The light ray from the source to a scatterer and the line of sight from the camera to the scatterer define a plane of incidence. We divide the airlight into two polarization components that are *parallel* and *perpendicular* to this plane, A^{\parallel} and A^{\perp} , respectively. The airlight degree of polarization is then

$$p \equiv (A^{\perp} - A^{\parallel})/A, \quad (4)$$

where

$$A = A^{\perp} + A^{\parallel} \quad (5)$$

is the total radiance due to airlight, given in Eq. (1).

The degree of polarization greatly varies as a func-

tion of the viewing and illumination directions, the density of the scattering particles, and their sizes. Depending on the size distribution of the scattering particles,^{17,18,38} the airlight is partially linearly polarized either perpendicular to the plane of incidence^{15,39} or parallel to it. When scattering is dominated by independent air molecules and small dust particles (Rayleigh scattering), $A^\perp \geq A^\parallel$. On the other hand, larger haze particles may cause $A^\perp < A^\parallel$. We use the following convention throughout this paper: To avoid confusion without loss of generality, we associate the parallel component notation (\parallel) with the minimum measured radiance at a pixel and perpendicular component notation (\perp) with the maximum radiance. We now explain the effectiveness of polarization in various haze and illumination conditions.

1. Trivial Case

Under special conditions, optical filtering alone is sufficient to remove haziness in images. For Rayleigh scattering,^{14,15,34,36,40} the degree of polarization, p , is $(\sin^2 \phi)/(1 + \cos^2 \phi)$, where ϕ is the scattering angle (the angle between the illumination ray and the line of sight). Only when the light source is normal to the viewing direction is the airlight totally polarized ($p = 1$) perpendicular to the plane of incidence. Thus it can be eliminated if the image is captured through a polarizing filter oriented parallel to the plane of incidence. Dehazing in this case is trivial, since it is achieved with optical filtering alone. This situation is quite restricted, since it occurs only when the aerosols are very small and when the Sun is in a favorable position.

2. General Case

In general, the airlight will not be completely polarized.³⁶ Thus a polarizing filter cannot remove the airlight on its own. For example, in Rayleigh scattering the degree of polarization p decreases as the direction of illumination deviates from 90° (relative to the viewing direction). Reduction of polarization is caused by scattering from large haze particles, which never completely polarize light. Moreover, airlight due to large haze particles may be polarized orthogonally to light scattered by air molecules^{17,18,38} causing partial annihilation of polarization. The degree of polarization p is also decreased by *depolarization*. This is caused by multiple scatterings from multiple directions: An illuminant of a scattering particle may be another particle in the air (e.g., a haze particle, a cloud drop, a molecule creating the skylight). Multiple scatterings^{14,15,39,40} are more probable when the density of scatterers is high (poorer visibility). To make matters more complicated, these mechanisms depend on the wavelength.^{14,15}

Fortunately, our algorithm *does not* require explicit modeling of the precise mechanisms of scattering. The method is based on the fact that even a partial polarization of the airlight can be exploited in post-processing to remove scattering effects. This degree of polarization needs to be significant enough to be

detected by the sensor. For this reason we concentrate in this paper on vision through haze, in which multiple scattering is much weaker than in fog. There are some weather conditions under which the algorithm may not be effective, as discussed in Section 8.

B. Direct Transmission Polarization

In addition to the presence of airlight, the scattering medium degrades images by attenuating the light emanating from scene objects. Let the object radiance be L^{object} in the absence of haze (scattering) between the scene and the viewer. When haze is present, as a ray of light progresses toward the viewer (Fig. 1), part of its energy is scattered in other directions, and a small portion of it may be absorbed. Thus the radiance sensed by the viewer is an attenuated fraction of L^{object} , called the *direct transmission*.⁷ As a function of the distance z the direct transmission is

$$D = L^{\text{object}} t(z), \quad (6)$$

where $t(z)$ is given in Eq. (2).

We make three approximations in this paper. First, we concentrate on the degradation due to the attenuation of the signal and due to the additive airlight. We do not deal with image blur. Second, we take single-scattering effects to be dominant over multiple-scattering effects, which cause image blur and reduce the degree of polarization along the line of sight. Finally, we assume that light emanating from scene objects has insignificant polarization.

It follows from the last assumption that the polarization of the direct transmission is insignificant. If the scattering particles have random orientations, then the directly transmitted light will not be polarized in any macroscopically preferable orientation. Hence the polarization state of the unscattered light does not change,^{39,40} although the radiance is attenuated.

The last assumption is invalid for specular surfaces. Nevertheless, when a specular object is far enough, its direct transmission makes a negligible contribution to the measured polarization. The reason is that the direct transmission decreases [Eq. (6)] whereas airlight¹⁰ increases [Eq. (1)] with distance. Thus airlight and its polarization dominate the measured light for distant objects. Hence the model becomes more accurate where it is needed most—for distant objects that are most affected by haze. This property is useful if we know the relative distances to the scene objects.

Recall that airlight is just the aggregation of light scattered by particles at various distances along the line of sight. Since the polarization of this light does not change along the line of sight^{39,40} p [Eq. (4)] does not depend on the distance.

The observations regarding the polarization described in Subsections 2.A and 2.B are unaffected by the exact dependence of t on z . The dominance of airlight polarization means that the polarizing filter




Fig. 2. At each point the minimum measured image irradiance as a function of α is I^{\parallel} . The maximum is I^{\perp} . The difference between these measurements is due to the difference between the airlight components A^{\parallel} , A^{\perp} . This difference is related to the unknown airlight A by the parameter p , which is the airlight degree of polarization. Without a polarizer the image irradiance is I^{total} , which is proportional to the sum of airlight and the unknown direct transmission.

modulates mainly the measured airlight but not the light originating from the objects. This is the key to subsequent calculations that remove the effects of haze.

3. Image Formation

The scene radiance is measured by the detector plane of the camera. The detected image irradiance is proportional to scene radiance. Since the proportionality depends on the optical system parameters and not on the weather effects, we treat the image irradiance and the scene radiance as equivalent. The overall radiance we sense is the incoherent sum of the airlight and the direct transmission. Without mounting a polarizer on the camera, the image irradiance is

$$I^{\text{total}} = D + A, \quad (7)$$

up to the said proportionality factor. It has been shown¹⁰ that except for rather close objects (for which $\beta z < 0.2$), I^{total} is typically dominated by the airlight and not by the direct transmission. Thus typically most of the light we measure is not attributed to the signal D , whose origin is L^{object} . This is reinforced by the fact that most terrestrial objects have a low albedo, further decreasing the signal.

When a linear polarizer is mounted on the camera, the image irradiance changes as a function of the polarizer orientation angle α . Figure 2 describes the irradiance at a single pixel. The irradiance is a cosine function of α . On average, the image irradiance is $I^{\text{total}}/2$. (We do not deal here with the global absorptivity of unpolarized light that is common in sheet linear polarizers. Although this absorptivity effects the image irradiance, it does not modulate the effects of haze.)

