

Recreating the Past

Duncan Brown Alan Chalmers Paul Debevec
Kate Devlin Philippe Martinez Greg Ward

Course #27, SIGGRAPH 2002
San Antonio, Texas, USA.
21 – 26 July 2002

Abstract

Recent developments in computer graphics and interactive techniques are providing powerful tools for modelling multi-dimensional aspects of data gathered by archaeologists. This course addresses the problems associated with reconstructing archaeological and heritage sites on computer and evaluating the realism of the resultant models. The crucial question considered is: are the results misleading and thus are we in fact misinterpreting the past?

We will never know precisely what was in the mind of our ancestors as they painted rock shelters in France 25 thousand years ago, or raised the pyramids in Egypt, or even purchased a particular brightly coloured pot during the Middle Ages. Recently archaeologists have been increasingly turning to computer graphics and interactive techniques to help interpret material preserved from ancient cultures. This course describes currently used state-of-the-art techniques for reconstructing archaeological sites and addresses the issues that still need to be resolved so that these techniques can indeed play a significant role in helping us understand the past.

The attendees should have an interest in “understanding the past” and a basic knowledge of computer graphics. No prior knowledge of laserscanning, lighting simulation or visual perception evaluation is assumed, although any knowledge will be an advantage.

The course covers the following topics: creating models of the sites, including laserscanning; very realistic lighting simulation; quantifying the realism of the results using human visual perception and psychophysical methods; valid interpretation of the results by the archaeologists and general public.

All topics are illustrated by case studies.

Course Schedule

Module 1 - Creating the Past

- 1.30** Introduction to Recreating the Past - Chalmers
 - some intuitive examples of applications
 - role of realism
 - our focus: understanding the past using computer graphics methods
- 1.50** Creating the Models - Debevec & Martinez
 - using all the evidence
 - laser scanning
 - recreating color and textures
- 2.40** Very Realistic Lighting Simulation - Ward & Brown
 - experimental archaeology to recreate the ancient fuel types
 - accurate simulation of light propagation
 - tone mapping operators to achieve meaningful results
- 3.15** Break

Module 2 - Interpreting the Past

- 3.30** Quantifying Realism - Chalmers
 - psychophysics: procedures for comparing real and synthetic images
 - fidelity assessment
 - case studies
- 4.00** Interpretation of the Models - Brown & Devlin
 - displaying the information
 - setting standards
 - interpreting the results
 - avoiding misinterpretation
 - developing new hypotheses
- 4.40** Conclusion & Summary - Chalmers & Martinez
- 5.00** Discussion and questions - All

About the authors



Alan Chalmers is a Reader in the Department of Computer Science at the University of Bristol, UK. He has published over 80 papers in journals and international conferences on very realistic graphics. He is currently Vice President of ACM SIGGRAPH. His research is investigating the use of very realistic graphics in the accurate visualisation of archaeological site reconstructions and techniques which may be used to reduced computation times without affecting the perceptual quality of the images.

Dept. of Computer Science, University of Bristol, Woodland Road, Bristol BS8 1UB, United Kingdom.

Email: Alan.Chalmers@bris.ac.uk

URL: <http://www.cs.bris.ac.uk/~alan/>



Kate Devlin is a research member of the Graphics Group in the Department of Computer Science at the University of Bristol, UK. Her undergraduate degree was in archaeology and she worked as a field archaeologist and site draughtsman for two years before studying for an MSc in Computer Science. Her research interests are realistic archaeological reconstructions of flame-lit environments and accurate display of high dynamic range scenes. Her personal interests lie in representation and interpretation of archaeological records.

Dept. of Computer Science, University of Bristol, Woodland Road, Bristol BS8 1UB, United Kingdom

Email: devlin@cs.bris.ac.uk

URL: <http://www.cs.bris.ac.uk/~devlin/>



Duncan Brown is the Curator of Archaeological Collections at Southampton City Heritage, UK; a visiting Research Fellow at the University of Southampton, UK; and a freelance medieval pottery specialist who has spent several years digging in Bucks, Northants, Wales, Sussex, Cheshire. He has worked on assemblages from Fyfield Down, Winchester, the Tower of London, Reading, the Mary Rose, Guernsey, Isle of Wight and elsewhere. Duncan specialises in pottery imported from mainland Europe and has taught ceramic analysis at Bournemouth University and archaeological theory at King Alfred's College. Duncan has been co-editor of the newsletter of the Society for Medieval Archaeology, Secretary of the Medieval Pottery Research Group, founder committee member and subsequently chair of the IFA Finds Group.

City Heritage Services, Southampton City Council, Southampton, United Kingdom.

Email: duncan@potboiler.fsnet.co.uk



Paul Debevec received his Ph.D. from UC Berkeley in 1996 and leads a computer graphics research group at the University of Southern California's Institute for Creative Technologies. His computer graphics research has concerned reconstructing real world environments from photographs, augmenting such models with computer-generated objects and humans, and creating animations using image-based rendering and global illumination. His films and art installations have featured photo-real reconstructions of the Rouen Cathedral, the UC Berkeley Campus, and St. Peter's Basilica and have contributed to the visual effects techniques seen in "The Matrix" and "X-Men". Debevec is a member of ACM SIGGRAPH, the Visual Effects Society, and the recipient of SIGGRAPH's 2001 Significant New Researcher Award.

Executive Producer, Graphics Research, USC Institute for Creative Technologies, 13274 Fiji Way, 5th Floor, Marina del Rey, CA 90292, USA.

Email: paul@debevec.org

URL: <http://www.debevec.org/>



Philippe Martinez has a Ph.D. in Egyptology and a research degree from Ecole du Louvre. After completing his studies he joined the staff of the Franco-Egyptian Center in Karnak and worked there as research assistant until 1992. During this period, he was in charge of on site documentation, specializing in the epigraphical survey and study of the dismantled limestone monuments of the Middle Kingdom and beginning of the 18th dynasty. In 1989, Philippe Martinez joined a team of computer engineers at France's Electricity R{D department, working on a 3D reconstruction of the ancient Egyptian temples of Karnak and Luxor. Philippe Martinez' major contribution was in the use of 3D modelling in archaeological hypothesis testing. Since 1999, Philippe Martinez has been a member of the ECHO Project (Egyptian Cultural Heritage Operation) as lead archaeologist. Using 3D reconstruction in the field of monumental archaeology, he is trying to work on the very idea of digital epigraphy in the hope of applying fast, reliable, robust and cost effective documentation techniques to the thousands of monuments actually disappearing in Egypt as well as in the rest of the world. His current mission takes him around the Mediterranean, in Italy, Tunisia and Egypt.

École Normale Supérieure, Paris, France.

Email: pmartine@ens.fr



Greg Ward (a.k.a. Greg Ward Larson) graduated in Physics from UC Berkeley in 1983 and earned a Masters in Computer Science from SF State University in 1985. Since 1985, he has worked in the field of light measurement, simulation, and rendering variously at the Berkeley National Lab, EPFL Switzerland, Silicon Graphics Inc., Shutterfly, and Exponent. He is author of the widely used Radiance package for lighting simulation and rendering.

1200 Dartmouth St., #C, Albany, CA 94706, USA.

Email: gward@lmi.net

Contents

1	Introduction	3
1.1	Cap Blanc	4
2	Creating the models	9
2.1	3D scanning in archaeological perspective	9
2.2	Affordable 3D scanning for archaeologists: the possibilities of structured light	12
2.3	Slides	15
2.4	A photometric approach to digitizing cultural artifacts: slides . . .	34
3	Very Realistic Lighting Simulation	51
3.1	Slides	52
4	Quantifying Realism	69
4.1	Luminaires	70
4.1.1	Creating a Realistic Flame	70
4.1.2	Converting Pixel Information to Radiance Files	75
4.1.3	Creation of a Radiance Scene	76
4.1.4	Creation of Final Images	77
4.1.5	Automation	78
4.2	Implementation	78

4.2.1	Flame to Film to Frames	78
4.2.2	Creating Sphere Data	80
4.2.3	Superimposing the Flame	87
4.3	Converting Luminaire Data	89
4.4	Validating Realism	90
5	Representation and Interpretation	95
5.1	Introduction	95
5.2	A brief history of archaeological illustration	96
5.2.1	Archaeological illustration: an overview	96
5.2.2	Case study: seeing Stonehenge	99
5.3	The idea of realism	101
5.3.1	Terms and concepts	103
5.3.2	The nature of archaeological data	106
5.3.3	Context	108
5.3.4	An established reality	109
5.4	Representing for a purpose	110
5.4.1	Representations for the archaeologist	110
5.4.2	Representations for the computer scientist	111
5.4.3	Representation as advertising	111
5.4.4	Representations for the public	112
5.4.5	Fit for purpose	113
5.5	Misinterpretation	113
5.5.1	Different outcomes from the same evidence	113
5.5.2	Seeing what we want to see	114
5.5.3	Reducing misinterpretation	115
5.6	Setting standards	115

<i>CONTENTS</i>	1
5.6.1 Metadata	115
5.6.2 Alternative representations	117
5.6.3 Preserving information	120
5.6.4 Standardisation	121
5.7 Developing new hypotheses	121
5.7.1 New ideas from light and colour perception	122
5.7.2 Case study: Medieval pottery	123
5.8 Summary	131
5.9 Slides	133
Bibliography	144
A Included papers	151

Chapter 1

Introduction

by Alan Chalmers, University of Bristol, UK.

Recent developments in computer graphics and interactive techniques are providing powerful tools for modelling multi-dimensional aspects of data gathered by archaeologists. This course address the problems associated with reconstructing archaeological and heritage sites on computer and evaluating the realism of the resultant models. The crucial question considered is: are the results misleading and thus are we in fact misinterpreting the past.

Archaeology provides an excellent opportunity for computer graphics to explore scientific methods and problems at a human scale and to introduce a cultural dimension which opens up avenues for new and alternative interpretations. For example, pottery is a very common find at the excavation of medieval sites. Archaeologists have been studying this material for many years and its means of manufacture, distribution and use are very well understood. With this sound basis provided by years of data acquisition, we can now begin to investigate less easily comprehended aspects of pottery. Medieval pottery was often very colourful, wonderfully texture and vibrantly decorated. Was this necessary because of the lighting conditions which prevailed in medieval society, or were the pots perfectly visible and the people simply wanted some colour in their otherwise dull lives? Computer graphics offers the possibility of exploring these questions.

In order to minimise misinterpretation of the archaeological evidence, the question is whether computer visualisation of archaeological sites should be very realistic including the accurate modelling of the prevalent illumination, the complex 3D

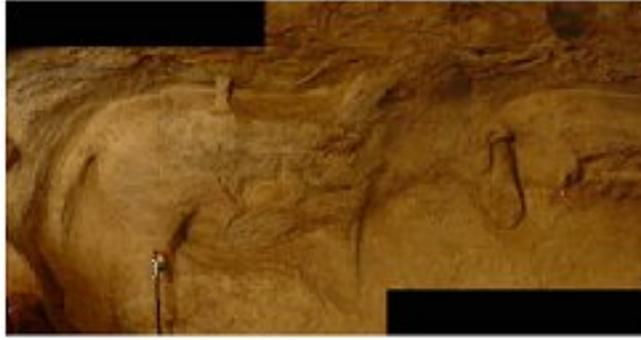


Figure 1.1: Part of the frieze from Cap Blanc

environment, and atmospheric factors such as smoke and dust. We will never know exactly why medieval pottery was so brightly coloured, but perhaps computer graphics and interactive techniques can establish a framework of possibilities within which archaeologists can hypothesise as to probable solutions.

1.1 Cap Blanc

As a way of illustrating the potential computer graphics has to offer archaeology we consider the prehistoric site of Cap Blanc. The rock shelter site of Cap Blanc, overlooking the Beaune valley in the Dordogne, contains perhaps the most dramatic and impressive example of Upper Palaeolithic haut relief carving. A frieze of horses, bison and deer, some overlain on other images, was carved some 25,000 years ago into the limestone as deeply as 45cms, covers 13m of the wall of the shelter. Since its discovery in 1909 by Raymond Peyrille several descriptions, sketches, and surveys of the frieze have been published, but they appear to be variable in their detail and accuracy .

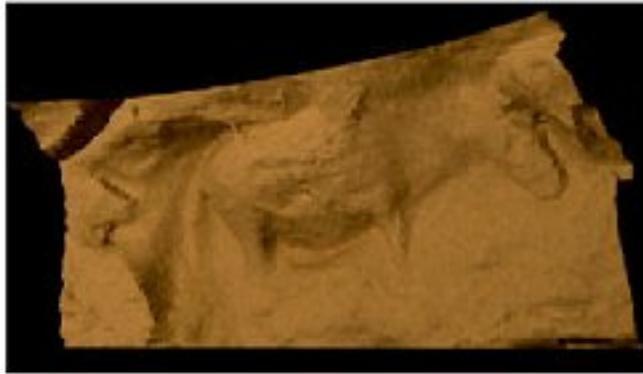
In 1999, a laser scan of was taken of part of the frieze (Figure 1.1) at 20mm precision [9]. It was obviously of utmost importance that an eye safe laser was used to ensure there was no damage at all to the site.

Figure 1.1 shows part of the frieze from Cap Blanc Some 55,000 points were obtained in two scans of the upper and lower part of the selected area, Figure 1.1 (the MDL laser scanner used did not have sufficient memory to store all the points in a single scan). These points were stitched together and converted into a triangular

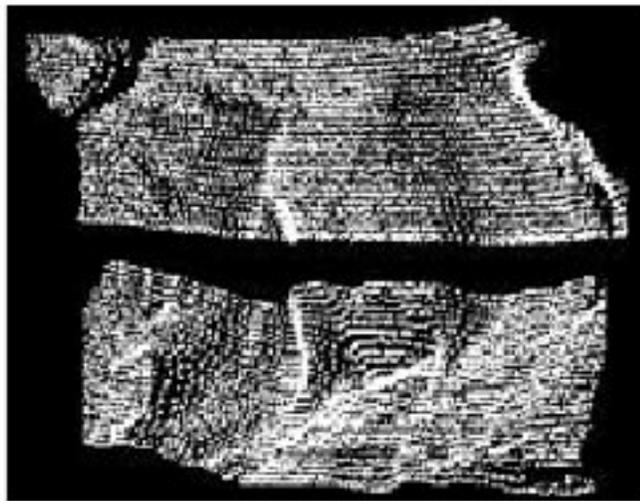
mesh.

Detailed photographs of the frieze were taken, each one of which included a standard rock art colour chart, just visible in Figure 1.1. As the exact spectral data for the colour chart is known, this enabled us to compensate for the lighting in the photograph and thus obtain approximate illumination-free textures to include with the wire-frame model. Images were rendered using Radiance. Figure 1.1(a) shows the horse illuminated by a simulated 55W incandescent bulb (as in a low-power floodlight), which is how visitors view the actual site today. In Figure 1.1(b) the horse is now illuminated by an animal fat tallow candle as it may have been viewed 25,000 years ago. As can be seen the difference between the two images is significant with the candle illumination giving a ‘warmer glow’ to the scene as well as more shadows.

This shows that it is important for archaeologists to view such art work under (simulated) original conditions rather than under modern lighting. (It is of course impossible to investigate these sensitive sites with real flame sources). For this site, we wanted to investigate whether the dynamic nature of flame, coupled with the careful use of three-dimensional structure, may have been used by our prehistoric ancestors to create animations in the cave art sites of France, 25,000 years ago. The shadows created by the moving flame do indeed appear to give the horse motion. We will of course never know for certain whether the artists of the Upper Palaeolithic were in fact creating animations 25,000 years ago, however the reconstructions do show that it certainly is possible. There is other intriguing evidence to support this hypothesis. As can be seen in the figures, the legs of the horse are not present in any detail. This has long been believed by archaeologists to be due to erosion, but if this is the case, why is the rest of the horse not equally eroded? Could it be that the legs were deliberately not carved in any detail to accentuate any motion, that is to create some form of motion blur? Furthermore traces of red ochre have been found on the carvings. It is interesting to speculate again whether the application of red ochre at key points on the horse’s anatomy may also have been used to enhance any motion effects.



(a) Clouds of points from the scan

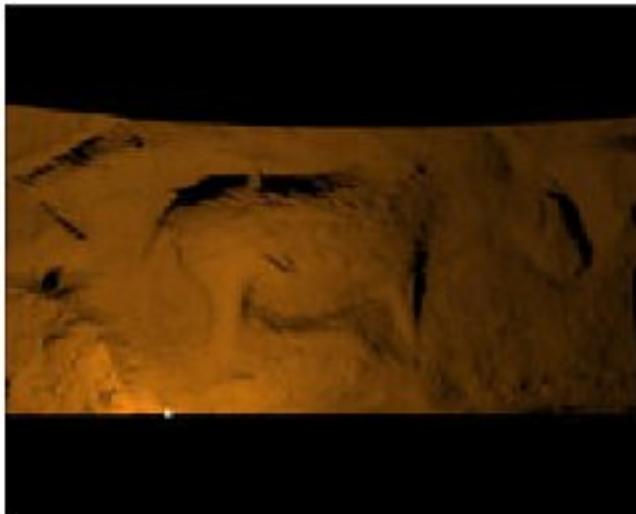


(b) The reconstructed horse

Figure 1.2: Cap Blanc horse: from point cloud to reconstruction



(a) 55W incandescent bulb



(b) animal fat candle

Figure 1.3: Lighting the reconstructed frieze

Chapter 2

Creating the models

by **Paul Debevec**, USC Institute for Creative Technologies, and **Philippe Martinez**, École Normale Supérieure.

2.1 3D scanning in archaeological perspective

For most people, archaeology means digging up things and putting them on a shelf in some kind of museum, putting back together the so-called material culture of a civilisation. But for the professional archaeologist or art historian, the main part of this process is spent on what we shall call bluntly ‘documentation’.

Very generally and simply put, it means the faithful and complete portraiture through any available mean and medium of the discovered artefacts in connection with their context of discovery. In most of the cases, this has been done most classically through verbal descriptions and drawings. The end of the 19th century brought a more consistently faithful tool with the advent of photography. However, the time as well as the equipment needed to pursue the goal of complete and accurate documentation of the ancient remains have contributed to raising its cost and lowering its effectiveness. Where actual publications, which are most of the time the only form of documentation available, would need to present the whole global and different aspects of the discoveries, they still tend to render a very verbal, and sometimes verbose if not partial, account of the archaeological context or monuments they concern.

Digital tools appeared during the last two decades tend slowly to make those difficulties disappear. Thus the total stations used for topographical surveys now enable us to get a fair idea of the setting of huge settlements. Databases tend to make the difficult art of dealing with thousands of very varied types of data at least possible, while CAD tools enable archaeologists who are not professional draughtsmen to create drawings that are sometimes a lot better than the simple sketches we were unfortunately tolerating with the slow disappearance of dedicated skills and staff. Very simple and effective technologies like QTVRs enable us also to get very realistic and somewhat complete visual descriptions of objects or settlements. However all these possibilities may turn also into the source of new problems when the same staff member tries to cover all the different needs of documentation and publication, while never taking — or simply never having — the time to learn the different skills that lie behind the interface of the software they use.

However, the appearance during the last decade of new 3D tools now enables us to at least try to get towards something like the global documentation archaeologists have been dreaming of over the last two centuries, regretting the incomplete documents left by their predecessors. Apart from the 3D modelling packages that enable an artistic or technical representation of the object or architecture or terrain studied, the recent availability of 3D scanning techniques tends to make it possible to try to obtain a very accurate and faithful representation of the archaeological remains, with the possibility to work on a model that is as close as possible to reality. These techniques enable us to take virtual casts of the objects without having to physically touch or risk harming them, while also avoiding a human intellectual reinterpretation interfering with the final portraiture of reality (even if this human interference cannot seem to be ever totally avoided during the complex processing of the original very accurate set of physically captured data).

The available techniques tend to diversify greatly over the years. However, the cost, fragility and difficulty of use of most of these technologies have also made them almost impossible to use directly on site by ‘impoverished’ archaeological research institutions and professionals. When they should be used on a daily basis to salvage information that is destroyed everyday by trowels and shovels, they are unfortunately reserved for luckier projects that try to lure sponsors or media representatives by electing a very well known and not really threatened masterpiece as their centre of interest. However, these technologies being clearly in the process of development, these facts might as well look like a small price to pay for their availability, growing robustness and cost effectiveness.

Different aspects of the studied subjects are to be taken into account: size, weight, material, colour, settings, artistic quality, fragility, reflective or non-reflective qualities of the material etc... These intrinsic qualities have to be listed to be able to choose the technology to best fit the aims and needs. If we examine the different technologies, it is clear that none of them can be seen as universal and that anyone envisaging the use of this approach has to deal with a pretty complex toolbox. The same thing can be said of the methodology to be employed, once one or more technologies have been singled out. There is no way to define a methodology that could be applied blindly and thoroughly. Each project will in fact bring the definition of a specific strategy according to the target and to the type of final representation that would seem to be the desired aim of the project.

To go one step further, the very idea of trying to use 3D scanning tools on cultural heritage and archaeological subjects has to be defined according to goals which are clearly complementary but can finally end up being almost contradictory during the actual performance of the project. A very big issue is the quality of the measurements gathered, and thus from them the quality of the final model. The archaeologist or art historian will naturally tend to ask for a very precise virtual cast with as high a resolution as possible. This stage has to be considered as the core of documentation and the file thus gathered could be seen as a digital archive that could enable its curators to replace the object if this one would end up being damaged or destroyed. But a high resolution scan turns out to be a painful thing to use once you try to turn into any kind of modelling and rendering. Thus the graphic specialists who take part in such projects would rather gather less-rich 3D scans that they could use simply as templates for a precise but somewhat reinterpreted representation. However it should be stressed from the outset that it seems always more important to try to get the best virtual cast as possible as an archive, to save as much of the available information as possible and to choose to resample the data to be able to use it for modelling.

Apart from this idea of cultural heritage neutral archiving, one has also to think of the perception we have of these objects to understand and take into account the need of precise 3D capture, modelling and rendering of cultural objects. Most of those are linked to some kind of artistic and aesthetic aspect that make it unique and result in a complex mixture of feeling in its perception by the human mind. To go through a very straightforward comparison, human body and faces have been the nightmare of computer graphics for years now, the problem remaining almost unsolved today. The reason for this does not lie in technical inabilities but maybe in our perception that is centered around what a human being's essence is — our

brain reacting very suspiciously to any ‘unnatural’ or not realistically true representation of a human being. It is striking to note that, on the contrary, phoney or caricatured, humanised animals are easily accepted by our mind. We can possibly draw the same line between any industrial object that does not have any specifics that make it unique, and a cultural artefact that is most of the time the result of a very specific craftsmanship (if not artistry) that comes to us modelled by human use and the passing of time. Capturing the soul and essence of these objects to be able to transmit it to the largest possible audience, away from the object itself, is the challenge facing anyone willing to use digital tools to document, understand and explain ancient cultural artefacts and monuments. The problematic subject of realistic representation is discussed in more detail in section 5.3.

2.2 Affordable 3D scanning for archaeologists: the possibilities of structured light

Among the existing technologies, the use of structured light might very well be the most promising, as it seems to be one that takes into accounts most of the real needs of the archaeologist.

This technology is based mainly on a optical approach that can be simplified as the projection on an object of one or more geometrical patterns or grids whose topology is very specifically known. The capture consist of digital images showing either the grid as projected on the object or/and an image of the object itself taken from the very same angle that shall be used as a texture. The grid projected on the object is of course going to be deformed by the shape of the object. It is this deformation that is in fact complexly processed by a software that extracts x , y , z , measurements for points of the projected grid or grids.

Different softwares are already available on the market today but research is still taking form. At present, the limitations of this technology reside mainly in the need for powerful projection devices that would enable the digital camera to capture a fine but dense grid on any kind of material and in the difficult, sometimes very bright, lighting conditions that would be encountered on site. All this tends to make structured light a technology best suited for small or medium size capture. However, it enables the processing of dense and precise measurements. The quality of the optics and the resolution of the camera are of course of the greatest importance, but the progress made in these aspects during the last few years make

2.2. AFFORDABLE 3D SCANNING FOR ARCHAEOLOGISTS: THE POSSIBILITIES OF STRUCTURED

it possible to think that the necessary parameters will be available in affordable quality cameras in the coming years. But aside from that, the required apparatus can be put together mainly from off-the-shelf tools.

To be able to deal with a portable, flexible device, it has to be constructed from different parts that need to be put back together for every use, and whose precise relationship is a condition to the quality of the capture. The capture itself can be done with either a still digital camera or a high resolution video camera. The choice is open but it also depends from the idea of using one or many different grids. Those can be projected either through the use of glass slides or through computer generated and displayed graphic files. In this case, it is interesting to use a digital projector that is directly linked to a laptop computer equipped with a powerful graphics card that will simultaneously control the display and of the capture.

These two devices, projector and camera, have to be mounted together on a stand (which could be any professional, stable, photographic tripod). However, the need to take the whole device onto the site makes it very difficult to ensure that the spatial relationship of projector and camera will always be the same, even approximately — something that is necessary to the actual processing of the data and a condition to the quality of the results. From this fact comes the need to predefine the neutral result of this spatial relationship through the use of calibration files that capture the image of the grid or grids on a known object (which can be either a plan or a box which is itself "configured" through a regular and known pattern applied to it). It has to be stressed that the greatest care has to be taken with this reference object, its deformation inevitably leading to flaws in the quality and homogeneity of the measurements captured. A calibration file has thus to be processed every time the spatial relationship between camera and projector is altered. However, if the setting is good enough to accommodate a whole set of objects or different parts of the same big object, then hundreds of captures can be made with the same fixed setting, with the possibility of having the whole rig on wheels to move along or around the studied objects. In this case, there is a need to keep the eye of the camera at the same distance and angle of the object to keep a reasonable homogeneity and quality to the captured data. Very simple devices can be used to this effect, like markings on the ground showing the different predefined positions of the camera, and a track enabling the operating staff to check the relative spatial position of the camera-projector rig compared to the surface of the object, using, for example, laser pointers to keep the alignment to the predefined track.

These precautions being taken, the capture device can be used to capture the object from different, complementary angles. The technology being optically based, it has to be kept in mind that we are in fact taking 3D photographs of complex objects. It is thus necessary to take multiple views of the same object to capture the complexity of its topology: either going around it, (having it turning in front of the camera if we are dealing with small portable artefacts), or panning in front of it from different directions and angles to cover as much as possible of its geometric specifics. This then leads to the difficult challenge to put these different views of the object together, once the original measured data is processed. This registration stage can be solved either by the use of specific targets present in the scene or on the object. However, most of the existing softwares enable us also to use a simple, robust and automatic fitting of the entities through common surfaces that are defined topologically or through the precise "mapping" of the pixels of the captured image. All this make it very important to have as much common ground between the captured image as possible to ensure a steady workflow and precise results during the post-processing of the data.

In comparison to the data gathered though metrically based 3D measurement and scanning that generate only x , y , z , text files, archiving is an issue to be taken seriously into account as the original raw data consists of a large number of high resolution image data, but the actual progress in the domain of storage makes it possible to deal with, with a certain confidence. It is also to be kept in mind that the images thus gathered are precious 'portraits' of the object which, used as texture maps, result in the elaboration of a very rich and natural rendering of complex artefacts.

We can thus describe this technique as very straightforward, cost and time effective once specific methodological precautions are applied during the original capturing stage of the work. Its portability, flexibility and robustness makes it a precious tool for those who are dealing with fragile objects or who want to capture as much of the information as possible that gets destroyed daily during excavations.

2.3 Slides

Creating the Models

SAN FRANCISCO
SIGGRAPH
2002



Paul Debevec
USC Institute for
Creative Technologies

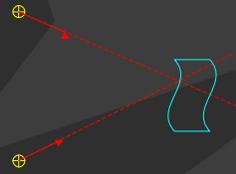
Philippe Martinez
Ecole Normale
Superieure

3D Scanning Overview

- Scanning determines 3d location of surface points
- Two general methods:
 - Triangulation
 - Time of Flight

Triangulation

- Two known positions
- Find direction to a point
- Intersect line segments



- **Human depth perception**
- **Measures direction**
- **Occlusion problems**

Time of Flight

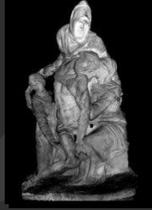
- **Send out a signal pulse**
- **Measure time for reflection to return**
- **Measures time**
- **No occlusion**



Recent 3D Scanning Projects



Electricite de France



IBM Watson



Stanford University



CNR Pisa

Low-cost, high speed, sculpture scanner – structured light approach



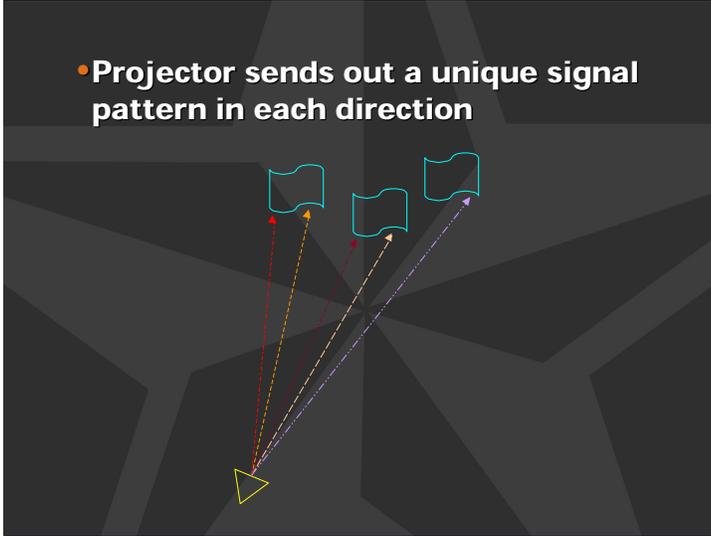
DLP Projector
(1024x768)



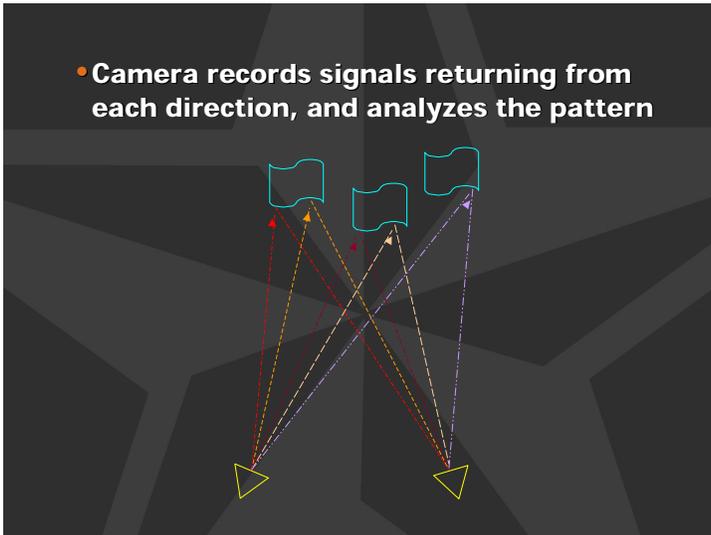
Digital Video
Camera (1024x1024)



- Projector sends out a unique signal pattern in each direction



- Camera records signals returning from each direction, and analyzes the pattern



Advantages/Disadvantages

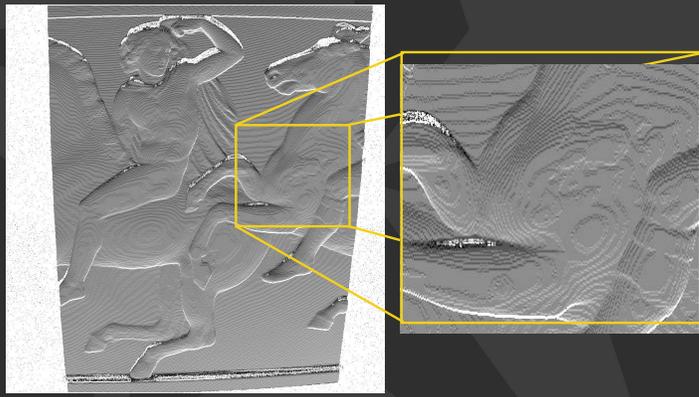
• Advantages

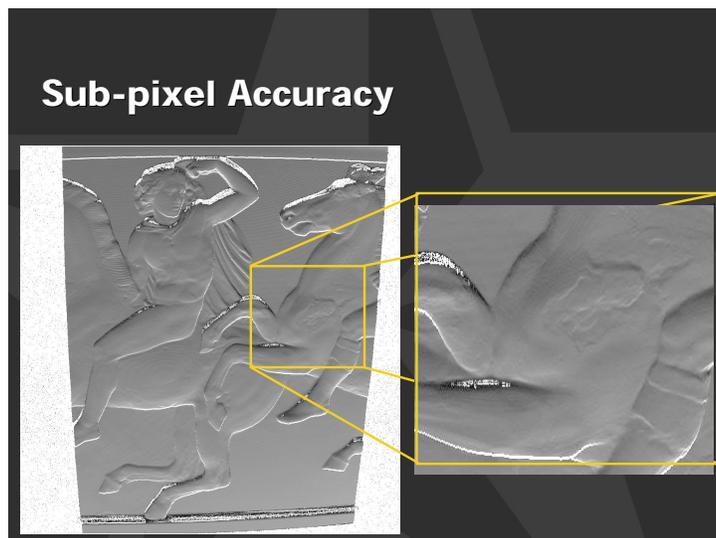
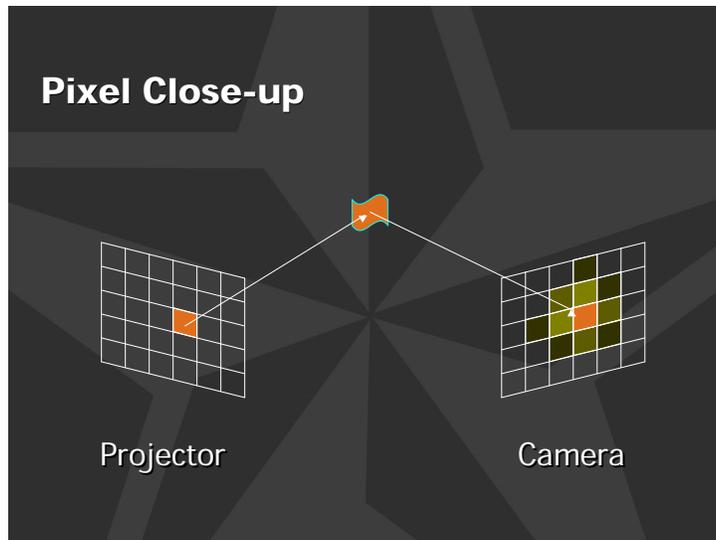
- Cheap (several thousand for the hardware)
- Fast (~15 seconds / 1,000,000 points)
- Accurate
- Can get textures at the same time

• Disadvantages

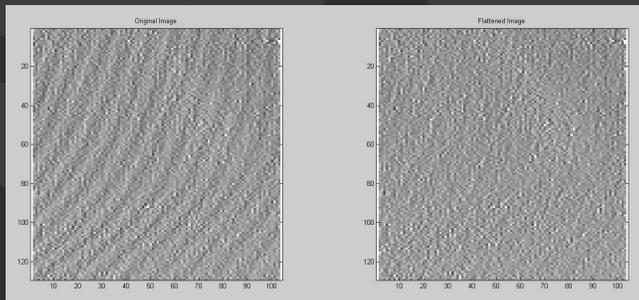
- Hard to scan large objects (more than 8 ft)
- Or in bright light
- Calibration is not easy

Pixel Accuracy





Sub-pixel Curve Modeling

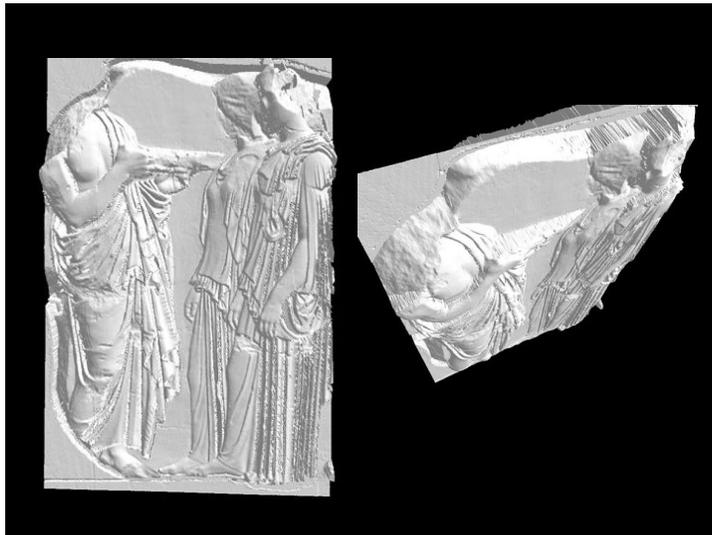


Chris Tchou Master's Thesis 2002

3D Scanning Parthenon Sculptures

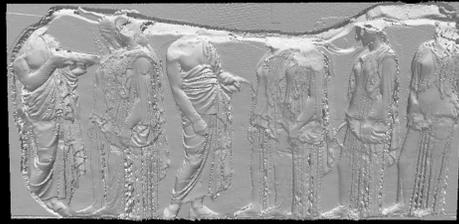


MODELS



Additional Evidence:

Drawings can provide additional lost information. How can this be incorporated?



3D Scan, 2001



Carrey Drawing, 1674

Scanning Casts:

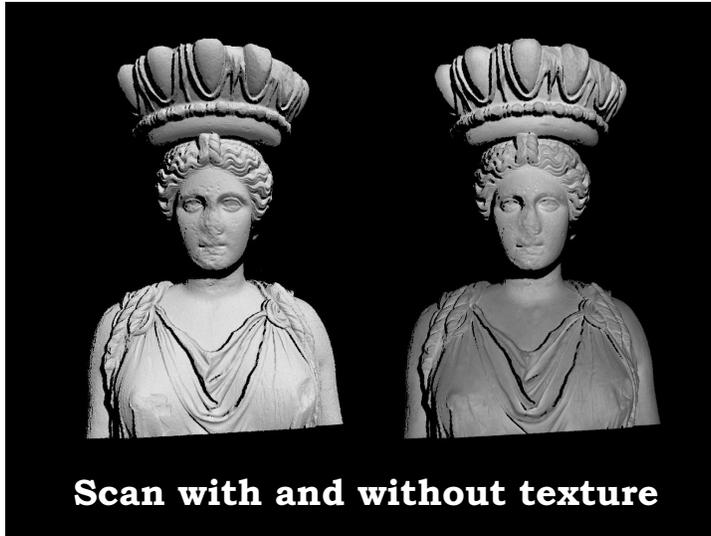
Sometimes in better condition than originals



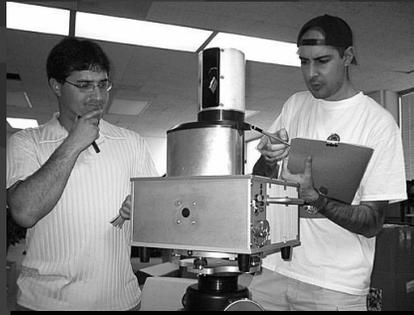
original

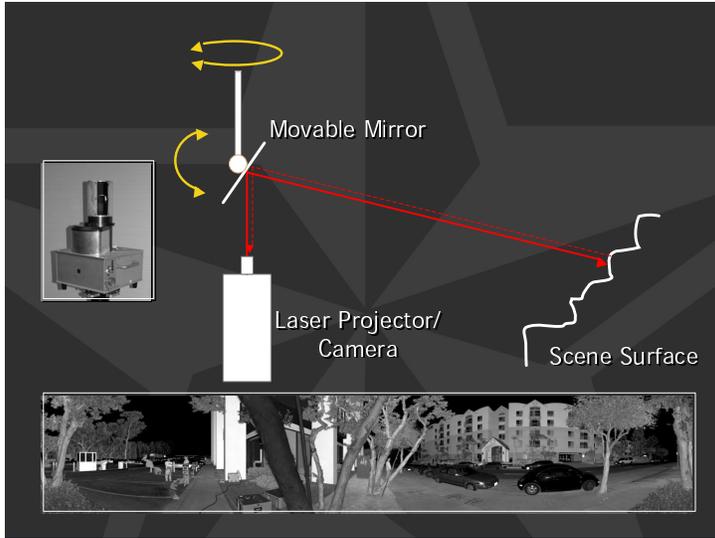


scan of cast



Scanning Environments



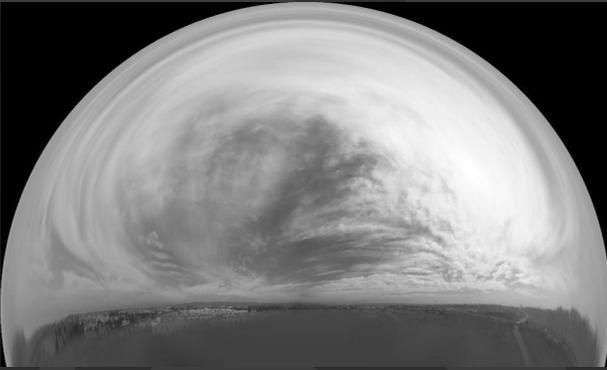


**Rendering Archaeological
Models with Global
Illumination and Image-Based
Lighting**

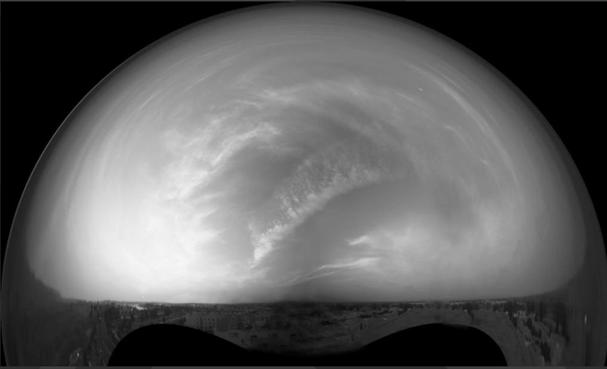
**Acquiring Real-World
Illumination**



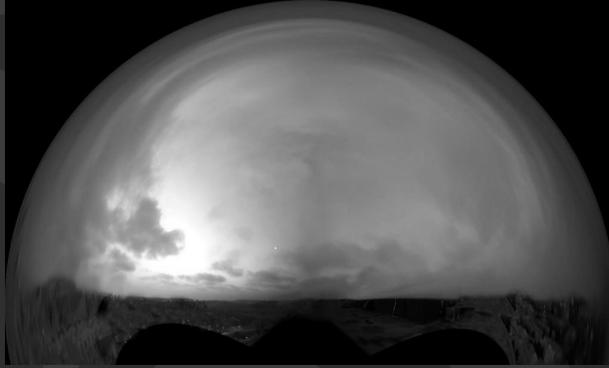
Outdoor Light Probes



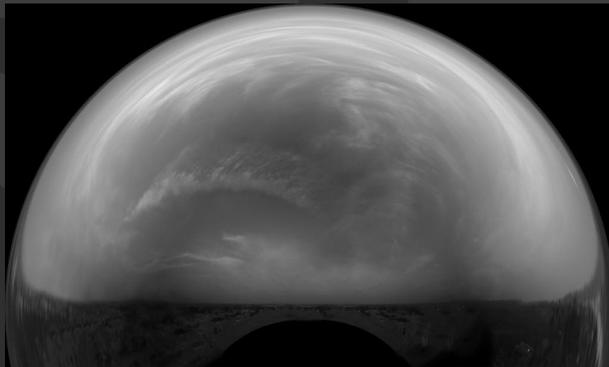
Outdoor Light Probes



Outdoor Light Probes



Outdoor Light Probes





Untextured Model rendered with real-world illumination



Lighting Concept Drawings by Mark Brownlow

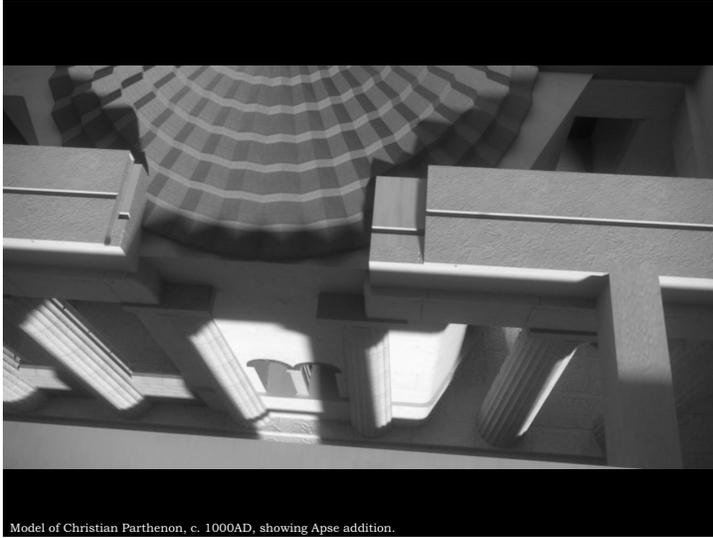
MODELS



Computer model of Parthenon, c. 1830, illuminated with image-based lighting, Arnold global illumination, depth of field



Model of contemporary Parthenon, illuminated by evening light of Marina del Rey, CA



Model of Christian Parthenon, c. 1000AD, showing Apse addition.



Computer model of the Duveen Gallery in the British Museum, site of many of the Parthenon sculptures.

MODELS



Rendering of a computer scan of a cast of West Panel II of the Parthenon frieze in the Basel Skulpturhalle.



Rendering of a computer scan of the head of a Caryatid cast scanned in the Basel Skulpturhalle.

Modeling and Animation

Brian Emerson
 Craig "X-Ray" Halperin
 Mark Brownlow
 Yikuong Chen
 Diane Suzuki
 Hiroyuki Matsuguma
 Jamie Waese
 Rippling Tsou
 Shivani Khanna
 Patrick Lee

Arnold Rendering Software

Marcos Fajardo

HDR Image Processing

Chris Tchou

Archaeological Consultant

Philippe Martinez

Sculpture Scanning

Chris Tchou
 Tim Hawkins
 Paul Debevec
 Philippe Martinez

Scanning Hardware

Tim Hawkins
 Chris Tchou
 Paul Debevec

Scanning Software

Chris Tchou
 Jonathan Cohen
 Fred Pighin

Video Editing

Paul Asplund

3D Scanning made possible by Tomas Lochman of the Basel Skulpturhalle, Jean-Luc Martinez of the Musee du Louvre, and with the support of TOPPAN Printing Co., Ltd.

2.4 A photometric approach to digitizing cultural artifacts: slides

ch to digitizing cultural

A Photometric Approach to Digitizing Cultural Artifacts

Paul Debevec
USC Institute for Creative Technologies



Tim Hawkins, Jonathan Cohen, and Paul Debevec. 2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage (VAST 2001). Glyfada, Greece, November 2001

 www.debevec.org/Research/LS2 

Current 3D Scanning Process

- Acquire Laser or Structured Light Scans from Different Angles
- Merge Scans into complete model
- Register photographs to determine object appearance or reflectance
- Project and merge photographs onto geometry



Deriving Object Appearance

- **Approach 1: Project photos onto object as texture maps**
 - Façade 96, Pulli 97, Miller98, Wood2000
 - Lighting is fixed
- **Approach 2: Perform reflectometry to estimate diffuse and specular components for each surface point**
 - Sato 97, Rushmeier 98, Marschner 99, Yu 99, Levoy 2000
 - Still not general

SIGGRAPH
2002

Deriving Textures

- **Approach 2: Derive lighting-independent textures**
 - Photograph surfaces under known illumination conditions
 - Estimate diffuse albedo and specular components for each surface
 - Disadvantages
 - difficult to control illumination
 - complex reflectance models require many observations
 - Models do not extend to translucency, subsurface effects

SIGGRAPH
2002

Strengths of current 3D Scanning

- **Works well for plaster and marble sculptures**
 - Well-defined geometric surfaces
 - Mostly diffuse reflectance
 - Limited mutual illumination
 - Limited self-occlusion



Stanford University



IBM Watson



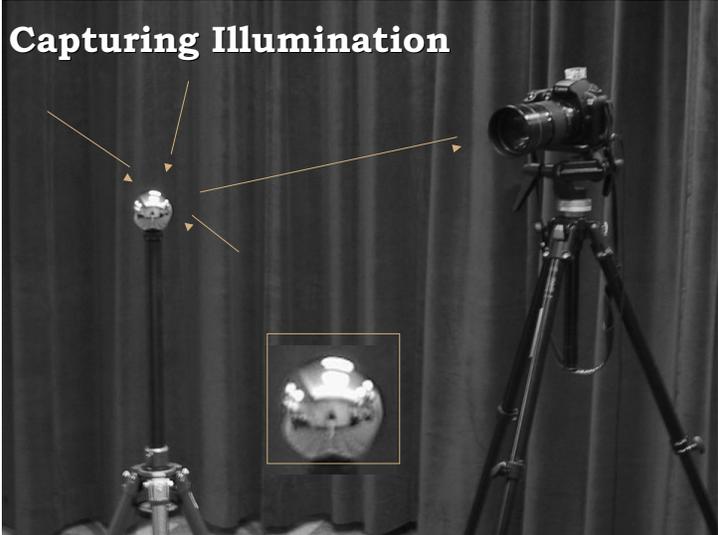
CNR Pisa

SIGGRAPH
2002

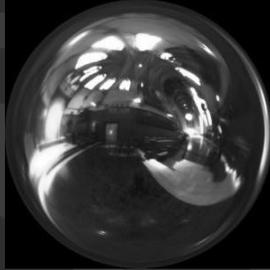
Challenging cases for current 3D scanning

- **Complex geometry**
 - Fur, plants
- **Non-diffuse reflectance**
 - Shininess (esp. spatially varying)
 - Anisotropy
- **Translucency**
 - Jade, gems
- **Interreflection**
 - Anything non-convex and light-colored

SIGGRAPH
2002



Captured Lighting Environments



SIGGRAPH
2002

Light Stage 1.0

Debevec, Hawkins,
Tchou, Duiker, Sarokin,
and Sagar. *Acquiring
the Reflectance Field
of a Human Face.*
SIGGRAPH 2000.

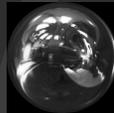


Sample Images

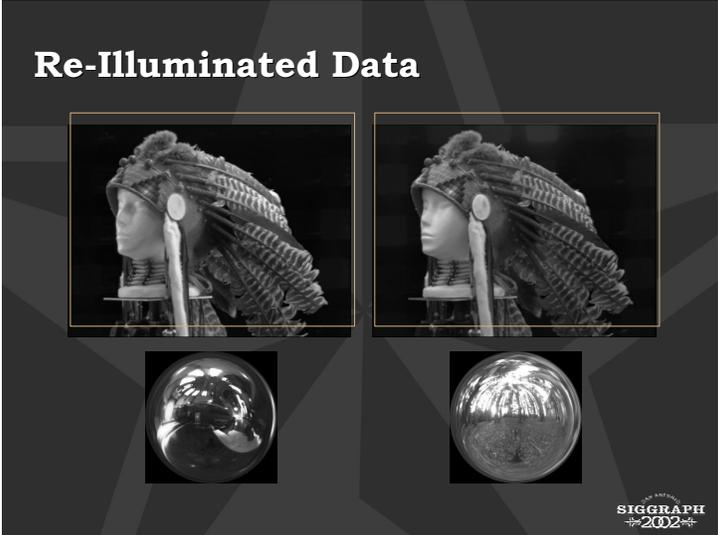
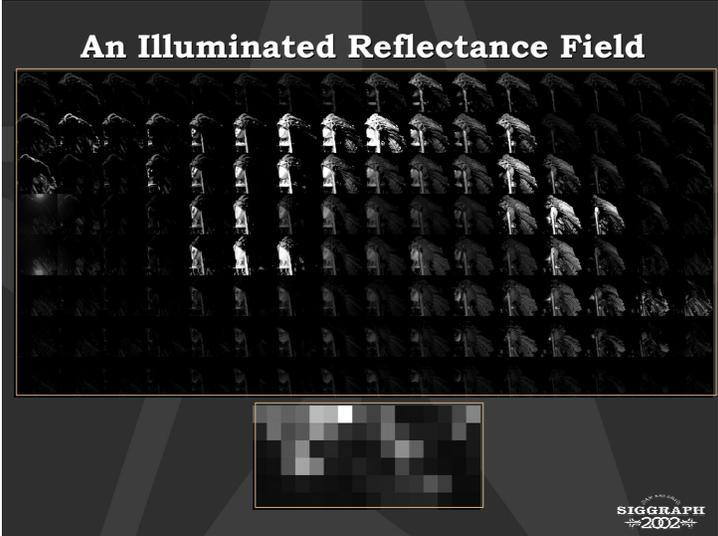


SIGGRAPH
2002

4D Reflectance Field Data



SIGGRAPH
2002



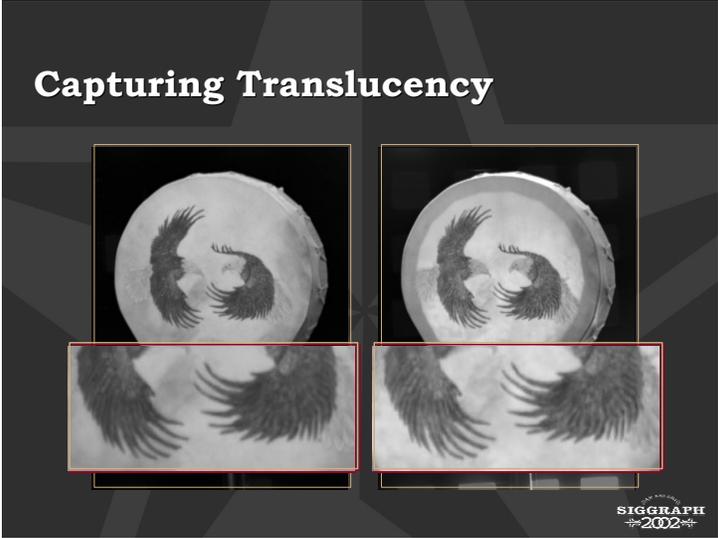
Capturing Translucency

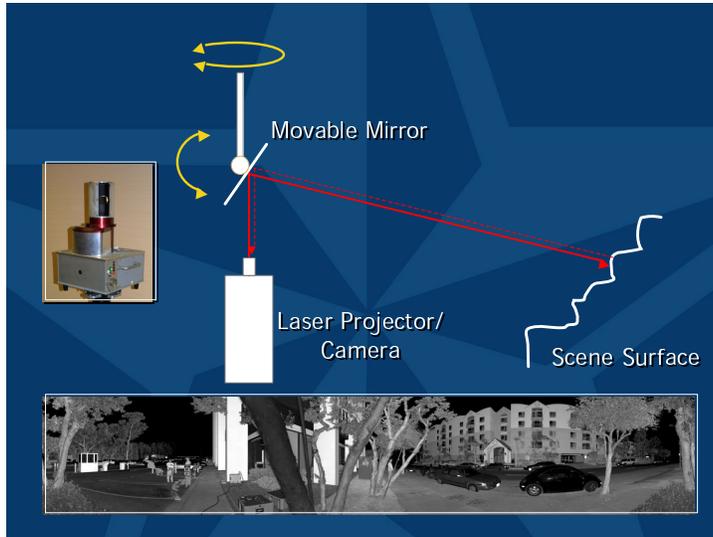


SIGGRAPH
2002



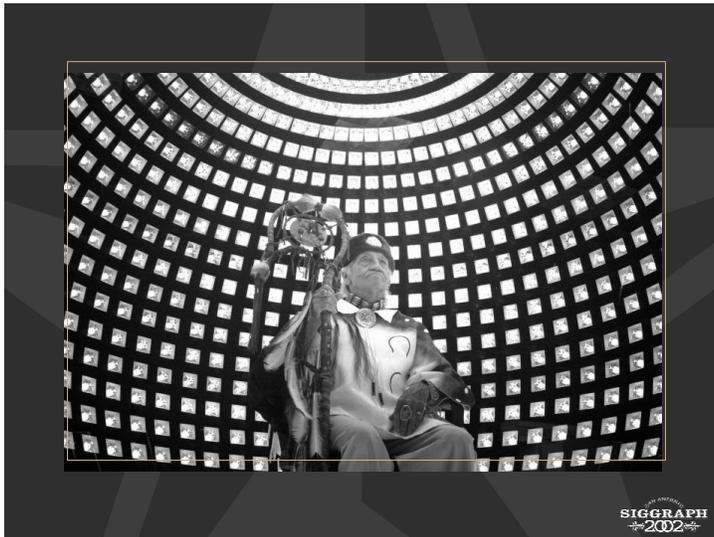
SIGGRAPH
2002

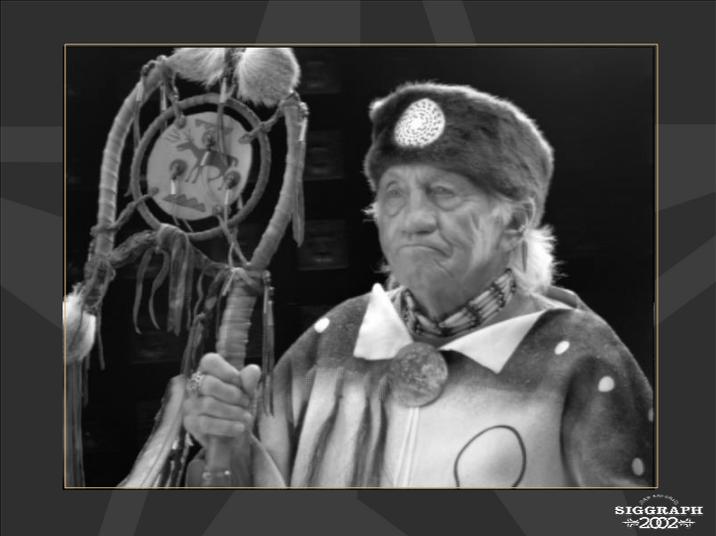




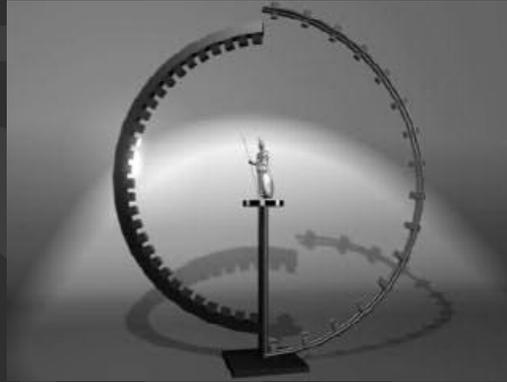


MODELS





Future work: Capturing from all viewpoints



SIGGRAPH
2002

View Interpolation



SIGGRAPH
2002

Conclusion: What do you do with the virtual object?

- View it from different angles
- Illuminate it differently
- Place it in virtual surroundings
- Compute its volume
- Manipulate it
- Feel its shape and materials
- Sense its smells
- Analyze its composition

SIGGRAPH
2002

Thanks!

- George Randall - collection owner
- Maya Martinez - production coordinator
- Chris Tchou - rendering programs
- Dan Maas and Chris Tchou - real-time demo
- Brian Emerson - 3D modeling
- Andrew Gardner - video editing
- Bobbie Halliday and Jamie Waese - production
- USC ICT, TOPPAN Printing, Inc. - support

SIGGRAPH
2002

Chapter 3

Very Realistic Lighting Simulation

by **Greg Ward**, Exponent.

3.1 Slides

The Three Steps of Lighting Simulation

1. Model Generation
2. Simulation / Rendering
3. Visualization

1. Model Generation

Determine the Accuracy Required

- Tape Measure or Laser Scanner?
- Macbeth Chart or Spectrophotometer?
- Aerial Photo or Site Survey?

The accuracy you require for your simulation depends on the application. It may not be as great as you think. If you only want to see what it would have looked like, you can often get by with relatively simple and inexpensive measurement techniques.

Tape Measure or Depth Scanner?

Tape Measure Requirements and Accuracy

- Digital camera and an assistant (both optional)
- Centimeter accuracy

Depth Scanner Requirements and Accuracy

- Scanner + computer + set-up + data reduction
- Millimeter accuracy

The modeling step following each measurement may require more or less time depending on the tools available and geometric complexity.

Tape Measure Example

- Accident Reconstruction (recent archaeology)
- Needed to know driver's eye height
- Photo with tape measure followed by computer modeling
- Centimeter accuracy



Use tape measure that is visible in photographs and/or record measurements on audio or in notebook. Photogrammetry may also be used in scenes containing straight lines.

Depth Scanner Example

- **Pietà Project**
www.research.ibm.com/pieta
- **Multi-baseline stereo camera with 5 lights**
- **Captured geometry and reflectance**
- **Sub-millimeter accuracy**



Technically involved procedure of measuring artifacts or sites with laser scanners requires painstaking set-up and tedious data reduction. The results can be quite impressive.

Macbeth Chart or Spectrophotometer?

Macbeth Chart Requirements and Accuracy

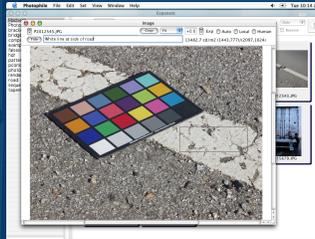
- ColorChecker chart and a digital camera
- Accurate to about 8 ΔE (1994 CIE Lab)

Spectrophotometer Requirements and Accuracy

- Hand-held spectrophotometer
- Accurate to about 1 ΔE

Macbeth Chart Example

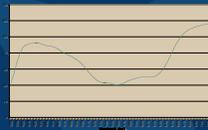
- Digital photo with ColorChecker under uniform illumination
- Compare points on image and interpolate
- Best to work with HDR image
- Accurate to $\sim 8 \Delta E$



Need to make sure you have diffuse illumination at a uniform angle and no highlights. Average over areas to avoid problems with noise in image and texture on object. Values are difficult to interpolate from standard image.

Spectrophotometer Example

- Commercial spectrophotometers run about \$5K
- Measure reflectance spectrum for simulation under any light source
- Accurate to $\sim 1 \Delta E$



Some spectrophotometers can separate diffuse and specular components. Most devices record reflectance at 10 nm increments over 400-700 nm range. Measuring textured materials is a problem.

Aerial Photo or Site Survey?

Aerial Requirements and Accuracy

- Satellite photos or (better) fly-over
- 1-10 meter accuracy, usually without elevation

Site Survey Requirements and Accuracy

- GPS or traditional surveying equipment
- 1-10 centimeter accuracy, with elevation

Aerial Photo Example

- Giza Pyramids
- Fly-over aerial photo shows positions of pyramids and tombs
- Requires perspective correction
- Accuracy is ~ 5 m



Taken from the website <http://sphinxtemple.virtualave.net>

Site Surveying

- Traditional instruments measure point-to-point
- GPS equipment measures absolute position
- Accuracy 1-10 cm



Taken from the website
<http://www.johann-sandra.com/surveying1.htm>

2. Simulation / Rendering

- *Radiance* Input Requirements
- Rendering Time and Accuracy
- Output Options

Radiance website is <http://radsite.lbl.gov/radiance>
Siggraph '94 paper is in the Appendix.

***Radiance* Input Requirements**

- *Radiance* is a physically-based lighting simulation and rendering tool from LBNL
- Takes geometry and RGB reflectances (also BRDFs, procedural textures, etc.)
- Prefiltering with source illuminant yields accurate colors ($\sim 2 \Delta E$)
- Give attention to source photometry□

See “Picture Perfect RGB Rendering Using Spectral Prefiltering and Sharp Color Primaries” on course notes

Example Radiance Image



Simulation of San Francisco air traffic control tower created to examine different shading devices and monitor equipment. The model was created from many measurements, including spectrophotometry and captured textures and patterns.

Rendering Time and Accuracy

- Diffuse interreflection and output resolution are the main parameters
- Many other parameters controlling time and accuracy of direct, specular, etc.
- User-friendly front-end program called “rad” is handy to control rendering

Computers are so much faster than when Radiance was written; it can now handle huge models with difficult lighting quite easily. Other methods such as Monte Carlo Path Tracing start to be competitive.

Rad Parameter Settings

Rad Variable Settings			
Variable Name	Interpretation	Legal Values	Affects
DETAIL	geometric detail	High, Med, Low	image sampling, irradiance value density
EXPOSURE	picture exposure	positive real	final picture brightness, ambient approximation
INDIRECT	importance of indirect diffuse contribution	0, 1, 2, ...	number of diffuse interreflections
PENUMBRAS	importance of soft shadows	True, False	source subdivision, source sampling, image plane sampling
QUALITY	rendering quality/accuracy	High, Med, Low	nearly everything
VARIABILITY	light distribution in the space	High, Med, Low	indirect irradiance interpolation, hemisphere sampling
ZONE	region of interest	Interior/Exterior keyword plus bounding box	irradiance value density, standard viewpoints

Rendering Quality Comparison



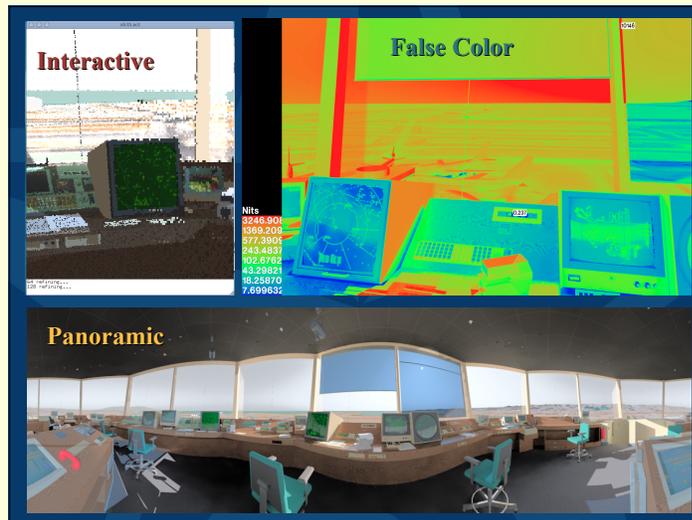
Low, Medium, and High quality rendering parameters as set by “rad.”

Output Options

- **Radiance picture format is gaining popularity in HDR imaging**
 - 4-byte RGBE uses common exponent per pixel
 - XYZE format for photometric images
 - Converts to and from other formats, including LogLuv TIFF
- **Interactive rendering for previewing**
- **Direct numerical output is also supported**

Usual output from Radiance is a picture, but numerical output and direct ray-tracing control is possible with the system using the rtrace program.

Interactive rendering is also supported with rview and rholo.



Interactive rendering allow us to preview our model and find good views.
 False color images allow us to analyze light levels.
 Panoramic images permit QTVR viewing of scene.

3. Visualization

- **Numerical Visualization**
 - False color images and plots
 - Visibility analysis - how well could people see?
- **Tone-mapping**
 - Visibility-matching tone operator
- **High Dynamic Range Display**

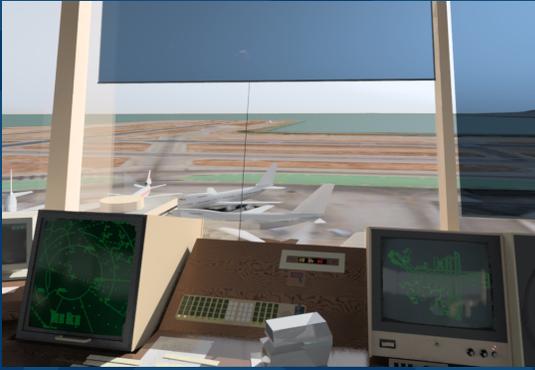
Visibility-matching tone operators include Larson et al '97 (in Appendix) and Pattanaik et al (Siggraph '98).

Tone-mapping Goal: Colorimetric



Cannot represent entire range on display, so we end up clamping.

Tone-mapping Goal: Optimize Contrast



Produces nice-looking result, but what does it mean?

Tone-mapping Goal: Match Visibility



Goal: If we can see it in the real world, we can see it on display and vice versa.

Operator Comparison



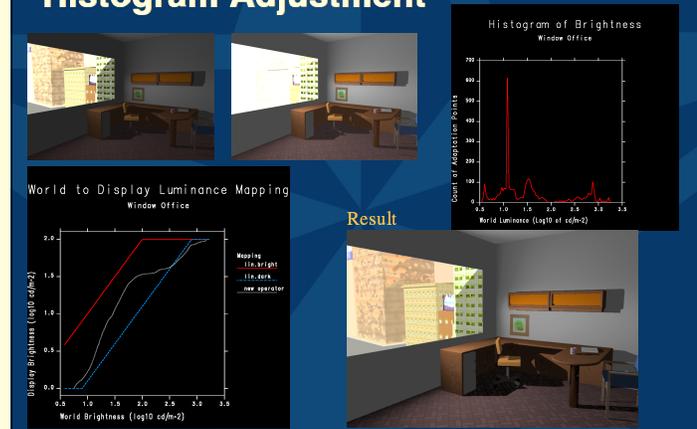
Clamping is obvious.
Differences in other operators more difficult to understand.
All examples here are global operators -- spatially varying operators are more interesting and potentially more powerful.

Visibility-matching Tone-mapping Approach

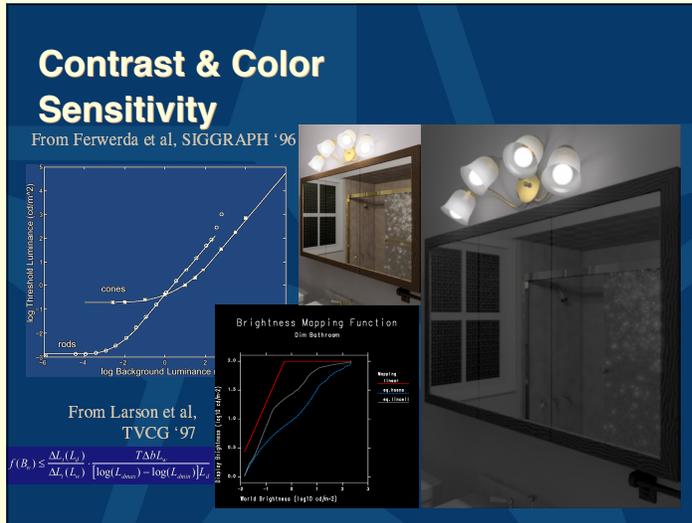
- Generate histogram of log luminance
- Redistribute luminance to fit output range
- Optionally simulate human visibility
 - match contrast sensitivity
 - scotopic and mesopic color sensitivity
 - disability (veiling) glare
 - loss of visual acuity in dim environments

Contrast and color adjustments are global over image; glare and acuity simulation are local.

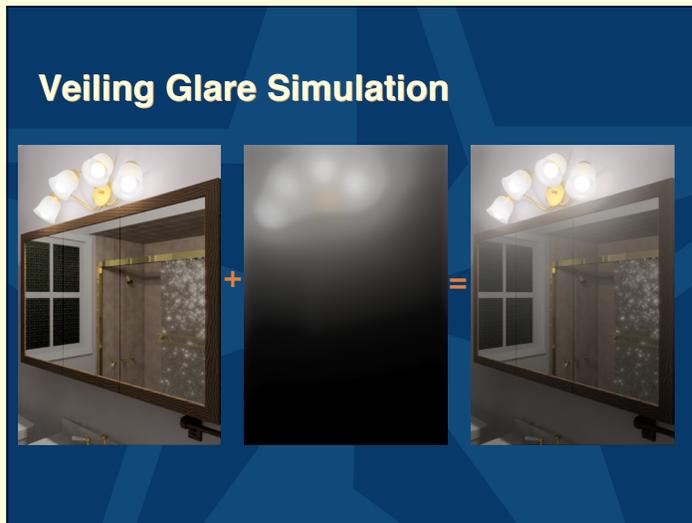
Histogram Adjustment



Dark and light colorimetric exposures vs. histogram adjustment.



Matching contrast and color sensitivity yields more representative visualization.



Glare is caused by scattering in the lens, iris, and aqueous humor, and results in veil cast on nearby photoreceptors.

High Dynamic Range Projection

- **TI Micro-mirror Device**
 - Good dynamic range, tunable gamut
 - Widely used for still projection systems
 - Already in trials for digital cinema
- **Silicon Light Machines Grating Light Valve**
 - Amazing dynamic range, widest gamut
 - Still in development
 - Promising for digital cinema

These projection systems are under development, and may provide wider gamuts and greater dynamic range than current LCD-based systems.

High Dynamic Range Display

- **UBC Structured Surface Physics Lab LCD-based HDR Display**
 - 1024x768 resolution
 - 60,000:1 dynamic range
 - 2,000 cd/m² maximum luminance
- **Next generation planned**
 - 2048x1536 resolution
 - 70,000:1 DR and 30,000 max. luminance

Working with Canadian research group to develop systems and applications for high dynamic range imaging.

High Dynamic Range Viewer



First prototype developed at LBNL 10 years ago.
Paper from PICS conference included on course CD-ROM.

Further Reference

- viz.cs.berkeley.edu/gwlarson
 - publication list with online links
 - LogLuv TIFF pages and images
- www.debevec.org
 - publication list with online links
 - *Radiance* RGBE images and light probes
- radsite.lbl.gov/radiance
 - *Radiance* rendering software and links

Chapter 4

Quantifying Realism

by **Alan Chalmers**, University of Bristol, UK.

The appearance of computer reconstructions of archaeological sites can often owe more of their form to the imagination of the artist than to the physics of the model. The rendered images may look ‘real’, but how can their validity be guaranteed? One approach is render images which are based on as much real-world data as is obtainable, rather than a preconceived aesthetic.

Much research has been undertaken into accurately modelling archaeological sites and reconstructing incomplete structures. Unfortunately, the luminaires used by standard modelling packages to render these scenes tend to be based on parameters for daylight, filament bulbs or cold fluorescent tubes, rather than lamp or candle light. Before the advent of modern lighting, illumination within ancient environments was dependent on daylight and flame. The lack of windows in ancient environments (certainly the case in caves) and later, glass if available was expensive, windows were inappropriate for defense etc., meant that even during daylight hours, some form of firelight was necessary for interior illumination. The fuel used for the fire directly affects the visual appearance of the scene. Furthermore, flames are not static and their flicker may create patterns and moving shadows that further affect how objects lit by these flames might look. Any realistic reconstruction of an archaeological site must take into account that these would have been flame-lit environments and thus the reconstructions should not only incorporate the accurate spectral profile of the fuel being burnt, but also the manner in which the flame may move over time.

4.1 Luminaires

“Command the Israelites to bring you clear oil of pressed olives for the light so that the lamps may be kept burning.” Exodus 27:20

This sections describes techniques which combine experimental archaeology and psychophysics, with computer graphics and vision in order to simulate flame lit environments to new levels of visual appearance and accuracy.

4.1.1 Creating a Realistic Flame

Modelling the shape of a flame mathematically is complex. An alternative approach is to incorporate video footage of a real flame into the virtual environment. Correctly included in the virtual environment the video provides a realistically shaped flame, while the illumination of the flame within the virtual environment can be computed from the size and position of this real flame. To capture the correct shape of a flame, and following this, the movement, the first step is to film a real flame. This is best done using a high quality digital video camera. This film can then be transferred onto a computer, and by using various techniques as described below, the flame from each frame can then be incorporated into the virtual environment.

‘Green Screen’ Technique

The simple technique of blue-screening, widely employed in the film industry, can be used to cut out an object from its background surroundings. Filming the flame against an evenly coloured, matt background enables thresholding of each frame to take place, which is used to identify and dismiss a background colour, effectively separating the flame from any unwanted parts of the scene. The background colour should be chosen so that it does not occur within the foreground. ‘Blue Screen’ uses thresholding techniques to achieve this. Thresholding is the process of identifying a range of colour, and changing all areas within an image within this colour range to another

specified colour. Simple, solid objects can easily be separated from a background, but a flame produces some of its own difficulties:

- It is useful to have static, even lighting on the background to simplify the thresholding process. Filming a flame clearly creates a problem in that it is itself a light source, and it is moving, disrupting otherwise static lighting, and producing lighting effects on the background.
- Parts of a flame may be translucent or partly transparent, so on a film the background colour may seep through into what we would identify as the flame. This is hard to avoid, but can be compensated for later, by deliberately seeping the background colour of the modelled scene into the flame.
- The efficiently burning part of the flame, around its base, is generally blue in colour, so rather than using a blue screen, another colour, such as green should be employed.

Once an object has been separated from its background, it can then be placed in front of other backgrounds. If used sensibly, the object can be blended into the new background to give the appearance that this is an unaltered, original scene.

Capturing the Flame

As discussed above, care has to be taken when filming in front of the green screen, to simplify the separation of the flame from its background. The green screen should be as evenly coloured as possible, and with as little shine as can be achieved. Simple things such as moving the flame as far from the screen as is possible help to achieve this. Figure 4.1 shows the required set up.

It is important to capture as much of the actual flame as possible, and as the melting of candle wax at the base of the flame creates a substantial depression, care has to be taken to keep the whole flame visible. The video is streamed onto a computer, and thresholding is used to separate the background. The digital video has to be broken down into individual frames before this can take place, and a suitable file

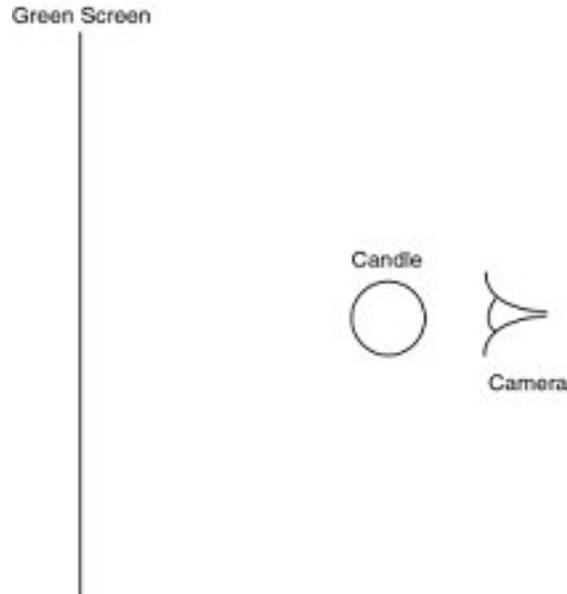


Figure 4.1: Green Screen filming

format used that will allow manipulation of the picture files. Common picture formats have different amounts of encoding to reduce their size, so they need to be decoded before any changes can take place. It is also important at this stage to attempt to keep as much accuracy as possible, so lossy formats such as JPEG should be avoided. In terms of frames per second, 25 fps will be used which is sufficient to provide smooth, flowing results. Correct thresholding should immediately identify the area in which the flame exists. One way by which the flame could be represented would be by using an enclosed edge map to give the outer line for the flame, in a picture file, or by a series of connected coordinates.

To ensure high fidelity reconstructions, rendering is carried out using Radiance [28]. Light sources in Radiance are best constructed from various predefined geometric shapes, such as cones, cylinders, and spheres. The simplest way of representing the shape of the flame is by using spheres.

This involves finding the top and the base of the flame and dividing the space in between into a number of segments, determined by how



Figure 4.2: Basic flame

many spheres will be used to model the flame. On each of the dividing lines, a sphere is created, the radius and centre point of which can be found by looking for the boundaries of the flame along the dividing line. This is achieved using the thresholded image. Each pixel on a line is tested, starting from the sides of the image and working in towards the flame. When a non-background colour is found, then this pixel represents the edge of the flame on one side. The same is done from the other side, until two co-ordinate positions are found, representing the flame boundaries on the particular line. Figure 4.3 shows the flame represented by just 3 spheres: it is clearly not modelling the shape of the flame accurately, shown by the large amounts of yellow 'flame' that are not covered by the green spheres.

Increasing the number of spheres should increase the accuracy of the representation. However, using more spheres leads to greater complexity in the rendering process, as each of the spheres will be an individual light source.

Note: Care must be taken if the spheres overlap each other as the manner in which Radiance implements the illum material can mean this leads to problems of self-shadowing amongst the spheres. This will need to be accounted for during the rendering process.

As Figures 4.3– 4.5 show, the theoretical benefit of adding more spheres to the representation slows quite quickly, the jump from 3 to 7 spheres being more useful in terms of percentage of the flame covered, than

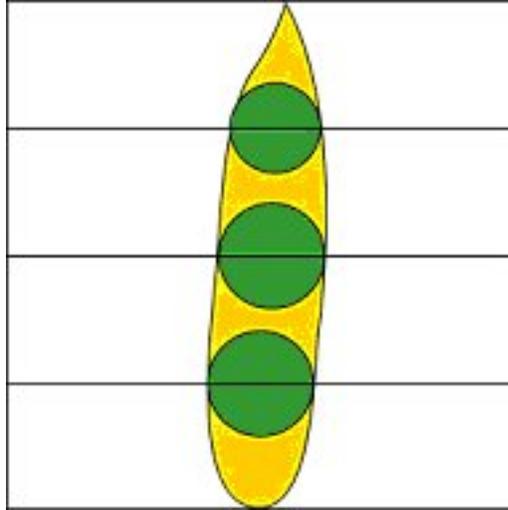


Figure 4.3: 3 sphere representation

raising from 7 to 15 spheres. Also, there is a large amount of duplication as more spheres are used, leading to an inefficient representation, which will result in a much longer rendering time. The above simple method for attaining sphere data is flawed when the flame is not vertical, overemphasising the size of the flame as shown Figure 4.6.