One of our goals is to decouple the airlight and

direct transmission. Since we assume that direct transmission is not polarized, its energy is evenly distributed between the polarization components. The variations due to the polarizer rotation are assumed to be mainly due to airlight. As seen in Fig. 2, when the polarizing filter is oriented such that the image irradiance is minimal ($\alpha = \theta_{\parallel}$), we measure

$$I^{\parallel} = D/2 + A^{\parallel}, \quad (8)$$

where [from Eqs. (4) and (5)]

$$A^{\parallel} = A(1 - p)/2. \quad (9)$$

This is the best state of the polarizer because here the image irradiance is the closest to the irradiance corresponding to the direct transmission (except for a factor of $1/2$). There is a difference between I^{\parallel} and $D/2$, because the airlight is not completely polarized ($A^{\parallel} \neq 0$).

In Section 4 we recover D by comparing two images taken with different orientations of the polarizer. For instance, one image can be I^{\parallel} , whereas the other,

$$I^{\perp} = D/2 + A^{\perp}, \quad (10)$$

is acquired when the filter is oriented perpendicular to θ_{\parallel} . From Eqs. (4) and (5),

$$A^{\perp} = A(1 + p)/2. \quad (11)$$

From Eqs. (5), (8), and (10),

$$I^{\text{total}} = I^{\parallel} + I^{\perp}. \quad (12)$$

Note that I^{\perp} is the worst state of the polarizer, because the airlight is enhanced relative to the direct transmission. To dehaze the image, we first have to remove the airlight A . The key step here is the estimation of p , the degree of polarization of airlight. As shown in Fig. 2, p relates the unknown airlight A to the difference between the image irradiances I^{\perp} and I^{\parallel} .

Acquisition of polarimetric images is easy and feasible at the video rate.^{34,41–44} However, for demonstration purposes we photographed the scene on a Fuji Sensia 100 slide film, using a common SLR (single-lens-reflex) (Cannon EOS-5) camera. The slides were scanned by a Nikon LS2000 35-mm film scanner. Before processing the images by the algorithm described in the following sections, we linearized the raw photographs to compensate for the system's radiometric response. This response was estimated from images of the Macbeth ColorChecker.⁴⁵ We modulated the polarization by mounting a standard linear polarizer on the camera lens.

We took images of a distant scene at two different orientations of a polarizer approximately corresponding to the perpendicular and the parallel airlight polarization components. The images are shown in Fig. 3. The raw images were *not* acquired in the trivial situation described in Subsection 2.A.1: The Sun was almost behind the camera (the Sun was in the south, and the picture was taken toward the north), whereas the haze was rather dense. For this



Fig. 3. Images of the polarization components corresponding to the minimal and the maximal radiances. Note that I^{\parallel} (the image of irradiance) has the best image contrast that optics alone can yield, and yet there is no significant improvement over the image of the worst polarization state.

reason, I^{\parallel} has only a slight improvement of image contrast relative to the contrast in I^{\perp} . Because of the partial polarization of the airlight, I^{\parallel} was lower than I^{\perp} . For clarity of display, the brightness of each of the photos shown in Fig. 3 is contrast stretched.

4. Dehazing Images

Both the attenuation [Eq. (6)] and the airlight [Eq. (1)] depend on the distance z of the scene point. The distance to the objects is *spatially varying*, since different image pixels (x, y) correspond to objects at different distances. Therefore compensation for the effects of haze is spatially varying. Our dehazing method automatically accounts for this spatial variation of scene depth.

For each image pixel we have three unknowns: the object radiance (without haze) L^{object} , the airlight A , and the transmittance $t(z)$. These determine the irradiance at each image pixel. The airlight is related to $t(z)$ by Eq. (1). Thus the number of unknowns per pixel is reduced to two. These unknowns can be estimated from two images taken at almost any gen-

eral unknown (but nondegenerate) orientations of the polarizing filter. We actually do not need the four measurements commonly used for full estimation of polarization. The reason for this is that the goal is not the estimation of polarization but the dehazing of the scene, and this can be done with two raw images. This is proved in Subappendix A.1. Nevertheless, the most-stable results are obtained if the algorithm is based on I^{\parallel} and I^{\perp} . Therefore we focus here on this case.

Let the raw images correspond to the estimated polarization components, \hat{I}^{\parallel} and \hat{I}^{\perp} . We assume that we have an estimate of the global parameters A_{∞} and p . Estimation of these parameters is described in Section 7. From Eqs. (8)–(11) it is seen that we can estimate the airlight of any point as

$$\hat{A} = (\hat{I}^{\perp} - \hat{I}^{\parallel})/p, \quad (13)$$

and the unpolarized image [Eq. (12)] as

$$\hat{I}^{\text{total}} = \hat{I}^{\parallel} + \hat{I}^{\perp}. \quad (14)$$

With Eq. (7) the estimated direct transmission is therefore

$$\hat{D} = \hat{I}^{\text{total}} - \hat{A}. \quad (15)$$

In this image the additive effect of the airlight is removed.

Recall that besides the addition of airlight, the haze attenuates the light coming from the object. The transmittance is estimated from Eqs. (1) and (13) as

$$\hat{t} = 1 - \hat{A}/A_{\infty}. \quad (16)$$

Thus we can compensate for the attenuation of the transmitted light. From Eqs. (6), (15), and (16) we obtain an estimate for the radiance that would have been sensed in the absence of atmospheric attenuation,

$$\hat{L}^{\text{object}} = \frac{\hat{I}^{\text{total}} - \hat{A}}{\hat{t}} = \frac{\hat{I}^{\text{total}} - \hat{A}}{1 - \hat{A}/A_{\infty}}. \quad (17)$$

\hat{L}^{object} is hence the dehazed image. Dehazing thus amounts to simple pointwise operations such as subtraction and division of corresponding pixel values. Since different image pixels are processed independently, the spatially varying inversion of the haze effects is implicit: \hat{A} is spatially varying, since it is larger for the more-distant objects.

We note that A_{∞} , p , and the extinction coefficient β are functions of the light wavelength λ . For Rayleigh scattering,^{39,40} $\beta \sim 1/\lambda^4$, so the airlight in moderate haze is typically bluish. To account for the wavelength dependence, it is best to analyze the images with high spectral resolution. Each wavelength band can be analyzed independently. In our experiments, though, we used the traditional coarse wideband red, green, and blue (RGB) sampling of the spectrum.

We applied the dehazing method to the images shown in Fig. 3, after estimating the parameters A_{∞} and p on the basis of sky measurements (for the parameter estimation, see Section 7 and in particular the stabilizing approach in Subsection 7.C). The resulting dehazed image is shown in Fig. 4(a). The brightness of the displayed image has the same scale as was used for displaying the best-polarized image I^{\parallel} in the bottom of Fig. 3. The contrast of features in the dehazed image is greatly improved relative to I^{\parallel} and I^{\perp} . Note, for instance, the distant mountain ridge (especially on the left), which is not visible in the raw images. Moreover, the algorithm removed the blue color bias, which existed in the raw images. This enables distinguishing the different vegetation types by hue. Although most of the scene is dehazed, the sky recovery is noisy, and there is a residual haziness at the more distant objects. These artifacts are explained in Subsection 7.C.