To deal with this, the centreline of the flame should be found, running from wick to tip. By using the normal to this line, a more accurate description can be produced. Figure 4.7 shows how with a rough red central line found for the flame, normals to this line can be used to divide the flame up in the same method as described above. This time, by working outwards from where the centre line intersects each normal, the boundaries for the flame can be found, and a sphere created.

To define a sphere in 3D space we need a centre point (x,y,z) and a radius. A flame is created using many of these sets of data, one for each sphere. This information then needs to be output to a text file. The information is currently in a state relative to pixels, and not to any real physical size. This means that using different resolution images of the same flame at a particular point in time, will give different data for the output spheres. This does not matter though, as the dimensions of each sphere, and the positions of all the spheres are

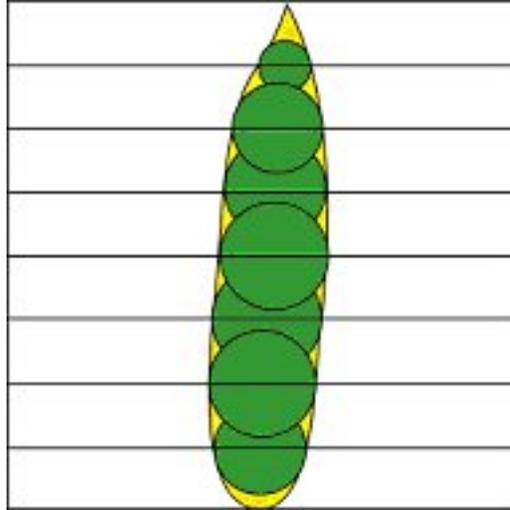


Figure 4.4: 7 sphere representation

relative. The spheres will need to be scaled suitably to fit in with the modelled environment. There are some problems with representing the flame in this manner. Firstly, it is quite likely that some of the spheres will be larger than the part of the flame they are attempting to model. This is particularly the case near the base and tip of the flame, or any part where the gradient of the edge is changing quickly. These bigger spheres mean slightly more light will be produced than is wanted. Another problem is the computational expense. Figure 4.7 demonstrates the difference between the 3, 7 and 15 sphere representations. With 15 spheres there is massive overlapping, and only very small parts of the original 3 and 7 spheres can be seen. It would be far better to find a way of converting the numerous spheres into 1 object, cutting out the duplication.

4.1.2 Converting Pixel Information to Radiance Files

The above method creates data on a pixel level, but to be used in Radiance, this data needs to be scaled and converted from raw data. Scaling can only be achieved by judgement and trial and error within the Radiance scene. A Radiance definition for a sphere is as follows:

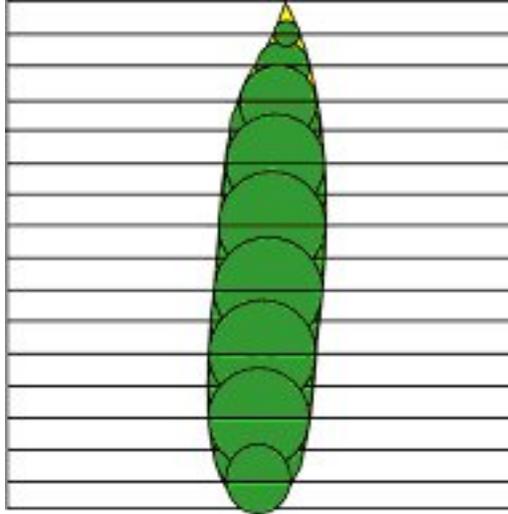


Figure 4.5: 15 sphere representation

```

candlelight sphere sphere1
0 0 4
X Y Z R

```

, where *candlelight* represents a previous declaration of the material type to be used, *sphere* describes the object to be created, and *sphere1* is the name of the new object. The next two lines show that there are no modifiers. *X*, *Y*, *Z* and *R* represent the (x,y,z) coordinates and the radius. The *candlelight* declaration will be used to describe the spectral properties of the light emitted from this sphere, in terms of a red, green and a blue component. A file containing the above description (with *X*, *Y*, *Z* and *R* replaced with suitable numbers) can be used in Radiance to create a single sphere. Many of the above descriptions will be used to create the flame.

4.1.3 Creation of a Radiance Scene

With a Radiance file created for each frame of flame, this now needs to be incorporated into a Radiance scene. The only specification for this part is for the scene to be interesting enough, and to be able to show the movement of the flame, and the light emitted from it. In

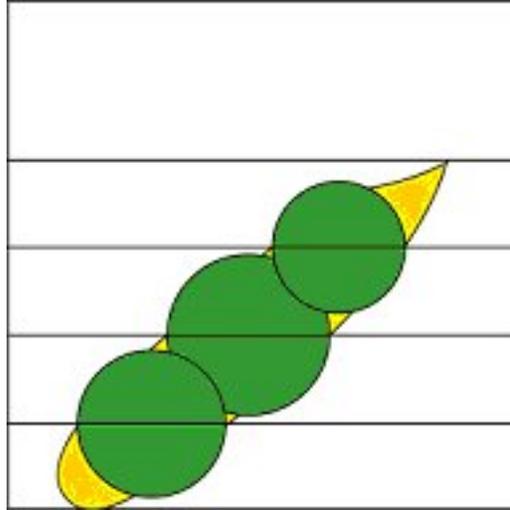


Figure 4.6: Oversized model

Radiance, there are several different materials that can be used as light sources. Most commonly used is *light*, but in this case *light* would be inappropriate, as when viewing the light source itself, it is visible. What is needed is an invisible light source, which is provided by the material type *illum*. When viewed directly the object made from *illum* is invisible, but light is still emitted from it. Accordingly, the *illum* material type will be needed to create the flames.

4.1.4 Creation of Final Images

With pictures rendered, it is now necessary for the original picture of the flame to be pasted in the right place on the scene. Some care needs to be taken here to accurately position the flame, and also to blend it into the scene, rather than simply sticking it on. As discussed earlier, parts of the flame may contain some of the old background colour within it, so this needs to be compensated for with the new background colour. With the flame replaced in all of the rendered frames, these can be played in sequence, at 25 fps, to create the finished animation. Following the above process should provide a path from the original film of a flame to another animation, with the same

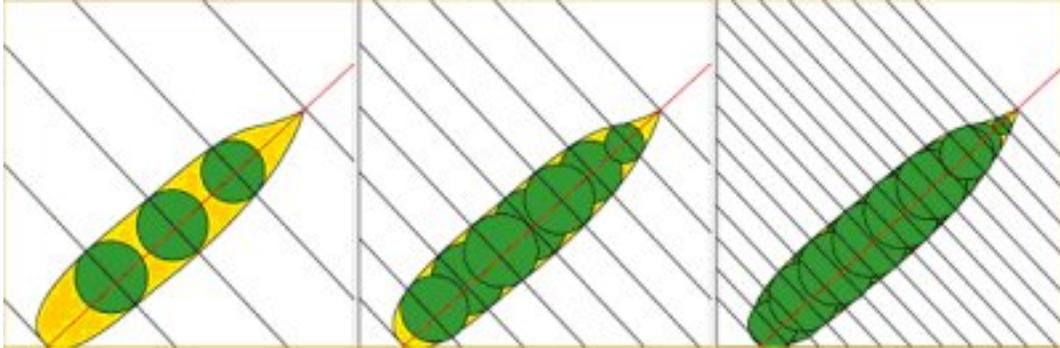


Figure 4.7: 3, 7 and 15 sphere representations

flame, but in a completely different, realistic environment.

4.1.5 Automation

It is very important that the whole process should be automated. When creating film at a rate of 25 fps, to get any reasonable amount of time out, many, many frames need to be produced. The programs for thresholding, creating flame data and replacing the flame, must work for all the frames at once, without needing to change any settings. If it works for the first frame, it must work for all the frames. For the rendering itself, this is not so important, as currently a complex scene may take many hours to render, so due to hardware constraints, this may have to be processed individually anyway.

4.2 Implementation

4.2.1 Flame to Film to Frames

After filming the candle flame in front of a green screen, a file format was chosen for which to convert the digital video. Binary .ppm (Poskanzer Portable Pixmap) files were chosen for several reasons. Firstly they are not lossy, and although in Windows environments ppm's are not well used, they are a standard on Unix, and can easily

be converted into Radiance pic files using *ra_ppm*, a built-in Radiance program. For instance,

```
ra\_ppm picture.pic  
picture.ppm
```

would convert the Radiance image (pic file), into a .ppm image. Another reason is that .ppms can be manipulated using standard Image Processing techniques, such as provided by the Image Processing Library (IPLIB). This provides C functions for decoding ppm's enabling them to be altered. The function *readppm* reads in a .ppm file and converts it into a pixmap, a 3D array. The first two indices represent pixel height and width. The third dimension/index of the array represents the red, green and blue parts of the image, stored with values of 0, 1, and 2 respectively. For instance, the statement

```
image[15][34][1] =  
240
```

would assign the green value of the pixel at (15,34) to 240. The values that can be contained in the array are unsigned chars, i.e. any value from 0 to 255. The video was streamed onto a computer and split into .bmp frames using Adobe Premiere, and then converted into .ppms using Paint Shop Pro. The frames must be ordered with a 3-digit sequence number, prior to the file extension. For instance the first file might be called *flame000.ppm*, so the second must be *flame001.ppm*, etc. This is a quick yet necessary process to turn the film into suitable media.

Figure 4.8 shows one frame from the original candle movie, shot against a green screen. The original film was recorded at a rate of 25 frames per second. This leads to a great deal of individual frames. Up to this starting point, other programs have been employed to turn the original footage into a series of images. It makes sense to use these packages to get to this early state, but from this point onwards it new programs need to be written. Programs (written in C) are needed that can take a large series of individual images, and quickly create data representing each of the flames within the image. Part of this also involves separating the flame from the green screen background. This



Figure 4.8: The Original Candle Flame



Figure 4.9: Get Flame Program Sequence

leaves a series of flames that can be used in a second program: code that blends in the individual flames onto the corresponding rendered scenes to create the final realistic images and animations.

4.2.2 Creating Sphere Data

At each stage in Figure 4.9, functions need to be written to carry out the specified task. The process starts with the original series of images. They are then reduced in size through cropping to remove unwanted background where the flame does not appear. The next stage is to threshold the image to separate the background from the flame itself. The crux of the process, creating the sphere data, is then undertaken. Finally, this is converted into Radiance files that can be used in the rendering process.



Figure 4.10: Original Flame Image to Cropped Flame

Cropping

The larger the picture files, the more work will need to be done, so it is important to get rid of unnecessary data at an early stage. It is also important to remove any candle or base that lies beneath the flame (Figure 4.10). Accordingly, two sets of coordinates are required, one specifying the top left corner of the cropped picture, and the other specifying the bottom right.

These coordinates are entered into the definitions section at the beginning of the program. As they are only entered once for the whole film, it is important to allow enough room so that the moving flame will always remain within these specified boundaries.

Thresholding

Thresholding replaces a specified range of colour values with a single colour. As shown in Figure 4.11, this has been used to separate the flame from its background, which has been turned to a single colour — in this case black. Some of the heat haze around the central flame has also been captured to add to the realism of the picture.

Changing the threshold settings can alter the amount of haze captured. For instance, if just the pure flame is required with no haze, by altering the inputs slightly, this can be achieved. These parameters are located



Figure 4.11: Cropped Flame to Thresholded Flame

at the definitions section at the start of the `getFlame` program file. Maximum and minimum values can be specified in each of the red, green and blue parts to separate those parts of the image that are to be removed, and those which are to be kept.

```

if(p[i][j][R] > RLOW && p[i][j][R] < RHIGH && p[i][j][G] > GLOW &&
    p[i][j][G] < GHIGH && p[i][j][B] > BLOW && p[i][j][B] < BHIGH &&
    p[i][j][G] > p[i][j][R]) {
q[i][j][R]= COLR;
q[i][j][G]= COLG;
q[i][j][B]= COLB;
}
else {
q[i][j][R]=p[i][j][R];
q[i][j][G]=p[i][j][G];
q[i][j][B]=p[i][j][B];
}

```

The code fragment above shows some thresholding. The pre-threshold image is stored in the pixmap p , and the new, thresholded image is stored in q . The code shows how one pixel of the original image is considered. If its red, green and blue values are within the designated colour values, specified by $RLOW$, $RHIGH$, $GLOW$, $GHIGH$, $BLOW$, and $BHIGH$, then it is considered an unwanted background area, and so replaced in the new image by another colour, specified by $COLR$, $COLG$, and $COLB$. If its colour is not within the boundaries, then the values of the original image are retained for the new image. To remove the green background used, suitable threshold values were:

```
#define RLOW 10
```

```
#define RHIGH 200
#define GLOW 80
#define GHIGH 230
#define BLOW 35
#define BHIGH 180
```

However, with different colour backgrounds, this can easily be altered.

Finding Spheres

The number of spheres used to model each flame can be specified at the start of the program using the *SPHERES* definition. This parameter can be used to change the balance between accuracy of the flame, and the speed of the rendering. A higher number of spheres leads to a better representation, but one that will take longer to render. The spheres are found using the method previously described. Some key code fragments demonstrating important parts of this process is as follows:

```
for(i=0;i<height;i++) {
    for(j=0;j<width;j++) {
        if(p[i][j][R] != COLR && p[i][j][G] != COLG &&
        p[i][j][B] != COLB) {
            tx = j;
            ty = i;
            i = height;
            j = width;
        }
    }
}
```

The above code shows a method to find the top of the flame, needed to work out a central line of the flame. It cycles through the thresholded image and finds the uppermost pixel that is not the same colour as the background colour (i.e. not *COLR*, *COLG*, *COLB*), and will therefore be the highpoint of the flame.

```

m2 = -(1/m);

divH = (by-ty)/(SPHERES+1.0);

divW = (bx-tx)/(SPHERES+1.0);

```

Once the centre line gradient has been found, the gradient of the normal, $m2$, can also be calculated. At this point the centre line can be divided up into the necessary number of chunks that is required. This is specified through the *SPHERES* definition. *divH* and *divW* represent the height and width jumps from one sphere to the next, bx , by , tx , ty , representing the top and base coordinates of the flame. It is now possible using the equation of the centre line, starting from the tip of the flame, to work down it, using *divH* and *divW* as the decrement values. Each time the decrement occurs a new sphere can be created.

```

for(j=levelW;j<width;j++) {
    y2 = m2 * j + c; /*Equation of normal*/
    y = round(y2);
    if(p[y][j][R] == COLR && p[y][j][G]
== COLG && p[y][j][B] == COLB){
        rightH = y;
        rightW = j;
        j = width;
    }
}

```

On each decrement, we start with a centre point of the sphere on the centre line. It is now necessary to find the edges of the flame on the normal to the centre line, and then calculate a more accurate centre point and radius for the sphere. The above code shows how the equation of the normal is used to find the pixel coordinates on which it lies. These values have to be rounded to keep them as pixel values. The code fragment shows how the program works outwards from the centre point to find the right hand side boundary of the flame. This point is stored in the *rightH* and *rightW* variables. The same is then done to find the left hand side boundary.

```

centreH2 = (rightH + leftH)/2.0;
centreW2 = (rightW + leftW)/2.0;
centreH = round(centreH2);
centreW = round(centreW2);

radius = (sqrt(pow((rightW-leftW),2)
+ pow((rightH-leftH),2)))/2.0;

```

The real centre point can then be determined, as well as the radius of the flame at that point, as the previous piece of code shows. So for each flame, a number of spheres have been defined as a triple; an (x,y) co-ordinate position centre point, and a radius. All the sphere data, for all the files is output into a single text file called “flameinfo.txt”. The number of files that are processed is specified by the *FILES* definition. The first line of this file contains information on the number of spheres used in each file and the number of files.

Figure 4.12 shows the ‘flameinfo’ output if *SPHERES* was set to 5 and *FILES* set to 4 (i.e. one sphere representation was created comprising of 5 spheres). The rest of the text shows the X , Y and Z co-ordinates of each sphere. The code also allows for extra images to be produced at this stage showing the centre points for each sphere, and the sphere edge boundaries, superimposed on the flame image. This is used as a tool to check the algorithm is working correctly. Figure 4.13 shows 2 such images. The green dots show the boundary of the flames. The red dots represent the predicted centre line, calculated by forming the line from tip to base. The blue dots show the actual centres at that particular point in the flame. By taking the blue dots as centre points and using the distance from this to the closest green point on the normal line as the radius, each of the spheres are formed.

Producing Radiance Scene Files

The next part of the program reads in the file ‘flameinfo.txt’, and creates the output, as Radiance scene descriptions. These are ordered in the same manner as the input files, going from *f000.rad* up to *fnnn.rad*.

```

5 4
-0.50000 66.00000 5.50000
-0.50000 54.50000 8.51469
-0.50000 43.00000 9.50000
-0.50000 31.50000 9.51315
0.50000 19.50000 8.51469
-1.00000 65.50000 6.02080
-0.50000 54.00000 8.50000
-0.50000 42.50000 9.51315
0.00000 31.00000 9.00000
0.50000 19.50000 8.51469
-1.00000 65.50000 6.02080
0.00000 54.00000 9.00000
0.00000 42.50000 10.01249
0.50000 31.00000 9.50000
0.50000 19.50000 8.51469
0.00000 65.50000 6.02080
0.00000 54.00000 9.00000
0.00000 42.50000 10.01249
0.50000 31.00000 9.50000
0.50000 19.50000 8.51469

```

Figure 4.12: File information

```

for(j=0 ; j<spheres ; j++) {
    fscanf(input, "%lf", &x);
    fscanf(input, "%lf", &z);
    fscanf(input, "%lf", &radius);
    fprintf(output, "candlelight sphere flame%d\n", (j+1));
    fprintf(output, "0\n0\n4  ");
    fprintf(output, "%.5lf 0  %.5lf  %.5lf\n\n",
            x/SCALEX, z/SCALEZ, radius/SCALER);
}

```

The data has been scaled so that it creates flames of the right size within Radiance, through the *SCALEX*, *SCALEZ* and *SCALER* defini-



Figure 4.13: Edge and Centre Markers for 7 and 20 Spheres

tions. This may need to be adapted after rendering a scene, to check that the flame is the correct size in the picture. It is hard to predict the relative size beforehand, so actually rendering the scene with the spheres as solid, non-light emitting objects, is perhaps the best way of scaling properly. It was found that at this pixel level, dividing the original data by a factor of 2000 led to flames being created at the right scale within Radiance.

Other Matters

Another value that must be specified at the start of the program is that of the centre point of the candle (i.e. the width value at the base of the wick). This should not change throughout, so only needs to be specified once. This makes it simpler to find the centre lines of the flames. The program should be written so that it may easily be adapted to suit others needs, when using other film with different characteristics.

4.2.3 Superimposing the Flame

With the rendering completed, a large sequence of images will have been produced. It now remains for the real flame corresponding to each image to be reinserted to create the final scenes. Another program needs to be written to handle this. As the images have retained the same sequential numbering system throughout, it is a simple task to match rendered images with the correct cut out flame. However

some groundwork must be done to ensure the flame is positioned correctly within the scene. A co-ordinate position must be specified, so that the flame is correctly placed at the point of the candle top. It is obvious when this is wrong, as it looks strange on the output, so trial and error is again the best way of achieving this. The actual flame picture may also need to be scaled at this point if the candle in the scene is further away or closer to the viewpoint, than the original flame was to the camera.

```

if(p[i-IMIN][j-JMIN][R] > 200 && p[i-IMIN][j-JMIN][G] > 150
&& p[i-IMIN][j-JMIN][B] > 100) {
    q[i][j][R] = p[i-IMIN][j-JMIN][R]; //catches the bright areas
    q[i][j][G] = p[i-IMIN][j-JMIN][G];
    q[i][j][B] = p[i-IMIN][j-JMIN][B];
}
else if(p[i-IMIN][j-JMIN][R] < 250 && p[i-IMIN][j-JMIN][G] < 130
&& p[i-IMIN][j-JMIN][B] < 90) {
    q[i][j][R] = p[i-IMIN][j-JMIN][R]; //catches the dark areas
    q[i][j][G] = p[i-IMIN][j-JMIN][G];
    q[i][j][B] = p[i-IMIN][j-JMIN][B];
}
else {
    q[i][j][R] = (p[i-IMIN][j-JMIN][R] + r[i][j][R])/2;
    //For the remainder (i.e. heat haze and blue bits)
    //colour ratio is even
    q[i][j][G] = (p[i-IMIN][j-JMIN][G] + r[i][j][G])/2;
    q[i][j][B] = (p[i-IMIN][j-JMIN][B] + r[i][j][B])/2;
}

```

The central parts of the flame are placed straight on. These are the areas where none of the background colour will come through, such as the brightest part of the flame and the darkest parts such as the wick. This is shown on the *if* and *if else* parts of the above code. The surrounding areas, like the heat haze produced by the flame and the blue at the base around the wick, are blended in with the background colour to produce a more realistic effect.

This is demonstrated in the Figure 4.14. The left image shows the flame placed straight on, with no blending techniques applied. This

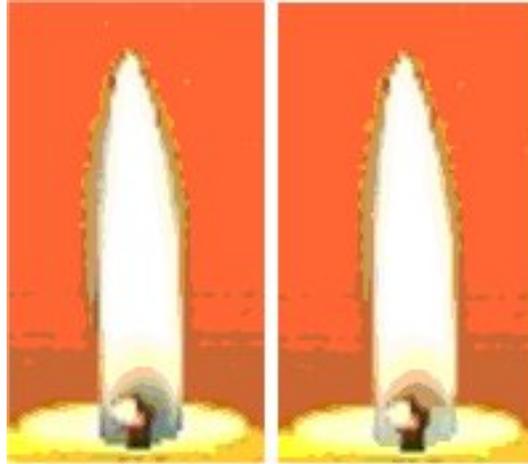


Figure 4.14: Different Ways of Blending the Flame

gives a harsher, unrealistic edge to the flame, with a solid haze. The right image shows the haze somewhat wispy, as would be expected. Also on the left image, the blue base of the flame does not depict the transparency that a real flame gives. The right image shows a truer image, with some of the background showing through the transparent blue base. The program allows different areas of the flame to be identified through a colour range, and each one can then be treated differently. They can be placed straight on, blended in to varying degrees, or not put on at all. Examples of the results can be seen in Figures 4.15– 4.16.

4.3 Converting Luminaire Data

The final problem to overcome is the conversion of luminaire data into RGB values. Because we wish to represent the results on computer displays we need to break the data down into spectral contributions for the three primary phosphor colours. When this process is performed, the detailed spectrum data from the spectroradiometer is merged into values representing the red, green and blue portions of the spectrum. It is essential that this conversion is calculated in a perceptually valid way, as defined by the CIE (Commission International de l'Eclairage)

1931 1-degree standard observer. This system specifies a perceived colour as a tristimulus value (i.e. three coordinates) indicating the luminance and chromaticity of a stimulus as it is perceived in a 1-degree field around the foveal centre.

Figure 4.17 shows the functions for the X , Y and Z channels. The Y channel measures the luminance of a source, and the X and Z channels measure the chromaticity. This information is more useful when broken down as follows: If we let $x = X/(X+Y+Z)$ and $y = Y/(X+Y+Z)$ Then we can calculate the exact colour values for the red, green and blue sections of the spectrum, disregarding luminance. For a canonical set of VDU phosphors; RED(x, y) = (0.64, 0.33) GREEN(x, y) = (0.29, 0.60) BLUE(x, y) = (0.15, 0.06). Many luminance/chrominance meters will record the coordinates of the CIE standard observer. Fortunately, Radiance comes with a program *rcalc* to convert the xyY coordinates to RGB values, using *xyz_rgb.cal*. It is worth re-iterating a couple of points. Firstly, the accurate modelling of luminaire colour values and temperature is usually unnecessary unless high measurement accuracy is required, because the image will appear unrealistically tinted. The perceived colour will always be shifted towards the white. Secondly, a reliance on extreme accuracy in these images will probably be unwise, because an RGB calculation is by definition an approximation of the colours present. If very accurate colour readings are required, we should calculate the convolution of the emission spectrum of the light source with the reflectance curve of the material under examination.

4.4 Validating Realism

The aim of realistic image synthesis is the creation of accurate, high quality imagery which faithfully represents a physical environment, the ultimate goal being to create images which are perceptually indistinguishable from an actual scene. Advances in image synthesis techniques allow us to simulate the distribution of light energy in a scene with great precision. Unfortunately, this does not ensure that the displayed image will have a high fidelity visual appearance. Reasons for this include the limited dynamic range of displays, any resid-

ual shortcomings of the rendering process, and the extent to which human vision encodes such departures from perfect physical realism. Conversely, along many parameters, the visual system has strong limitations, and ignoring these leads to an over specification of accuracy beyond what can be seen on a given display system. This gives rise to unnecessary computational expenses. It is increasingly important to provide quantitative data on the fidelity of rendered images. This can be done either by developing computational metrics which aim to predict the degree of fidelity, or to carry out psychophysical investigations into the degree of similarity between the original and rendered images.

Techniques to compare real and synthetic images, identify important visual system characteristics and thus produce benefits to the graphics community such as being able to reduce rendering times significantly, have been the subject of two previous courses at SIGGRAPH: “Seeing is Believing: Reality Perception in Modeling, Rendering and Animation” [12] and Image Quality Metrics [11].

McNamara et al.’s paper, “High Fidelity Image Synthesis” [30] which discusses many of the issues is included in Appendix A, as is Devlin and Chalmers’ application of the above methods, “Realistic Visualisation of the Pompeii Frescoes” [16].

section*Acknowledgements We would like to thank Jean Archambeau, the owner of Cap Blanc, for his permission to work at the site and his interest and Francesco d’ Errico, Kate Robson Brown, Ian Roberts, Chris Green, Natasha Chick and Michael Hall for their input to this work. Many thanks also to the Bristol/Bordeaux Twinning Association and the ALLIANCE / British Council programme (Action integre franco-britannique) for their financial support.

Much of this work first appeared in:

- I. Roberts, *Realistic modelling of flame*, A. Chalmers (advisor), BSc Hons Thesis, University of Bristol, May 2001.
- C. Green *The visualisation of ancient lighting conditions*, A. Chalmers (advisor), BSc Hons Thesis, University of Bristol, May 1999.



Figure 4.15: Array of gently Flickering Flame



Figure 4.16: Other candlelit environments

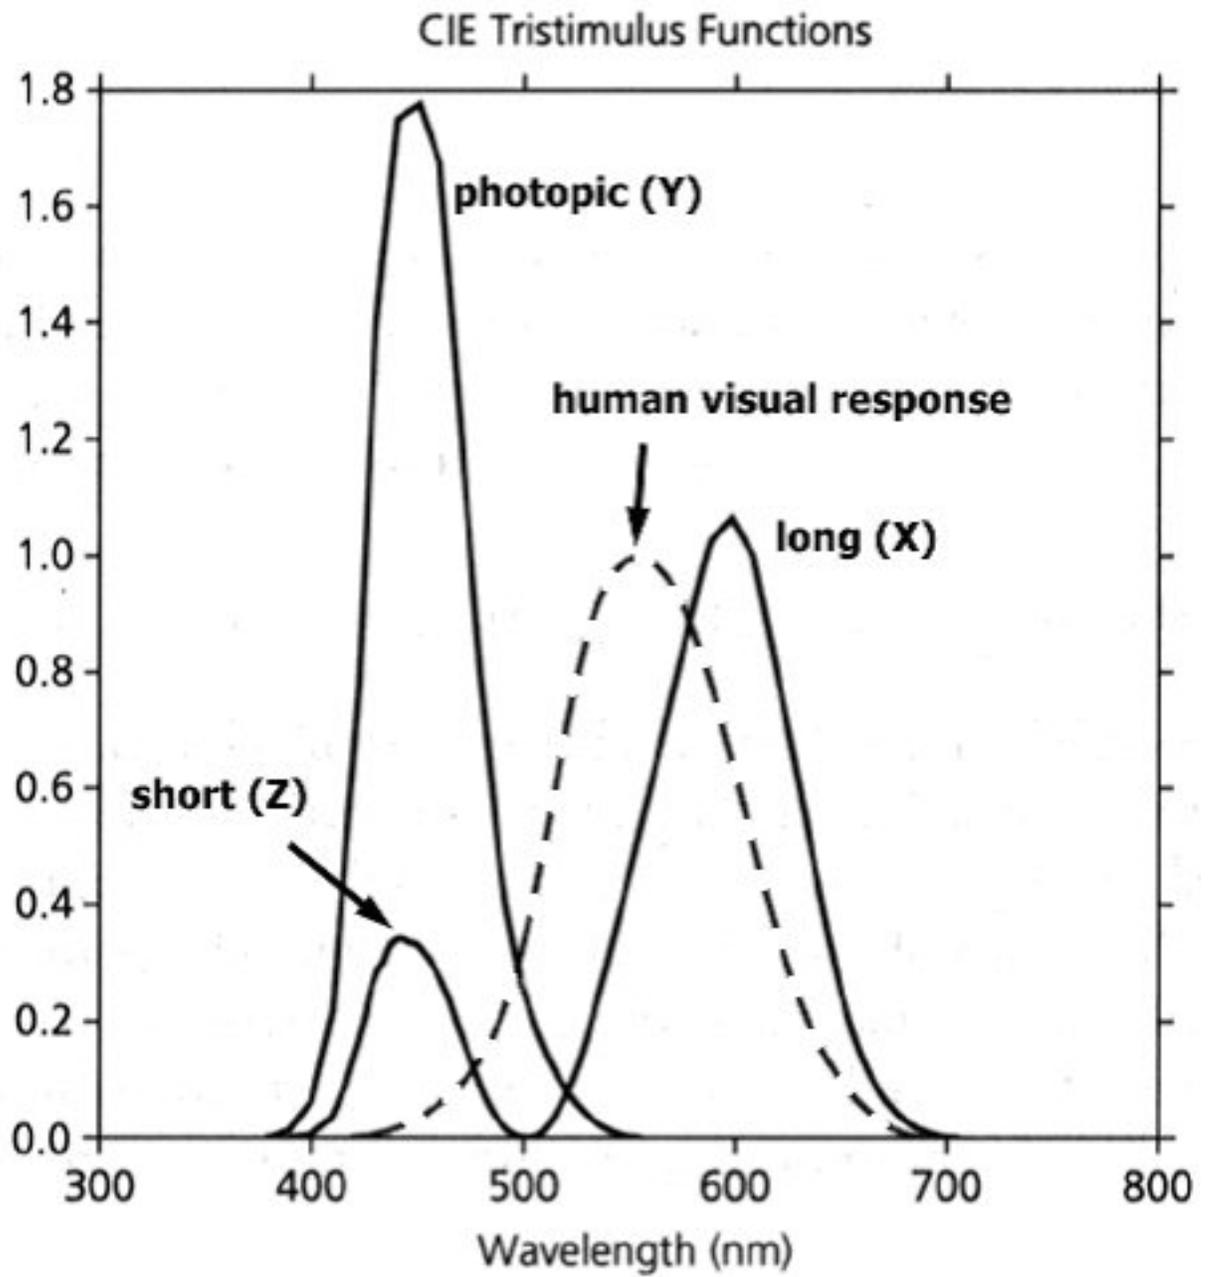


Figure 4.17: The CIE tristimulus curves

Chapter 5

Representation and Interpretation

by **Kate Devlin**, University of Bristol, UK.

5.1 Introduction

The idea of representing a past environment in the form of an interpretive illustration is by no means a new one. While reconstruction drawings date back to the beginnings of the discipline itself, computer graphics has enabled new dimensions — quite literally — to be added to visual representations of archaeological sites. The archaeological community has embraced this technology, finding in it new methods of presenting the past to a public whose demands grow increasingly sophisticated. With the advent of ‘media archaeology’ in the form of television documentaries and World Wide Web presentations, and the increasing use of audio-visual displays in museums and interpretive centres, three-dimensional computer graphics provide an aesthetic and convenient means of enhancing an archaeological experience, allowing us a glimpse of the past that might otherwise be difficult to appreciate.

To date, however, the emphasis has been on using three-dimensional computer graphics for display purposes, with interpretation and re-

search taking second place to the need for media representations. The current trend for artistic conception and photo-realism in reconstructions is not enough to benefit the archaeological community, and for the archaeologist to use the computer-generated environments as a research tool, stricter controls are necessary. It is only when we can explain and quantify the accuracy of the generated image that it can be used for an interpretative purpose.

5.2 A brief history of archaeological illustration

“Without drawing or designing the Study of Antiquities or any other Science is lame and imperfect.” William Stukeley, 1717.

This section provides a brief overview of the development of archaeological illustration, from depictions of monuments in historical manuscripts through to the immersive Virtual Reality worlds created today, outlining the main developments in society and technology that permitted advances to be made. This provides a context for the aspects of representation that will be discussed later on in this chapter.

5.2.1 Archaeological illustration: an overview

Illustrations depicting archaeological sites date back to the medieval period, as far back as the recorded interest in antiquities itself. The Renaissance heralded a rediscovery of classical Antiquity and marked a new approach to knowledge that in the UK, for example, was manifested in the foundation of the national academy of science — the Royal Society — in 1660 [34]. This led in turn to the regular meetings of the Society of Antiquaries of London from 1717 onwards, with the aim of “the encouragement, advancement and furtherance of the study and knowledge of the antiquities and history of this and other countries”. Systematic illustration of excavated artefacts took

hold, and by the mid-nineteenth century archaeology (rather than antiquarianism) became established as a discipline. Advances in metal-plate engraving for the printing of illustrations and development in wood engraving and lithography permitted the creation of detailed and intricate work.

When the pioneering archaeologist Pitt-Rivers published his excavation reports in the late nineteenth century his illustrations came closer to the standard of work required today than to that of his contemporaries. The advent of photography brought a greater choice of mediums in the twentieth century, and following the First World War the potential of aerial photography was realised, as was the importance of distribution maps for analytical purposes [1]. The possibility of printing good quality photographic reproductions meant that the 'realistic' drawing of the site could be discarded, yet the archaeologist today works with stylised two-dimensional site plans. Stripped of its subjectivity, the drawing in the form of a plan is still a means of explicitly conveying information.

Computer Added Design (CAD) was first developed towards the end of the 1970s, initially two-dimensional but overtaken by three-dimensional packages in the second half of the 1980s. CAD packages have found a market in the archaeological community where they provide a robust, user-friendly means of linking site plans and section drawings to create a complete three-dimensional record of an excavation [15]. Visualisation projects in the UK originated in the late 1980s with Woodward and Bowyer's reconstruction of the Temple Precinct from Roman Bath, a project that inspired a succession of similar applications elsewhere [32].

Other forms of computer applications were subsequently explored. The use of GIS (Geographics Information System) was first seen in archaeology in 1986 in the US. GIS provides a method of combining spatial data and textual information for the purpose of landscape visualisation and analysis, and is in widespread use today. In terms of three-dimensional representations, VRML (Virtual Reality Modelling Language) was recognised as an international standard in 1997, providing a scene description language that permits the viewing and manipulating of a 3D 'world'. The low cost, ease of use and portability of VRML led to its adoption for interactive explorations of archaeo-



Figure 5.1: INSITE project reconstruction

logical representations, although interactivity is at the cost of realism. Conversely, the move towards photo-realism in computer graphics inspired ‘accurate reconstructions’ of sites based on excavation reports or standing remains. However, the idea of portraying a subjective interpretation as ‘real’ is no less fanciful than the antiquarians’ paintings of the past. Approaches have been taken to quantify the realism of the images offered, such as the INSITE project at the University of Bristol which sought to accurately simulate the light distribution in a scene using the spectral value of the original light source (Figure 5.1) [10].

Having moved into the twenty-first century, multi-sensory and mixed reality applications now provide further ways to present heritage information. Shaderlamps — a method of graphically animating physical real world models with projectors (Figure 5.2.1) has brought three-dimensional computer graphics and animation outside of the computer monitor [35].

From headmounted displays to total immersion in graphics CAVE Automatic Virtual Environments (CAVE), Virtual Reality has been embraced as a new way of imparting archaeological representations [3]. Nowadays, museums and heritage centre come replete with audio-visual displays, yielding to the demands of an increasingly sophisticated public. The all-pervading mediums of television and Internet have brought archaeology into our homes. The study of the past has become high-tech and sexy — slick, speedy and visually stunning presentation is expected.

The future of archaeological representation is limited only by tech-



(a) Shaderlamps set-up

(b) Resulting model

Figure 5.2: Shaderlamps: illuminating with projectors. *By permission of R. Raskar.*

nology. As computer power continues to grow and hardware and software fall in price, new applications can and will be found. We are moving into the realm of multi-sensory VR experiences — acoustic rendering is a growing area of interest, and touch and smell have been integrated into museum displays. The promising area of augmented reality can integrate archaeology into our everyday experiences. So far, the demand is there, and the thrill of the subject is as relevant today as it was to the first antiquarians of the Renaissance. With the current media fascination for all things archaeological perhaps we should ask ourselves how far we can go with excavation plans and records before it all becomes *infra dig*?

5.2.2 Case study: seeing Stonehenge

A good example of the changing nature of the representation of archaeological sites is that of Stonehenge in Wiltshire, England. It is depicted in a fourteenth century manuscript (albeit as an illustration of a legend first appearing around 1136 of the magician Merlin building the monument, Figure 5.3), but the essential form of the henge is apparent in this and also in a contemporary manuscript where the circle of stones is shown as a square due to the layout of the page.



Figure 5.3: Fourteenth century manuscript depicting Merlin erecting Stonehenge (MS Egerton 3208 f.30r). *By permission of the British Library.*

Paintings in the sixteenth and seventeenth centuries, presumably based on verbal accounts due to some inaccuracies, show an increasing awareness in antiquarian studies. A ‘reconstruction’ of the monument as a purported Roman Temple, *Stonehenge Restored*, by Inigo Jones in 1655 portrayed a completed and orderly drawing in full architectural splendour - the way he thought it should have been. John Aubrey made the first plan of the site in 1666.

In William Stukeley’s topographical recording of Stonehenge, published in 1740 at a time when the Society of Antiquaries was well-established, *scientific* visualisations of a sort (for his imagination ran unfettered in some of his representations) were truly underway, such were his intentions to carry out the work for a distinctly archaeological purpose. Stonehenge remained in the eye of the artist, however, and Turner’s watercolour of 1829 is a no-holds-barred, intimidating portrayal of storm-lashed stones, a howling dog and an unfortunate shepherd struck dead by lightning. Constable also painted the monument in a similarly apocalyptic manner in 1836, continuing the tradition of romance and drama that so inspired earlier depictions. These nineteenth century paintings may have called artistic licence into play, but they are nonetheless imbued with the approach seen several hundred years before, and if we are to refer to the latter as a form of

archaeological representation, then on the same grounds we should not exclude the former.

The twentieth century produced standardised archaeological excavation plans of Stonehenge, symbols and conventions demarcating the relevant features. A proliferation of drawings, paintings, photographs and postcards added to the purely archaeological record. The site never once wavered in popularity, with a three-dimensional replica of the monument even appearing, tragically not-to-scale, in 1984 rock ‘mockumentary’ *This Is Spinal Tap*. As VRML grew in popularity, so did a rash of virtual Stonehenge models created by enthusiasts the world over, its potential and popularity easy to comprehend. It is fair to say, though, that none came close in sheer scale and complexity than that developed by Virtual Presence Ltd. co-sponsored by Intel and English Heritage. Using source data such as photogrammetry, GIS data, site plans, excavation reports and astronomical maps, a detailed interactive model was created. (Figure 5.2.2 shows screenshots from this project.)

Stonehenge will undoubtedly continue to prove inspirational, and we should look forward to future interpretations using the visualisation techniques that are developing today. From the earliest recorded drawings to an interactive replica, each interpretation stimulates interest in the monument, and that — above all — is surely what we should be trying to achieve.

5.3 The idea of realism

“Even where both the light rays and the momentary external conditions are the same, the preceding train of visual experience, together with the information gathered from all sources, can make a vast difference in what is seen. If not even the former conditions are the same, duplication of light rays is no more likely to result in identical perception than is duplication of the conditions if the light rays differ...the behaviour of light sanctions neither our usual nor any other way of rendering space; and perspective



(a) Close-up of trilithon



(b) Aerial view



(c) Sunrise



(d) Sunrise (detail)

Figure 5.4: VR Stonehenge screenshots. *By permission of by Virtual Presence Ltd.*

provides no absolute or independent standard of fidelity.”
Nelson Goodman, 1976.

This section discusses issues with interpreting images that are intended to be realistic, and defines the key terms and concepts in this area. It briefly outlines the problems of defining what we mean when we say an image is real. It queries the amount of authenticity we can expect from a virtual representation of an archaeological site given limitations in both image generation and archaeological data, and should illustrate the need for placing representations into a context in order for them to be useful.

Because this topic covers some theoretical approaches to archaeological representation, an example is helpful. Suppose we have a house that was inhabited during the epidemic of Bubonic Plague that swept through Europe in the Fourteenth century. Archaeologists excavate the house and it is decided to use it as the basis for a museum on life in that time. We will use this example to illustrate issues of particular relevance to representation.

5.3.1 Terms and concepts

Reconstruction vs. representation

To begin, we must find a term to define the end-product of our work. The idea of classifying a generated image as a *reconstruction* — as such computer-generated past environments are often described — is slightly misleading. Although the geometry of the synthesised scene can be based on site plans and other captured data, any gaps in the information cannot be portrayed with justification. Where information is lacking, conjecture takes over. Where a construction is the original real-world scene, a reconstruction implies an objective rebuilding of this based on the material remains. By contrast, a *representation* more accurately describes one particular interpretation, to which there may be many and varied alternatives.

If we consider our example of the Plague house, we can rebuild it objectively, either physically, on paper, or on computer, using the same dimensions and materials. This is a reconstruction. When we recreate

the house and portray it *as we think it was in that time*, where gaps in our knowledge are filled in with educated guesses, it becomes a representation. In this chapter we will therefore use the term ‘representation’ when referring to our computer-generated archaeological scenes.

Defining realism

The human visual system is far too complex and unwieldy to currently know enough about its processes. It operates as a whole; we cannot yet replicate that on a computer. If we present humans with the same information the visual system can process it in the same manner, but on a higher (personal) level there is subjectivity. Before we produce an image and claim it is ‘realistic’, we must first define what we mean by ‘realism’. To say an image is *physically* real - i.e. the scene geometry is identical to the scene in the real world and the light distribution is accurately simulated - is concrete and understood, but to say it is *perceptually* real is to use a term that needs to be quantified.

There exists a wide range of philosophical literature and discussion on the nature of realism, and each tends to state that our perception of a scene is subject to some factor: association, connotation, contextual/comparative views, semantic meaning, Platonic idealism, expectation and experience, discrimination vs. association, recognition over time/exposure, sensory appeal, familiarity and recognition, and so on, and so forth. . . With so many apparent influences it becomes difficult to know how to approach a definition of realism. If the viewer brings their own experiences and subjectivity to an interpretation then each viewer must have their own version of reality. By trying create reality by faithfully copying something we are immediately limited by the inability to specify what it is that we are copying.

How do we decide if an image is real unless we compare it with something that we have defined as being real? What, therefore, constitutes realism of representation? Goodman suggests that it is “the probability of confusing the representation with the represented. . . how far the picture and object, under conditions of observation appropriate to each, give rise to the same responses and expectations” [19]. Certainly, in perceptually realistic graphics, the idea is to evoke the same

response from a generated scene as we might have to a real-world scene. Admittedly, it is an explanation that can be questioned given the subjective nature of personal response, but it does suffice to avoid becoming entrenched in tangled rhetoric.

What is in a name? Virtual Reality vs. hyperreality

Much of the above can be (and indeed has been) applied to the area of Virtual Reality (VR). *Virtual Reality* (a phrase which, in itself, seems contradictory) describes a environment generated by computer software. The use of VR as a term to describe the creation of realistic past environments comes under scrutiny and faces suggestions that the term ‘hyperreality’ is perhaps a better explanation for the representation of past environments [17, 3]. The term *hyperreality*, which is still evolving in definition, was brought into use by Baudrillard’s discussion of simulation: “it is the generation by models of a real without origin or reality: a hyperreal” [5]. The term describes a functioning copy of something that never existed in the first place. Spinal Tap — the mock rock group mentioned in Section 5.2 is an example of this. They were a fictional band, played by actors, featuring in a spoof documentary-style fictional film, yet the success of the film led to them producing records and touring despite the fact that they weren’t a real band and never actually existed. Now apply this term to computer-generated representations of archaeological sites: we can interpret the evidence and create a site that never actually existed, but which has become real to us because we have generated it.

The idea of *more real than real* can be seen in the concept of living museums such as folk parks, where actors dressed in period costume portray the day-to-day life of that time. They represent a reality, but it is a reality taken out of context, where certain parts of the experience are emphasised and heightened, with the customer experiencing a conjectural world as a real one. Applying the example of the Plague house, if there were actors portraying the daily life of someone during the time the house was in use then the situation arises where the real (the remains of the house) and the fake (the actors, reproduction furniture and artefacts) co-exist. Historical information has been displayed alongside duplications and fakes in order to make people view

it as being more authentic.

This hyperreality is no more unusual than escapism in a novel or film — we seek out something that is more exciting, more dramatic and more memorable than mundane life. For a museum visitor, this may not necessarily be such a bad thing. If we are stimulated, does it matter if it is simulated? This is where the purpose of the representation comes into play — each need has a niche. To encourage people to think how life might have been in the past, it may be enough to offer them a lesser degree of reality. For the archaeologist to test and explore new hypotheses, more accuracy may be required.

Virtual Reality, virtuality, hyperreality — it is the underlying concept rather than the name that matters. Reality can be viewed as a singularity — no matter how much we try to emulate it we will always fall on the side of either hypo- or hyperreality. In spite of this, a lack of authenticity does not mean representations are useless. We must be aware of what our motivation and purpose is (see Section 5.4). Also, providing we determine what we are trying to achieve and place it in a broader context, we do not need to spend too much time worrying over nomenclature, but such a discussion helps to highlight the issues in defining realism.

5.3.2 The nature of archaeological data

One of the first discussions of the idea of Virtual Archaeology categorised the VR model as a *replacement* — a duplicate — for an original, with the fidelity of the model dictated by the dataset from which it was created [36]. However, the above section indicates that this does not necessarily guarantee realism. Also, given the wide-ranging aspects of archaeological evidence, it is improbable that we can display it all at once in any meaningful way. Gillings drew attention to the problematic nature of a *tangible referent*, the idea of an existing reality to which we can compare our model [17].

There are several issues that stem from the fact that we wish to portray this tangible referent. The first problem is that the site may no longer exist in the form in which we wish to represent it: walls tumble down, buildings crumble; paintings fade. The second problem is

the diachronous nature of archaeological sites. In general, an archaeological site consists of material remains that have accumulated over a period of time; it may be very short or may span over hundreds over years. Either way, it exists *over a period of time*; it cannot be isolated to one single instance. How, therefore, do we display this? A three-dimensional spatial representation depicts a single moment in time. To show the site evolving in a temporal manner we must add a fourth dimension — a timeline.

By way of example, suppose that the museum at the Plague house depicts life during an ordinary day in the fourteenth century. Presumably the people will have been living there for some time — the house has been built, it is full of furniture and useful items, and the people are carrying out their usual daily tasks. However, an archaeological record of that site does not focus on one particular day in the fourteenth century. The archaeologists may have sifted through twenty-first century topsoil, through twentieth century discarded trash, backwards through the waste of the preceding centuries until they reach the remains of the Plague house. These are carefully recorded, but the excavation does not necessarily stop there. They can go further back in time, and uncover older archaeology lying beneath the house. They may keep excavating until they reach the subsoil where no more archaeology is to be found. From all of this comes an incredibly detailed excavation record spanning centuries, and yet only one typical fourteenth century day is shown to the public.

It is tempting to say that a tangible original never actually existed, that there is no single, all-encompassing form that can be represented. Aside for the purely temporal, a site has many different aspects and can be viewed in many different ways.

Having said that, the process of archaeology itself is materially and culturally selective — human influence exerts itself at every level. Sites are excavated in a manner suited to each excavation director, and may vary from one trench to the next. The archaeological record depends on how the site has been dug. The interpretation of this record depends on the interpreter. The reality “exists in the interface between the site and the excavator or analysis. It is the information itself” [33].

5.3.3 Context

The context of time has already been mentioned, but this is merely one aspect that needs to be considered. Social context is important, and is something that is lacking in many computer representations due to the absence of human figures in scenes. The addition of a human presence in a representation conflicts with achieving physical reality as it requires the addition of extraneous, conjectural material. Human figures feature in many drawn archaeological illustrations, providing not only a social element, but a convenient measure of scale. However, in synthesised scenes, if we try to make our representations as hands-off as possible in the hope that this will preserve a greater degree of authenticity, adding evidence of human habitation that cannot be inferred from the archaeological record results in artistic interpretation. This is an area where a balance needs to be found. Social context is at the expense of physical accuracy. It is a dilemma that again calls purpose and motivation into question. If we want to safeguard a degree of realism then our virtual worlds are barren and unpopulated. Conversely, how can we identify with such a sterile scene? How can we view an empty room as being real when it was probably furnished and populated in the past?

This problem is not unique to computer representations, and the idea of living museums are, once more, testament to that. Archaeology is the study of material remains, but it was people who made and used these. If we neglect them, we exclude a facet of reality. If we include them, we stand accused of conjecture.

Another interesting factor is the perspective from which we view the represented scene. Often, it is the form of a fly-by where the user zooms over and around the virtual site for a panoramic view. This is not a real view if we consider that the inhabitants of the fourteenth century did not have helicopters, and so would never have seen their home from this viewpoint. Nonetheless, it is a useful form of display, not only for the public to get an overall sight, but also for the archaeologist to establish spatial relationships. Conversely, when working from an eye-level view, it is important to remember that our tall and healthy bodies of today may not bear great resemblance to a medieval peasant, so a decision has to be made as to whether we are at our

present eye-level or the eye-level of the original inhabitant.

Emotionally, it is difficult to convey realism — our response to a representation of a house during the time of Plague is unlikely to be the same as our response to the original environment. There is no danger, no fear, no worry or hope or loss in a safe representation. We do not have the threat of death and disease hanging over us when we gaze at an image or walk through a museum conveying such times. We remain distant and detached because we are not experiencing this original reality. The long list of factors that may shape our personal interpretations comes into play when we consider that we are subconsciously influenced by a wide range of experiences, unique to each individual. In a way, we look at things with our own archaeology of past events shaping each interpretation.

However, although the image may not display context, there is a way of providing it. Computer-generated scenes offer means of putting objects and places into context in a way that museums with their shelves of artefacts cannot. By providing a method of determining contextual information, representations are afforded a way of explaining the processes that led to their creation. Section 5.6 will discuss this in more detail.

5.3.4 An established reality

If we are to use computer graphics for predictive purposes - for sites that no longer exist (e.g. archaeology) or those that have not yet been built (e.g. architectural simulations) there is no real-world scene with which to compare our representation. Since we do not have a tangible original, we must try to establish that we have made the scene as close a representation as possible by using some form of quantification. If we are worried about artistic interpretation and conjecture then a method of providing information about the virtual image needs to be employed. The degree of realism we require depends on purpose — on the questions that we ask — rather than being a blanket-term for a impossible standard to which we try to adhere.

5.4 Representing for a purpose

Representations of archaeological sites can be created for a number of different reasons, and this section aims to identify some of the key reasons and show how the intended purpose of an image can influence the ways in which it is represented, and thus interpreted. It is also important to realise that even if an image is not useful in an archaeological respect it does not make it useless altogether, as some of the cases below will illustrate.

5.4.1 Representations for the archaeologist

It is likely that the primary purpose for representations created for use by an archaeologist will be for research. While their own personal interest might warrant the creation of images for aesthetic reasons, computer-generated representations provide the archaeologist with a means for exploring the past and testing new hypotheses in a safe and controlled manner (see Section 5.7).

Establishing spatial relationships

Three-dimensional computer graphics provide a convenient means for the archaeologist to gain spatial awareness of a two-dimensional site record and to establish spatial relationships within that environment. For the visualisation of the site layout and the distribution of artefacts a CAD wireframe model or a GIS model may well be sufficient. Realism is not a priority, but the ability to quickly and easily navigate the virtual world is essential. This is the most objective form of representation as only the details from the excavation record are likely to be displayed.

Investigating new hypotheses

A computer-generated representation of the archaeological evidence provides the archaeologist with the chance to manipulate variables in

the virtual environment in a way that cannot be done in the real environment. This is explained in greater detail in Section 5.7. The main thrust here is that the archaeologist is choosing to emphasise their own ideas, so subjectivity is a factor. The end product is intended to either prove or disprove a particular hypotheses — the archaeologist may not necessarily be interested in exploring other avenues of thought. The representation may therefore have a tendency to be biased towards supporting the archaeologist's views. This is not a bad thing; indeed, it is the very nature of archaeology to provide new theories, and this gives the archaeologist a chance to visualise their ideas where beforehand they may have only written about them.

5.4.2 Representations for the computer scientist

There is the possibility that the archaeological representation has been created for the purpose of providing an application for the use of new computer graphics techniques. In this instance there is the danger that the archaeological aspects of the representation may be neglected, with the emphasis placed on demonstrating advances in computer graphics research [40]. The image may well be meaningful, but the motivation behind it might limit its usefulness for the purposes of archaeological research. However, it could prove highly desirable as a demonstration of cutting-edge graphics.

5.4.3 Representation as advertising

Cynical as it might seem, there is always the chance that visualisation might have been carried out for the primary reason of advertising. Such projects in the past have been the results of sponsorship by large commercial organisations, concentrating on archaeology as an area rich in media attention and public interest and likely to provide public relations points [32]. In these cases, research was not the main aim of the project, but it certainly did not harm the archaeologists who benefited from the money and resources. However, with an emphasis on advertising and less direct control from the archaeological side of things, and with well-known archaeological examples (rather than

unexplored datasets) being chosen as subjects for the work, there was little new insight to be gained.

5.4.4 Representations for the public

Archaeological representations aimed at the public fall into two broad areas: education and entertainment. Increasingly, the boundary between the two is blurring as the idea of making education enjoyable — or making entertainment educational — takes hold.

Education

The move towards increased use of audio-visual displays for heritage purposes has led to a demand for high-quality, comprehensive representations that inspire as well as instruct. Computer-generated representations designed for this purpose run the risk of providing a model that is visually appealing at the expense of being informative, but in a museum context the primary information (in the form of the museum displays) reinforce the didactic objectives of the virtual environment.

This is an area where representations can vary greatly depending on the demographics of the intended audience. Nonetheless, the main objective is to provide information to the non-specialist.

Entertainment

The growth of media archaeology — television and the Internet being the main examples — has led to representations based on the fact that they look stunning and slick and attract the attention of the viewer, thus pushing up the ratings or the hit-counter. Indeed, it is becoming rare nowadays to watch an archeological documentary on television that does not contain some sort of virtual representation (generally touted in the television listings as “state-of-the-art computer graphics”). These programmes may run the gamut from the educational documentary to the sensationalist suppositions, but they work from the same motivation — attract as many viewers as possible. Like the case for advertising mentioned above, this is not always particularly

useful, but it does provoke interest in the subject and can therefore be considered as indirectly educational.

5.4.5 Fit for purpose

The above cases highlight the need for clearly establishing the motivation behind creating a representation. As we expect, representations must be tailored towards their intended user. This also gives us greater insight into the path that led to the specific interpretation portrayed.

5.5 Misinterpretation

“In presenting a very visual and solid model of the past there is a danger that techniques of visualization will be used to present a single politically correct view of the past, and will deny the public the right to think for themselves.”
Miller and Richards, 1994.

Given the problems in archaeological representation with defining realism, the motivation of the work and the inclusion of inference, the possibility of misinterpretation is high. As Miller and Richards remarked in 1994, “there is little, if any quality control for computer graphics and they are not subject to the same intense peer review as scientific papers” [32].

5.5.1 Different outcomes from the same evidence

One way of demonstrating the potential of myriad interpretations derived from a single dataset is to ask a group of people to independently represent the same scene. A project of this nature was undertaken by Hodgson in 2001. He advertised on the web site of the Association of Archaeological Illustrators and Surveyors in the UK¹ appealing to people interested in creating a representation of a site. Each participant was given the same design brief containing information about the

¹<http://http://www.aais.org.uk/>

site (Dewlish Roman villa), the purpose of the end product (a popular publication about the site for the “interested layman”), photographs, plans and sketches from the excavation report, and a choice of any medium. Additionally, participants were informed that “the illustrator is not bound to use only that material which is contained in the brief. Use of reference material for costume, furnishings, implements etc. is perfectly valid; illustrators are only asked to keep a record of any additional sources used for the ‘debriefing’ questionnaire” [25]. At the time of print, Hodgson was compiling and analysing the completed representations. The intended outcome is an overall picture of “how the reconstruction process functions” [25], taking into consideration the content of the completed images, the illustrator’s responses to the questionnaire, and the audience’s reaction. We await the results with interest. Projects of this type are an excellent way of highlighting how many interpretations can be derived from the one source, and are useful in emphasising to the public that a single representation cannot be taken for granted.

5.5.2 Seeing what we want to see

Conversely to the above, given the previous discussion in Section 5.3 where the viewer brings their own experiences to their viewing of an image, there is the chance that the synthesised scene becomes a type of Rorschach test where the viewer projects their thoughts onto the image and sees it as it was never actually intended to look. It is not just the public who are at risk of this. The archaeologist may focus on an insignificant detail and magnify its importance, losing the overall impact of the image (and the idea of alternatives) by concentrating all their attention on a singular aspect. Likewise, a computer scientist may marvel over a great graphics effect, which to them is of the utmost importance to the scene (a scenario not unlike graphics researchers watching animated feature films and focusing all their attention on how amazing the rendered hair looks, rather than on the storyline).

5.5.3 Reducing misinterpretation

Since the chances of misinterpretation are significant, a method of reducing it must be implemented. Information about the dataset and the interpretative process is key, and informing the viewer of the methods and decisions taken to produce a synthesised scene allows them to place the image in context. Section 5.6 discusses how this can be achieved.

5.6 Setting standards

The previous sections have demonstrated how misinterpretation of an image might arise. By incorporating information pertaining to a VR world or computer-generated image, each synthesised scene can be analysed and compared, allowing the viewer to determine information regarding decisions over representations of incomplete geometry and inclusion of artefacts, and also information image attributes such as rendering quality or resolution. If a level of standardisation can be reached, representations of past environments can be of more use to both the archaeologists and the public. If alternatives are provided, or if the image is treated with the same scrutiny as a documented source, we can limit the danger of awarding graphics more influence than they should actually hold. This section discusses the application of metadata, tagging and standardisation to computer graphics as a means of increasing contextual information and reducing misinterpretation.

5.6.1 Metadata

To avoid misinterpretation due to a lack of information and to place an archaeological representation in its appropriate context, some form of description of the dataset and the decisions involved in creating the representation is required. Metadata — data about data, or information about information — is one method of providing this. The idea of providing metadata has flourished with the advent of the Internet, although the phrase has been around since the 1960s. Metadata exists in familiar forms, such as a bibliography that tells us when a book

was written and by whom it was published, or on a map where we can check the scale and the date of survey. The idea of metadata is not dissimilar to treating the image in the same manner as a documentary source. Any source used for historical/archaeological evidence is subject to scrutiny, and the nature of the document is always questioned — who wrote it, when, where, why and in what context. Questioning a computer-generated image the same way makes good sense, and metadata provides the answers.

There are two issues about metadata that need to be standardised: the format of the metadata (the syntax, file format, etc.) and the actual metadata required, which is subject-area specific. A number of standards for metadata exist, and it has proved especially useful in archaeology when documenting and cataloguing material for digital archives. Applications and standards for metadata for digital archives are well-documented (metadata about metadata!) and published guides (such as “Creating Digital Resources for the Visual Arts: Standards and Good Practice” [21], published by the UK’s Arts and Humanities Data Service) seek to make digital archives more accessible.

Syntax

Standards for metadata may be highly specific, such as MARC (Machine Readable Catalogue) used for a traditional library or the FGDC (Federal Geographic Data Committee), or may be simple and in widespread use, such as the Dublin Core standard [26], an international and interdisciplinary initiative, which is commonly used to describe online resources on the Internet. This is an attempt to standardise the information located in the META tags within the HEAD tags of an HTML (HyperText Markup Language) file, thus hopefully returning more comprehensive and less random information from an Internet search. There are fifteen core elements - structured categories such as title, author, date, source, etc.

Metadata for archaeological purposes

Metadata can be represented in several different syntaxes, and Dublin Core is applicable to virtually any file format. As mentioned above, HTML is one commonly known form, but XML (eXtensible Markup Language) is perhaps a better-suited format for our purposes, allowing the creator to define the actual semantics [13]. Like HTML, XML uses tags and attributes to delimit data, but unlike HTML it is not a fixed format, so elements can be customised to suit the user's purpose. A customised markup application can therefore be created for exchanging information in a particular subject area, and can the same information can be read across different operating systems.

Work on online image retrieval has focused attention on the need for metadata for images, but it is the setting of information standards for VR that is perhaps more useful. The AHDS has determined the Core metadata required for VR (Table 5.1) and has also compiled a checklist of information that should be included when documenting virtual reality models (Table 5.2) [42].

5.6.2 Alternative representations

In addition to the provision of metadata, providing multiple interpretations of archaeological data is another means of drawing attention to the subjective aspects of representation. In order to provide a method of portraying the tentative nature of interpretations, Roberts and Ryan identified four modes of operation that allow the generation of different views of a site representation [40]. Each of these four modes allow alternative interpretations of the same data to be represented in a VRML world:

- *Require New* hinges on the archaeologist providing multiple instances of the world, including the description of possibilities and arrangements that could have existed. The client then downloads each model they require.
- *Switch Change* has multiple interpretations in one document, and the client (who requests this single document) can switch between various configurations.

INFORMATION TYPE	SCOPE NOTE
Title	The name of the bubbleworld, panorama or virtual reality model.
Survey index	The identification number/code used internally for the project.
Description	A brief summary of the main aims and objectives of the project for which the model was developed and a summary description of the model itself.
Language	An indication of the program language(s) in which interactions take place in the virtual reality model, e.g. Javascript.
Type	The type of resource, e.g. three-dimensional model, interactive resource, collaborative virtual environment.
Format	The data format of the resource, e.g. VRML 97.
Subject	Keywords indexing the subject content of the model. If possible these can be drawn from existing documentation standards, e.g. for archaeology the English Heritage Thesaurus of Monument Types, the mda Archaeological Objects Thesaurus. If a local documentation standard is used a copy should be included with the data set.
Temporal Coverage	The time period covered by the virtual reality model.
Spatial Coverage	Where the model relates to a real world location, give the current and contemporary name(s) of the country, region, county, town or village covered by the model and map co-ordinates (e.g. in the UK national grid).
Administrative area	Where appropriate give the District/County/Unitary Authority in which the model lies.
Country	Where appropriate, give the country to which the model relates.
Date	The dates of the first and last day of the virtual reality modelling project.
Creator	The name(s), address(es) and roles of the creator(s), compiler(s), funding agencies, or other bodies or people intellectually responsible for the model.
Publisher	List details of any organisation that has published the model including the URL of on-line resources.
Depositor	The name, address and role of the organisation or individual(s) who deposited the data related to the virtual reality model.
Related archives	References to the original material from which the model was derived in whole or in part from published or unpublished sources, whether printed or digital. Give details of where the sources are located and how they are identified there (e.g. by file name or accession number).
Copyright	A description of any known copyrights held on the source material.

Table 5.1: AHDS Core metadata for virtual reality models, from <http://ads.ahds.ac.uk/project/goodguides/vr/appendix2.html>.

INFORMATION TYPE		INFORMATION TYPE	
Project documentation		Application development	
Project name	✓	Model type	✓
Survey index	✓	Application platform	✓
Description	✓	Hardware platform	✓
Bibliographic references	✓	Authoring software	✓
Subject keywords	✓	3D drawing tools	✓
Spatial coverage	✓	3D scanners	✓
Administrative area	✓	Animation scripts	✓
Country	✓	Sound clips	✓
Date	✓	Image format	✓
Subject: discipline	✓		
Subject: type	✓	Delivery platform	
Subject: period	✓	Operating system	✓
Creator	✓	Browser	✓
Client	✓	Plug-in / viewer	✓
Funding body	✓	Scripting language	✓
Depositor	✓	Hardware platform	✓
Primary archives	✓	Network connection	✓
Related archives	✓	Target frame rate	✓
Copyright	✓		
Target Audience		Description of archive	
Audience	✓	List of all file names	✓
Mediator	✓	Explanation of codes used in file names	✓
Standard	✓	Description of file formats	✓
Interactivity type	✓	List of codes used in files	✓
Interactivity level	✓	Date of last modification	✓
Typical Time	✓		

Table 5.2: A check-list for documenting virtual reality models
<http://ads.ahds.ac.uk/project/goodguides/vr/appendix2.html>.

- *Functional Change* — a world is generated with objects and characteristics that can be easily altered (for example, building height) by means of a control panel.
- *Program Run* works on the basis of providing a version of the data and program to the client who can run it on their own machine and generate different views and models.

This idea of providing alternatives is nothing new — archaeological illustrators have often raised similar issues about drawings, with pleas for “more partial and avowedly tentative reconstructions, with more alternatives offered when, as so often, the evidence is indecisive” [43]. The ability to identify which parts of a scene are conjectural and which come from the excavation record would undoubtedly be useful. A tool for switching on and off attributes of a scene, or for rearranging furniture, or altering lighting allows greater flexibility for the archaeologist’s investigations, and a chance for the public to see how the decision process works regarding interpretation.

5.6.3 Preserving information

Another important aspect is to store the both the synthesised scenes and the associated information in a relevant technical form. A digital resource is of no use if it cannot be read, and changes in technology should be considered. An example of digital mismanagement is the BBC’s (British Broadcasting Corporation) Domesday Project of the 1980s. The original Domesday Book was a survey commissioned by William the Conqueror in 1085AD. Upon completion in August 1086 it contained records for 13418 settlements in England, providing a detailed record of life in Britain in the eleventh century. The BBC launched a project involving schoolchildren all over the UK, asking them to record their life and lands 900 years after the original Domesday Book. The information gathered by the schoolchildren was stored on 12 inch video disks which today can no longer be read as the technology is obsolete. By contrast, the original eleventh century Domesday book is still legible (albeit in handwritten Latin). This serves as a warning that advances in technology should be considered and storage of information should be carefully planned and monitored.

5.6.4 Standardisation

In a subject such as archaeology where alternative explanations may be equally plausible, the option to move between a number of different interpretations is undoubtedly important. A form of standardisation is most desirable, but rather impractical given the diverse scope of the subject. In the same way that the archaeological evidence on an excavation needs to be recorded as thoroughly as possible, so too does the process used to create the computer-generated representations so that all the factors might be displayed, allowing the user to make up their own mind based on the supporting material. If we strive to provide information about the underlying decisions taken in the creation of our work then our virtual worlds have the potential of being meaningful, useful pieces of information.

5.7 Developing new hypotheses

² As well as providing a good means of teaching the public about past environments, computer representations of archaeological sites can also be of use to the archaeologist. The immediate advantage is the ability to provide a spatial framework. For an archaeologist used to working with two-dimensional site plans which must be correlated with section drawings and level readings, a three-dimensional view of their site provides a framework for determining spatial relationships. This might well prove useful for establishing hypotheses about function of certain areas, or artefacts found therein.

However, computer graphics also offers the chance to explore another aspect of past environments that does not appear in the archaeological record: light.

²This section is by Duncan Brown, Southampton City Heritage, UK. The papers that form the basis of this discussion are included in the Appendix.

5.7.1 New ideas from light and colour perception

Light is something we take for granted, that we create through the movement of a switch. Light is not something archaeologists can recover or record and consequently its importance is rarely considered in interpretations of past ways of living. Light is fundamentally important, yet at present it is rarely considered as a medium for comprehending how people behaved or how they perceived their environment. The hard archaeological evidence is unrevealing. Windows remain as evidence for the way of introducing light; lamps, lanterns and candlesticks for the means of its creation. Yet there is a more subtle way of approaching this problem and that is to look at the objects we find, the clues to the ways in which buildings were decorated, in an effort to understand what sort of environment past peoples created for themselves. Their perceptions can be revealed by the appearance of the objects they used.

In searching for insights into the ways past people perceived their environments and the objects they interacted with it is important to consider how they illuminated their lives. Even more crucial is the relationship between light and colour.

We are not used to looking at objects in conditions where light sources are dim, generally of a red cast and also moving. Nor indeed, once they have been removed from the ground, are we used to viewing artefacts only in daylight. In our age, the variations in the provision of light are beyond our immediate experience and understanding. The ways in which we view, perceive and understand objects is governed by current lighting methods (electric light and large windows) but in order to understand how objects were viewed and understood in the past we must consider how they were illuminated.

It is essential that any system for reconstructing and visualising ancient environments must be as accurate as possible and flexible, allowing archaeologists to alter the scene parameters in order to investigate different hypotheses concerning the structure and contents of a site [10].



Figure 5.5: Examples of medieval pottery

5.7.2 Case study: Medieval pottery

The basic premise is that the colours of medieval pots are related to the lighting conditions that medieval people were accustomed to. Some pots are brightly coloured and highly decorated, others are dull (Figure 5.5). This is related in part to vessel function but must also reflect the intended place of use and thus variations in lighting conditions. Furthermore, pottery colours may actually reflect a typical absence, rather than presence of light, or at least lighting at much lower levels than we are used to. This case study therefore considers the ways in which medieval interiors were illuminated and how lighting conditions might affect the ways in which objects were perceived and designed.

Provision of light

As has already been inferred, archaeological evidence for the creation of light is relatively rare. The ceramic evidence is perhaps most commonplace on excavations and a variety of lamps and candleholders are known throughout the medieval period. These are generally portable types and are consequently small in size. It is difficult

to envisage such objects being used to illuminate whole rooms. It is known, of course, that torches were extensively used and wall-brackets for these survive. Another source of light, and possibly a very important one, must have been the fires and braziers that were used to heat rooms. On this evidence it seems that the medieval interior must have been a flickering, smoke-beset world and perhaps this explains the bright colours used to decorate medieval objects and indeed rooms. The evidence for those is drawn primarily from manuscript illuminations and paintings, where furniture is often shown brightly painted, and walls display rich hangings.

It may, however, be the case that most of these things were not meant to be seen at their best in artificial light. The hours of daylight regulated pre-industrial life and provided the fundamental means of illumination. In the present day, sunlight is almost irrelevant to the conducting of our lives; houses, offices, shops, factories are almost all permanently lit by artificial means. In the medieval period the sun was probably viewed as the only constant source of light and it is in the architecture that the best evidence for lighting can be found. Big windows provided lots of light but given the limited availability of window-glass they also brought draughts. Windows also created a security risk, as is shown most obviously in castles. Indeed the largest windows may be found in ecclesiastical buildings, those which one might presume to have been least threatened. This provokes the thought that the provision of light on through such grand openings was as much a signal of devotion to God as the building of the entire edifice. Light and holiness seem almost to be related. The way light was used within a church or cathedral may also, perhaps, reflect the controlling aspects of ecclesiastical architecture.

Medieval pottery, however, was used more frequently in a domestic environment and the windows in houses are therefore more pertinent to this discussion. Here, window size might be related to status. The windows in surviving English peasant houses are generally small, leading to the conclusion that warmth was more important than light, perhaps because the rural lifestyle was in any case governed by the rising and the setting of the sun. Rural manor houses, and the homes of late medieval yeomen farmers exhibited larger windows, as did medieval town houses. In all instances the window provided light for

the carrying out of daily activities such as weaving and sewing. In towns, where jettied upper stories often brought houses within a few feet of each other, windows had to be larger but they also allowed those sitting at them to communicate with their neighbours, as well as people in the street. The window therefore played an important social role as well as a domestic one. It is clear that different dwellings gave different lighting conditions. One might therefore expect the most brightly coloured objects to be associated with the best lit settings and this is, to some extent true. The most highly decorated types of pottery, for instance, are not found at sites of the lowest status. However, that does not tell us how such objects were perceived by those who used them but simply how much light might have been available to see them by. It is an understanding of medieval perception that is being sought here and the next section considers ways of looking for that, if not necessarily finding it.

Colour in the medieval period

The pottery used in England throughout the medieval period was mostly earthenware and may be divided into white-firing and red-firing types. Pots made from white-firing clays appear white to buff in colour but were usually covered, completely or partially in a lead glaze. Clear lead glazes give a yellow to amber appearance to white pottery but the addition of copper creates green, which can vary from dark through olive to bright, apple hues. In general, the most brightly coloured medieval pottery was made from white-firing clays. Vessels made from iron-rich, red-firing clays range in colour from brick red to dull brown, when fired in oxidising conditions. Clear lead glazes can give a red-orange appearance to red wares while lead with copper glazes produce a variety of greens generally less vibrant or consistent in colour than those seen on white earthenware. Red clays fired in an oxygen-free atmosphere range in colour from grey through to black and in such conditions an ordinary lead glaze will turn slightly green or greenish-clear. In short, medieval pots were given a variety of colours, including white, red, brown, yellow, orange, green, grey and black. The appearance of vessels was also enhanced, and the colour effect slightly altered, by decoration such as painted lines, coloured

slips and applied clays.

So what? may be the first question that springs to mind after reading that. What, if anything, does this information tell us about the medieval period? It is possible to quantify medieval pottery assemblages by colour, but is that a useful thing to do? Did the colours of medieval pots have any meaning?

The range of colours produced in pottery at least tells us something about ceramic technology. Aspects of kiln construction and management, clay exploitation and mineral use are all illuminated through an understanding of how pots were made in the colours that they were. Also revealed is the importance of colour in medieval society. The most colourful and highly decorated vessels were jugs, normally identified as table ware, used in the serving of liquids, especially wine. It is possible that these were used at high profile dinners. The dullest vessels, unglazed and undecorated, were those used in low-visibility activities such as cooking. Colour was therefore an important element in the culture of display that prevailed in medieval society, at least in its upper echelons. In these terms it is useful, therefore, to consider quantifying pottery assemblages by colour, as the quantities of highly decorated and plain pots might be related to status: high numbers of gaudy vessels might indicate a household where presentation was important.

It may be possible to put meanings to colours but one might be wary of doing so in the case of pottery. After all, the range of colours is partly determined by the available technology, in terms of glaze and mineral use and kiln structure. Such a view, of course, excludes the element of human decision-making and if we accept that medieval potters gave pots the colours their customers wanted then we must ask why it was that those colours appealed. An analysis of colour symbolism may not appear very helpful, however. The most common pottery was coloured red-brown. In medieval art, brown was associated with mourning and red with God, while green, the most common colour for glazed pottery, signified the holy spirit and therefore also the bishops. Yellow was regarded as a substitute for gold, which denoted heaven [6].

The next step, of course, is to ask whether medieval potters were fully

aware of the symbolism in art. The same colours have meanings in folk-lore that are not entirely consistent with those of medieval Christianity. Brown is of course the colour of the earth and its association with mourning is perhaps related to the notion that humans are earthly beings whom, in death, revert to clay. Red is the colour of blood and, perhaps more pertinently with regard to pottery, of fire, and in its positive manifestation signifies life, love, warmth and passion. Its negative aspect is that of destruction and war. Green was also related to life, through its association with plants and the shoots of springtime as well as water. It thus symbolised hope and longevity [6]. There is more to go on here, and it is easy to create associations between pots of certain colours and the offering, presumably through their contents, of valued qualities such as life, warmth and a long life. Yellow is more of a problem, however, as in the middle ages it was the colour of envy. It may be safer, therefore, to look upon it as a substitute for gold, and here is the crux of the matter — it is we who are now doing the looking, thus creating our own understandings of the colours we see.

Medieval pottery

Issues of colour and perception in pottery have already been raised and it is clear that we must understand the relationships between pottery and those who used it. In the first place, it is not always clear whom those users were, or at least it is clear that in many medieval households several different users were involved. In the townhouse of a wealthy burgess, such as that which provides the setting for the case study set out below, there was a hierarchy of individuals and their roles. It seems reasonable therefore, that different types of vessel were also fitted into that scheme. At the lowest level, ceramically, we may place vessels used in food preparation. Medieval cooking pots were simple forms, cheaply made and cheap to buy. Their use would have been confined to the kitchen and scullery areas and cooking pots would probably not have appeared at high table or been considered for display. These vessels were of unglazed earthenware, usually buff, red-brown or grey in colour. Unglazed, or partially glazed, jugs were also produced and these may have been used in the same areas of the

house. It is most likely that domestic servants used pottery bought for storage and cooking and, as far as the householder was concerned, the appearance of that pottery may have been largely irrelevant. There seems, therefore, to be little perceptual influence in the acquisition and use of kitchenwares and this is borne out by the fact that in most excavated medieval assemblages cooking pots and jars are invariably of local origin [7] and have a uniform appearance.

Jugs, however, the medieval vessel type we commonly characterise as tableware, were derived from a wider variety of sources, even at humble farmsteads and issues of taste and display must be considered. Rich glazes and elaborate decoration suggest use in public ways, certainly outside the confines of the kitchen area. A penchant for display is known to have permeated most levels of medieval society and ceramics would have been a cheap way of assuaging such a requirement. It is important to recognise that pottery could never compete with glass or metal as a medium for showing off but it is easy to identify highly decorated vessels as intended for use at table. Here, aspects of colour become more important perhaps. The use of glaze allows a greater degree of consistency in the colouring of pottery and it must be assumed that this was done deliberately to appeal to the tastes of prospective customers. If, therefore, a domestic assemblage yields tableware of a certain range of hues, then it seems reasonable to suggest that those were the colours that the users preferred.

The conclusion, therefore, is that the jug is a suitable focus for research into colour and perception in medieval pottery, at least at this initial stage. Jugs present the widest range of shapes, decorative designs and colours of any ceramic form in medieval assemblages; were probably used in a greater variety of domestic situations and might to a higher degree reflect the tastes of consumers.

Putting things in context

The first question that needs to be answered in this consideration of colour and perception must be this; what did medieval pots look like in their original setting? That question has been addressed by research conducted at the Department of Computer Science at Bristol University, where a computer model of the hall of a medieval town house has



Figure 5.6: Photograph of the Medieval Merchant's House, Southampton, UK.

been constructed. The model is based on the Medieval Merchant's House museum in Southampton, a half-timbered structure renovated by English Heritage as accurately as possible to represent a 13th century dwelling of some economic status (Figure 5.6).

Computer modelling offers a flexible approach to investigating hypotheses regarding colour and lighting. It is not possible to occupy the actual building, light a fire, position candles and torches and record the results and if it were it is still not feasible to remove the existing fireplace and chimney-breast. In a computer-generated environment we are able to remove extraneous structures, change the size and colour of wall-hangings, even change the colours of the walls. We can increase or diminish the size of the hearth and alter its position and we can move other illuminants around as we wish, placing them on furniture or anywhere on the walls. It may also be possible to recreate vision defects, such as shortsightedness, that were left uncorrected in the 13th century.

Figure 5.7 shows a view of the hall with the jugs placed on the main table. This scene is an accurate depiction of the actual building, the wall hanging and the furniture are pieces currently on display. The three jugs are a local redware baluster, red-brown in colour with a partial greenish-clear glaze, a Dorset white ware jug with an overall yellow glaze, decorated with dark brown applied vertical stripes and lines of pellets and a Saintonge white ware vessel with an overall bright green glaze. They are contemporary, dateable to around 1270-1300 and it is possible that similar vessels could have been in use in a

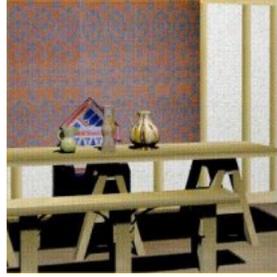


Figure 5.7: Computer-generated model of the hall of the Medieval Merchant's House, French Street, Southampton, lit by generic approximation of daylight (created by Ann McNamara, now of Trinity College, Dublin, Ireland).

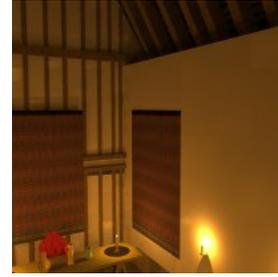


Figure 5.8: Computer-generated model of the hall of the Medieval Merchant's House, French Street, Southampton, lit by candlelight (created by Patrick Ledda, University of Bristol, UK).

Southampton household at the same time. The rendering in Figure 5.7 is lit with generic lighting, comparable with daylight, rather than a specifically modelled light source. Figure 5.8 shows the same scene from a viewpoint on the gallery above the hall, lit with candles.

These computer-generated images (Figures 5.7 and 5.8) serve to illustrate the value of this approach. This work has also confirmed the necessity of understanding how objects looked in their original contexts if we seek any insight into past perceptions. It is already apparent that many pots would have looked brightest when lit from above. Only the top half of some jugs are glazed and decorated, and this is perhaps indicative of how they were illuminated in use, perhaps by daylight through windows or perhaps from torches hung on walls. The purpose of this research is the revelation of detail such as this and there is great potential to go deeper into medieval ways of living.