As another example consider the images shown in Figs. 5(a) and 5(b). Here \hat{I}^{\parallel} and \hat{I}^{\perp} were calculated from images taken at several orientations of the po-

larizer, as is commonly done in polarimetric imaging.^{29,31} This experiment was conducted in conditions far from the trivial case: The haze was dense (visibility of a few kilometers), and the contrast in the parallel component was only slightly better than in the perpendicular component (all displayed images are linearly contrast-stretched versions of the raw images). The dehazed image is shown in Fig. 5(c). We obtain a significant improvement of contrast and color: The green forest is visible in the distant scene, whereas in the raw images that area looks like grayish-blue noise. The colors of the red bricks and roofs of the distant buildings are also restored.

5. Range Map of the Scene

A. Range Estimation

As a byproduct of the dehazing process we get an estimate of the range (depth) map of the scene. The estimation exploits the dependence of the transmittance t on the distance z . Note that t is always a monotonically decreasing function of z , and thus Eq. (16) immediately indicates the distance ordering of objects in the scene. Assuming the extinction coefficient to be distance invariant, $t(z) = \exp(-\beta z)$. Then, from Eq. (16),

$$\widehat{\beta z} = -\ln[1 - \hat{A}(x, y)/A_{\infty}]. \quad (18)$$

Note that the distance z is estimated as a function of (x, y) up to a global scale factor, which is the unknown extinction coefficient β .

Recall that we get an independent estimated range map for each color channel: $\widehat{\beta_r z}$, $\widehat{\beta_g z}$, and $\widehat{\beta_b z}$, where the subscripts r, g, b denote the three color channels. These maps should differ only in their scale. This scale is set by the ratios of their scalar extinction coefficients, β_r , β_g , and β_b . We combine the range maps into a single, average one:

$$\overline{\beta z}(x, y) \equiv [\widehat{\beta_r z}(x, y) + \widehat{\beta_g z}(x, y) + \widehat{\beta_b z}(x, y)]/3. \quad (19)$$

Note that there may be more optimal methods of combinations, such as a weighted average, in which the weights depend on the noise in each channel. The estimation may be further improved if it is based on narrow spectral bands, rather than on RGB.

We derived the range map of the scene corresponding to Fig. 3 as a byproduct of dehazing. To make the depth map more robust to noise, we median filtered \hat{A} . The range map is shown in Fig. 6. The darker points correspond to more distant objects. The map is qualitatively consistent with the scene, for example, indicating the close tree in the foreground and the gradual increase of range in the background.

The range map can be used to render the scene from viewpoints other than the one used during acquisition. We texture map the dehazed image on the range map (surface) and then look at the texture-



Fig. 4. (a) The dehazed image has much better contrast and color than the optically filtered image, especially in the distant regions of the scene (compare with Fig. 3). (b) and (c), As described in Section 5, we estimate the range map of the scene. We use it to render the dehazed scene from different perspectives, as if the viewer descends. Note the occlusion of the background by the foreground tree on the right. Note also the distant mountains occluded by the closer ridge.

mapped surface from various viewpoints. Examples for this application are shown in Figs. 4(b) and 4(c). Here the appearance of the scene is shown from two different viewpoints. The images are rendered by rotation of the textured surface by 22° and 31° relative to normal viewing direction shown at the top of Fig. 4. This creates the impression that the viewer gradually descends relative to the acquisition position.

One may see that the valley with the agricultural segments is planar, since straight lines on it remain quite straight when viewed from different directions. It can clearly be seen in Fig. 4 how the large tree on the right-hand side of foreground occludes the buildings in the valley behind it and the mountain range bounding the valley. In addition, one may see that there are consecutive mountain ranges (the close one, which bounds the valley, is greener, whereas the far-

ther ones are more pale as explained in Subsection 7.C). As the viewer descends, the farther ridges become occluded by the close ridge, and as the viewer ascends, the distant ridges gradually appear.

B. Range Accuracy

We now analyze the accuracy of range estimation on the basis of scattering. Let the uncertainty (scaled by the camera's dynamic range) in the measurement of \hat{I}^{\parallel} or \hat{I}^{\perp} be $\sigma = 2^{-b}$. Here b is the number of bits per pixel of the camera, assuming that quantization is the dominant noise source and that the radiometric response is linear. The uncertainty of the estimated depth is then

$$\delta z \approx \frac{1}{\beta} \exp(\beta z - b \ln 2) \frac{1}{pA_{\infty}}. \quad (20)$$



Fig. 5. Photograph with the best contrast that optics alone can give (a) is almost as poor as the worst polarization state (b). The dehazed image (c) has much better contrast and color, especially in the distant regions of the scene (note the green forest and the red roofs).

The range uncertainty grows exponentially with the distance. Beyond some distance z_{\max} , range estimation is too uncertain to be reliably considered. Nevertheless, relation (20) shows that the degrading effect due to growth of distance z can be compensated by a proportional increase of the number of bits b . Thus, $z_{\max} \propto b$. For example, a 12-bit camera will be able to probe 50% farther into the scene than an 8-bit camera.

One may note that haze can actually be useful for estimating distances, since without scattering the perception of depth due to airlight is lost. Indeed, $\delta z \rightarrow \infty$ when $\beta \rightarrow 0$, i.e., when attenuation is weak. On

the other hand, $\delta z \rightarrow \infty$ also when $\beta \rightarrow \infty$, that is, when attenuation is too strong (e.g., in fog). The optimal extinction coefficient minimizing δz is $\beta = 1/z$. For example, estimates of object ranges around the distance of 3 km are most accurate on a day when the effective attenuation distance due to haze is 3 km [$\beta = 1/3 \text{ (km}^{-1}\text{)}$].

The estimation is prone to significant error for close objects reflecting significantly polarized light such as specular objects. As discussed in Subsection 2.B, the degree of polarization of the distant, hazy objects is small relative to the airlight. This may not be true for close objects. Figure 7 shows




Fig. 6. Range map of the scene shown in Fig. 3, estimated as a byproduct of the dehazing algorithm. The farther the object, the darker the shading.

the range map obtained for the scene shown in Fig. 5. The map is qualitatively consistent with the scene, for instance, indicating the close buildings and the distant ridge. Yet a significant partial polarization was observed in some surfaces on the close buildings, especially those directly lit by the Sun. In Fig. 7 this manifests in a dark shading of the points corresponding to these objects (rather than a bright shade). Note also that haze homogeneity is the basis for Eq. (18); thus range estimation becomes less accurate when there are significant spatial variations of the haze.

To conclude this section, using haze and polarization to estimate a range map of the scene is prone to several sources of inaccuracy. Nevertheless, it enables passive estimation of very long distances, without resorting to geometric methods (triangulation), which are prone to matching problems. Instead of geometry, the use of photometry enables the rough estimation of the large distances from a single viewpoint. The limit of the estimated range depends on the visibility. We demonstrated ranging in a scene with visible objects tens of kilometers away, but the maximum range can be much larger, depending on β and b .




Fig. 7. Range map of the scene shown in Fig. 5, estimated as a byproduct of the dehazing algorithm. Some surfaces of close objects are wrongly marked as distant ones as a result of their high degree of polarization.

6. Information about the Aerosols

In Section 4 we showed that on the basis of as few as two images we can dehaze the imaged scene. Now we will show that with the same raw images we can extract information related to the atmospheric particles that degrade the scene visibility.

Consider the range maps of each color channel, which were described in Section 5. Averaging over the image pixels, we define scalars corresponding to each color channel:

$$s_r = \frac{\sum_{x,y} \widehat{\beta_r z}(x, y)}{\sum_{x,y} \widehat{\beta z}(x, y)}, \quad (21)$$

$$s_g = \frac{\sum_{x,y} \widehat{\beta_g z}(x, y)}{\sum_{x,y} \widehat{\beta z}(x, y)}, \quad (22)$$

$$s_b = \frac{\sum_{x,y} \widehat{\beta_b z}(x, y)}{\sum_{x,y} \widehat{\beta z}(x, y)}. \quad (23)$$

These scalars express the extinction coefficients of the atmosphere, in each of the color channels, up to a single scale factor. This result is valuable because the relative scattering coefficients are determined by the size of the scattering particles.^{2,15,39} Assuming that scattering is the dominant process of attenuation, these ratios provide rough indication about the distribution of the particles' size. This information may be used in conjunction with other methods for estimating the particle size from spectral and polarization information.¹⁶ It may be incorporated into models that make explicit physical analysis of atmospheric scattering, as well as in applications of ecological monitoring. As with image dehazing and range estimation, this application would be more accurate if narrow spectral bands were used, rather than RGB.

In the experiment based on the images shown in Fig. 5, we obtained

$$\begin{pmatrix} s_r \\ s_g \\ s_b \end{pmatrix} = \begin{pmatrix} 0.26 \\ 0.32 \\ 0.42 \end{pmatrix}, \quad (24)$$

which means that the scattering in the blue band is ~60% stronger than the scattering in the red band. Had the dominant particles been small enough to obey Rayleigh's $1/\lambda^4$ rule, we might have expected that the scattering of the blue wavelengths were much stronger relative to the red wavelengths. It is difficult to pinpoint the exact relative strength when very broadband spectral channels are used. So to get a rough estimate, if we take 450 and 650 nm as typical mid-band wavelengths for blue and red, respectively, then in Rayleigh scattering we may expect the blue scattering to be (an order of) 300% stronger

than the red one. Therefore, as expected, the experiment was conducted under conditions in which scattering was dominated by particles that do not fit into the Rayleigh model. Therefore, even if the Sun had been perpendicular to the viewing direction (the trivial case for Rayleigh scattering), light would not have been sufficiently polarized to enable dehazing by optical filtering alone.

7. Estimating A_∞ and p

To dehaze the image with Eqs. (13), (15), and (17), we need an estimate of the global parameters A_∞ and p . In Ref. 10 it was concluded that human vision can correct for effects of aerial perspective on the basis of context. We use this observation as the way to obtain A_∞ and p . The image context we mainly rely on is the sky. This section discusses estimation of these parameters. The experimental results were based on the principles described in Subsections 7.A–7.C.

A. Sky Measurements at the Horizon

Since $t(z) \rightarrow 0$ as $z \rightarrow \infty$, we get

$$I^{\text{total}} = L^{\text{object}}t(z) + A_\infty[1 - t(z)] \rightarrow A_\infty. \quad (25)$$

The degree of polarization of the measured scene (i.e., the direct transmission combined with airlight) is

$$\hat{P}(x, y) = \frac{\Delta I(x, y)}{I^{\text{total}}(x, y)}, \quad (26)$$

where

$$\Delta I(x, y) \equiv [\hat{I}^\perp(x, y) - \hat{I}^\parallel(x, y)]. \quad (27)$$

As $z \rightarrow \infty$,

$$\hat{P}(x, y) \rightarrow \frac{A_\infty^\perp - A_\infty^\parallel}{A_\infty^\perp + A_\infty^\parallel} = p. \quad (28)$$

We can measure these parameters directly from the images. We can use points that are seen through enough haze such that their direct transmission is practically zero. Such points are not always available, so we use some heuristics based on context (as in human vision¹⁰) to estimate these parameters. The most direct way is to measure patches of the sky at the horizon:

$$\hat{A}_\infty = I^{\text{total}}(\text{sky}), \quad \hat{p} = \frac{\Delta I(\text{sky})}{I^{\text{total}}(\text{sky})}. \quad (29)$$

As an example, in the experiment corresponding to the images shown in Fig. 5, the average measured values of \hat{p} with Eq. (29) were

$$\begin{pmatrix} \hat{p}_r \\ \hat{p}_g \\ \hat{p}_b \end{pmatrix} \approx \begin{pmatrix} 0.28 \\ 0.25 \\ 0.22 \end{pmatrix}. \quad (30)$$

Note that $\hat{p}_r > \hat{p}_g > \hat{p}_b$. This is consistent with the literature^{14,15}: In haze, long wavelengths (red) undergo less multiple scattering and therefore maintain

a higher degree of polarization, compared with short (blue) wavelengths. Since p depends on the size and density of the scatterers,¹⁵ the estimation of p in the different spectral bands may provide additional information about the aerosols in the scene.¹⁶

Note that if the horizon is cloudy and $t(\infty) \neq 0$, then the measured light is due not only to haze airlight but also to the object (cloud). Thus this method will be erroneous, and we may need to apply the method described in Subsection 7.D.

B. Spatial Variability

Although we treat the parameters A_∞ and p as global, they may vary across the field of view. The sky (horizon) radiance A_∞ depends on the angular scattering coefficient $\beta(\phi)$, as explained in Subappendix A.2. Therefore it depends on the position of the Sun relative to the viewing direction.¹² For instance, because of strong forward scattering and backscattering, A_∞ will usually be larger when the Sun is in front or behind the camera. Also, p depends on the position of the Sun relative to the viewing direction.

The spatial variations of A_∞ and p across the field of view are much slower than the typical variations in radiance that are due to texture. Thus we can account for the horizontal variation of A_∞ and p by sparsely sampling the sky radiance at the horizon across the field of view and then interpolating the values, using a smoothly varying function. In the experiments we performed we estimated \hat{A}_∞ and \hat{p} by sparsely measuring the sky values above the distant ridges across the images. We then fit a second-order polynomial to the measurements.

Vertical variations are more complicated. Sky measurements can change as a function of altitude. Haze density can change significantly as a function of altitude¹¹ within the first few vertical kilometers of atmosphere,² and even within a few hundred meters. Moreover, even on a clear day the sky radiance and polarization change as a function of the viewing direction relative to the zenith.^{38,46} The plane of polarization may change by 90° above the solar and antisolar horizons when the Sun is low. Thus one has to be careful when applying this method in a large vertical field of view. A direct implication is that measuring \hat{A}_∞ and \hat{p} from a sky patch at high elevation angle will usually be error prone. For this reason we measure the sky values close to the horizon. We also applied a stabilizing bias, described next.