In searching for insights into the ways past people perceived their environments and the objects they interacted with it is important to consider how they illuminated their lives. Even more crucial is the relationship between light and colour. Colours will change in appearance according to the types of light source present; yellow, for instance is especially affected by the RGB factor of certain illuminants. The recreation of medieval lighting conditions is therefore seen as a vi-

tal step in comprehending attitudes to colour, and eventually perhaps, shape and decoration. If there is any symbolic meaning in the use of colour on pottery then this might be revealed through an exploration of medieval perception, through the recreation of a medieval environment. The modelling of a realistic environment through the application to computer graphics of computer science and psychophysics, is perceived to be the most far-reaching and flexible way of exploring human perceptions in the past.

This research into colour and light has shown how easy it is for our own preconceptions to intrude into the ways we view archaeological objects or sites. Our research is intended to reveal more about medieval perceptions by investigating the context of colour. This has led to a re-evaluation of how we ourselves look at pottery, the chosen medium for our research. One result of our inquiries may be to suggest new, or additional ways of recording ceramic assemblages. The aim is to find methods of analysis that could take us beyond those typical questions of chronology and provenance. If the aim of archaeology is to provide insights into the lives of past individuals, communities and cultures then we need to show a greater respect for the things they have left behind and attempt more refined ways of understanding them.

This project is an attractive mix of archaeological intuition and philosophy together with hard science. The proposition is even more exciting when viewed as a voyage of discovery which is bound to open up new lines of enquiry and thought.

5.8 Summary

This chapter has discussed the problems of representation and interpretation of archaeological datasets. Visualising a past environment is fraught with difficulties from the outset. The archaeological record itself can be subjective, the level of realism must be defined, motivation and audience should be established, the potential for misinterpretation needs to be minimised, and standards have to be set. If the archaeological information is placed in context and the decision-making process is outlined then the synthesised scenes can become

meaningful and useful. It is only then that we can confidently begin to explore and investigate these virtual environments as a means of understanding the past.

5.9 Slides



Representation and Interpretation

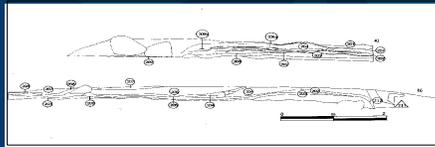
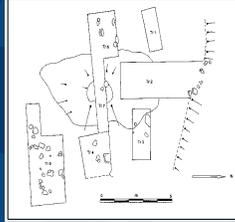
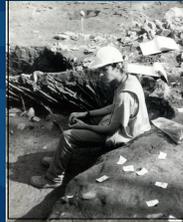


- Archaeological illustration
- The idea of realism
- Representing for a purpose
- Misinterpretation
- Setting standards
- Developing new hypotheses

The slide has a dark blue background with a subtle geometric pattern. The title "Representation and Interpretation" is in a white, sans-serif font. The SIGGRAPH 2002 logo is positioned in the top right corner. The list of topics is in a white, sans-serif font, with each item preceded by a small orange circle.

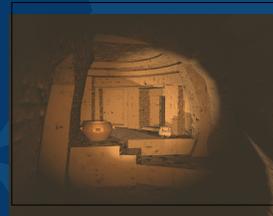
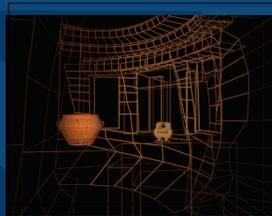
Recording sites

SAN ANTONIO
SIGGRAPH
2002



Visualising the data

SAN ANTONIO
SIGGRAPH
2002



From dataset wireframe rendering

Case Study: Depicting Stonehenge

SAN ANTONIO
SIGGRAPH
2002

Fourteenth century through to the present day



Terms and Concepts

SAN ANTONIO
SIGGRAPH
2002

- Reconstruction – an objective rebuilding
- Representation – a subjective interpretation
- What do we mean by 'realism'?
- Virtual Reality vs. hyperreality

Defining realism

SAN ANTONIO
SIGGRAPH
2002

- Everyone has their own version of reality - we bring our own years of experience to all that we view.
- **Perceptual realism:** when a generated scene evokes the same response as the original scene?

The tangible referent

SAN ANTONIO
SIGGRAPH
2002

We try to emulate a 'tangible referent'.

PROBLEM: we do not have a reality with which to compare our scenes.

How do we choose which aspects of a multi-faceted site to represent?

Context

SAN ANTONIO
SIGGRAPH
2002

Representations must be placed in context:

- Temporal
- Social
- Emotional



Representing for a purpose

SAN ANTONIO
SIGGRAPH
2002

- FOR THE ARCHAEOLOGIST – establishing spatial relationships, investigating new hypotheses
- FOR THE COMPUTER SCIENTIST – new graphics techniques
- FOR ADVERTISING – PR for companies
- FOR THE PUBLIC – educational, entertainment

Misinterpretation

SAN ANTONIO
SIGGRAPH
2002

- A single site can be interpreted in many different ways (e.g. Dewlish Roman villa project).
- We see what we want to see – our synthesised scenes work in the same way as Rorschach inkblot tests.
- The decision-making process that led to a particular interpretation needs to be documented.

Setting standards

SAN ANTONIO
SIGGRAPH
2002

- **METADATA** – information about information. Need produce a standard format and standard input.
- **PROVIDING ALTERNATIVES** – offering more than a single interpretation. Need ways to present alternative representations.
- **PRESERVING INFORMATION** – keeping data accessible. Beware the advances of technology!

Developing new hypotheses

SAN ANTONIO
SIGGRAPH
2002

Computer graphics offers us a chance to test ideas in a safe and controlled manner.

We are no longer limited to visualising data in two dimensions.

We can recreate past environments and change the variables.

Making it meaningful

SAN ANTONIO
SIGGRAPH
2002

We must consider:

- The questions that we want answered
- Context
- Purpose
- Target audience
- Supporting information

...in order for our images to be useful.

Case study: Medieval pottery

SAN ANTONIO
SIGGRAPH
2002

The premise is that the colours of medieval pots are related to the lighting conditions of that time.



Case study: Medieval pottery

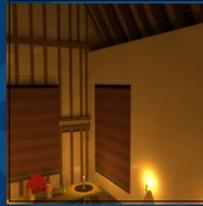
SAN ANTONIO
SIGGRAPH
2002



Medieval merchant's house, Southampton, UK.

Case study: Medieval pottery

SAN ANTONIO
SIGGRAPH
2002



Computer representations

Summary

SAN ANTONIO
SIGGRAPH
2002

In order for our computer-generated representations to be useful, additional information about them needs to be provided.

It is only when our synthesised scenes have been placed in context and are fully documented that they become meaningful.



Acknowledgements

Many thanks to Bob Stone, John Hodgson, Ramesh Raskar and Northern Archaeological Consultancy (UK) for their helpful input and permission to reproduce images.



Further information

<http://www.cs.bris.ac.uk/Research/Graphics/>

Acknowledgements

Many thanks indeed to Bob Stone, John Hodgson and Ramesh Raskar for providing some of the images and information referred to in this

chapter, and for doing so in such willing and helpful ways.

Bibliography

- [1] L. Adkins and R. A. Adkins. *Archaeological Illustration*. Cambridge University Press, Cambridge, UK, 1989.
- [2] R. Arnheim. *Art and Visual Perception*. University of California Press, Berkeley, CA, second edition, 1974.
- [3] J. A. Barcelo, M. Forte, and D. H. Sanders. *Virtual Reality in Archaeology*. Archeopress, Oxford, UK, 2000.
- [4] J. Bateman. Immediate realities: an anthropology of computer visualisation in archaeology. *Internet Archaeology*, 8, 2000.
- [5] J. Baudrillard. *Simulacra and Simulations*. University of Michigan Press, 1994.
- [6] U. Becker. *The Continuum Encyclopedia of Symbols*. Continuum, New York, NY, 2000.
- [7] D. H. Brown. Pots from houses. *Medieval Ceramics*, 1997.
- [8] D. H. Brown and A. Chalmers. Light, perception and medieval pottery. Unpublished.
- [9] K. A. Robson Brown, A. G. Chalmers, T. Saigol, C. Green, and F. d'Errico. An automated laser scan survey of the upper palaeolithic rock shelter of cap blanc. *Journal of Archaeological Science*, 28:283–289, 2001.
- [10] A. Chalmers, S. Stoddart, J. Tidmus, and R. Miles. Insite: an interactive visualisation system for archaeological sites. In J. Huggett and N. Ryan, editors, *Computer Applications and Quantitative Methods in Archaeology 1994*, pages 225–228. BAR International Series 600, Archeopress, 1995.

- [11] A. G. Chalmers, A. McNamara, S. Daly, K. Myszkowski, and T. Troscianko. Image quality metrics. In *SIGGRAPH 2000 Course #44*. ACM SIGGRAPH, July 2000.
- [12] A. G. Chalmers, A. McNamara, S. Daly, K. Myszkowski, and T. Troscianko. Seeing is believing: Reality perception in modeling, rendering and animation. In *SIGGRAPH 2001 Course #21*. ACM SIGGRAPH, August 2001.
- [13] World Wide Web Consortium. Extensible markup language (xml). URL <http://w3.org/XML/>.
- [14] G. Currie. *Image and Mind*. Cambridge University Press, Cambridge, UK, 1995.
- [15] R. Daniels. The need for the solid modelling of structure in the archaeology of buildings. *Internet Archaeology*, 2, 1997.
- [16] K. Devlin and A. Chalmers. Realistic visualisation of the pompeii frescoes. In Alan Chalmers and Vali Lalioti, editors, *AFRIGRAPH 2001*, pages 43–47. ACM SIGGRAPH, November 2001.
- [17] M. Gillings. Engaging place: a framework for the integration and realisation of virtual-reality approaches in archaeology. In L. Dingwall, S. Exon, V. Gaffney, S. Laflin, and M. van Leusen, editors, *Computer Applications and Quantitative Methods in Archaeology 1997*, pages 187–200. BAR International Series 750, Archeopress, 1999.
- [18] E. H. Gombrich. *The Image and the Eye*. Phaidon Press Limited, London, UK, 1999.
- [19] N. Goodman. *Languages of Art*. Hackett Publishing Company, Indianapolis, Indiana, second edition, 1976.
- [20] R. L. Gregory. *Eye and Brain*. Oxford University Press, Oxford, UK, fifth edition, 1998.
- [21] C. Grout, P. Purdy, J. Rymer, K. Youngs, J. Williams, A. Lock, and D. Brickley. *Creating Digital Resources for the Visual Arts: Standards and Good Practice*. Oxbow Books, Oxford, UK, 2000.
- [22] H. Eiteljorg. The compelling computer image - a double-edged sword. *Internet Archaeology*, 8, 2000.

- [23] T. Hawkins, J. Cohen, and P. Debevec. Photometric approach to digitizing cultural artifacts. In *2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage (VAST 2001)*, 2001.
- [24] J. Hodgson. Style and content: The effects of style on archaeological reconstruction. *Graphic Archaeology*, 1997.
- [25] J. Hodgson. Dewlish roman villa: Design brief. Design brief for reconstruction project, 2001.
- [26] Dublin Core Metadata Initiative. The dublin core metadata initiative. URL <http://dublincore.org>.
- [27] J. Kantner. Realism vs. reality: Creating virtual reconstructions of prehistoric architecture. In *Virtual Reality in Archaeology*. Archeopress, Oxford, UK, 2000.
- [28] G. Ward Larson and R. Shakespeare. *Rendering with Radiance: The Art and Science of Lighting Visualization*. Morgan Kaufmann, San Francisco, CA, 1998.
- [29] A. McNamara, A. Chalmers, and D. Brown. Light and the culture of medieval pottery. In *Proceedings of the International Conference on Medieval Archaeology*, pages 207–219, October 1997.
- [30] A. McNamara, A. Chalmers, T. Troscianko, and I. Gilchrist. Comparing real and synthetic scenes using human judgements of lightness. In *Proceedings of the 11th Eurographics Rendering Workshop*, pages 207–219. Springer Verlag, June 2000.
- [31] A. McNamara, A. Chalmers, T. Troscianko, and E. Reinhard. Fidelity of graphics reconstructions: A psychophysical investigation. In *Proceedings of the 9th Eurographics Rendering Workshop*, pages 237–246. Springer Verlag, June 1998.
- [32] P. Miller and J. Richards. The good, the bad, and the downright misleading: archaeological adoption of computer visualisation. In J. Huggett and N. Ryan, editors, *Computer Applications and Quantitative Methods in Archaeology 1994*, pages 19–22. BAR International Series 600, Archeopress, 1995.

- [33] B. Molyneaux. From virtuality to actuality: the archaeological site simulation environment. In *Archaeology and the Information Age*. Routledge, London, UK, 1992.
- [34] S. Piggott. *Antiquity Depicted*. Thames and Hudson, London, UK, 1978.
- [35] R. Raskar, G. Welch, K. Low, and D. Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. In S. J. Gortler and K. Myszkowski, editors, *Rendering Techniques 2001*, pages 89–102. Springer-Verlag, 2001.
- [36] P. Reilly. Towards a virtual archaeology. In K. Lockyear and S. Rahtz, editors, *Computer Applications and Quantitative Methods in Archaeology 1990*, pages 133–140. BAR International Series 565, Archeopress, 1991.
- [37] P. Reilly. Three-dimensional modelling and primary archaeological data. In *Archaeology and the Information Age*. Routledge, London, UK, 1992.
- [38] P. Reilly and S. Rahtz. *Archaeology and the Information Age*. Routledge, London, UK, 1992.
- [39] C. S. Rhyne. Computer images for research, teaching, and publication in art history and related disciplines. *An International Journal of Documentation*, XII:19–51, 1995.
- [40] J. C. Roberts and N. Ryan. Alternative archaeological representations within virtual worlds. In Richard Bowden, editor, *Proceedings of the 4th UK Virtual Reality Specialist Interest Group Conference*, pages 179–188, Uxbridge, Middlesex, November 1997.
- [41] N. Ryan. Computer based visualisation of the past: technical ‘realism’ and historical credibility. In P. Main T. Higgins and J. Lang, editors, *Imaging the past: electronic imaging and computer graphics in museums and archaeology*, number 114 in Occasional Papers, pages 95–108. The British Museum, London, November 1996.
- [42] Archaeology Data Service. Creating and using virtual reality: a guide for the arts and humanities. URL <http://ads.ahds.ac.uk/project/goodguides/vr/appendix2.html>.

- [43] J. T. Smith. The validity of inference from archaeological evidence. In P. J. Drury, editor, *Structural Reconstruction*, pages 7–19, Oxford, UK, 1982. BAR International Series 110, Archaeopress.
- [44] D. Spicer. Computer graphics and the perception of archaeological information: Lies, damned statistics and...graphics! In C. L. N. Ruggles and S. P. Q. Rahtz, editors, *Computer Applications and Quantitative Methods in Archaeology 1987*, pages 187–200. BAR International Series 393, Archeopress, 1988.
- [45] T. Troscianko, A. McNamara, and A. Chalmers. Measures of lightness constancy as an index of the perceptual fidelity of computer graphics. In *European Conference on Visual Perception 1998, Perception Vol 27 Supplement*, pages 25–25. Pion Ltd, August 1998.
- [46] G. Ward and E. Eydberg-Vileshin. Picture perfect rgb rendering using spectral prefiltering and sharp color primaries. In *(to appear) Eurographics 2002*, September 2002.

Appendix A

Included papers

LIGHT AND THE CULTURE OF COLOUR IN MEDIEVAL POTTERY

by Duncan H. Brown, Alan Chalmers and Ann MacNamara

Light is something we take for granted, that we create through the movement of a switch. Light is not something archaeologists can recover or record and consequently its importance is rarely considered in interpretations of past ways of living. Light is something that we, as students of the past, need to understand, for the introduction or the creation of light, and its use, facilitated the activities, rituals and lives of our ancestor societies. Light is fundamentally important, yet at present it is rarely considered as a medium for comprehending how people behaved or how they perceived their environment. The hard archaeological evidence is unrevealing. Windows remain as evidence for the way of introducing light; lamps, lanterns and candlesticks for the means of its creation. Yet there is a more subtle way of approaching this problem and that is to look at the objects we find, the clues to the ways in which buildings were decorated, in an effort to understand what sort of environment past peoples created for themselves. Their perceptions can be revealed by the appearance of the objects they used.

This paper is specifically concerned with medieval pottery because that is my particular specialism but the philosophy behind this discussion should lend itself to the study of any type of object of any date. The basic premise is that the colours of medieval pots are related to the lighting conditions that medieval people were accustomed to. Some pots are brightly coloured and highly decorated, others are dull. This is related in part to vessel function but must also reflect the intended place of use and thus variations in lighting conditions. Furthermore, pottery colours may actually reflect a typical *absence*, rather than presence of light, or at least lighting at much lower levels than we are used to. This short paper therefore considers the ways in which medieval interiors were illuminated and how lighting conditions might affect the ways in which objects were perceived and designed. The basis of this discussion is a computer science research project which places medieval pots into simulated environments and introduces different types of illumination. This project is in its early stages but this is seen as an opportunity to show how even these preliminary developments can lead to new lines of enquiry of medieval ceramics. This paper therefore presents a philosophical discussion rather than hard evidence but in doing so hopefully suggests new lines of enquiry and thought.

MEDIEVAL LIGHTING

As has already been inferred, archaeological evidence for the creation of light is relatively rare. The ceramic evidence is perhaps most commonplace on excavations and a variety of lamps and candle-holders are known throughout the medieval period. These are generally portable types and are consequently small in size. It is difficult to envisage such objects being used to illuminate whole rooms. It is known, of course, that torches were extensively used and wall-brackets for these survive. Another source of light, and possibly a very important one, must have been the fires and braziers that were used to heat rooms. On this evidence it seems that the medieval interior must have been a flickering, smoke-beset world and perhaps this explains the bright colours used to decorate medieval objects and indeed rooms. The evidence for those is drawn primarily from manuscript illuminations and paintings, where furniture is often shown brightly painted, and walls display rich hangings.

It may, however, be the case that most of these things were not meant to be seen at their best in artificial light. The hours of daylight regulated pre-industrial life and provided the fundamental means of illumination. In the present day, sunlight is almost irrelevant to the conducting of our lives; houses, offices, shops, factories are almost all permanently lit by artificial means. In the medieval period the sun was probably viewed as the only constant source of light and it is in the architecture that the best evidence for lighting can be found. Big windows provided lots of light but given the limited availability of window-glass they also brought draughts. Windows also created a security risk, as is shown most obviously in castles. Indeed the largest windows may be found in ecclesiastical buildings, those which one might presume to have been least threatened. This provokes the thought that the provision of light on through such grand openings was as much a signal of devotion to God as the building of the entire edifice. Light and holiness seem almost to be related. The way light was used within a church or cathedral may also, perhaps, reflect the controlling aspects of ecclesiastical architecture. Medieval pottery, however, was used more frequently in a domestic environment and the windows in houses are therefore more pertinent to this discussion. Here, window size might be related to status. The windows in surviving English peasant houses are generally small, leading to the conclusion that warmth was more important than light, perhaps because the rural lifestyle was in any case governed by the rising and the setting of the sun. Rural manor houses, and the homes of late medieval yeomen farmers exhibited larger windows, as did medieval town houses. In all instances the window provided light for the carrying out of daily activities such as weaving and sewing. In towns, where jettied upper stories often brought houses within a few feet of each other, windows had to be larger but they also allowed those sitting at them to communicate with their neighbours, as well as people in the street. The window therefore played an important social role as well as a domestic one. It is clear that different dwellings gave different lighting conditions. One might therefore expect the most brightly coloured objects to be associated with the best lit settings and this is, to some extent true. The most highly decorated types of pottery, for instance, are not found at sites of the lowest status. However, that does not tell us how such objects were perceived by those who used them but simply how much light might have been available to see them by. It is an understanding of medieval perception that is being sought here and the next section considers ways of looking for that, if not necessarily finding it.

LIGHT AND COLOUR

It was in the thirteenth and fourteenth centuries, the high medieval period, that English pottery was most elaborately decorated and given the brightest colours. There were dark grey pots, brown, red-brown, brick red, orange, pink and white ones, and those were not glazed. Medieval lead glazes tend either to draw their own colour from that of the clay beneath, thus a clear glaze on a white body appears bright yellow; or are themselves coloured by means of additives, so a bright green is created with the addition of copper. External glazing was almost certainly a form of decoration and it is hard to find any symbolic distinctions between colours in a medium which is inherently prosaic. At the same time, pottery played an increasingly important role in domestic life from the 12th century onwards and it seems to have been consumed at a high rate. The exuberant ceramic forms of the high medieval period coincided with many other cultural developments and at present it is the intention to concentrate on those types. The purpose is to examine the relationship between lighting conditions and the

consumption of pots of a particular appearance by means of trying to find out how those pots might have looked in different settings.

PHOTO-REALISTIC VISUALISATION

Those settings are being created on a computer at the University of Bristol. A selection of pots of different shapes and colours are being computer-modelled and different environments and lighting conditions are being computer-simulated. The value of using a computer is the speed at which one can alter environments. Initially, it is necessary to simulate and measure the optical environment prevailing at the time. This will be achieved by using a controlled test environment, and simulating different conditions by filling it with smoke and dust. Optical measurements of the scene will be made, including replicas of medieval pots, under these murky conditions, using a hyper-spectral camera system developed at Bristol University for the purpose of scene analysis. Similar measurements will also be made of a restored medieval house in Southampton. The next goal is to be able to represent the scene on a high-quality stereoscopic computer display device. This will allow convenient inspection of the likely appearance of artefacts under conditions prevailing at their time of use, and it will be possible to examine the scene from different viewpoints. The fidelity of the reproduction process will be assessed by comparing human visual performance both in the original room and on the computer screen; if the optic arrays are similar in both cases then visual perception and discrimination would follow identical functions. Thus, it will be possible to use an operational measure of human vision to ensure the computer is representing the original scene with accuracy. Finally, the computer-based virtual environment will be used to assess the effects of smoke, dust and other particles on the appearance of a pot.

It is essential that any system for reconstructing and visualising ancient environments must be as accurate as possible and flexible, allowing archaeologists to alter the scene parameters in order to investigate different hypotheses concerning the structure and contents of a site (Chalmers and Stoddart, 1994). Over the last decade, computer graphics techniques have shown an astonishing increase in functionality and performance. Real-time generation of images has been made possible by implementing projective display algorithms in specialised hardware. With the latest graphical hardware systems it is now possible to walk through virtual environments and scenes and perform tasks within these environments. Although the image quality of projective methods is good enough for spatial impressions and for interaction in a virtual environment, it is not sophisticated enough for realistic lighting simulation. Only by simulating the physical propagation of light in the environment can something approaching photo-realism be achieved.

In computer graphics, the illumination at any point in a scene can be determined by solution of the *rendering equation* (Kayija, 1986). Unfortunately, the general form of this equation involves a complex integral over the entire environment and, as such, photo-realistic computer graphics techniques are only able to approximate the solution. Of those currently in use, the particle tracing method is able to approximate most closely all the lighting effects in a closed environment (Pattanaik, 1993). The particle tracing model follows the path of photons as they are emitted from the surface of the light sources and uses the reflected particle flux given by a large number of these particles per unit time as a measure of the illumination of points in the environment. In this way, the particle tracing method is able to simulate direct as well as indirect reflection, diffuse and specular reflection, and the effects of participating

media such as flame, smoke, dust and fog (Lafortune and Williams, 1996). All these effects are essential for a physically accurate lighting simulation.

CONCLUSION

It may seem that computer-simulation will never approach the authenticity of finding a medieval house, filling it with pottery and lighting fires, torches and lamps. However, the time and resources needed for doing that are very restrictive. On a computer it is possible to change the shape of a room, the colours of the walls, the colours of the pots, and the quantity of light and smoke, with relative ease. That facility will open up new lines of enquiry. The project is in its early stages but it is revealing that so far observation shows that the appearance of the objects in a room changes with the colours of the walls and the angle of the light source. The constancy of the light source will also have an effect, hence the development of a flickering light and the introduction of smoke and dust particles.

This project is an attractive mix of archaeological intuition and philosophy together with hard science. The proposition is even more exciting when viewed as a voyage of discovery which is bound to open up new lines of enquiry and thought.

BIBLIOGRAPHY

Chalmers, AG and Stoddart, SKF, 1994, Photo-realistic graphics for visualising archaeological adoption of computer visualisation. In Hugget J and Ryan N, 'Computer applications and quantitative methods in archaeology' pp19-22, Tempus Reparatum, Glasgow.

Kajiya, JT, 1986, The Rendering Equation, in ACM Computer Graphics, 20 (4), pp143-150.

Lafortune, E and Williams, Y, 1996, Rendering participating media with bidirectional path tracing, in Pueyo X and Schroeder, P, 'Seventh Eurographics Workshop on Rendering', pp 92-101, Oporto.

LIGHT, COLOUR, PERCEPTION AND MEDIEVAL POTTERY

Duncan H. Brown and Alan Chalmers

This paper begins with a discussion of some of the issues that need to be acknowledged by archaeologists studying the use and perception of colour in past societies. This is followed by a case study that, it is hoped, further illuminates the subject, but is something of an interim statement on a continuing research project, where computer graphics are used to re-create medieval lighting effects. The purpose is to reach some understanding of how objects, specifically pottery vessels, might have appeared in a medieval environment. The context given to the whole of this paper is therefore that of medieval England, although the general principles should be applicable throughout archaeology.

LIGHT

It is difficult to appreciate colour without understanding light, for without light there is no colour. In attempting to understand any significance particular colours might have had in the past, it is therefore necessary to overcome some of our own preconceptions. We, after all, are used to illuminating our environment simply by flicking a switch, but the types of light cast by electric sources are very different to those experienced by our medieval predecessors. Each different light source has its own spectral profile, some lights are more red, for instance, and others more blue. This is visible to us as 'warm' or 'cold' light. Undergraduate research by Natasha Chick, of the Department of Computer Science at Bristol University, compared the spectral properties of various light sources by converting spectroradiometer data to RGB format and taking readings from a MacBeth colour chart. At the yellow square of the chart an animal fat candle gives RGB readings of 0.759% red, 0.24% green and 0.001% blue. A 55-watt bulb, by comparison, shows as 0.524% red, 0.345% green and 0.131% blue (Chick, 2000). This demonstrates that an electric bulb gives a much greater blue cast than a tallow flame, which is more red. The results are shown visually in figures 1 and 2. Light sources also differ widely in the amount of light produced, so that there are bright or dim lights. Flames were the usual means of introducing light in the past, and these have the extra dimension of flicker, thus creating patterns and moving shadows that further affect how objects might have looked.

We are not used to looking at objects in conditions where light sources are dim, generally of a red cast and also moving. Nor indeed, once they have been removed from the ground, are we used to viewing artefacts only in daylight. The role of sunlight would seem less important when studying objects that were mainly, as far as we understand, used indoors, but in this we might be mistaken. Medieval daily life was governed by the rising and setting of the sun, as documents such as books of hours demonstrate. The creation of light was a costly business, especially through the use of candles, which were beyond the means of many, and most people would not have stayed up much beyond

nightfall. Some medieval buildings, such as churches and shops, were constructed with large windows that were designed to make the most of sunlight, allowing visitors a good view of what was being offered for there is no doubt that most of the public rituals or transactions enacted at these places took place during daylight hours. Others, such as castles and rural dwellings, contained small windows that served to provide illumination without compromising the requirements of protection, either from people or the elements. The provision of daylight was equally crucial in these places, however, and in the rural context at least, once the sun had gone activities and movement were restricted. In any event, the quality of daylight available must have affected the appearance of the objects used indoors and again, in our well-glazied age, the variations in the provision of light are beyond our immediate experience and understanding. The ways in which we view, perceive and understand objects is governed by current lighting methods (electric light and large windows) but in order to understand how objects were viewed and understood in the past we must consider how they were illuminated.

COLOUR

The pottery used in England throughout the medieval period was mostly earthenware and may be divided into white-firing and red-firing types. Pots made from white-firing clays appear white to buff in colour but were usually covered, completely or partially in a lead glaze. Clear lead glazes give a yellow to amber appearance to white pottery but the addition of copper creates green, which can vary from dark through olive to bright, apple hues. In general, the most brightly coloured medieval pottery was made from white-firing clays. Vessels made from iron-rich, red-firing clays range in colour from brick red to dull brown, when fired in oxidising conditions. Clear lead glazes can give a red-orange appearance to red wares while lead with copper glazes produce a variety of greens generally less vibrant or consistent in colour than those seen on white earthenware. Red clays fired in an oxygen-free atmosphere range in colour from grey through to black and in such conditions an ordinary lead glaze will turn slightly green or greenish-clear. In short, medieval pots were given a variety of colours, including white, red, brown, yellow, orange, green, grey and black. The appearance of vessels was also enhanced, and the colour effect slightly altered, by decoration such as painted lines, coloured slips and applied clays.

So what? may be the first question that springs to mind after reading that. What, if anything, does this information tell us about the medieval period? It is possible to quantify medieval pottery assemblages by colour, but is that a useful thing to do? Did the colours of medieval pots have any meaning?

The range of colours produced in pottery at least tells us something about ceramic technology. Aspects of kiln construction and management, clay exploitation and mineral use are all illuminated through an understanding of how pots were made in the colours that they were. Also revealed is the importance of colour in medieval society. The most colourful and highly decorated vessels were jugs, normally identified as table ware, used in the serving of liquids, especially

wine. It is possible that these were used at high profile dinners. The dullest vessels, unglazed and undecorated, were those used in low-visibility activities such as cooking. Colour was therefore an important element in the culture of display that prevailed in medieval society, at least in its upper echelons. In these terms it is useful, therefore, to consider quantifying pottery assemblages by colour, as the quantities of highly decorated and plain pots might be related to status: high numbers of gaudy vessels might indicate a household where presentation was important.

It may be possible to put meanings to colours but one might be wary of doing so in the case of pottery. After all, the range of colours is partly determined by the available technology, in terms of glaze and mineral use and kiln structure. Such a view, of course, excludes the element of human decision-making and if we accept that medieval potters gave pots the colours their customers wanted then we must ask why it was that those colours appealed. An analysis of colour symbolism may not appear very helpful, however. The most common pottery was coloured red-brown. In medieval art, brown was associated with mourning and red with God, while green, the most common colour for glazed pottery, signified the holy spirit and therefore also the bishops (Becker, 2000). Yellow was regarded as a substitute for gold, which denoted heaven (ibid).

The next step, of course, is to ask whether medieval potters were fully aware of the symbolism in art. The same colours have meanings in folk-lore that are not entirely consistent with those of medieval Christianity. Brown is of course the colour of the earth (ibid) and its association with mourning is perhaps related to the notion that humans are earthly beings whom, in death, revert to clay. Red is the colour of blood and, perhaps more pertinently with regard to pottery, of fire, and in its positive manifestation signifies life, love, warmth and passion. Its negative aspect is that of destruction and war (ibid). Green was also related to life, through its association with plants and the shoots of springtime as well as water. It thus symbolised hope and longevity (ibid). There is more to go on here, and it is easy to create associations between pots of certain colours and the offering, presumably through their contents, of valued qualities such as life, warmth and a long life. Yellow is more of a problem, however, as in the middle ages it was the colour of envy. It may be safer, therefore, to look upon it as a substitute for gold, and here is the crux of the matter. For it is we who are now doing the looking, thus creating our own understandings of the colours we see.

PERCEPTION

There is no guarantee that the green pot we now look upon was also understood to be green by medieval people. Or to put it another way, although they may have known it to be green, they might have appreciated it as something else. There is no doubt that colour can transmit mood and meaning, and medieval folk may have read their own particular meanings into pots of certain colours. The use of green, for instance, might have signified its quality and its suitability for certain functions, furthermore, the application of slipped decoration might even have indicated those uses to which a pot should *not* be put. We come, here, to a

further issue. If there is a gap between that which we perceive and what medieval people understood, there was also a gap between what the creator of a pot understood and the perceptions of the consumer. This may be true, and perhaps more readily recognised, in terms of function as well as appearance. A pot that the potter considered to be an excellent wine jug may have been bought for use as a piss-pot. A simple example of differences in perception may be that of a vessel made in Muslim Spain and imported into Christian England. It is likely that any significance or message in the decoration became irrelevant in the transfer from one culture to another. The same may therefore be true of colour.

This discussion, however, is focused on the perceptions of the users, rather than the makers, of medieval pottery and here again we must recognise how discrepancies might creep in. It is unlikely, at certain levels of medieval society, that domestic goods were acquired directly by those householders who kept servants. The pots we recover from their dwellings may not, therefore, accurately reflect the tastes of the principal occupants. It is, however, reasonable to assume that householders would ensure that they used objects they liked the look of. Once we acknowledge taste as a factor in deciding which objects to use then we raise the question of perception for no colour, shape or decorative design is actually worse or better, it is only perceived to be so. Cultural, social, political and economic attitudes will be brought to bear in determining what is acceptable, and may be expressed in the symbolism of artefact appearance and use and domestic ritual. Here again, our own preconceptions will interfere. The framework for our own perceptions is quite different from that for medieval society, which will remain, at best, difficult to comprehend. Certain clues survive, and the significance of colour in medieval culture is one of them. There is no physical difference, between humans in the present and the medieval past, in the actual mechanisms for perceiving, the eye and brain, so changes in our perceptions must be explained culturally and emotionally rather than clinically.

POTTERY

Issues of colour and perception in pottery have already been raised and it is clear that we must understand the relationships between pottery and those who used it. In the first place, it is not always clear whom those users were, or at least it is clear that in many medieval households several different users were involved. In the townhouse of a wealthy burgess, such as that which provides the setting for the case study set out below, there was a hierarchy of individuals and their roles. It seems reasonable therefore, that different types of vessel were also fitted into that scheme. At the lowest level, ceramically, we may place vessels used in food preparation. Medieval cooking pots were simple forms, cheaply made and cheap to buy. Their use would have been confined to the kitchen and scullery areas and cooking pots would probably not have appeared at high table or been considered for display. These vessels were of unglazed earthenware, usually buff, red-brown or grey in colour. Unglazed, or partially glazed, jugs were also produced and these may have been used in the same areas of the house. It is most likely that domestic servants used pottery bought for storage and cooking and, as far as the

householder was concerned, the appearance of that pottery may have been largely irrelevant. There seems, therefore, to be little perceptual influence in the acquisition and use of kitchenwares and this is borne out by the fact that in most excavated medieval assemblages cooking pots and jars are invariably of local origin (Brown 1997) and have a uniform appearance.

Jugs, however, the medieval vessel type we commonly characterise as tableware, were derived from a wider variety of sources, even at humble farmsteads (ibid) and issues of taste and display must be considered. Rich glazes and elaborate decoration suggest use in public ways, certainly outside the confines of the kitchen area. A penchant for display is known to have permeated most levels of medieval society and ceramics would have been a cheap way of assuaging such a requirement. It is important to recognise that pottery could never compete with glass or metal as a medium for showing off but it is easy to identify highly decorated vessels as intended for use at table. Here, aspects of colour become more important perhaps. The use of glaze allows a greater degree of consistency in the colouring of pottery and it must be assumed that this was done deliberately to appeal to the tastes of prospective customers. If, therefore, a domestic assemblage yields tableware of a certain range of hues, then it seems reasonable to suggest that those were the colours that the users preferred.

The conclusion, therefore, is that the jug is a suitable focus for research into colour and perception in medieval pottery, at least at this initial stage. Jugs present the widest range of shapes, decorative designs and colours of any ceramic form in medieval assemblages; were probably used in a greater variety of domestic situations and might to a higher degree reflect the tastes of consumers.

CONTEXT

The first question that needs to be answered in this consideration of colour and perception must be this; what did medieval pots look like in their original setting? That question has been addressed by research conducted at the Department of Computer Science at Bristol University, where a computer model of the hall of a medieval town house has been constructed. The model is based on the Medieval Merchant's House museum in Southampton, a half-timbered structure renovated by English Heritage as accurately as possible to represent a 13th century dwelling of some economic status. The modelling was carried out by Ann McNamara as part of her post-graduate research into visual perception and computer graphics. Three 13th century jugs were chosen from the archaeology collections of Southampton City Council and these have been modelled, not to a very high degree of accuracy as yet, and inserted into the environment. Realistic renderings of medieval-type light sources have not yet been introduced, although Natasha Chick's research now makes that possible.

Figure 3 shows a view of the hall with the jugs placed on the main table. This scene is an accurate depiction of the actual building, the wall hanging and the furniture are pieces currently on display. The three jugs are a local redware

baluster, red-brown in colour with a partial greenish-clear glaze, a Dorset white ware jug with an overall yellow glaze, decorated with dark brown applied vertical stripes and lines of pellets and a Saintonge white ware vessel with an overall bright green glaze. They are contemporary, dateable to around 1270-1300 and it is possible that similar vessels could have been in use in a Southampton household at the same time. The rendering in figure 3 is lit with generic lighting, comparable with daylight, rather than a specifically modelled light source. Figure 4 shows the same scene lit from a central hearth with two candles on the table. The later fireplace and chimney shown in figure 3, which is still in place in the actual building, has been removed from the rendering in figure 4, thus giving a more accurate model of the 13th century hall. The next steps will be to illuminate the scene with accurately rendered light sources, such as tallow flames, then to introduce atmospheric pollutants, such as smoke and dust, that will effect the behaviour of the light. More accurate models of the pots will probably be best achieved through the use of a high-resolution laser scanner.

Once a computer model has been generated, however, it is necessary to ascertain that the viewer's perception of what is visible on a monitor equates with how an actual scene would be perceived. In other words, how 'real' is a computer-generated environment? The aim of realistic image synthesis is the creation of accurate, high quality imagery that faithfully represents a physical environment, the ultimate goal being to create images that are perceptually indistinguishable from an actual scene. Advances in image synthesis techniques allow us to simulate the distribution of light energy in a scene with great precision. This does not, unfortunately, ensure that the displayed image will have a high fidelity visual appearance. Reasons for this include the limited dynamic range of displays, any residual shortcomings of the rendering process, and the extent to which human vision encodes such departures from perfect physical realism.

Computer graphics techniques are increasingly being used to reconstruct and visualise features of cultural heritage sites that may otherwise be difficult to appreciate. While this new perspective may enhance our understanding of the environments in which our ancestors lived, if we are to avoid misleading impressions of a site, then the computer generated images should not only look "real", but there should be a quantifiable metric of image fidelity by which this "realism" can be measured. Computational image quality metrics may be used to provide quantitative data on the fidelity of rendered images; however, if such images are to be used to investigate visual perception of environments by humans then they must include an understanding of the features of the Human Visual System (HVS). The HVS comprises many complex mechanisms that work in conjunction with each other, making it necessary to consider the HVS as a whole rather than study each function independently. Psychophysical experiments can be used to investigate the human perception of scenes. Subjects are presented with a (random) selection of images produced on a viewing device and the same real world scenario. Interactive responses and detailed questionnaires are used to evaluate the perceptual quality of the rendered images. Outcomes from these psychophysical experiments then feed

back to further improve the rendering process. Preliminary results, based on work Ann McNamara carried out at Bristol, show that there is a close correlation between perceptions of actual and computer-generated scenes (McNamara et al, 2000).

We are a long way from a satisfactory model but this exercise has already thrown up a number of issues. Computer modelling offers a far more flexible approach. It is not possible to occupy the actual building, light a fire, position candles and torches and record the results and if it were it is still not feasible to remove the existing fireplace and chimney-breast. In a computer-generated environment we are able to remove extraneous structures, change the size and colour of wall-hangings, even change the colours of the walls. We can increase or diminish the size of the hearth and alter its position and we can move other illuminants around as we wish, placing them on furniture or anywhere on the walls. It may also be possible to recreate vision defects, such as shortsightedness, that were left uncorrected in the 13th century. These computer-generated images are far from ideal but figures 3 and 4 serve to illustrate the value of this approach. This work has also confirmed the necessity of understanding how objects looked in their original contexts if we seek any insight into past perceptions. It is already apparent that many pots would have looked brightest when lit from above. Only the top half of some jugs are glazed and decorated, and this is perhaps indicative of how they were illuminated in use, perhaps by daylight through windows or perhaps from torches hung on walls. The purpose of this research is the revelation of detail such as this and there is great potential to go deeper into medieval ways of living.

CONCLUSIONS

The research described above has a long way to go yet, but certain points have already been raised and are perhaps worth summarising here. In searching for insights into the ways past people perceived their environments and the objects they interacted with it is important to consider how they illuminated their lives. Even more crucial is the relationship between light and colour. Colours will change in appearance according to the types of light source present; yellow, for instance is especially affected by the RGB factor of certain illuminants. The recreation of medieval lighting conditions is therefore seen as a vital step in comprehending attitudes to colour, and eventually perhaps, shape and decoration. If there is any symbolic meaning in the use of colour on pottery then this might be revealed through an exploration of medieval perception, through the recreation of a medieval environment. The modelling of a 'realistic' environment through the application to computer graphics of computer science and psychophysics, is perceived to be the most far-reaching and flexible way of exploring human perceptions in the past. Here, however, the ways in which a computer-generated image might be perceived also need to be understood.

A broader issue, therefore, is how all of us read our experiences. This research into colour and light has shown how easy it is for our own preconceptions to intrude into the ways we view archaeological objects or sites. It

is rare for archaeologists to attempt to see things differently (ask the protesters at Sea Henge) and the ways in which we recover details of the past, through the recording of demonstrably secure data such as composition or dimensions, is designed to facilitate research into human activity rather than human thought. We all take too much for granted, certainly in the ways we consume but also in the ways we interpret. Our research is intended to reveal more about medieval perceptions by investigating the context of colour. This has led to a re-evaluation of how we ourselves look at pottery, the chosen medium for our research. One result of our inquiries may be to suggest new, or additional ways of recording ceramic assemblages. The aim is to find methods of analysis that could take us beyond those typical questions of chronology and provenance. If the aim of archaeology is to provide insights into the lives of past individuals, communities and cultures then we need to show a greater respect for the things they have left behind and attempt more refined ways of understanding them. How do archaeologists perceive the past and objects from the past, as archaeologists or as human beings; and why might there be a difference between the two? If nothing else, our research has raised such questions, and may help to answer them.

BIBLIOGRAPHY

Becker, U, 2000, 'The Continuum Encyclopedia of Symbols' Continuum

Brown, DH, 1997, 'Pots from Houses', *Medieval Ceramics* 21, 83-94

Chick, NJ, 2000, 'Realistic visualisation of medieval sites', Final year project dissertation, University of Bristol.

McNamara A., Chalmers A.G., Troscianko T., Gilchrist I. 'High Fidelity Image Synthesis', 11th Eurographics Workshop on Rendering, Brno, June 2000.

CAPTIONS

Figure 1: Computer-generated environment with MacBeth chart, lit by tallow flame (created by Natasha Chick).

Figure 2: Computer-generated environment with MacBeth chart, lit by 55-watt electric light (created by Natasha Chick).

Figure 3: Computer-generated model of the hall of the Medieval Merchant's House, French Street, Southampton, lit by generic approximation of daylight (created by Ann McNamara, now of Trinity College, Dublin).

Figure 4: Computer-generated model of the hall of the Medieval Merchant's House, French Street, Southampton, lit by a generic approximation of firelight (created by Ann McNamara, now of Trinity College, Dublin).

The following paper is a reprint of:

A Photometric Approach to Digitizing Cultural Artifacts

Tim Hawkins Jonathan Cohen Paul Debevec

University of Southern California Institute for Creative Technologies

Published at:

2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage

David Arnold, Alan Chalmers, and Dieter Fellner, co-chairs

Glyfada, Greece, November 2001.

<http://www.eg.org/events/VAST2001/>

A Photometric Approach to Digitizing Cultural Artifacts

Tim Hawkins

Jonathan Cohen

Paul Debevec

University of Southern California Institute for Creative Technologies¹

ABSTRACT

In this paper we present a photometry-based approach to the digital documentation of cultural artifacts. Rather than representing an artifact as a geometric model with spatially varying reflectance properties, we instead propose directly representing the artifact in terms of its *reflectance field* – the manner in which it transforms light into images. The principal device employed in our technique is a computer-controlled lighting apparatus which quickly illuminates an artifact from an exhaustive set of incident illumination directions and a set of digital video cameras which record the artifact's appearance under these forms of illumination. From this database of recorded images, we compute linear combinations of the captured images to synthetically illuminate the object under arbitrary forms of complex incident illumination, correctly capturing the effects of specular reflection, subsurface scattering, self-shadowing, mutual illumination, and complex BRDF's often present in cultural artifacts. We also describe a computer application that allows users to realistically and interactively relight digitized artifacts.

Categories and subject descriptors: I.2.10 [Artificial Intelligence]: Vision and Scene Understanding - *intensity, color, photometry and thresholding*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *color, shading, shadowing, and texture*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *radiosity*; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture - *radiometry, reflectance, scanning*; I.4.8 [Image Processing]: Scene Analysis - *photometry, range data, sensor fusion*. **Additional Key Words and Phrases:** image-based modeling, rendering, and lighting.

1 Introduction

Creating realistic computer models of cultural artifacts can aid their study by remote parties as well as serve as an improved record of the artifact for archival purposes. While a photograph faithfully records an artifact's appearance from a single point of view in a particular lighting environment, it can be far more informative to be able to see the artifact from any angle and in any form of lighting - allowing a scholar to give it greater scrutiny as well as the general observer to see the artifact in its natural environment.

The standard computer model of a cultural artifact consists of a geometric surface model covered by a texture map which represents the artifact's spatially varying diffuse reflectance properties. Sometimes, a specular component is added to render the shininess, often set to a single representative value for the object. An artifact's geometry is most commonly acquired using active sensing such as laser scanning or structured light, and the texture maps are usually constructed using a simplified form of reflectometry: lighting the object from a known direction with a calibrated light source, and then, using an estimate the object's surface orientation, determining the spatially varying albedo of the object. This basic approach has produced numerous high-quality records and render-

ings of cultural artifacts such as statues, vases, and interior environments.

Unfortunately, this traditional approach is difficult to apply to a large class of cultural artifacts – ones that exhibit complex reflectance properties such as anisotropy or iridescence, ones that exhibit significant self-shadowing or mutual illumination, ones that exhibit significant subsurface reflection, objects that are highly specular or translucent, and objects with intricate surface geometry. We can illustrate these difficulties with several examples:

- A fur headband would be difficult to digitize since it does not have a well-defined surface – the stripe of a laser scanner or video projector would scatter through the fur rather than drape over it, and the surface reconstruction algorithm would have difficulty reconstructing a surface from the data. Even if the fur headband's geometry could be captured, it could be impractical to represent the geometry of tens of thousands of individual hair segments, and to subsequently realistically render the effects of this fur being illuminated. Similar challenges exist for digitizing many types of clothing as well as for human hair.
- A small jade sculpture would exhibit significant subsurface scattering - light hitting it from behind would cause it light up from the front, and light striking the front would penetrate the surface a considerable distance. This could complicate the range scanning process, but would pose an even more significant problem for reflectometry: current techniques neither estimate nor represent subsurface scattering properties of an object. A computer rendering of the jade sculpture might more closely resemble painted green rock than luminous jade.
- An intricately carved ivory sculpture with complex internal geometry would be challenging to digitize due to both significant self-shadowing and mutual illumination. Scanning its geometry would be complicated by the self-shadowing; narrow crevices are difficult to record with triangulation-based scanning methods in which each surface point must be visible to both the source of light and the image sensor. Furthermore, the complexity of the geometry could complicate the range map merging process, which works best when there are relatively coherent surface sections. Mutual illumination – the fact that light will bounce between surfaces inside the object's reflective concavities – would complicate the reflectometry stage since it becomes very difficult to reliably control the illumination incident on any particular surface. The significant mutual illumination will also complicate rendering, requiring expensive global illumination algorithms to produce images that correctly replicate the appearance of the original object.
- A polished silver necklace, encrusted with rubies, diamonds, and abalone, would also be difficult to digitize. The low diffuse reflectance of the silver and the internal reflections and refractions in the rubies and diamonds would make geometry capture with structured light or laser scanning impractical. It would be hard to measure the reflectance of the reflective surfaces since most reflectometry techniques are better suited to materials with a significant diffuse reflection component. The translucence of the rubies and diamonds would make modeling their reflection characteristics difficult, and the reflectance

¹USC Institute for Creative Technologies, 13274 Fiji Way 5th floor, Marina del Rey, CA, 90292. Email: {timh, debevec}@ict.usc.edu.

of the iridescent abalone would be too complex to represent with most currently available reflectance models.

In summary, a large class of cultural artifacts exhibiting complex geometric and reflectance properties cannot be effectively digitized using currently available techniques. This poses a significant problem for the application of computer graphics to cultural heritage, as many of the materials and designs used by craftspeople and artisans are specifically chosen to have complex geometry or to reflect light in interesting ways.

In this paper, we show that an alternative approach based on capturing *reflectance fields* [4] of cultural artifacts can acquire, represent, and render any of the above artifacts just as easily as it could a clay jug or granite statue, with simple acquisition and rendering processes, and can produce photorealistic results. The technique is data intensive, requiring thousands of photographs of the artifact, and as currently applied allows only for relatively low-resolution renderings. As such we discuss its advantages and disadvantages over current techniques as well as potential hybrid methods.

In our proposed technique the artifact is photographed under a dense sampling of incident illumination directions from a dense array of camera viewpoints. The device we use to acquire this dataset is a *light stage* consisting of a semicircular rotating arm with an array of strobe lights capable of illuminating an object placed at its center from up to a thousand different directions covering the entire sphere of incident illumination. Images taken with digital cameras are compressed together into a single multi-dimensional dataset comprising the object's *reflectance field* [4], which characterizes how the object transforms incident illumination into radiant imagery. Renderings of the object can then be created under any form of illumination - such as the light in a forest, a cathedral, a museum, or in a candlelit hut, by taking linear combinations of the images in the reflectance field dataset. No geometric model of the object is required, and the resulting renderings capture the full complexity of the object's interaction with light, including self-shadowing, mutual illumination, subsurface scattering, and translucency, as well as non-lambertian diffuse and anisotropic specular reflection.

In this paper we apply this process to a number of Native American cultural artifacts including an otter fur headband, a feathered headdress, an animal-skin drum, and several pieces of neckwear and clothing. We show these artifacts illuminated by several real-world natural lighting environments, and describe a software program for interactively re-illuminating artifacts in real time. We also describe the reflectance field capture process and the equipment involved. We also propose using view interpolation to extrapolate a discrete set of original viewpoints to of an artifact to arbitrary novel viewpoints in conjunction with the reflectance field capture technique.

The central contribution of this paper is to demonstrate that the reflectance field capture technique, introduced in the context of rendering human faces in [4], provides particular advantages over current digitization techniques when applied to cultural artifacts. Since artifacts can have stronger specular components than human skin, we also show that capturing shiny artifacts requires the acquisition of high dynamic range [5] reflectance field image data. We present an interactive program for visualizing re-illuminated reflectance fields, and suggest how view interpolation could be used to continuously vary the viewpoint of an artifact from a discretely sampled set of viewpoints. Finally, we suggest improvements to the light stage apparatus specifically for acquiring complete view-independent models of cultural artifacts.

2 Background and Related Work

Current leading techniques for digitizing three-dimensional cultural artifacts involve acquiring multiple range scans of the artifact, as-

sembling them into a complete geometric surface model, and then using a form of reflectometry to derive lighting-independent texture maps for the artifact's surfaces.

Range scans are acquired most commonly through laser-stripe scanning, by projecting patterns of light from a video projector, or through illumination-assisted stereo correspondence. Individual range scans can be aligned to each other using Iterated Closest Points (ICP) algorithm [23] and merged using either polygon zipping [19] or volumetric range merging [2]. A sampling of recent projects which have used these techniques to derive geometric models of cultural artifacts are IBM Watson's Florentine Pieta Project [17], Stanford's Digital Michelangelo Project [11], Electricit de France's Cosquer cave (1994) and Colossus of Ptolemy (1997) projects, the Canadian National Research Council's museum artifact scanning work [1], and work to scan vases [16] at the Istituto di Elaborazione dell Informazione in Pisa.

Reconstructing an artifact's appearance - not just its geometry - is the second stage of the digitizing process. Some techniques (e.g. [6], [22]) directly project photographs of the object onto the artifact's geometry to form diffuse texture maps; this technique has the advantage that the renderings will exhibit realistic, pre-rendered shading effects including self-shadowing and mutual illumination, but have the disadvantage that the lighting is fixed according to the conditions in the original imagery. For environments it is sometimes acceptable to have static lighting; for artifacts static lighting is generally less acceptable since the directions of incident illumination on an artifact change as the artifact is rotated and there is often a desire to visualize an artifact as it would be seen in different illumination environments.

Deriving lighting-independent texture maps for artifacts has been done in several projects. One set of techniques lights an object from one or more directions and uses the geometric model to estimate the surface's diffuse albedo for all points on the surface. [13] solved for the spatially varying diffuse reflectance properties across a diffuse object using different lighting and observation directions. Related techniques are used in [17] and [11] to derive illumination-independent texture maps for marble statues. [1] uses the intensity return of its collimated tri-colored laser scanner to derive color lighting-independent texture maps. Recovering specular object properties has been investigated in [18].

These current techniques have produced excellent 3D models of artifacts with well-defined surface geometry and generally diffuse reflectance characteristics. However, these current techniques for model acquisition are difficult or impossible to apply for a large class of artifacts exhibiting complex surface microstructure, spatially varying specular reflection, complex BRDFs, translucency, and subsurface scattering. As a result, artifacts featuring silver, gold, glass, fur, cloth, jewels, jade, leaves, or feathers are very challenging to accurately digitize and to convincingly render.

Recent work [4] building upon related image-based rendering techniques [9, 15, 21] presented a technique for creating relightable computer models of human faces without explicitly modeling their geometric or reflectance properties. Instead, the face was illuminated from a dense array of incident illumination and a set of digital images were captured to represent the face's reflectance field. The images from the face's reflectance field were then combined together in order to produce images of the face under any form of illumination, including lighting environments captured from the real world as in [3]. [12] applied this technique in rendering cultural artifacts exhibiting diffuse reflectance properties. In this paper, we show that this technique can be applied to the digitization of cultural artifacts exhibiting any geometric or reflectance properties including those that are traditionally difficult to model and render. We discuss issues which arise in applying these techniques to capturing cultural artifacts, including the need to acquire image data sets using high dynamic range photography

[5] to properly capture and render specular highlights. We furthermore describe an interactive lighting tool that allows artifacts to be re-illuminated by a user in real time, and propose image-based rendering techniques that will allow an artifact to be manipulated in 3D as well as being arbitrarily illuminated. In this work we use a collection of Native American clothing and jewelry to demonstrate the possibilities of the technique.

3 Dataset Acquisition



Figure 1: **The Light Stage** illuminates an artifact (center) from a dense array of incident illumination directions as its appearance is recorded by high-speed digital video cameras (one can be seen in the upper left). This quarter-second photographic exposure shows several of the lights on at once although only one strobe light flashes at any given time.

The data acquisition apparatus used in this work consists of a semicircular arm three meters in diameter that rotates about a vertical axis through its endpoints. Attached to the arm are twenty-seven evenly spaced xenon strobe lights, which fire sequentially at up to 200 Hz as the arm rotates around the subject. The arm position and strobe lights are computer-controlled allowing the strobes to synchronize with high-speed video cameras. We have used two models of high-speed video cameras in this work. The first is the UniQ Vision UC-610; it has a single image sensor of 660 by 494 pixels and can run asynchronously at up to 110 frames per second, producing digital output to the computer. The second camera is a Sony DXC-9000, which has separate 640×480 image sensors for red, green, and blue channels and which runs progressively at 60 frames per second. The Sony camera produces sharper images and more vibrant colors, although it has analog video output which must be digitized by the computer.

Figure 1 shows the light stage capturing the reflectance field of a Lakota Native American headdress; the capture process takes approximately fifteen seconds to acquire 1,728 images of the artifact. The 1,728 images are arranged as an array of 64 directions of longitude corresponding to the rotation of the arm and 27 directions of latitude corresponding to the individual strobe lights. An alternative low-cost light stage apparatus presented in [4] consists of a single traditional light source on a two-axis rotation mechanism that spirals around the artifact, beginning at the north pole and spiraling down along the surface of a sphere to the south. This alternative device is less expensive to construct, but the device we present in this current paper allows the lights to be positioned with greater precision and repeatability and the datasets to be acquired significantly more rapidly, and with much less work.

Figure 2 shows a captured reflectance field dataset of the headdress. Images near the top of the figure are illuminated from above; images in the center of the figure are illuminated from straight forward, and to the right of the center are illuminated from the right. Images at the far left and right of the figure are actually illuminated from various directions *behind* the artifact; these images are important to capture since they show light grazes along the sides of the artifact and shines through the artifact’s translucent areas. The actual dataset taken by the apparatus is considerably higher resolution; the figure shows just every fourth image in both the horizontal and vertical dimensions.

For artifacts exhibiting relatively shiny reflectance, it becomes necessary to record the reflectance field dataset in *high dynamic range* – with imagery that can capture a greater ratio between light and dark areas than single exposures from the video cameras. For this we can employ a high dynamic range image acquisition method as in [5] to combine both under- and over-exposed images taken from the same viewpoint in the same lighting into images that capture the full dynamic range of the artifact under that lighting. To take these images, we capture the reflectance field dataset more than once; we first record it normally and then next time place neutral density filters over the video camera lenses in order to reduce the exposure of the imagery. In these subsequent passes, the images are darkened to the point that the specular highlights will be properly imaged without saturating the image sensor. In this work, we use 3-stop neutral density filters which reduce the exposure by a factor of eight. If the specular highlights still saturate the image, the exposure can be further reduced by adding additional neutral density filters or setting the cameras to a smaller aperture. In Section 5 we show an example of using this procedure to faithfully reproduce specular reflections in a synthetically illuminated artifact.

4 Illuminating Reflectance Fields of Artifacts

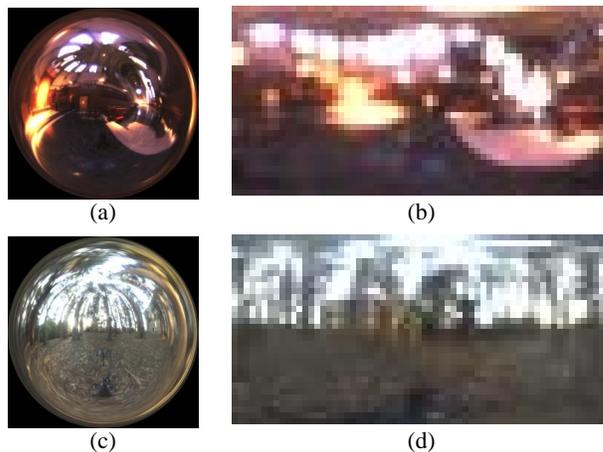


Figure 3: **Real-World Lighting Environments** (a) A light probe [3] image records the illumination in San Francisco’s Grace Cathedral (b) The probe image is resampled into a latitude-longitude format having the same coordinate system and resolution as the captured reflectance field in Figure 2. (c) A light probe image recording the incident illumination in a Eucalyptus grove. (d) The resampled version of (c). The incident illumination datasets were recorded by taking omnidirectional high dynamic range images of real lighting environments.

The motivation for our approach is to be able to realistically show digitized artifacts illuminated by any desired form of illumination. This ability can allow scholars to study how an artifact responds to light and how it may have appeared to people of the corre-

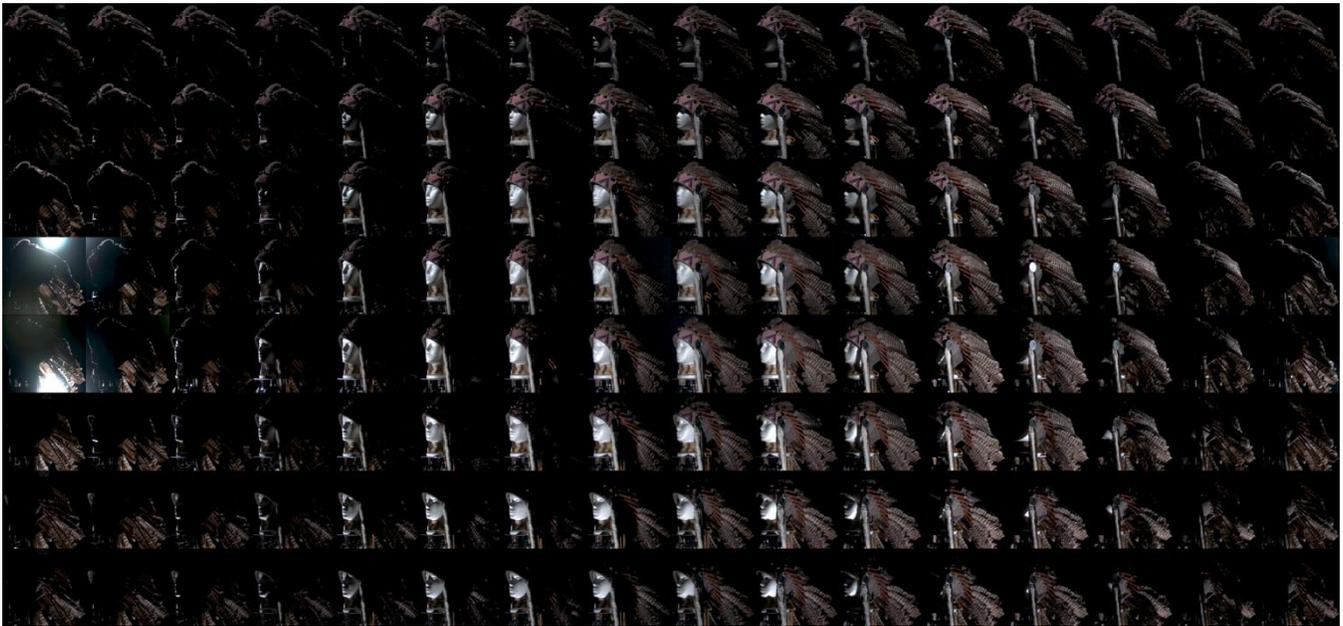


Figure 2: **A Reflectance Field Dataset** This mosaic of images of the headdress shows a sampling of the 1,728 images acquired in a light stage capture session (the original 64×27 dataset is shown as a 16×8 dataset.) The dataset shows the headdress illuminated from all possible directions of incident illumination.

sponding culture in the artifact’s natural environments. It can also allow the artifact to be realistically integrated into virtual cultural recreations or virtual museums, illuminating it with the specific illumination present in any given virtual environment.

To re-illuminate the artifacts, we employ the reflectance field illumination process originally applied for relighting human faces in [4]. First, an image of an incident illumination environment is captured or rendered; for this, the light probe technique presented in [3] can be employed. In this technique, a series of differently exposed images of a mirrored ball are combined to produce a high dynamic range omnidirectional image that measures the color and intensity of the illumination arriving from every direction in the environment. The two light probe images used in this paper are shown in Figure 3.

To create a rendering of an artifact as it would appear in such a sampled lighting environment, the light probe image is resampled into the same coordinate space and resolution as the artifact’s reflectance field; in our work this is a 64×27 image in a latitude - longitude coordinate system. The two images on the right of Figure 3 show the Grace Cathedral and Eucalyptus Grove lighting environments resampled into this coordinate system and resolution.

The next step is to multiply the red, green, and blue color channels of each of the reflectance field dataset images by the red, green, and blue colors of the corresponding pixel of the resampled lighting environment. For example, suppose that the pixel in the lighting environment corresponding to light coming directly from the right of the artifact is bright yellow. Then the reflectance field image illuminated from this direction in the light stage will be scaled so that it too is correspondingly bright and yellow. Thus, each image in the dataset becomes an accurate rendering of how the artifact would appear if illuminated by just its corresponding direction of light in the environment. The illuminated reflectance field dataset for the headdress is shown in Fig. 4.

The final step is to sum all of the images in the illuminated reflectance field dataset, producing a final rendered image showing the artifact as it would appear as illuminated by the entire sampled lighting environment at once. This procedure works because of the additive nature of light [9]: if an artifact is illuminated by

two sources of light and photographed separately as illuminated by each, then the sum of these two images will show what the artifact will look like as if illuminated by both sources at once. This assumes that the cameras taking these images are radiometrically calibrated; i.e. that the pixel values in each image are proportional to the amount of light received by the image sensor; we use the method of [5] in order to perform this calibration.

Figure 6 shows the result of illuminating the reflectance field dataset of the headdress by the two lighting environments in Figure 3 as well as a user-constructed lighting environment made with the interactive relighting tool described in Figure 5.

We note that this illumination technique yields results quite close to the physically correct answer of how the artifact would appear in the given light. Since the rendered image is a linear combination of real images of the artifact, the rendering will exhibit correct real-world illumination effects including those of anisotropic specularities, iridescence, self-shadowing, mutual illumination, subsurface scattering, translucency, and complex surface microstructure; and thus is a faithful rendering of what the artifact would look like in the specified environments. This is notable given that there is no modeling of the artifact’s surface geometry nor any steps to derive reflectance data for the artifact.

It should be mentioned that the technique will not yield perfect results in all cases. Scenes with very concentrated light sources should produce very sharp shadows on the artifact; such shadows will become slightly blurred using this technique since the reflectance fields are acquired at a finite resolution. The technique may also produce incorrect results if the spectrum of either the illumination or of the artifact’s reflectance has significant spectral peaks or valleys; carrying out the calculation solely on trichromatic RGB pixel values may fail to yield precise color balance in the renderings in such cases. Finally, this technique assumes that the artifact is illuminated by an even field of illumination; additional data acquisition and rendering would be required to show an artifact as it would be illuminated by dappled light or in partial shadow, or in close proximity to other artifacts or light sources¹.

¹This additional data could be recorded by using pixel-addressable video

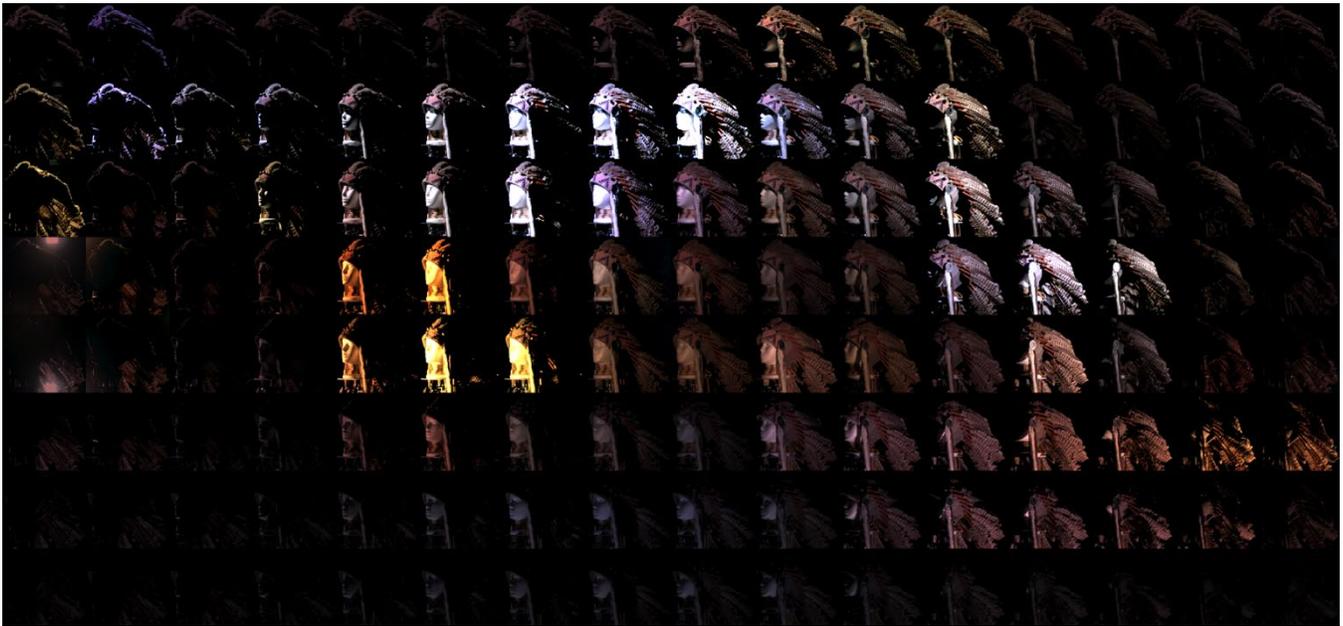


Figure 4: **Lighting a Reflectance Field Dataset** A reflectance field dataset is illuminated by coloring each image in the dataset according to the color and intensity of the illumination coming from the corresponding direction in a sampled lighting environment. The images in this figure are colored according to the illumination captured in Grace Cathedral (see Figure 3 using the light probe technique in [3]). A final image of the artifact as illuminated by the environment is obtained by adding together all of the transformed images in the dataset; such an image can be seen in Figure 6.

4.1 Interactive Reflectance Field Illumination

We have written an interactive computer program that implements our relighting technique in real time on a standard PC equipped with an OpenGL graphics card. Screen snapshots from this program are shown in Figure 5. The program operates in two modes; in the first, the user can choose from a variety of captured lighting environments with which to illuminate the artifact, and can rotate the environment about the y-axis to see the light from the environment reflect differently off of the artifact. In the second mode, the user can construct the lighting environment by hand using a number of light sources. For each source, the user selects its intensity, color, and direction, as well as whether the light is a hard source such as a point light or a soft source such as an area light. As the user moves the lights, the appearance of the artifacts updates interactively at over twenty frames per second on contemporary mid-range PCs. The demo uses compressed versions of the reflectance field datasets in order to compute the renderings; this degrades the quality of the renderings slightly but makes it possible to process the quantity of data necessary at interactive rates. We have found that being able to interactively control realistic illumination significantly helps a user sense both the geometric and material properties of an artifact.

In the work so far we only render the artifact from a static viewpoint. In the next section we describe how the artifact can be rendered from novel viewpoints in addition to being rendered under arbitrary illumination.

5 Results

We have applied the photometric digitization technique to a variety of cultural artifacts chosen to exhibit complex geometry and reflectance properties that would be difficult to capture using traditional digitization techniques. Some descriptive information on the artifacts themselves is presented in Section 8.

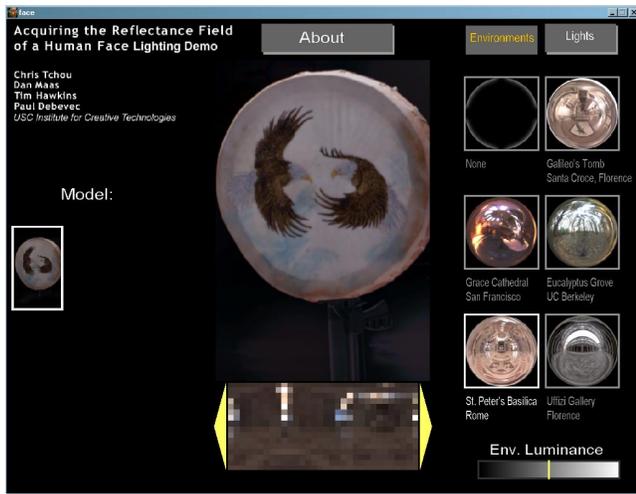
projectors, rather than uniform light sources, to illuminate the artifact.

Figure 6 shows a Lakota headdress synthetically illuminated by light captured from the Grace cathedral and Eucalyptus grove environments in Figure 3, as well as a user-specified lighting environment. Despite the headdress' complex and in places fuzzy geometry and its generally complex reflectance properties, the headdress appears realistically illuminated by each environment. A rendering of the headdress from a novel viewing angle is shown in Figure 10.

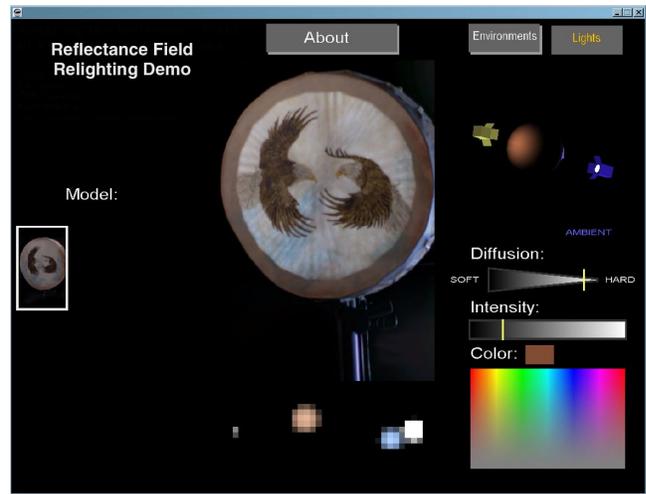
Figure 8 shows an otter skin cap and an animal bone choker necklace being captured and then illuminated by the Grace cathedral environment. The figure illustrates the need to capture the reflectance fields of shiny artifacts in high dynamic range - using multiple varying exposures to properly record both albedo reflection and specular highlights. Image (c) was illuminated using only standard single-exposure imagery in which highlight values were clipped, and yields a diffuse appearance to the reflective abalone necklace decoration. Image (d) used high dynamic range images to illuminate the artifact and thus properly replicates the shininess of the abalone.

Figure 7 shows two original light stage images and two synthetically illuminated renderings of a flat drum. Since the skin of the drum is thin and translucent, it becomes significantly illuminated even when lit from behind. Since this effect is captured in the reflectance field, the drum head will properly light up when illuminated from behind by a sampled illumination environment such as by the bright yellowish altar in the Grace cathedral lighting environment.

Figure 9 shows an approach to photometrically recording jewelry and clothing by having them be worn by an actual person. The light stage is outfitted with a chair and the reflectance field of the subject wearing the clothing is captured. Capturing clothing and jewelry in this manner can help us to better understand its physical and artistic design, and greater context is provided when the clothing is worn by a member of the culture from which the artifacts originate.



(a)



(b)

Figure 5: **An Interactive Illumination Program** The program seen above allows reflectance fields of cultural artifacts to be interactively illuminated with (a) different sampled illumination environments or (b) arbitrary user-specified illumination.



(a)



(b)



(c)



(d)

Figure 6: **Synthetically Illuminating an Artifact** (a) One of 1,728 original images in the reflectance field dataset of a headdress. (b) The headdress synthetically illuminated by the environmental lighting captured in Grace cathedral in Figure 3. (c) The headdress synthetically illuminated by the environmental lighting captured in a eucalyptus grove in Figure 3. (d) The headdress synthetically illuminated by a user-constructed lighting environment using the interactive relighting program seen in Figure 5. The renderings exhibit all of the artifact's complex properties of specularly, anisotropic reflection, translucency, and mutual illumination; such effects are usually challenging to model, represent, and render using currently available techniques for artifact digitization.

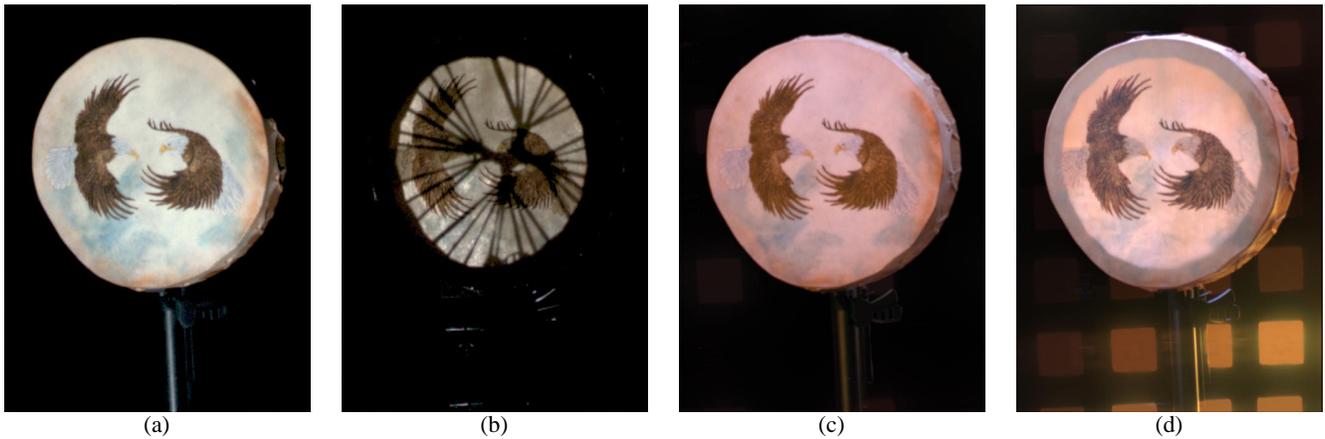


Figure 7: Capturing Translucency (a) The reflectance field of a Flat Drum is captured in the light stage. (b) When the drum is lit from behind, the front of the drum lights up due to the translucency of the drum head. The shadow of the strings that tighten the drum is also visible. (c) This image of the reflectance field of the drum illuminated by the Grace cathedral lighting environment in Figure 3 does not reveal the translucency since the environment is situated so that more light strikes the front of the drum than the back. (d) Rotating the lighting environment so that more of the light comes from behind the drum reveals the drum's translucency. A traditionally scanned 3D model of the drum with a diffuse texture map would not be able to reproduce this effect.

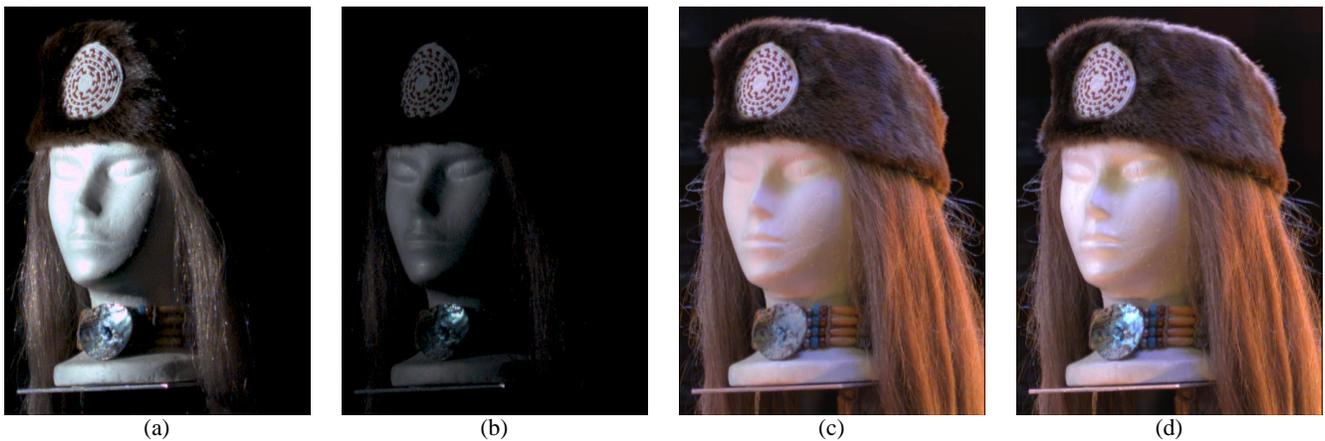


Figure 8: The Need for High Dynamic Range Datasets (a) a dataset image of an otter headband and an abalone shell necklace. The bright reflection in the necklace is too bright to be captured accurately by the video camera; the pixel values in the highlight have been clipped to the maximum pixel value. (b) A dataset image taken on a second pass of the light stage with a 3-stop neutral density filter placed in front of the camera lens. While most of the image is too dark to be useful, the specular highlight in the shell is properly imaged within the range of sensitivity of the camera. The two images can be combined into a single high dynamic range image as in [5]. Relighting the dataset with the method described in Section 4 will produce correct results only if the high dynamic range dataset is used. (c) A rendering of the artifacts using only the low dynamic range imagery captured from a single pass of the light stage as in (a); the abalone necklace piece appears to made of a diffuse plaster-like material. (d) A rendering of the artifacts using high dynamic range imagery captured from multiple passes of the light stage faithfully produces the bright iridescent specular reflection in the necklace.



(a)



(b)



(c)



(d)



(e)

Figure 9: **Capturing Jewelry and Clothing** The light stage can have a chair mounted within it that allows the appearance of traditional clothing and jewelry to be captured as worn by a person, in this case by a member of the particular culture of the artifacts' origin. (a) and (b) show tribal elder George Randall being photometrically recorded in the light stage wearing the otter fur hat, a deerskin Ghostdance shirt, a hair pipe necklace, and holding a prayer staff (Section 8 provides some more detailed information about the artifacts.) Both images used an extended shutter speed of ten seconds to capture an image of all of the lighting directions being illuminated at once. (c) and (d) show close-ups of Randall in the Grace Cathedral and a manually specified lighting environment. (e) shows a wider view of Randall with a prayer staff synthetically re-illuminated in the Grace Cathedral lighting environment.

6 Future Work

A benefit of the technique presented in this paper is that no geometric model of the artifact need be supplied or acquired. However, the technique as so far presented does not allow the viewpoint of the object to be changed, which is the principal way to give a sense of the geometry and reflectance of an object. One technique to render a reflectance field of an artifact from novel viewpoints would be to acquire a geometric model of the artifact, texture map the model with illuminated renderings of the artifact, and then render the texture-mapped model from different points of view. In this paper, we wish to model artifacts whose geometry can not easily be scanned with currently available 3D scanning techniques. As a result, we propose an image-based approach to rendering artifacts from novel viewpoints.