C. Stabilizing the Recovery

Even if all the measurements were perfect, we believe that avoiding a complete inversion of the image formation is beneficial from a perceptual-aesthetic point of view. There are two reasons for this. First, attempting to remove the atmospheric scattering from the sky itself means that the daylight sky should be removed. In a noiseless image the result would be dark deep-space sky in a clear daylight image! This is perceptually annoying. The sky appearance is further degraded by unstable amplification of noise,

since the denominator in Eq. (17) approaches zero when $\hat{A} \rightarrow A_\infty$. Second, even on the clearest day, we still encounter airlight because of scattering by the air molecules, which manifests in a bluish tint in the distance. Images lacking any aerial perspective,¹² in which far objects are the same as close ones, look strange and artificial.

If the dehazed images are meant for human inspection, it is preferable to avoid these phenomena. This is easily achieved if we somewhat bias \hat{p} by multiplying it by a factor ϵ , such that $1 \leq \epsilon \leq 1/\hat{p}$:

$$\hat{p} \rightarrow \epsilon \hat{p}. \quad (31)$$

If $\epsilon = 1$, we get Eq. (17) as the solution to the inverse problem. In the other extreme we may set $\epsilon \hat{p} = 1$ (recall that $p \leq 1$), and then the algorithm behaves as if the airlight were completely polarized: From Eqs. (13) and (15) $D(x, y) = \hat{I}^{\parallel}(x, y)$ when $p = 1$; hence the unprocessed best-polarized image component is used as the value for the direct transmission. For all intermediate values of ϵ the inverse solution (17) is moderated by means of weighting it with the raw image. Using a value of ϵ slightly larger than 1 leaves a residual haziness (originating from the raw image) that grows with the distance, thus making the image consistent with naked-eye experience. It can also be shown that when $\epsilon > 1$ the estimated dehazed sky value is $\hat{L}^{\text{object}}(\text{sky}) = A_\infty$; i.e., it retains its raw unpolarized sky value, automatically. The noise in the estimated sky value is

$$\delta \hat{L}^{\text{object}}(\text{sky}) = \sqrt{2}\sigma(1 - 1/\epsilon)^{-1}, \quad (32)$$

where σ is the noise standard deviation of the raw image components. It can be seen that noise is amplified in the sky by the dehazing process, but this amplification decreases with the increase of ϵ .

This bias is also beneficial to counter effects of error stemming from inaccurate estimation of p and A_∞ . As described in Subsection 7.B, we may expect such inaccuracy if it is based on sky measurements. From Eqs. (13)–(15), (26), and (27),

$$\hat{D}(x, y) = \frac{\hat{I}^{\text{total}}(x, y)}{\hat{p}} [\hat{p} - \hat{P}(x, y)]. \quad (33)$$

If our estimate of the airlight degree of polarization is too low ($\hat{p} < p$), then negative values can appear in the image of the direct transmission. This is especially relevant to distant objects, because $P(x, y) \rightarrow p$ when $z \rightarrow \infty$. Biasing \hat{p} with Eq. (31) reduces the occurrence of such problems. In addition, Eq. (31) better conditions the compensation for attenuation in case \hat{A}_∞ is inaccurate [see Eqs. (13), (16), and (17)].

In the experiments we performed, we first estimated \hat{A}_∞ and \hat{p} as described in Subsections 7.A and 7.B. Indeed, that heuristic method resulted in a slight error, and many of the resulting pixels of \hat{D} and \hat{L}^{object} had negative values, especially in the distant areas. In addition, the sky was quite noisy. So, to get the dehazing results presented in this paper, we fine tuned \hat{p} by increasing its values globally

by a few percent ($\epsilon = 1.09$). We thus gave up a few percent of the full inversion, and in return almost all the occurrences of negative values were eliminated and the sky regained a natural color (a tolerable noise is still present). As can be seen in Fig. 4, distant objects have a residual bluish haziness (aerial perspective¹²), making the recovered scene look perceptually acceptable.

D. Unavailable Sky Measurements

When direct sky measurements are not possible to obtain, we need to get the context for estimating p and A_∞ from nonsky regions of the scene. We now show that such an estimate can be obtained from *a priori* information about the identity of several scene points with similar but unknown radiance L^{object} , had the haze been absent. For instance, this is possible if parts of the scene had been observed on a clear day so that some objects are known to have corresponding colors. From Eqs. (13), (17), and (27),

$$\hat{I}^{\text{total}}(x, y) = L^{\text{object}} + \left(\frac{1}{p} - \frac{L^{\text{object}}}{pA_\infty} \right) \Delta I(x, y). \quad (34)$$

Assume that we know of a set of (at least two) pixels (x_k, y_k) for which the object radiance in the absence of haze is the same, $L^{\text{object}} = L_1^{\text{object}}$, but their distances from the viewer are different. For instance, the set of pixels can correspond to similar bushes at different distances. The value L_1^{object} does not have to be known in advance, as we will see in the following. Because of the differing depths, these points (x_k, y_k) will have different values of $\hat{I}^{\text{total}}(k)$ and $\Delta I(k)$. For all these points, however,

$$C_1 \equiv \left(\frac{1}{p} - \frac{L_1^{\text{object}}}{pA_\infty} \right) \quad (35)$$

is constant. Thus $\hat{I}^{\text{total}}(k)$ as a function of $\Delta I(k)$ forms a straight line,

$$\hat{I}^{\text{total}}(k) = L_1^{\text{object}} + C_1 \Delta I(k), \quad (36)$$

whose intercept on the \hat{I}^{total} axis is the radiance value L_1^{object} . Therefore knowing points (x_k, y_k) that have corresponding radiances, we can estimate their corresponding dehazed radiance by fitting a line to the measured $\hat{I}^{\text{total}}(k)$ and $\Delta I(k)$. The slope of the fitted line is C_1 .

Now that we know L_1^{object} and C_1 , we may rewrite Eq. (35) as

$$p = \frac{1}{C_1} - (L_1^{\text{object}}/C_1) \frac{1}{A_\infty}. \quad (37)$$

Thus the unknown p and A_∞ are constrained to lie on a line in the $(p, 1/A_\infty)$ plane. The line is defined by the already estimated L_1^{object} and C_1 . Now we can look at a different set of (at least two) pixels (x_n, y_n) , which in the absence of scattering effects have the same radiance L_2^{object} , where $L_2^{\text{object}} \neq L_1^{\text{object}}$. Once again, analogous to Eqs. (35) and (36), L_2^{object} and the corresponding C_2 are estimated if the pixels (x_n, y_n)

correspond to scene points at different distances. This supplies another line constraint,

$$p = \frac{1}{C_2} - (L_2^{\text{object}}/C_2) \frac{1}{A_\infty}. \quad (38)$$

The intersection of these lines (37, 38) yields the estimated values for p and A_∞ . For $N > 2$ sets of pixels corresponding to unknown dehazed object radiances L_l^{object} where $l = 1, \dots, N$, the estimation of p and A_∞ becomes more robust. The minimum, however, is two sets of two points. Note that for identifying the sets of pixels, this method requires some user interaction, as in the estimation using sky measurement.