The Light Field [10] and Lumigraph [8] works suggest acquiring a two-dimensional dataset of images from different viewpoints situated on a viewing surface around an object and then composing new viewpoints by sampling similar pixel rays found in existing images. Since this is done entirely with images, no three-dimensional model of the object is required. However, to produce sharp renderings from novel viewpoints the spacing of the image viewpoints must be very fine, which for our purposes would multiply

the already great quantity of data required to represent an object's reflectance field from just one viewpoint. In order to reduce the amount of data required, a practical suggestion would be to acquire more widely spaced viewpoints and geometrically interpolate between them using view interpolation [20].

To show the potential merit of such an approach, we produced two rendered views (Figures 10(a) and 10(c)) of the headdress made from reflectance fields acquired by two simultaneously running digital video cameras. From these two static viewpoints, we used manually provided image correspondences to produce an intermediate viewpoint as seen in Figure 10(b). Ideally, one would like to find such correspondences manually using an optic flow technique. While finding such correspondences can be challenging [14], we note that the problem may be easier in our case since we have many pictures under different lighting conditions of the object seen from the two viewpoints. Thus, matching could be performed on a much higher-dimensional space of image pixels (under all the different lighting conditionals) to help disambiguate correspondences.

Figure 11 shows a possible augmentation of the light stage apparatus that would capture such datasets. Instead of just a few cameras, a linear array of sixteen cameras aimed at the subject is distributed along a second semicircular arm. Each time the arm

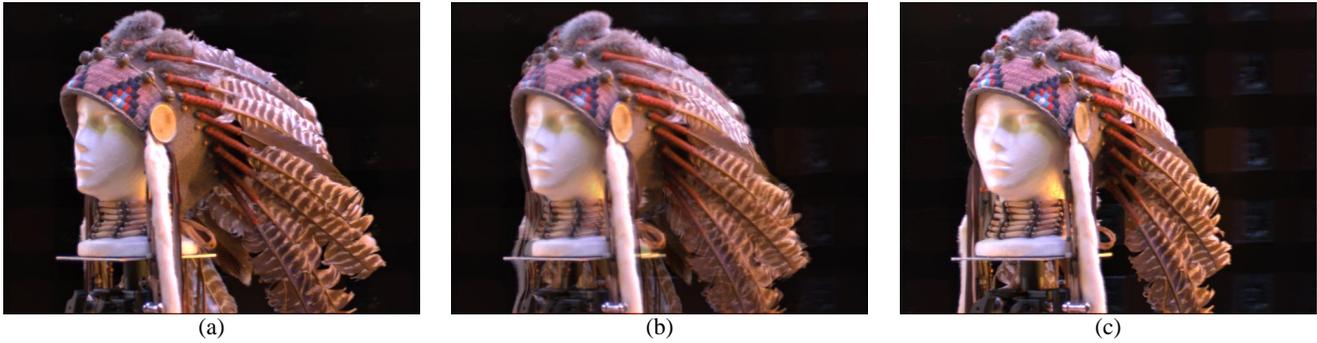


Figure 10: **Rendering the Artifact from Arbitrary Viewpoints (a) and (c)** Renderings of the artifact illuminated by the Grace Cathedral environment from different viewing directions. **(b)** An illuminated rendering of the artifact from an intermediate novel viewpoint made using the image-based rendering technique of view interpolation [20]. Using this technique, the artifact can be rendered from any viewpoint if reflectance field data is captured from a sufficiently dense sampling of discrete viewing directions. An augmented light stage that would automate such a capture process is shown in 11.

of lights goes around, a reflectance field of the artifact is captured from a variety of latitudes but from just one longitudinal direction. To capture views spaced out in longitude as well, there are two possibilities. The first is that the array of cameras would rotate in fixed angular increments after each rotation of the lights; the other is that the artifact itself would be rotated by a motion-controlled platform after each lighting pass.



Figure 11: **An Evolved Light Stage** includes a stationary array of cameras, seen at right, that record the artifact from a one-dimensional array of directions, and a motorized platform that can rotate the artifact a set number of degrees for every rotation of the light stage arm. The device would capture a complete light field of the artifact for every direction of incident illumination, allowing artifacts to be rendered from any viewpoint as well as illuminated from any direction.

In this manner, we could effectively capture an entire light field [10, 8] of the artifact as illuminated from every possible incident illumination direction, and from such a dataset could render the artifact from arbitrary angles and in arbitrary illumination, all without having geometric information for the artifact. For the moment, we leave this acquisition process for future work.

Another important avenue for future work will be to capture higher spectral resolution of the light reflected from the object as illuminated by the light stage. Treating reflectance and incident illumination with just three spectral bands (R, G, and B) is an approximation that can create color rendition problems when either the object's reflectance or the illuminant has a complex spectrum. One

technique for this would be to use a monochrome camera equipped with a multispectral filter wheel to capture an object's reflectance field at a variety of spectral bands. Another technique would be to place a Liquid Crystal Tunable Filter (LCTF) [7] in front of the camera to choose the reflectance bands. Combined with multispectral measurements of incident illumination (taken, for example, by placing an LCTF in front of the camera imaging a mirrored ball), such datasets could yield significantly more accurate renditions of artifacts under novel illumination conditions (such as under fire-light from a torch or fluorescent light in a museum), even if the target rendering space remains just (R,G,B).

7 Conclusion

In this paper we have presented an alternative technique for photometrically acquiring computer graphics models of real-world cultural artifacts. In this work we acquire reflectance fields of the artifacts – image datasets that directly measure how an artifact transforms incident illumination into radiant imagery – rather than surface geometry and texture maps. The method allows the artifacts to be rendered under arbitrary illumination conditions, including image-based illumination sampled from the real world. In this work we have focused on rendering artifacts from the same viewpoints from which their imagery is captured, but have shown that image-based rendering techniques can allow rendering from novel viewpoints to be performed as well. We demonstrated realistic illuminated renderings of a variety of cultural artifacts which would be challenging to model, represent, and render using current digitization techniques. It is our hope that this line of research will eventually help yield practical methods for digitizing photometrically accurate models of cultural artifacts, as well as provide insight into improving current techniques.

8 About the Artifacts

The artifacts featured in this paper are from the private collection of George Randall, a tribal elder from the White Earth Chippewa Reservation in Minnesota.

The headdress in Figure 6 is a classic war bonnet from the Lakota tribe made around 1920 at the Rosemont Reservation in South Dakota near Little Big Horn. The main head band is made from a Union Army blanket, and the headdress features bead work, metal bells, white and black ermine fur, colored ribbons, and turkey feathers augmented with tufts of rabbit fur at the tips. All the materials are organic except for the bells and mirrors, which were typical trade goods.

The drum in Figures 7 and 5 is a ceremonial flat drum made of

stretched elk stomach held together by leather lacing. The painted design shows two eagles, or *hanbli*, a prominent icon in the tribe's narrative tradition.

The otter skin hat in Figure 8 was made in California from a freshwater otter skin from Canada. The hat is approximately twenty years old and is typical of a style of headdress in use for approximately two hundred years.

The choker necklace in Figure 8 is also approximately twenty years old and is made from "hair pipe" (hollowed out bone from small animals), trade beads, and an abalone shell at the front.

The Ghostdance shirt in Figure 9 is sewn from baby deer skin. The front features six strands of horse hair wrapped with red thread, as well as three horse shoe prints indicating that the wearer owns three horses. The back of the shirt features three ermine tails attached to abalone shell circles.

These styles of clothes and jewelry are typical of Native American tribes spanning from the Northeast to Wisconsin such as the Seneca, Mohawk and Wampanoag.

Acknowledgements

We would like to express our gratitude to George Randall for making himself and his collection of artifacts available for this project and our great thanks to Maya Martinez-Smith for organizing his visit and making this collaboration possible. We thank Chris Tchou for his extensive contributions to the software employed in creating the renderings and both Chris Tchou and Dan Maas for writing the interactive reflectance field visualization program. We also wish to thank Brian Emerson for his modeling for figure 11, Andy Gardner for helping with the video editing, and Bobbie Halliday and Jamie Waese for their invaluable production assistance. This work was supported by the University of Southern California, the United States Army, and TOPPAN Printing Co., Inc. and does not necessarily reflect any corresponding positions or policies and no official endorsement should be inferred.

References

- [1] BARIBEAU, R., COURNOYER, L., GODIN, G., AND RIOUX, M. Colour three-dimensional modelling of museum objects. *Imaging the Past, Electronic Imaging and Computer Graphics in Museum and Archaeology* (1996), 199–209.
- [2] CURLESS, B., AND LEVOY, M. A volumetric method for building complex models from range images. In *SIGGRAPH 96* (1996), pp. 303–312.
- [3] DEBEVEC, P. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH 98* (July 1998).
- [4] DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. Acquiring the reflectance field of a human face. *Proceedings of SIGGRAPH 2000* (July 2000), 145–156. ISBN 1-58113-208-5.
- [5] DEBEVEC, P. E., AND MALIK, J. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97* (August 1997), pp. 369–378.
- [6] DEBEVEC, P. E., TAYLOR, C. J., AND MALIK, J. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *SIGGRAPH 96* (August 1996), pp. 11–20.
- [7] GAT, N. Real-time multi- and hyper-spectral imaging for remote sensing and machine vision: an overview. In *Proc. 1998 ASAE Annual International Mtg.* (Orlando, Florida, July 1998).
- [8] GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. The Lumigraph. In *SIGGRAPH 96* (1996), pp. 43–54.
- [9] HAEBERLI, P. Synthetic lighting for photography. Available at <http://www.sgi.com/grafica/synth/index.html>, January 1992.
- [10] LEVOY, M., AND HANRAHAN, P. Light field rendering. In *SIGGRAPH 96* (1996), pp. 31–42.
- [11] LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. The digital michelangelo project: 3d scanning of large statues. *Proceedings of SIGGRAPH 2000* (July 2000), 131–144. ISBN 1-58113-208-5.
- [12] MALZBENDER, T., GELB, D., AND WOLTERS, H. Polynomial texture maps. *Proceedings of SIGGRAPH 2001* (August 2001), 519–528. ISBN 1-58113-292-1.
- [13] MARSCHNER, S. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, August 1998.
- [14] MCMILLAN, L., AND BISHOP, G. Plenoptic Modeling: An image-based rendering system. In *SIGGRAPH 95* (1995).
- [15] NIMEROFF, J. S., SIMONCELLI, E., AND DORSEY, J. Efficient rendering of naturally illuminated environments. *Fifth Eurographics Workshop on Rendering* (June 1994), 359–373.
- [16] ROCCHINI, C., CIGNONI, P., AND MONTANI, C. Multiple textures stitching and blending on 3d objects. *Eurographics Rendering Workshop 1999* (June 1999). Held in Granada, Spain.
- [17] RUSHMEIER, H., BERNARDINI, F., MITTLEMAN, J., AND TAUBIN, G. Acquiring input for rendering at appropriate levels of detail: Digitizing a pieta. *Eurographics Rendering Workshop 1998* (June 1998), 81–92. ISBN 3-211-83213-0. Held in Vienna, Austria.
- [18] SATO, Y., WHEELER, M. D., AND IKEUCHI, K. Object shape and reflectance modeling from observation. In *SIGGRAPH 97* (1997), pp. 379–387.
- [19] TURK, G., AND LEVOY, M. Zippered polygon meshes from range images. In *SIGGRAPH 94* (1994), pp. 311–318.
- [20] WILLIAMS, L., AND CHEN, E. View interpolation for image synthesis. In *SIGGRAPH 93* (1993).
- [21] WONG, T.-T., HENG, P.-A., OR, S.-H., AND NG, W.-Y. Image-based rendering with controllable illumination. *Eurographics Rendering Workshop 1997* (June 1997), 13–22.
- [22] WOOD, D. N., AZUMA, D. I., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D. H., AND STUETZLE, W. Surface light fields for 3d photography. *Proceedings of SIGGRAPH 2000* (July 2000), 287–296. ISBN 1-58113-208-5.
- [23] Y.CHEN, AND MEDIONI, G. Object modeling from multiple range images. *Image and Vision Computing* 10, 3 (April 1992), 145–155.

Realistic Visualisation of the Pompeii Frescoes

Kate Devlin and Alan Chalmers

University of Bristol
Merchant Venturers Building
Woodland Road
BRISTOL BS8 1UB

devlin@cs.bris.ac.uk

Alan.Chalmers@bris.ac.uk

ABSTRACT

Three dimensional computer reconstruction provides us with a means of visualising past environments, allowing us a glimpse of the past that might otherwise be difficult to appreciate. Many of the images generated for this purpose are photorealistic, but no attempt has been made to ensure they are physically and perceptually valid. We are attempting to rectify these inadequacies through the use of accurate lighting simulation. By determining the appropriate spectral data of the original light sources and using them to illuminate a scene, the viewer can perceive a site and its artefacts in close approximation to the original environment. The richly decorated and well-preserved frescoes of the House of the Vettii in Pompeii have been chosen as a subject for the implementation of this study. This paper describes how, by using photographic records, modelling packages and luminaire values from a spectroradiometer, a three dimensional model can be created and then rendered in a lighting visualisation system to provide us with images that go beyond photorealistic, accurately simulating light behaviour and allowing us a physically and perceptually valid view of the reconstructed site. A method for capturing real flame and incorporating it in a virtual scene is also discussed, with the intention of recreating the movement of a flame in an animated scene.

KEYWORDS

Computer graphics, reconstructions, archaeology, visualization, visual perception.

1. INTRODUCTION

The application of computer graphics to the field of archaeology is becoming more commonplace. From providing the archaeologist with an aid to interpretation to giving the public an animated glimpse of the past, the use of realistic graphics

provides a powerful tool for modelling multi-dimensional aspects of archaeological data. Sites and artefacts can be reconstructed and visualised in 3D space, providing a safe and controlled method of studying past environments. This new perspective may enhance our understanding of the conditions in which our ancestors lived and worked. To date, however, limitations have been imposed with regard to the validity of these reconstructions [1]. The concept of realistic image synthesis centres on generating scenes with an authentic visual appearance. The modelled scene should not only be physically correct but also perceptually equivalent to the real scene it portrays [10] and this research seeks to address the problems encountered in the realistic simulation of archaeological sites.

Today our world is lit by bright and steady light, but past societies relied on daylight and flame for illumination. There is well-documented archaeological evidence for the use of flame; the presence of hearths, the remains of lamps and historical documentation where it exists all provide a source of information regarding the use of artificial light. If we consider our perception of the world we inhabit today and compare the modern lighting with that of the past, it is evident that there are significant differences in how it appears [2]. It would seem, therefore, that the photo-realistic site reconstructions often produced are flawed in regard to lighting conditions. Although they may look "real" their validity cannot be guaranteed as no attempt has been made to use physically accurate values for ancient light sources and surface reflectance. They owe more to an artist's imagination than to an interpretation based on numerical simulation. The commonly used software packages base the lighting conditions on daylight, fluorescent light or filament bulbs and not the lamp and candlelight that would have been used in the past. In some cases the reconstructions are lit with physically impossible lighting values. Our perception of past environments should consider the lighting conditions of that time - the use of natural daylight and the use of flame in a variety of forms. The different fuel types of each of these sources will affect our perception of a scene, and this needs to be taken into account [6]. Our perception of colour is affected by the amount and nature of light reaching the eye, so by simulating the behaviour of the appropriate type of light in an environment it should be possible to demonstrate how it may have looked in the past. The goal is to produce images that recreate accurately the visual appearance of an environment illuminated by flame.

2. LUMINAIRES

The luminosity of flame is due to glowing particles of solids in laminar flux, the colour of which is primarily related to the emission from incandescent carbon particles. A typical fuel/air wick flame consists of three distinct zones: the inner core, the blue intermediate zone and the outer cone [5]. The different zones of the flame produce different emissions depending on the fuel type and environment conditions.

Previous work on modelling flame has focussed on large-scale flames such as fires [4], fireballs and explosions [15] [13] or more generic flames [17][18][19]. Inakage introduced a simplified candle flame model [7], which Raczkowski extended to incorporate the dynamic nature of the flame [14]. In this study we are interested in the perception of objects illuminated by different fuel types.

2.1 Building the luminaires

The acquisition of valid experimental data is of vital importance as the material used may have had a significant influence on the perception of the ancient environment. In consultation with the Department of Archaeology at the University of Bristol, various light sources were recreated. These included processed beeswax candles, tallow candles (of vegetable origin), unrefined beeswax candles, a selection of reeds coated in vegetable tallow, a rendered animal fat lamp, and an olive oil lamp.



Figure 1. Experimental archaeology: reconstructing ancient light sources

The appropriate sources for this project were judged to be olive oil lamps, the most readily available fuel type for that area. Water was added to some of the lamps to keep them cool whilst being carried and to stop the oil from sticking. Salt was added to others to make the oil burn for longer.

Detailed spectral data was gathered using a spectroradiometer, allowing us to measure the absolute value of the spectral characteristics without making physical contact with the flame. This device can measure the emission spectrum of a light source from 380nm to 760nm, in 5nm wavelength increments, giving an accurate breakdown of the emission lines generated by the combustion of a particular fuel. The measurements were performed in a completely dark room and the device was aimed at a board coated with a 99% optically

pure Eastman Kodak Standard white powder, which diffusely reflects the aggregate incident light. Ten readings were taken for each lamp type and an average was calculated. This data can be used to create an RGB colour model for use in rendering the scene.

3. THE POMPEII FREScoes

The House of the Vettii in Pompeii is one of the best-preserved and decorated buildings in this World Heritage site, and is the most frequently visited building in Pompeii [12]. The rich colours and extensive use of artistic techniques such as *trompe l'oeil*, along with its magnificent state of preservation draws millions of visitors through its rooms each year. However, the impact of time and tourism on such a site has led to serious deterioration. Computer reconstruction of the House of the Vettii allows us to view it as it might have been when it was in use before the eruption of Vesuvius in AD 79.

The room chosen for the study was an *oecus*, or reception room, which opens onto a colonnaded garden. The high quantity of red and yellow pigments in this room was of specific interest to our study. Its three walls are richly decorated by intricate frescoes in the IV Style, also termed the “illusionist style” (Figure 2). Descriptions of how frescoes are created appear in Classical literature. It involves the application of colour pigments to wet plaster so that the plaster and the paint are merged and dry together, creating a permanent and vivid display. The fact that the House of the Vettii was immediately and sympathetically restored around the frescoes has meant that the paint colours have been well preserved.

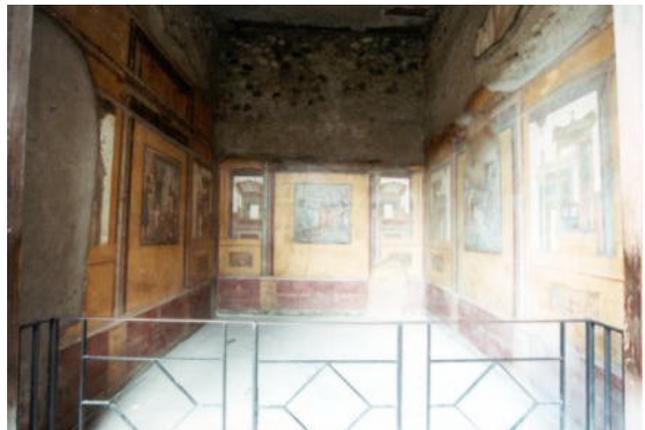


Figure 2. The room as it appears today

The frescoes in the room were recorded photographically. A colour chart was included at either side of each photo to permit calibration, identify illumination levels and allow any gradient in the light to be calculated. Using a scale plan of the room a 3D model was generated.

3.1 Converting the luminaire data

The resulting luminaire data obtained from the spectroradiometer was then converted to RGB values to enable display on a computer monitor. It is essential, when converting the detailed spectrum data from the spectroradiometer into values representing the red, green and blue portions of the spectrum, that this conversion is calculated in a perceptually valid way, as defined by the CIE (Commission International de l'Eclairage) 1931 1-degree standard observer¹.

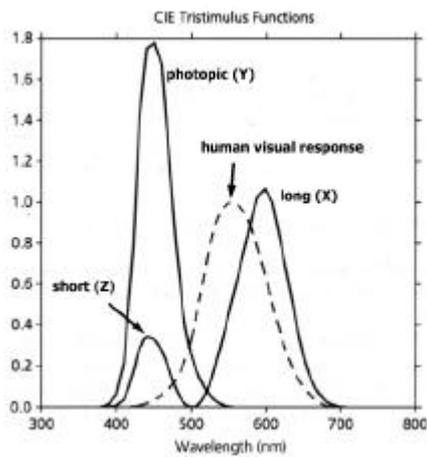


Figure 3. CIE Tristimulus functions

Figure 3 shows the functions for the X, Y and Z channels. The Y channel measures the luminance of a source, and the X and Z channels measure the chromaticity. This information is more useful when broken down as follows:

$$\text{If we let } x = \frac{X}{(X + Y + Z)} \text{ and } y = \frac{Y}{(X + Y + Z)}$$

then we can calculate the exact colour values for the red, green and blue sections of the spectrum, disregarding luminance. For a canonical set of VDU phosphors: RED(x,y) = (0.64, 0.33); GREEN(x,y) = (0.29, 0.60); and BLUE(x,y) = (0.15, 0.06) [22].

Radiance [22], a lighting visualization system, was used to render the images. It contains a function (*rcalc*) to convert the xyY coordinates to RGB values. Radiance then takes these RGB values and accurately simulates the associated light source behaviour in a modelled scene.

¹ This system specifies a perceived colour as a tristimulus value indicating the luminance and chromaticity of a stimulus as it is perceived in a 1-degree field around the foveal centre.

3.2 Modelling the flame

The type of flames that we are interested in are categorised as diffusion wick flames, in which heat transfer from the flame cause a steady production of flammable vapour. Observations from the reconstructed light sources showed that all of the flames were small and (ignoring the effect of air turbulence for now) fairly steady.

Rather than attempting to model the shape of the flame accurately for this initial stage of the project, video footage of the real flames from the reconstructed light sources was processed using computer vision techniques, the shape of the flame extracted and the real flame incorporated in the virtual scene. To capture the flame a 'blue screen' technique was used. This simple technique is widely employed in the film industry, and is used to cut an object from its background surroundings. Filming the flame against an evenly coloured, matt background enables thresholding of each frame, which is used to identify and dismiss a background colour, effectively separating the flame from any unwanted parts of the scene. The background colour should be chosen so that it does not occur within the foreground. As the intermediate zone of the flame is generally blue in colour, a green screen rather than a blue screen should be used (Figure 4).



Figure 4: Video still of the flame against a green screen

Thresholding is the process of identifying a range of colour, and changing all areas within an image that fall into this range to another specified colour. Simple, solid objects can easily be separated from a background, but a flame produces some of its own difficulties. It is useful to have static and even lighting on the background to simplify the thresholding process. Filming a flame creates a problem in that it is itself a light source and may therefore disrupt otherwise static lighting, producing unwanted background lighting effects. Parts of a flame may be translucent or partly transparent, so seepage of the background colour into what we identify as the flame may occur. This is difficult to avoid, but we can compensate for this at a later stage by deliberately seeping the background colour of the modelled scene into the flame. By splitting the video stream into separate frames an animated sequence can be achieved.

Once an object has been separated from its original background, it can then be placed in a new scene. If used sensibly, the object can be blended into the new background to give the appearance that this is an unaltered, original scene.

The illumination of the flame in the environment was achieved by approximating the shape of the flame by a series of *illum* spheres [22], as shown in figure 5(a) and included in a virtual scene 5(b). The material type *illum* is an invisible light source. When viewed directly, the object made from *illum* is invisible but it still emits light. The number and size of the spheres can of course be varied to achieve a better “fit” to the shape of the flame for each frame of the video sequence. Once the images have been rendered the original picture of the flame is pasted into the scene. Some care needs to be taken here to accurately position the original picture of the flame and blend it into the scene.



Figure 5: (a) Simple luminaire model (b) real flame in virtual environment

A set of programs was written to take a sequence of flame pictures as input and to output data representing the flame, to any desired level of accuracy [16]. This provides an efficient method of incorporating the real flame in a synthetic image.

3.3 Changes in perception in flame-lit conditions

Conversion of the spectral data to RGB does of course lead to an approximation of the colours present. In future, for more accuracy we will need to consider calculating the convolution of the emission spectrum of the light source with the reflectance curve of the material under examination. However, even with the approximation, significant perceptual differences related to fuel type are achieved. These simulations can be validated with real scenes [11]. Figure 6 shows a test scene including a MacBeth colour chart illuminated with (a) a 55W electric bulb, (b) an olive oil fuel. The difference in fuel type has a discernible effect on the appearance of the MacBeth chart. The apparent differences indicate that it is important for archaeologists to view such

artwork under (simulated) original conditions rather than under modern lighting. It is of course impossible to investigate these sensitive sites with real flame sources.



**Figure 6: The effect of different fuel types
(a) modern (b) olive oil**

The initial results of this on the Pompeii frescoes can be seen below. It is noticeable that the lamp-lit scenes (Figures 7b – 7d) can be perceived as warmer in appearance when compared to the modern light (Fig. 7a), with the yellow and red pigments particularly well emphasised. The appearance of the three-dimensional *trompe l’oeil* art is also influenced.

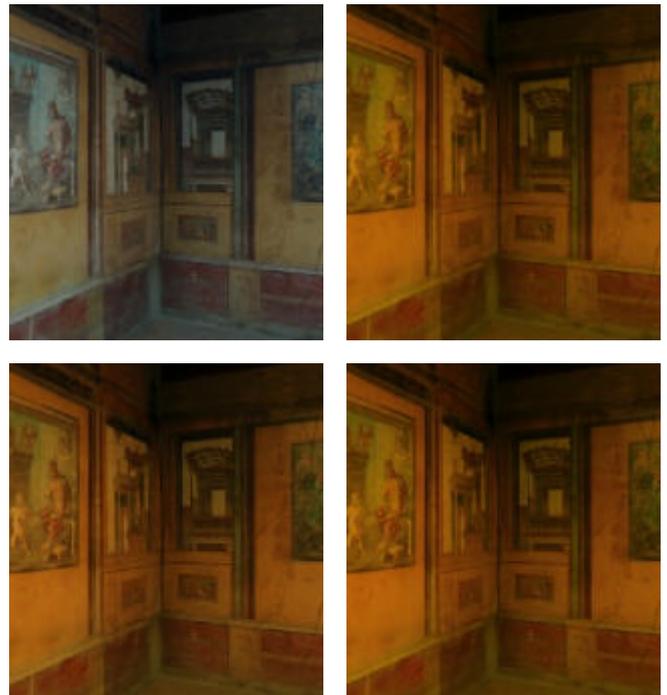


Figure 7. Clockwise from top left: (a) modern lighting (b) olive oil lamp (c) olive oil lamp with salt (d) olive oil lamp with water

These are preliminary images only, and current work involves the addition of models of appropriate Roman furniture and artefacts to the scene. This will not only create a more realistic scene, but will allow archaeologists to investigate the appearance of objects under their original lighting conditions. It is important to remember that a reconstruction is only one glimpse of many valid interpretations. Various configurations of lamps are also being modelled to provide a number of possible scenarios.

4. FUTURE WORK

In flame-lit environments the range of light can vary greatly over short distances. Human vision ranges over nine orders of magnitude whereas the dynamic range of most display devices covers only two, thus some form of compensation is required to map the light-dependent way we will view a scene [20][3][21].

The ultimate aim of realistic graphics is the creation of images that provoke the same response and sensation as a viewer would have to a real scene, i.e. the images are physically and perceptually accurate when compared to reality. Given that the aim is to create an environment that can be perceived as real, future work involving tone mapping will allow us to gain more perceptual accuracy of the scene. Furthermore, the use of an eyetracking device to measure involuntary eye movement will enable us to define which areas of the room are emphasised under different lighting conditions, and will give us an insight into the effectiveness of the three-dimensional paint techniques employed in the frescoes. Above all this, the need to establish a metric for realism through comparison with reality is essential if the images are to be of full use [11], and work in the future will attempt to quantify how "real" our reconstructions actually are.

5. CONCLUSIONS

To date, the work has provided us with a means of viewing a reconstructed site under its original lighting, and allows us to place a real flame in a virtual environment. The method of incorporating a real flame in a rendered image also provides for movement by means of a sequence of frames, so that an animation of the scene can have a dynamic flame inserted in it.

This method of visualization of past environments provides a safe and controlled manner in which the archaeologist can test hypotheses regarding perception and purpose of colour in decoration and artefacts. Computer generated imagery indistinguishable from the real physical environment will be of substantial benefit to the archaeological community, and this research is one method of moving beyond the current trend of photo-realistic graphics into physically and perceptually realistic scenes which are ultimately of greater use to those investigating our past.

6. ACKNOWLEDGMENTS

Many thanks to Ian Roberts for his work regarding flame modelling, and Helen Legge for her input in Italy. We would also like to thank the Defence Science and Technology Laboratory

(formerly DERA), and in particular Marilyn Gilmore, for their help and financial assistance.

7. REFERENCES

- [1] Barcelo, J.A., Forte, M. and Sanders, D.H., eds. *Virtual Reality in Archaeology*. (2000) ArchaeoPress.
- [2] Chalmers A., Green C. and Hall M. "Firelight: Graphics and Archaeology", Electronic Theatre, SIGGRAPH 00, New Orleans, July 2000.
- [3] Ferwerda, J. A., Pattanaik, S., Shirley, P. and Greenberg, D. P. "A Model of Visual Adaptation for Realistic Image Synthesis". *Proceedings of ACM SIGGRAPH 96* (August 1996) Addison Wesley, pp. 249-258.
- [4] Gardner G.Y., "Modeling Amorphous Natural Features", in *SIGGRAPH 94 Course Notes 22* (1994).
- [5] Gaydon, A.G. and Wolfard, H.G. *Flames: Their Structure, Radiation and Temperature*. (1979) Chapman and Hall.
- [6] Green, C. *The Visualisation of Ancient Lighting Conditions*. Project Thesis submitted in support of the degree of Bachelor of Science in Computer Science, University of Bristol (1999).
- [7] Inakage M., "A Simple Model of Flames", Computer Graphics Around the World, ed. T.S.Chua, T.L.Kunii, *Proceedings of Computer Graphics International*, (1990) Springer-Verlag, pp. 71-81.
- [8] Mavrodineanu R. et al., "Analytical Flame Spectroscopy, Selected Topics". (1970) MacMillan.
- [9] McNamara A., Chalmers A., Troscianko T. and Gilchrist I. "Comparing real & synthetic scenes using human judgements of lightness". In B. Peroche and H. Rushmeier, editors, *Rendering Techniques 2000*, Springer Wien.
- [10] McNamara, A., Chalmers, A., Troscianko, T. and Reinhard, E., "Fidelity of Graphics Reconstructions: A Psychophysical Investigation". *Proceedings of the 9th Eurographics Workshop on Rendering* (June 1998) Springer Verlag, pp. 237 - 246.
- [11] McNamara, A. and Chalmers, A., *Image Quality Metrics*, Image Quality Metrics Course Notes, SIGGRAPH 00, (July 2000).
- [12] Nappo, S., *Pompeii: Guide to the Lost City*. (1998) Weidenfeld and Nicolson.
- [13] Perlin K., Hoffert E.M., "Hypertexture", *Proceedings of ACM SIGGRAPH 89* (1989) pp. 253-262, 1989.
- [14] Raczkowski J. "Visual Simulation and Animation of a laminar Candle Flame". *International Conference on Image Processing and Computer Graphics*, (1996) Poland.
- [15] Reeves W.T., "Particle Systems - A Technique for Modeling a Class of Fuzzy Objects". *Proceedings of ACM SIGGRAPH 83* (1983) pp. 359-376.

- [16] Roberts, I. *Modelling Realistic Flame*. Project Thesis submitted in support of the degree of Bachelor of Science in Computer Science, University of Bristol (2001).
- [17] Rushmeier, Holly E. "Rendering Participating Media: Problems and Solutions from Application Areas". *Proceedings of the Fifth Eurographics Workshop on Rendering* (June 1995) Springer-Verlag.
- [18] Sakas G., "Cloud modeling for visual simulators", in G. von Bally and H.I. Bjelkhagen, editors, *Optics for protection of man and environment against natural and technological disasters* (1993) Elsevier Science Publishers B.V., pp.323-333.
- [19] Stam J., Fiume E., "Turbulent Wind Fields for Gaseous Phenomena". *Proceedings of SIGGRAPH 93* (1993) pp.369-376.
- [20] Tumblin, J. and Rushmeier, H., "Tone Reproduction for Realistic Images", *IEEE Computer Graphics and Applications* (November 1993) 13(6), pp. 42 – 48.
- [21] Ward Larson, G., Rushmeier, H. and Piatko, C. "A Visibility Matching Tone Operator for High Dynamic Range Scenes", *IEEE Transactions on Visualization and Computer Graphics* 3 (1997) no. 4, pp. 291 – 306.
- [22] Ward Larson, G. and Shakespeare, R., *Rendering with RADIANCE: The art and science of lighting simulation*. (1998) Morgan Kaufman.

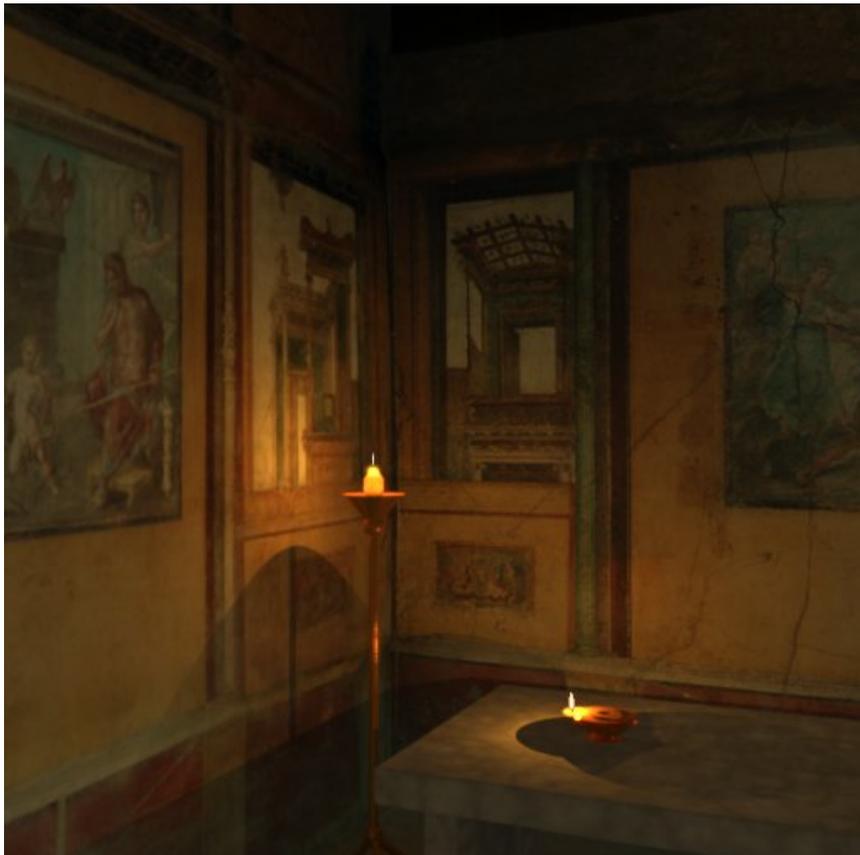
Computer Representation of the House of the Vettii, Pompeii



Under modern (55w) light



Under olive oil lamp



Under olive oil lamp, with furniture to show shadow effects

Comparing Real & Synthetic Scenes using Human Judgements of Lightness

Ann McNamara Alan Chalmers
Department of Computer Science

Tom Troscianko Iain Gilchrist
Department of Experimental Psychology
University of Bristol Bristol
mcnamara@cs.bris.ac.uk

Abstract. Increased application of computer graphics in areas which demand high levels of realism has made it necessary to examine the manner in which images are evaluated and validated. In this paper, we explore the need for including the human observer in any process which attempts to quantify the level of realism achieved by the rendering process, from measurement to display. We introduce a framework for measuring the perceptual equivalence (from a lightness perception point of view) between a real scene and a computer simulation of the same scene. Because this framework is based on psychophysical experiments, results are produced through study of vision from a *human* rather than a *machine* vision point of view. This framework can then be used to evaluate, validate and compare rendering techniques.

1 Introduction

The aim of realistic image synthesis is the creation of accurate, high quality imagery which faithfully represents a physical environment, the ultimate goal being to create images which are perceptually indistinguishable from an actual scene. Rendering systems are now capable of accurately simulating the distribution of light in an environment. However, physical accuracy does not ensure that the displayed images will have authentic visual appearance. Reliable image quality assessments are necessary for the evaluation of realistic images synthesis algorithms. Typically the quality of an image synthesis method is evaluated using numerical techniques which attempt to quantify fidelity using image to image comparisons (often comparisons are made with a photograph of the scene that the image is intended to depict).

Several image quality metrics have been developed whose goals are to predict the *visible* differences between a pair of images. It is well established that simple approaches, such as mean squared error (MSE), do not provide meaningful measures of image fidelity, more sophisticated techniques are necessary. As image quality assessments should correspond to assessments made by humans, a better understanding of features of the **H**uman **V**isual **S**ystem (HVS) should lead to more effective comparisons, which in turn will steer image synthesis algorithms to produce more realistic, reliable images. Any feature of an image not visible to a human is not worth computing. Results from psychophysical experiments can reveal limitations of the HVS. However, problems arise when trying to incorporate such results into computer graphics algorithms. This is due to the fact that, often, experiments are designed to explore a single dimension of the HVS at a time under laboratory conditions. The HVS comprises

many complex mechanisms, which rather than function independently, often work in conjunction with each other, making it more sensible to examine the HVS as a whole. Rather than attempting to reuse results from previous psychophysical experiments, new experiments are needed which examine the complex response HVS as a *whole* instead of trying to isolate features for individual investigations. In this work we study the ability of the HVS to perceive albedo and the impact of rendering quality on *this task*. Rather than deal with atomic aspects of perception, this study examines a complete task in a more realistic setting.

Human judgements of lightness are compared in real scenes, and synthetic images. Correspondence between these judgements is then used as an indication of the fidelity of the synthetic image.

1.1 Lightness Perception

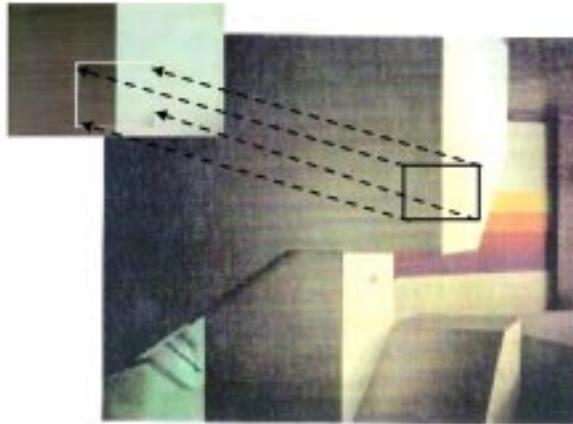


Fig. 1. Importance of depth perception for lightness constancy

Lightness is apparent reflectance, brightness is apparent intensity of the illuminant. Reflectance is the proportion of light falling on an object that is reflected to the eye of the observer. Reflectance (albedo) is constant, the perception of lightness depends of reflectance [1]. Gilchrist [8] showed that the perception of the degree of “lightness” of a surface patch (i.e. whether it is white, gray or black) is greatly affected by the perceived distance and orientation of the surface in question, as well as the perceived illumination falling on the surface - where the latter were experimentally manipulated through a variety of cues such as occlusion, or perspective.

Perception of the lightness of patches varying in reflectance may thus be a suitable candidate for the choice of visual task. It is simple to perform, and it is known that lightness constancy depends on the successful perception of lighting and the 3D structure of a scene, for example figure 1. When viewed in isolation the patches on the top left hand corner appear to be of different luminance. However, when examined in the context of the entire scene, it can be seen that the patches have been cut from the edge of the stairwell, and is perceived as an edge where the entire stairwell has the same luminance. Eliminating the depth cues means the patches are perceived as different, demonstrating

the dependency of lightness perception on the correct perception of three dimensional structure, [10]. As the key features of any scene are illumination, geometry and depth, the task of lightness matching encapsulates all three key characteristics into one task. This task is particularly suited to this experimental framework, apart from being simple to perform it also allows excellent control over experimental stimuli. Subsequent sections describe an experimental framework, with such a lightness matching task at the core, to allow human observers to compare real and synthetic scenes.