8. Discussion

We have shown that physics-based image analysis that follows acquisition of polarization-filtered images can remove visual effects of haze. Although it is based on some approximations, this approach proved to be effective in dehazing, when the problem could not be solved by optics alone. The method is quick and does not require temporal changes in weather conditions. In addition to the dehazed image, the method also yields information about scene structure and the atmospheric particles. These results can form the basis for useful tools in photography and remote sensing.

Our method is based on the partial polarization of airlight. Therefore its stability will decrease as the airlight degree of polarization decreases. For instance, the method may be less effective when the illumination is less directional (overcast skies). We expect it to have just a limited effect, or even fail, in cases of strong depolarization, as occurs in fog. Nevertheless, with more-exact scattering models, such as those that include multiple scattering, this research may be extended to complicated weather conditions and perhaps to other scattering media (e.g., underwater environments³⁴ and tissues).

Appendix A

1. Dehazing with Two Arbitrary Images

In Sections 4 and 5 we used estimates of I^\parallel and I^\perp in the dehazing algorithm. We now show that in theory the method can work on the basis of two images taken through any nondegenerate polarization orientations. Let θ_\parallel be the orientation of the polarizer for best transmission of the component parallel to the plane of incidence (Fig. 1). For a general orientation α the observed airlight is

$$A(\alpha) = A\{1 - p \cos[2(\alpha - \theta_\parallel)]\}/2, \quad (A1)$$

which coincides with Eqs. (9) and (11) if $\alpha = \theta_\parallel$, $\theta_\parallel + 90^\circ$. Assume that we take two images of the scene with arbitrary orientations of the polarizer, $\alpha_1 \neq \alpha_2$. Because the direct transmission is unaffected by the

polarizer orientation, the images are

$$I_1 = D/2 + A(\alpha_1), \quad (A2)$$

$$I_2 = D/2 + A(\alpha_2). \quad (A3)$$

Let us define an effective airlight

$$A_{\text{effective}} \equiv A(\alpha_1) + A(\alpha_2), \quad (A4)$$

with an effective degree of polarization

$$p_{\text{effective}} \equiv \frac{A(\alpha_2) - A(\alpha_1)}{A_{\text{effective}}}, \quad (A5)$$

where we set $A(\alpha_2) \geq A(\alpha_1)$, without loss of generality. We also define an effective unfiltered image

$$I_{\text{effective}}^{\text{total}} \equiv I_1 + I_2 = D + A_{\text{effective}}. \quad (A6)$$

It can easily be shown that $A_{\text{effective}}$ is proportional to the actual airlight,

$$A_{\text{effective}} = fA = fA_\infty[1 - t(z)] = A_\infty^{\text{effective}}[1 - t(z)], \quad (A7)$$

where $A_\infty^{\text{effective}}$ is the effective airlight at infinity (the horizon). The proportion factor f is

$$f = 1 - p \cos(\alpha_1 + \alpha_2 - 2\theta_\parallel) \cos(\alpha_1 - \alpha_2). \quad (A8)$$

Since we do not know θ_\parallel on the basis of two arbitrary polarizer angles, f is unknown.

Assume now that we have estimates of the parameters $p_{\text{effective}}$ and $A_\infty^{\text{effective}}$. These parameters can be estimated by measurement of the image irradiances I_1 and I_2 at the sky, similar to the way described in Section 7. Then we estimate the effective airlight at each point,

$$\hat{A}_{\text{effective}} = \frac{I_2 - I_1}{p_{\text{effective}}}. \quad (A9)$$

From Eq. (A6) the estimated direct transmission based on the raw images I_1 and I_2 is

$$\hat{D} = I_{\text{effective}}^{\text{total}} - \hat{A}_{\text{effective}}. \quad (A10)$$

From Eq. (A7) the estimated transmittance is

$$\hat{t} = 1 - \frac{\hat{A}_{\text{effective}}}{A_\infty^{\text{effective}}}. \quad (A11)$$

Thus the dehazed image is

$$\hat{L}_{\text{object}} = \frac{I_{\text{effective}}^{\text{total}} - \hat{A}_{\text{effective}}}{1 - \hat{A}_{\text{effective}}/A_\infty^{\text{effective}}}. \quad (A12)$$

We can check the stability of using an arbitrary pair of images. It is easy to show that

$$p_{\text{effective}} = \frac{Ap}{A_{\text{effective}}} \sin(\alpha_1 + \alpha_2 - 2\theta_\parallel) \sin(\alpha_2 - \alpha_1). \quad (A13)$$

Equation (A9) becomes unstable when $p_{\text{effective}} \rightarrow 0$. Besides the obvious case in which $p = 0$, this happens when

$$(\alpha_1 + \alpha_2)/2 = \theta_{\parallel}, \theta_{\parallel} + 90^\circ. \quad (\text{A14})$$

This is expected because the acquired images are equal if taken on symmetric angles relative to the extrema of the cosine in Eq. (A1). Therefore changing the orientation from α_1 to α_2 is degenerate. Except for these singular cases, dehazing is possible with two images. The best stability of dehazing is achieved when $p_{\text{effective}}$ is maximum, that is, when $\alpha = \theta_{\parallel}, \theta_{\parallel} + 90^\circ$. Therefore we focus here on dehazing based on I^{\parallel} and I^{\perp} .

By rotating the polarizer to achieve an extremum of the image irradiance or contrast, it is often easy to detect visually the states corresponding to I^{\parallel} and I^{\perp} . However, it is easier and more accurate to estimate these components with 3 or more images taken through different general orientations of the polarizer. This is a common practice in polarization imaging, as detailed in Refs. 29, 31, 34, 43, 47, and 48.

2. Inhomogeneous Attenuation: Model and Recovery

This subsection describes the image-formation model and the dehazing method when the attenuation varies along the line of sight. It is given here mainly to make the paper self-contained for readers unfamiliar with radiative transfer. When light propagating toward the camera passes through an infinitesimal layer of scattering media, some percentage of it is lost as a result of scattering to other directions and as a result of absorption. For a layer of thickness dz' , the direct transmission change dD is given^{7,36} by the differential equation

$$\frac{dD(z')}{D} = -\beta(z')dz'. \quad (\text{A15})$$

$\beta(z')$ is the extinction coefficient at depth z' . We obtain it by integrating the angular scattering coefficient $\beta(\phi, z')$ over all scattering angles ϕ and adding the absorption coefficient, if absorption exists. When we integrate Eq. (A15) over the distance z from the object, the measured transmitted light is $D = L^{\text{object}}t(z)$, where $L^{\text{object}} = D|_{z=0}$, and the atmospheric transmittance is

$$t(z) = \exp\left[-\int_0^z \beta(z')dz'\right]. \quad (\text{A16})$$

In the special case when $\beta(z') = \beta$ independently of the distance, we obtain Eq. (3).

Now let us derive the expression for airlight. Consider a layer of scattering media, of infinitesimal depth dz'' illuminated by a light source at an arbitrary direction (say, the Sun). Part of this light is scattered *toward* the camera. The radiance of the scattered light is proportional to the angular scattering coefficient $\beta(\phi, z'')$. Note⁷ that $\beta(\phi, z'') \propto \beta(z'')$. This is because light scattered toward a certain di-

rection is a fraction of the total amount of light removed from the incident beam by scattering in all directions and by absorption. The scattered light radiance is also proportional to both the illumination irradiance and to dz'' . Thus we may conclude that the ambient light scattered by this layer toward the camera is given by $\kappa\beta(z'')dz''$, where κ encapsulates the illumination irradiance and the proportion of light scattered in the direction of the camera, relative to the total scattering (and absorption). This expression also describes the case in which the layer is illuminated by a distribution of source directions.⁷

Once this light has been directed toward the camera, it undergoes attenuation on its way, as dictated by $t(z'')$. Eventually, the airlight from the above-mentioned layer is

$$dA(z'') = \kappa\beta(z'')dz'' \exp\left[-\int_0^{z''} \beta(z')dz'\right]. \quad (\text{A17})$$

The total airlight radiance (path radiance) is obtained by means of integrating the airlight contribution from all layers⁴⁹:

$$\begin{aligned} A(z) &= \int_0^z dA(z'') = -\kappa\left\{\exp\left[-\int_0^{z''} \beta(z')dz'\right]-1\right\}_0^z \\ &= \kappa[1 - t(z)]. \end{aligned} \quad (\text{A18})$$

The airlight of a scene point at infinity is

$$A_{\infty} = \kappa[1 - t(\infty)]. \quad (\text{A19})$$

Therefore the airlight is

$$A = \frac{A_{\infty}}{1 - t(\infty)}[1 - t(z)]. \quad (\text{A20})$$

In the homogeneous haze model [Eq. (3)] we have $t(\infty) = 0$, that is, object points that are far enough are completely attenuated. If we apply the assumption of total attenuation for objects at infinity also to inhomogeneous haze, then we set $t(\infty) = 0$ in Eq. (A20). We then obtain

$$A = A_{\infty}[1 - t(z)], \quad (\text{A21})$$

as in Eq. (1). This is the situation assumed throughout the paper.

It is interesting to examine how the solution is influenced when $t(\infty)$ is unknown and is not zero. Assume that we perform the recovery in the same way. First, we estimate \hat{p} and \hat{A}_{∞} by sampling the image of the sky, as in Subsection 7.A. This time,

$$\Delta I(\text{sky}) = A_{\infty}p, \quad (\text{A22})$$

while

$$I^{\text{total}}(\text{sky}) = L^{\text{object}}(\text{sky})t(\infty) + A_{\infty}. \quad (\text{A23})$$

Fortunately, we can set $L^{\text{object}}(\text{sky}) = 0$. The reason for this is that at night, when there is no airlight ($A_{\infty} = 0$), the sky is dark [$I^{\text{total}}(\text{sky}) = 0$]. Therefore,

when airlight exists, $I^{\text{total}}(\text{sky}) = A_{\infty}$. Hence we can safely use Eqs. (29) to estimate \hat{A}_{∞} and \hat{p} .

Now that we have the parameters, we can look at the scene-dehazing equation. It can be shown that the estimation of the direct transmission \hat{D} with Eq. (15) and \hat{p} is correct. Therefore we can remove the additive airlight, using the same procedure as with $t(\infty) = 0$. To complete the dehazing process, we need to compensate for the attenuation. Similar to Eqs. (15) and (17),

$$\hat{L}^{\text{object}} = \hat{D}/\hat{t}. \quad (\text{A24})$$

This time [see Eq. (A20)],

$$\hat{t} = 1 - \frac{\hat{A}}{A_{\infty}} [1 - t(\infty)], \quad (\text{A25})$$

which is somewhat different from Eq. (16). If $t(\infty)$ is unknown and cannot be neglected, then our estimation of the transmittance with Eq. (16) is biased toward a value lower than the true one. This will lead to some overamplification (brightening) of the image radiance corresponding to distant objects in the dehazed image. The overamplification of distant objects is reduced when we bias \hat{p} by a factor of $\epsilon > 1$ as in relation (31); this bias reduces \hat{A} , thereby increasing \hat{t} in Eq. (16) in the same way as the factor $[1 - t(\infty)]$ does in Eq. (A25).

Partial and preliminary results were presented in Ref. 47. We thank the reviewers for their helpful remarks. This study was supported in part by a grant from the Defence Advance Research Projects Agency (DARPA) Human Identification at a Distance program, contract N00014-00-1-0916; by National Science Foundation (NSF) award IIS-99-87979; and by the Morin Foundation.

References and Note

1. L. Grewe and R. R. Brooks, "Atmospheric attenuation reduction through multi-sensor fusion," in *Sensor Fusion: Architectures, Algorithms, and Applications II*, B. V. Dasarthy, ed., Proc. SPIE **3376**, 102–109 (1998).
2. N. S. Kopeika, *A System Engineering Approach to Imaging* (SPIE, Bellingham, Wash., 1998), pp. 446–452.
3. J. P. Oakley and B. L. Satherley, "Improving image quality in poor visibility conditions using a physical model for contrast degradation," IEEE Trans. Imag. Proc. **7**, 167–179 (1998).
4. K. Tan and J. P. Oakley, "Physics-based approach to color image enhancement in poor visibility conditions," J. Opt. Soc. Am. A **18**, 2460–2467 (2001).
5. S. G. Narasimhan and S. K. Nayar, "Chromatic framework for vision in bad weather," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 2000), Vol. I, pp. 598–605.
6. S. G. Narasimhan and S. K. Nayar, "Removing weather effects from monochrome images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 2001), Vol. II, pp. 186–193.
7. S. K. Nayar and S. G. Narasimhan, "Vision in bad weather," in *Proceedings of the IEEE International Conference on Computer Vision* (Institute of Electrical and Electronics Engineers, New York, 1999), pp. 820–827.
8. P. S. Pencikowski, "Low-cost vehicle-mounted enhanced vision system comprised of a laser illuminator and range-gated camera," in *Enhanced and Synthetic Vision*, J. G. Verly, ed., Proc. SPIE **2736**, 222–227 (1996).
9. B. T. Sweet and C. L. Tiana, "Image processing and fusion for landing guidance," in *Enhanced and Synthetic Vision*, J. G. Verly, ed., Proc. SPIE **2736**, 84–95 (1996).
10. R. C. Henry, S. Mahadev, S. Urquijo, and D. Chitwood, "Color perception through atmospheric haze," J. Opt. Soc. Am. A **17**, 831–835 (2000).
11. D. K. Lynch, "Step brightness changes of distant mountain ridges and their perception," Appl. Opt. **30**, 3508–3513 (1991).
12. S. D. Gedzelman, "Atmospheric optics in art," Appl. Opt. **30**, 3514–3522 (1991).
13. F. Cozman and E. Krotkov, "Depth from scattering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 1997), pp. 801–806.
14. W. A. Shurcliff and S. S. Ballard, *Polarized Light* (Van Nostrand, Princeton, N.J., 1964), pp. 98–103.
15. G. P. Können, *Polarized Light in Nature* (Cambridge University, Cambridge, UK, 1985), pp. 1–10, 29–54, 60–62, 131–137, 144–145.
16. B. Cairns, B. E. Carlson, A. A. Lacis, and E. E. Russell, "An analysis of ground-based polarimetric sky radiance measurements," in *Polarization: Measurement, Analysis, and Remote Sensing*, D. H. Goldstein and R. A. Chipman, eds., Proc. SPIE **3121**, 382–393 (1997).
17. K. L. Coulson, "Polarization of light in the natural environment," in *Polarization Considerations for Optical Systems II*, R. A. Chipman, ed., Proc. SPIE **1166**, 2–10 (1989).
18. S. J. Hitzfelder, G. N. Plass, and G. W. Kattawar, "Radiation in the earth's atmosphere: its radiance, polarization, and ellipticity," Appl. Opt. **15**, 2489–2500 (1976).
19. D. K. Lynch and P. Schwartz, "Rainbows and fogbows," Appl. Opt. **30**, 3415–3420 (1991).
20. M. S. Quinby-Hunt, L. L. Erskine, and A. J. Hunt, "Polarized light scattering by aerosols in the marine atmospheric boundary layer," Appl. Opt. **36**, 5168–5184 (1997).
21. M. J. Raković, G. W. Kattawar, M. Mehrübeoğlu, B. D. Cameron, L. V. Wang, S. Rastegar, and G. L. Côté, "Light back-scattering polarization patterns from turbid media: theory and experiment," Appl. Opt. **38**, 3399–3408 (1999).
22. D. B. Chenault and J. L. Pezzaniti, "Polarization imaging through scattering media," in *Polarization Analysis, Measurement, and Remote Sensing III*, D. B. Chenault, M. J. Guggin, W. G. Egan, and D. H. Goldstein, eds., Proc. SPIE **4133**, 124–133 (2000).
23. L. J. Denes, M. Gottlieb, B. Kaminsky, and P. Metes, "AOTF polarization difference imaging," in *27th AIPR Workshop: Advances in Computer-Assisted Recognition*, R. J. Mericsko, ed., Proc. SPIE **3584**, 106–115 (1998).
24. O. Emile, F. Bretenaker, and A. Le Floch, "Rotating polarization imaging in turbid media," Opt. Lett. **21**, 1706–1708 (1996).
25. X. Gan, S. P. Schilders, and Min Gu, "Image enhancement through turbid media under a microscope by use of polarization gating method," J. Opt. Soc. Am. A **16**, 2177–2184 (1999).
26. H. Horinaka, K. Hashimoto, K. Wada, T. Umeda, and Y. Cho, "Optical CT imaging in highly scattering media by extraction of photons preserving initial polarization," in *International Symposium on Polarization Analysis and Applications to Device Technology*, T. Yoshizawa and H. Yokota, eds., Proc. SPIE **2873**, 54–57 (1996).
27. M. P. Rowe, E. N. Pugh Jr., J. S. Tyo, and N. Engheta, "Polarization-difference imaging: a biologically inspired

- technique for observation through scattering media," *Opt. Lett.* **20**, 608–610 (1995).
28. J. G. Walker, P. C. Y. Chang, and K. I. Hopcraft, "Visibility depth improvement in active polarization imaging in scattering media," *Appl. Opt.* **39**, 4933–4941 (2000).
29. Y. Y. Schechner, J. Shamir, and N. Kiryati, "Polarization and statistical analysis of scenes containing a semireflector," *J. Opt. Soc. Am. A* **17**, 276–284 (2000).
30. H. Farid and E. H. Adelson, "Separating reflections from images by use of independent component analysis," *J. Opt. Soc. Am. A* **16**, 2136–2145 (1999).
31. S. K. Nayar, X. S. Fang, and T. Boulton, "Separation of reflection components using color and polarization," *Int. J. Comput. Vision* **21**, 163–186 (1997).
32. S. Rahmann and N. Canterakis, "Reconstruction of specular surfaces using polarization imaging," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 2001), Vol. 1, pp. 149–155.
33. M. Saito, Y. Sato, K. Ikeuchi, and H. Kashiwagi, "Measurement of surface orientations of transparent objects by use of polarization in highlight," *J. Opt. Soc. Am. A* **16**, 2286–2293 (1999).
34. L. B. Wolff, "Polarization vision: a new sensory approach to image understanding," *Image Vision Comput.* **15**, 81–93 (1997).
35. C. F. Bohren and A. B. Fraser, "At what altitude does the horizon cease to be visible?" *Am. J. Phys.* **54**, 222–227 (1986).
36. E. J. McCartney, *Optics of the Atmosphere: Scattering by Molecules and Particles* (Wiley, New York, 1976).
37. J. S. Tyo, M. P. Rowe, E. N. Pugh Jr., and N. Engheta, "Target detection in optically scattering media by polarization-difference imaging," *Appl. Opt.* **35**, 1855–1870 (1996).
38. R. L. Lee Jr., "Digital imaging of clear-sky polarization," *Appl. Opt.* **37**, 1465–1476 (1998).
39. E. Hecht, *Optics*, 3rd ed. (Addison-Wesley, New York, 1998), pp. 340–342.
40. S. Chandrasekhar, *Radiative Transfer* (Dover, New York, 1960), pp. 24–37, 280–284.
41. M. Ben-Ezra, "Segmentation with invisible keying signal," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 2000), Vol. 1, pp. 32–37.
42. T. Prosch, D. Hennings, and E. Raschke, "Video polarimetry: a new imaging technique in atmospheric science," *Appl. Opt.* **22**, 1360–1363 (1983).
43. A. M. Shutov, "Videopolarimeters," *Sov. J. Opt. Technol.* **60**, 295–301 (1993).
44. L. B. Wolff, "Polarization camera for computer vision with a beam splitter," *J. Opt. Soc. Am. A* **11**, 2935–2945 (1994).
45. A. S. Glassner, *Principles of Digital Image Synthesis* (Morgan Kaufmann, San Francisco, Calif., 1995), Appen. G.4.
46. R. L. Lee Jr., "Horizon brightness revisited: measurements and a model of clear-sky radiances," *Appl. Opt.* **33**, 4620–4628 (1994).
47. Y. Y. Schechner, S. G. Narasimhan, and S. K. Nayar, "Instant dehazing of images using polarization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 2001), Vol. 1, pp. 325–332.
48. J. E. Solomon, "Polarization imaging," *Appl. Opt.* **20**, 1537–1544 (1981).
49. For the calculation of the path radiance integral, we assume κ to be distance invariant. This is because typically most of the light in the scene comes from the Sun and sky and thus does not change much along the line of sight. Moreover, we assume that multiple scattering (which effects the angular scattering distribution) is dominated by single scattering.