The remainder of this paper is divided into the following sections. In Section 2, we describe previous research. In Section 3, we describe the steps taken to build the experiment in order to facilitate easy human comparison between real and synthetic scene, we also discuss the actual organisation of participants in terms of scheduling. Section 4 describes the experiment, the results are presented in section 5 and finally, conclusions are drawn in section 6.

2 Previous Work

Models of visual processing enable the development of perceptually based error metrics for rendering algorithms that will reduce the computational demands of rendering while preserving the visual fidelity of the rendered images. Much research investigating this issue is under way.

Using a simple five sided cube as their test environment Meyer et al [13] presented an approach to image synthesis comprising separate physical and perceptual modules. They chose diffusely reflecting materials to build a physical test environment. Each module is verified using experimental techniques. The test environment was placed in a small dark room. Radiometric values predicted using a radiosity lighting simulation of a basic environment are compared to physical measurements of radiant flux densities in the real environment. Then the results of the radiosity calculations are transformed to the RGB values for display, following the principles of colour science.

Measurements of irradiation were made at 25 locations in the plane of the open face for comparison with the simulations. Results show that irradiation is greatest near the centre of the open side of the cube. This area provides the best view of the light source and other walls. The calculated values are much higher than the measurements. In summary, there is good agreement between the radiometric measurements and the predictions of the lighting model. Meyer et al. then proceeded by transforming the validated simulated value to values displayable on a television monitor. A group of twenty experimental participants were asked to differentiate between real environment and the displayed image, both of which were viewed through the back of a view camera. They were asked which of the images was the real scene. Nine out of the twenty participants (45%) indicated that the simulated image was actually the real scene, i.e. selected the wrong answer, revealing that observers were simply guessing. Although participants considered the overall match and colour match to be good, some weaknesses were cited in the sharpness of the shadows (a consequence of the discretisation in the simulation) and in the brightness of the ceiling panel (a consequence of the directional characteristics of the light source). The overall agreement lends strong support to the perceptual validity of the simulation and display process.

Rushmeier et al. [15] used perceptually based metrics to compare image quality to a captured image of the scene being represented. The image comparison metrics were derived from [4],[6], [11]. Each is based on ideas taken from image compression techniques. The goal of this work was to obtain results from comparing two images using these models that were large if large differences between the images exist, and small

when they are almost the same. These suggested metrics include some basic characteristics of human vision described in image compression literature. First, within a broad band of luminance, the visual system senses relative rather than absolute luminances. For this reason a metric should account for luminance variations, not absolute values. Second, the response of the visual system is non-linear. The perceived “brightness” or “lightness” is a non-linear function of luminance. The particular non-linear relationship is not well established and is likely to depend on complex issues such as perceived lighting and 3-D geometry. Third, the sensitivity of the eye depends on the spatial frequency of luminance variations. The perceptual metrics derived were used to compare images in a manner that roughly corresponds to subjective human vision, in particular the Daly model performed very well.

The **Visible Difference Predictor (VDP)** is a perceptually based image quality metric proposed by Daly [4]. Myskowski [14] realised the VDP had many potential applications in realistic image synthesis. He completed a comprehensive validation and calibration of VDP response via human psychophysical experiments. Then, he used the VDP local error metric to steer decision making in adaptive mesh subdivision, and isolated regions of interest for more intensive global illumination computations. The VDP was tested to determine how close VDP predictions come to subjective reports of visible differences between images by designing two human psychophysical experiments. Results from these experiments showed a good correspondence with VDP results for shadow and lighting pattern masking and in comparison of the perceived quality of images generated as subsequent stages of indirect lighting solutions.

McNamara et al [12] built an experimental framework to facilitate human comparison between real and synthetic scene. They ran a series of psychophysical experiments in which human observers were asked to compare regions of a real physical scene with regions of the computer generated representation of that scene. The comparison involved lightness judgements in both the generated image and the real scene. Results from these experiments showed that the visual response to the real scene and a high fidelity rendered image was similar. The work presented in this paper extends this work to investigate comparisons using three dimensional objects as targets, rather than simple regions. This allows us to examine scene characteristics such as shadow, object occlusion and depth perception.

3 Experimental Design

This section outlines the steps involved in building a well articulated scene containing three dimensional objects placed within a custom built environment to evoke certain perceptual cues such as lightness constancy, depth perception and the perception of shadows. Measurements of virtual environments are often inaccurate. For some applications¹ such estimation of input may be appropriate. However, for these purposes an accurate description of the environment is essential to avoid introducing errors at such an early stage. Also, once the global illumination calculations have been computed, it is important to display the resulting image in the correct manner while taking into account the limitations of the display device. As we are interested in comparing different rendering engines, it is vital that we minimise errors in the model and display stages, this means then that any errors arising can be attributed to the rendering technique employed to calculate the image. This study required an experimental set-up comprised of a real

¹The level of realism required is generally application dependent. In some situations a high level of realism is not required, for example games, educational techniques and graphics for web design.



Fig. 2. The test environment showing real environment and computer image.

environment and a computer representation of that three dimensional environment. The measurements required for this study, the equipment used to record them are described herein, along with the rendering process employed to generate the physical stimuli.

3.1 The Real Scene

The test environment was a five sided box shown in figure 2. Several objects that were placed within the box for examination. All interior surfaces of the box were painted with white matt house paint. To accommodate the three dimensional objects, custom paints were mixed, using precise ratios to serve as the basis for materials in the scene. To ensure correct, accurate ratios were achieved, 30ml syringes were used to mix paint in parts as shown in Table 1. The spectral reflectance of the paints were measured using a TOPCON-100 spectroradiometer, these values were transformed to RGB tristimulus values following [16].

Appearance	% White	Reflectance	Patch#	Patch Reflectance
Black	0	0.0471	0	.0494
Dark Gray	10	0.0483	0	.0494
Dark Gray	20	0.0635	2	.0668
Dark Gray	30	0.0779	4	.0832
Dark Gray	40	0.0962	6	.1012
Dark Gray	50	0.1133	7	.1120
Gray	60	0.1383	9	.1224
Gray	70	0.1611	14	.1680
Light Gray	80	0.2002	15	.2259
Light Gray	90	0.3286	19	.3392
Light Gray	95	0.4202	23	.4349
Almost White	97.5	0.5292	26	.5512
Almost White	98.25	0.5312	26	.5512
White	100	0.8795	29	.8795

Table 1. Paint Reflectance along with Reflectance of Corresponding Patch

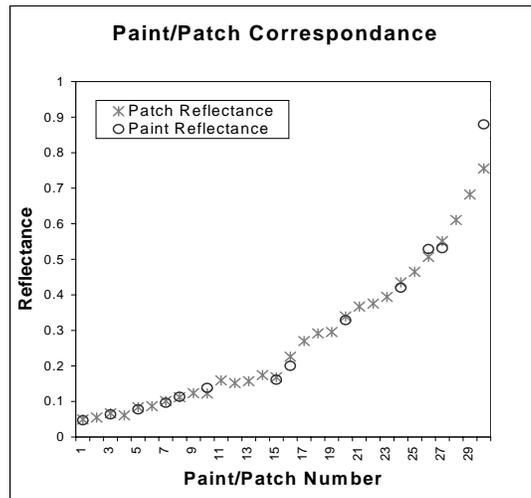


Fig. 3. Correspondence of Patches to Paints

As in [12] a small, front-silvered, high quality mirror was incorporated into the set up to allow the viewing conditions to facilitate alternation between the two settings, viewing of the original scene or viewing of the modelled scene on the computer monitor. When the optical mirror was in position, subjects viewed the original scene. In the absence of the optical mirror the computer representation of the original scene was viewed. The angular subtenses of the two displays were equalised, and the fact that the display monitor had to be closer to the subject for this to occur, was allowed for by the inclusion of a +2 diopter lens in its optical path; the lens equated the optical distances of the two displays.

3.2 Illumination

The light source consisted of a 24 volt quartz halogen bulb mounted on optical bench fittings at the top of the test environment. This was supplied by a stabilised 10 amp DC power supply, stable to 30 parts per million in current. The light shone through a 70 mm by 115 mm opening at the top of the enclosure. Black masks, constructed of matt cardboard sheets, were placed framing the screen and the open wall of the enclosure, a separate black cardboard sheet was used to define the eye position. An aperture in this mask was used to enforce monocular vision, since the VDU display did not permit stereoscopic viewing.

3.3 The Graphical Representations

Ten images were considered for comparison to the real scene, they are listed here along with the aims that we hoped to achieve from the comparison.

1. **Photograph:** Comparison to a photograph is needed to enable us to evaluate our method to more traditional image comparison metrics. The reasoning behind this is that most current techniques compare to “reality” by comparing to a captured

image. We wanted to see if this is equivalent to comparing to a real physical environment and so included a photograph, taken with a digital camera, as one of our test images.

2. **Radiance: 2 Ambient Bounces:** A Radiance [17] image generated using 2 ambient bounces is generally considered to be a high quality image. Here we wanted to determine if 2 ambient bounces gives a similar perceptual impression to an 8 ambient bounce image which is more compute intensive.
3. **Radiance: 8 Ambient Bounces:** We wanted to investigate if there was a marked difference using a Radiance image generated using 8 ambient bounces, as this involves considerably more compute time, and might not be necessary i.e. may not provide any more perceptual information than an image rendered using 2 ambient bounces.
4. **Radiance: 8 Ambient Bounces BRIGHT:** This image had its brightness increased manually to see if this affected perception. The brightness was doubled (i.e. the intensity of each pixel was multiplied by 2) to see what, if any effect this had on the perception of the image.
5. **Radiance: Default:** Image generated with the default Radiance parameters. This would determine whether extra compute time makes a significant difference. The default image renders in a very short time, however ambient bounces of light are absent, we wanted to compare this to imagery where interreflections were catered for.
6. **Radiance: Controlled Errors in Estimate Reflectance Values:** The RGB values for the materials were set to equal values to see what difference, if any, this made compared to using measured values. A poor perceptual response to this image would confirm our suspicion that material properties must be carefully quantified if an accurate result is required. This comparison, and the next, was to demonstrate the importance of using exact measurements rather than estimations for material values.
7. **Radiance: Controlled Errors in Estimate of Light Source:** The RGB values for the light source were set to equal values to see what difference this made compared to using measured values. This experiment will show the necessity of measuring emission properties of sources in an environment if an accuracy is the aim.
8. **Radiance: Tone Mapped:** We wanted to investigate the difference tone mapping would make to our test image. Tone mapping transforms the radiance values computed by the rendering engine to values displayable on a display device in a manner that preserves the *subjective* impression of the scene. The Tone Mapping Operator (TMO) used here was introduced by Ferwerda et al. [5]. Although the image examined does not have a very high dynamic range, we were interested to see the effects tone mapping would have on image perception.
9. **Renderpark: Raytraced:** This was a very noisy image generated using stochastic raytracing. This experiment was designed to see how under-sampling would affect perception. Here the effect of under-sampling is exaggerated but might give insights in to how much undersampling a rendering engine can "get away with" without affecting perceptual performance.
10. **Renderpark: Radiosity:** Finally, to investigate the effects of meshing in a radiosity solution, a poorly meshed radiosity image was used. We wanted to demonstrate the importance of using an accurate meshing strategy when employing radiosity techniques.

These images are shown in the accompanying colour plate.

The media used for stimulus presentation was a gamma corrected 20-inch monitor with the following phosphor chromaticity coordinates:

$$\begin{array}{cccc} x_r = 0.6044 & x_g = 0.2808 & x_b = .1520 & x_w = 0.2786 \\ y_r = 0.3434 & y_g = 0.6016 & y_b = .0660 & y_w = 0.3020 \end{array}$$

4 Experiment

Eighteen observers participated in the experiment, and were naive of the purpose of the experiment. All had normal or corrected-to-normal vision. Both condition order and trial order were fully randomised across subjects and conditions. Participants were given clear instructions.

4.1 Training on Munsell Chips

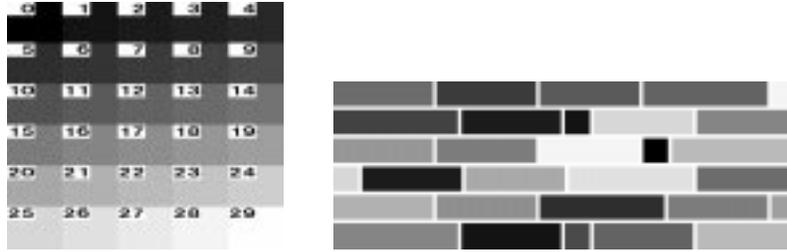


Fig. 4. Patch arrangement used to train participants with Reference Chart)

In [12], the task involved matching regions to a control chart which meant observers had to look away from the scene under examination to choose a match. Moving between scene and chart may affect adaptation to the scene in question, also the view point is not fixed, for this reason we decided to *train* participants on the control patches first. Once trained on the patches participants could then recall the match from memory. Training was conducted as follows. Observers were asked to select, from a numbered grid of 30 achromatic Munsell chips presented on a white background, a sample to match a second unnumbered grid (figure 4) simultaneously displayed on the same background, under constant illumination. The unnumbered grid comprised 60 chips. At the start of each experiment participants were presented with two grids, one an ordered numbered regular grid the other an unordered unnumbered irregular grid comprising one or more of the chips from the numbered grid. Both charts were hung on the wall approximately one meter from the participant. Each participant was asked to match the chips on the unnumbered grid to one of the chips on the numbered grid on the left. In other words they were to pick a numbered square on the left and place it right next to the grid on the right which in the grid would match it exactly. This is done in a random manner, a laser pointer² was used to point to the unnumbered chip under examination. Then the numbered chart was removed, and the unnumbered chart replaced by a similar chart but one where the chips had a different order. Participants repeated the task, this time working from memory to recall the number each chip would match to. The results of this training exercise are graphed in figure 5. The graph on the left shows the average

²non-invasive medium

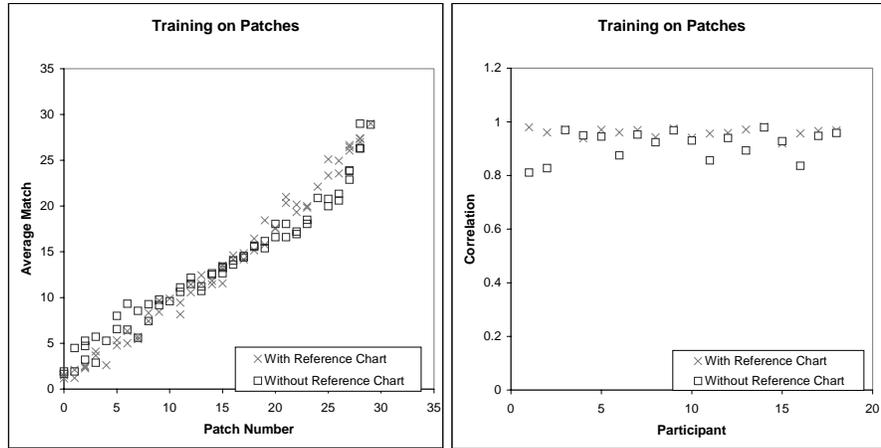


Fig. 5. Average of Matching to Training Patches with and without the reference chart shown on the right along with the Average Correlation for both cases on the left

match across 18 subjects, both with the reference chart and without the reference chart. The graph on the right shows the average correlation. This correlation gives an indication of the extent to which two sets of data are linearly related. A values close to 1 indicates a strong relationship, while a value of 0 signifies there is no linear relationship. A correlation of 1 would result if the participant matched each unnumbered patch to its corresponding numbered patch, in reality this is not the case and some small errors are made, what we need to determine is if the errors made when matching from memory i.e. without the chart are about the same size as the errors made with the reference chart in place. The correlation value when matching the patches with the chart in place is 0.96, and when matching from memory the result is 0.92, indicating a very small difference of 0.04 between the two conditions. From this small difference we can conclude that participants are *just as good* at matching the patches without the reference chart in place. Thus, this training paradigm proved to be reliable and stable. This has the dual benefit of speeding up the time taken per condition, as well as ensuring participants do not need to move their gaze from image to chart, thus eliminating any influence due to adaptation.

4.2 Matching to Images

Each participant was presented with a series of images, in a random order, one of which was the real environment. Participants were not explicitly informed which image was the physical environment. The images presented were the real scene, the photograph and the 9 rendered images. There were 17 different objects in the test environment, subjects were also asked to match the 5 sides of the environment (floor, ceiling, left

wall, back wall and right wall) giving a total of 21 matches. The paints used on the objects match to the training patches as shown in graph 3, and detailed in table 3.1. Participants were asked to judge the lightness of target objects in a random manner.

We chose this particular task - that of matching materials in the scene against a display of originals - because the task has a number of attractive features. First, Gilchrist [9, 7] has shown that the perception of lightness (the perceptual correlate of reflectance) is strongly dependent on the human visual system's rendition of both illumination and 3-D geometry. These are key features of perception of any scene and are in themselves complex attributes. However, the simple matching procedure used here depends critically on the correct representation of the above parameters. Therefore, the task should be sensitive to any mismatch between the original and the rendered scene. Secondly, the matching procedure is a standard psychophysical task and allows excellent control over the stimulus and the subject's response. The task chosen here corresponds closely to the methodology of Gilchrist [2, 9, 7] which permits simple measures (of lightness) to be made at locations in complex scenes. Ultimately, the task was chosen to be simple while also being sensitive to perceptual distortions in the scene.

5 Results

Results for each participant were recorded and analysed independently. The value (or gray level) chosen by each participant in the real scene was compared with the values chosen in the rendered image. For a rendered image to be a faithful reproduction, the values in both cases should be closely related. To examine this relationship we carried out a linear correlation for each subject. This correlation gives an indication of the extent to which two sets of data are linearly related. A values close to 1 indicates a strong relationship, whilst a value of 0 signifies there is no linear relationship. A correlation of 1 would result if the participant chose exactly the same gray level for each object in the real scene and rendered image. Correlation values are shown in table 2, and graphed as shown in the colour plate, the graph on the right shows these values averaged.

To examine *the pattern* of these correlations across participants we carried out ANalysis Of VAriance (ANOVA). ANOVA is a powerful set of procedures used for testing significance where two or more conditions are used, here 10 conditions were examined [3]. A *repeated measures within subjects* ANOVA was used. There was a significant effect of condition:

$$F(9, 153) = 80.3; p < .001$$

This equation can be read as follows, the F statistic equals 80.3, with 9 degrees of freedom (10 images), 153 degrees of freedom for the error term (calculated as a function of image combinations). The P value indicates the probability that these differences occur by *chance*. This is a repeated measures within subjects analysis of variance as each subject performed each condition.

This means there are statistically reliable differences between the conditions. This is to be expected as some images were deliberately selected for variation in quality.

The ANOVA showed there are significant differences in perception across images. Further analyses were carried out to investigate where these differences occur. These analyses took the form of a paired comparison t-test. Here we took the correlation between the real scene and the photograph, and compared it to the correlation of the real scene to the other images. Results from the correlations are shown in the following table.

Image	Mean Correlation with REAL
Photograph	.8918
* 2 Ambient Bounces	.843
8 Ambient Bounces	.884
Brightened 8 Ambient Bounces	.865
* Default	.337
* Controlled Error Materials	.692
Tone Mapped	.879
Controlled Error Illumination	.862
* Raytraced	.505
* Radiosity	.830

Table 2. Comparison of Rendered Images to Real Environment

A star in the table indicates a statistically significant difference, reflecting a reliable decrement in quality when compared to the photograph. The significant t values were as follows:

Two Ambient Bounces: ($t(17) = 3.11; p < .01$)

Default Image: ($t(17) = 12.4; p < .001$)

Guessed Materials Image: ($t(17) = 10.7; p < .001$)

Raytraced Image: ($t(17) = 9.36; p < .001$)

Radiosity Image: ($t(17) = 3.00; p < .01$)

The t statistic equals (take Two Ambient Bounces as an example) 3.11, with 17 degrees of freedom (18 participants). The probability, p of this distribution happening by chance is less than 0.01. This means that while there are some small differences between the results of matching to the photograph and matching to other images, these differences are not significant.

In summary, our results show that there is evidence that the 2 Ambient Bounces image, the Default image, the Controlled Error Materials image, the Raytraced image and the Radiosity image are perceptually degraded compared to the photograph. However, there is no evidence that the others images in this study are perceptually inferior to the photograph. From this we can conclude that the 8 Ambient Bounces image, the Brightened 8 Ambient Bounces image, the Tone Mapped image and the Controlled Error Illumination image are of the same perceptual quality as a photograph of the real scene.

6 Conclusions

We have introduced a method for measuring the perceptual equivalence between a real scene and a computer simulation of the same scene, from a lightness matching point of view. Because this model is based on psychophysical experiments, results are produced through study of vision from a human rather than a machine vision point of view.

By conducting a series of experiments, based on the psychophysics of lightness perception, we can estimate how much alike a rendered image is to the original scene. Results show that given a real scene and a faithful representation of that scene, the visual response function in both cases is similar.

Because the complexity of human perception and the computational expensive rendering algorithms that exist today, future work should focus on developing efficient methods from which resultant graphical representations of scenes yield the same perceptual effects as the original scene. To achieve this the full gamut of colour perception, as opposed to simply lightness, must be considered by introducing scenes of increasing complexity.

References

1. E. H. Adelson, *Lightness Perception and Lightness Illusions*, 339–351, MIT Press, 1999, pp. 339–351.
2. J. Cataliotti and A. Gilchrist, *Local and global processes in lightness perception*, *Perception and Psychophysics*, vol. 57(2), 1995, pp. 125–135.
3. H. Coolican, *Research methods and statistics in psychology*, Hodder and Stoughton, Oxford, 1999.
4. S. Daly, *The visible difference predictor: an algorithm for the assessment of image fidelity*, In A. B. Watson Editor, *Digital Images and Human Vision*, MIT Press, 1993, pp. 179–206.
5. J. A. Ferwerda, S.N. Pattanaik, P. Shirley, and D. P. Greenberg, *A model of visual adaptation for realistic image synthesis*, *Computer Graphics* **30** (1996), no. Annual Conference Series, 249–258.
6. J. Gervais, Jr. L.O. Harvey, and J.O. Roberts, *Identification confusions among letters of the alphabet*, *Journal of Experimental Psychology: Human Perception and Performance*, vol. 10(5), 1984, pp. 655–666.
7. A. Gilchrist, *Lightness contrast and failures of lightness constancy: a common explanation*, *Perception and Psychophysics*, vol. 43(5), 1988, pp. 125–135.
8. A. Gilchrist, S. Delman, and A. Jacobsen, *The classification and integration of edges as critical to the perception of reflectance and illumination*, *Perception and Psychophysics* **33** (1983), no. 5, 425–436.
9. A. Gilchrist and A. Jacobsen, *Perception of lightness and illumination in a world of one reflectance*, *Perception* **13** (1984), 5–19.
10. A. L. Gilchrist, *The perception of surface blacks and whites*, *Scientific American* **240** (1979), no. 3, 88–97.
11. J. L. Mannos and D. J. Sakrison, *The effects of a visual criterion on the encoding of images*, *IEEE Transactions on Information Theory* **IT-20** (1974), no. 4, 525–536.
12. A. McNamara, A. Chalmers, T. Troscianko, and E. Reinhard, *Fidelity of graphics reconstructions: A psychophysical investigation*, *Proceedings of the 9th Eurographics Rendering Workshop*, Springer Verlag, June 1998, pp. 237–246.
13. G. W. Meyer, H. E. Rushmeier, M. F. Cohen, D. P. Greenberg, and K. E. Torrance, *An Experimental Evaluation of Computer Graphics Imagery*, *ACM Transactions on Graphics* **5** (1986), no. 1, 30–50.
14. K. Myszkowski, *The visible differences predictor: Applications to global illumination problems*, *Rendering Techniques '98* (Proceedings of Eurographics Rendering Workshop '98) (New York, NY) (G. Drettakis and N. Max, eds.), Springer Wien, 1998, pp. 233–236.
15. H. Rushmeier, G. Ward, C. Piatko, P. Sanders, and B. Rust, *Comparing real and synthetic images: Some ideas about metrics*, *Eurographics Rendering Workshop 1995*, Eurographics, June 1995.
16. D. Travis, *Effective color displays*, Academic Press, 1991.
17. G. J. Ward, *The RADIANCE lighting simulation and rendering system*, *Proceedings of SIGGRAPH '94* (Orlando, Florida, July 24–29, 1994) (Andrew Glassner, ed.), *Computer Graphics Proceedings, Annual Conference Series*, ACM SIGGRAPH, ACM Press, July 1994, ISBN 0-89791-667-0, pp. 459–472.

The RADIANCE Lighting Simulation and Rendering System

Gregory J. Ward

Lighting Group
Building Technologies Program
Lawrence Berkeley Laboratory
(e-mail: GJWard@Lbl.Gov)

ABSTRACT

This paper describes a physically-based rendering system tailored to the demands of lighting design and architecture. The simulation uses a light-backwards ray-tracing method with extensions to efficiently solve the rendering equation under most conditions. This includes specular, diffuse and directional-diffuse reflection and transmission in any combination to any level in any environment, including complicated, curved geometries. The simulation blends deterministic and stochastic ray-tracing techniques to achieve the best balance between speed and accuracy in its local and global illumination methods. Some of the more interesting techniques are outlined, with references to more detailed descriptions elsewhere. Finally, examples are given of successful applications of this free software by others.

CR Categories: I.3.3 [Computer Graphics]: Picture/image generation - *Display algorithms*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *Shading*.

Additional Keywords and Phrases: lighting simulation, Monte Carlo, physically-based rendering, radiosity, ray-tracing.

1. Introduction

Despite voluminous research in global illumination and radiosity over the past decade, few practical applications have surfaced in the fields that stand the most to benefit: architecture and lighting design. Most designers who use rendering software employ it in a purely illustrative fashion to show geometry and style, not to predict lighting or true appearance. The designers cannot be blamed for this; rendering systems that promote flash over content have been the mainstay of the graphics industry for years, and the shortcuts employed are well-understood by the software community and well-supported by the hardware manufacturers.

Why has radiosity not yet taken off in the rendering market? Perhaps not enough time has passed since its introduction to the graphics community a decade ago [8]. After all, it took ray-tracing nearly that long to become a mainstream, commercial rendering technique. Another possibility is that the method is too compute-intensive for most applications, or that it simply does not fulfill enough people's needs. For example, most radiosity systems are not well automated, and do not permit general reflectance models or curved surfaces. If we are unable to garner support even from the principal beneficiaries, designers, what does that say of our chances with the rest of the user community?

Acceptance of physically-based rendering is bound to improve[†], but researchers must first demonstrate the real-life applicability of their techniques. There have been few notable successes in applying radiosity to the needs of practicing designers [6]. While much research has focused on improving efficiency of the basic radiosity method, problems associated with more realistic, complicated geometries have only recently gotten the attention they deserve [2,19,22]. For whatever reason, it appears that radiosity has yet to fulfill its promise, and it is time to reexamine this technique in light of real-world applications and other alternatives for solving the rendering equation [10].

There are three closely related challenges to physically-based rendering for architecture and lighting design: accuracy, generality and practicality. The first challenge is that the calculation must be accurate; it must compute absolute values in physical units with reasonable certainty. Although recent research in global illumination has studied sources of calculation error [1,20], few researchers bother to compute in physical lighting units, and even fewer have compared their results to physical experiments [15]. No matter how good the theory is, accuracy claims for simulation must ultimately be backed up with comparisons to what is being simulated. The second challenge is that a rendering program must be general. It is not necessary to simulate every physical lighting phenomenon, but it is important to do enough that the unsolvable rendering problems are either unimportant or truly exceptional. The third challenge for any rendering system is that it be practical. This includes a broad spectrum of requirements, from being reliable (i.e. debugged and tested) to being application-friendly, to producing good results in a reasonable time. All three of the above challenges must be met if a physically-based rendering package is to succeed, and all three must be treated with equal importance.

Radiance is the name of a rendering system developed by the author over the past nine years at the Lawrence Berkeley Laboratory (LBL) in California and the Ecole Polytechnique Federale de Lausanne (EPFL) in Switzerland. It began as a study in ray-tracing algorithms, and after demonstrating its potential for saving energy through better lighting design, acquired funding from the U.S. Department of Energy and later from the Swiss government. The first free software release was in 1989, and since then it has found an increasing number of users in the research and design community. Although it has never been a commercial product, *Radiance* has benefited enor-

[†]The term "physically-based rendering" is used throughout the paper to refer to rendering techniques based on physical principles of light behavior for local and global illumination. The term "simulation" is more general, referring to any algorithm that mimics a physical process.

mously from the existence of an enthusiastic, active and growing user base, which has provided invaluable debugging help and stress-testing of the software. In fact, most of the enhancements made to the system were the outcome of real or perceived user requirements. This is in contrast to the much of the research community, which tends to respond to intriguing problems before it responds to critical ones. Nine years of user-stimulated software evolution gives us the confidence to claim we have a rendering system that goes a long way towards satisfying the needs of the design community. Further evidence has been provided by the two or three design companies who have abandoned their own in-house software (some of which cost over a million dollars to develop) in favor of *Radiance*.

In this paper, we describe the *Radiance* system design goals, followed with the principal techniques used to meet these goals. We follow with examples of how users have applied *Radiance* to specific problems, followed by conclusions and ideas for future directions.

2. System Design Goals

The original goals for the *Radiance* system were modest, or so we thought. The idea was to produce an accurate tool for lighting simulation and visualization based on ray-tracing. Although the initial results were promising, we soon learned that there was much more to getting the simulation right than plugging proper values and units into a standard ray-tracing algorithm. We needed to overcome some basic shortcomings. The main shortcoming of conventional ray-tracing is that diffuse interreflection between surfaces is approximated by a uniform "ambient" term. For many scenes, this is a poor approximation, even if the ambient term is assigned correctly. Other difficulties arise in treating light distribution from large sources such as windows, skylights, and large fixtures. Finally, reflections of lights from mirrors and other secondary sources are problematic. These problems, which we will cover in some detail later, arose from the consideration of our system design goals, given below.

The principal design goals of *Radiance* were to:

1. Ensure accurate calculation of luminance
2. Model both electric light and daylight
3. Support a variety of reflectance models
4. Support complicated geometry
5. Take unmodified input from CAD systems

These goals reflect many years of experience in architectural lighting simulation; some of them are physically-motivated, others are user-motivated. All of them must be met before a lighting simulation tool can be of significant value to a designer.

2.1. Ensure Accurate Calculation of Luminance

Accuracy is one of the key challenges in physically-based rendering, and luminance (or the more general "spectral radiance") is probably the most versatile unit in lighting. Photometric units such as luminance are measured in terms of visible radiation, and radiometric units such as radiance are measured in terms of power (energy/time). Luminance represents the quantity of visible radiation passing through a point in a given direction, measured in lumens/steradian/meter² in SI units. Radiance is the radiometric equivalent of luminance, measured in watts/steradian/meter². Spectral radiance simply adds a dependence on wavelength to this. Luminance and spectral radiance are most closely related to a pixel, which is what the eye actually "sees." From this single unit, all other lighting metrics can be derived. Illuminance, for example, is the integral of luminance over a projected hemisphere (lumens/meter² or "lux" in SI units). Luminous intensity and luminous flux follow similar derivations. By computing the most basic lighting unit, our simulation will adapt more readily to new applications.

To assure that a simulation delivers on its promise, it is essential that the program undergo periodic validation. In our case, this means comparing luminance values predicted by *Radiance* to measurements of physical models. An initial validation was completed in 1989 by Grynberg [9], and subsequent validations by ourselves and others confirm that the values are getting better and not worse [14].

2.2. Model Both Electric Light and Daylight

In order to be general, a lighting calculation must include all significant sources of illumination. Daylight simulation is of particular interest to architects, since the design of the building facade and to a lesser degree the interior depends on daylight considerations.

Initially, *Radiance* was designed to model electric light in interior spaces. With the addition of algorithms for modeling diffuse interreflection [25], the software became more accurate and capable of simulating daylight (both sun and sky contributions) for building interiors and exteriors. The role of daylight simulation in *Radiance* was given new importance when the software was chosen by the International Energy Agency (IEA) for its daylight modeling task* [4].

2.3. Support a Variety of Reflectance Models

Luminance is a directional quantity, and its value is strongly determined by a material's reflectance/transmittance distribution function. If luminance is calculated using a Lambertian (i.e. diffuse) assumption, specular highlights and reflections are ignored and the result can easily be wrong by a hundred times or more. We cannot afford to lose directional information if we hope to use our simulation to evaluate visual performance, visual comfort and aesthetics.

A global illumination program is only as general as its local illumination model. The standard model of ambient plus diffuse plus Phong specular is not good enough for realistic image synthesis. *Radiance* includes the ability to model arbitrary reflectance and transmittance functions, and we have also taken empirical measurements of materials and modeled them successfully in our system [29].

2.4. Support Complicated Geometry

A lighting simulation of an empty room is not very interesting, nor is it very informative. The contents of a room must be included if light transfer is to be calculated correctly. Also, it is difficult for humans to evaluate aesthetics based on visualizations of empty spaces. Furniture, shadows and other details provide the visual cues a person needs to understand the lighting of a space. Modeling exteriors is even more challenging, often requiring hundreds of thousands of surfaces.

Although we leave the definition of "complicated geometry" somewhat loose, including it as a goal means that we shall not limit the geometric modeling capability of our simulation in any fundamental way. To be practical, data structure size should grow linearly (at worst) with geometric complexity, and there should be no built-in limit as to the number of surfaces. To be accurate, we shall support a variety of surface primitives, also ensuring our models are as memory-efficient as possible. To be general, we shall provide N-sided polygons and a mechanism for interpolating surface normals, so any reasonable shape may be represented. Finally, computation time should have a sublinear relationship to the number of surfaces so that the user does not pay an unreasonable price for accurate modeling.

*The IEA is a consortium of researchers from developed nations cooperatively seeking alternative energy sources and ways of improving energy efficiency in their countries.

2.5. Take Unmodified Input from CAD Systems

If we are to model complicated geometry, we must have a practical means to enter these models into our simulation. The creation of a complicated geometric model is probably the most difficult task facing the user. It is imperative that the user be allowed every means to simplify this task, including advanced CAD systems and input devices. If our simulation limits this process in any way, its value is diminished.

Therefore, to the greatest degree possible, we must accept input geometry from any CAD environment. This is perhaps the most difficult of the goals we have outlined, as the detail and quality of CAD models varies widely. Many CAD systems and users produce only 2D or wireframe models, which are next to useless for simulation. Other CAD systems, capable of producing true 3D geometric models, cannot label the component surfaces and associate the material information necessary for an accurate lighting simulation. These systems require a certain degree of user intervention and post-processing to complete the model. Even the most advanced CAD systems, which produce accurate 3D models with associated surface data, do not break surfaces into meshes suitable for a radiosity calculation. The missing information must either be added by the user, inferred from the model, or the need for it must be eliminated. In our case, we eliminate this need by using something other than a radiosity (i.e. finite element) algorithm.

CAD translators have been written for *AutoCAD*, *GDS*, *ArchiCAD*, *DesignWorkshop*, *StrataStudio*, *Wavefront*, and *Architrion*, among others. None of these translators requires special intervention by the user to reorient surface normals, eliminate T-vertices, or mesh surfaces. The only requirement is that surfaces must somehow be associated with a layer or identifier that indicates their material type.

3. Approach

We have outlined the goals for our rendering system and linked them back to the three key challenges of accuracy, generality and practicality. Let us now explore some of the techniques we have found helpful in meeting these goals and challenges.

We start with a basic description of the problem we are solving and how we go about solving it in section 3.1, followed by specific solution techniques in sections 3.2 to 3.5. Sections 3.6 to 3.9 present some important optimizations, and section 3.10 describes the overall implementation and use of the system.

3.1. Hybrid Deterministic/Stochastic Ray Tracing

Essentially, *Radiance* uses ray-tracing in a recursive evaluation of the following integral equation at each surface point:

$$L_r(\theta_r, \phi_r) = L_e(\theta_r, \phi_r) + \int_0^{2\pi} \int_0^\pi L_i(\theta_i, \phi_i) \rho_{bd}(\theta_i, \phi_i; \theta_r, \phi_r) |\cos\theta_i| \sin\theta_i d\theta_i d\phi_i \quad (1)$$

where:

- θ is the polar angle measured from the surface normal
- ϕ is the azimuthal angle measured about the surface normal
- $L_e(\theta_r, \phi_r)$ is the emitted radiance (watts/steradian/meter² in SI units)
- $L_r(\theta_r, \phi_r)$ is the reflected radiance
- $L_i(\theta_i, \phi_i)$ is the incident radiance
- $\rho_{bd}(\theta_i, \phi_i; \theta_r, \phi_r)$ is the bidirectional reflectance-transmittance distribution function (steradian⁻¹)

This equation is essentially Kajiyama's rendering equation [10] with the notion of energy transfer between two points replaced by energy passing through a point in a specific direction (i.e. the

definition of radiance). This formula has been documented many times, going back before the standard definition of ρ_{bd} [16]. Its generality and simplicity provide the best foundation for building a lighting simulation.

This formulation of the rendering problem is a natural for ray tracing because it gives outgoing radiance in terms of incoming radiance over the projected sphere, without any explicit mention of the model geometry. The only thing to consider at any one time is the light interaction with a specific surface point, and how best to compute this integral from spawned ray values. Thus, no restrictions are placed on the number or shape of surfaces or surface elements, and discretization (meshing) of the scene is unnecessary and even irrelevant.

Although it is possible to approximate a solution to this equation using uniform stochastic sampling (i.e. Monte Carlo), the convergence under most conditions is so slow that such a solution is impractical. For example, a simple outdoor scene with a ground plane, a brick and the sun would take days to compute using naive Monte Carlo simply because the sun is so small (0.5° of arc) in comparison to the rest of the sky. It would take many thousands of samples per pixel to properly integrate light coming from such a concentrated source.

The key to fast convergence is in deciding *what* to sample by removing those parts of the integral we can compute deterministically and gauging the importance of the rest so as to maximize the payback from our ray calculations. In the case of the outdoor scene just described, we would want to consider the sun as an important contribution to be sampled separately, thus removing the biggest source of variance from our integral. Instead of relying on random samples over the hemisphere, we send a single sample ray towards the sun, and if it arrives unobstructed, we use a deterministic calculation of the total solar contribution based on the known size and luminosity of the sun as a whole. We are making the assumption that the sun is not partially occluded, but such an assumption would only be in error within the penumbra of a solar shadow region, and we know these regions to represent a very small portion of our scene.

Light sources cause peaks in the incident radiance distribution, $L_i(\theta_i, \phi_i)$. Directional reflection and transmission cause peaks in the scattering function, ρ_{bd} . This will occur for reflective materials near the mirror angle, and in the refracted direction of dielectric surfaces (e.g. glass). Removing such peak reflection and transmission angles by sending separate samples reduces the variance of our integral at a comparatively modest cost. This approach was introduced at the same time as ray-tracing by Whitted [31]. Further improvements were made by adding stochastic sampling to the deterministic source and specular calculations by Cook in the first real linking of stochastic and deterministic techniques [5]. *Radiance* employs a tightly coupled source and specular calculation, described in [29].

3.2. Cached Indirect Irradiances for Diffuse Interreflection

No matter how successful we are at removing the specular reflections and direct illumination from the integral (1), the cost of determining the remaining diffuse indirect contributions is too great to recalculate at every pixel because this requires tracing hundreds of rays to reduce the variance to tolerable levels. Therefore, most ray-tracing calculations ignore diffuse interreflection between surfaces, using a constant "ambient" term to replace the missing energy.

Part of the reason a constant ambient value has been accepted for so long (other than the cost of replacing it) is that diffuse interreflection changes only gradually over surfaces. Thus, the contrast-sensitive eye usually does not object to the loss of subtle shading that accompanies an ambient approximation. However, the inaccuracies that result *are* a problem if one wants to know light levels or see the effects of daylight or

indirect lighting systems.

Since indirect lighting changes gradually over surfaces, it should be possible to spread out this influence over many pixels to obtain a result that is smooth and accurate at a modest sampling cost. This is exactly what we have done in *Radiance*. The original method for computing and using cached irradiance values [25] has been enhanced using gradient information [28].

The basic idea is to perform a full evaluation of Equation (1) for indirect diffuse contributions only as needed, caching and interpolating these values over each surface. Direct and specular components are still computed on a per-pixel basis, but hemispherical sampling occurs less frequently. This gives us a good estimate of the indirect diffuse contribution when we need it by sending more samples than we would be able to afford for a pixel-independent calculation. The approach is effectively similar to finite element methods that subdivide surfaces into patches, calculate accurate illumination at one point on each patch and interpolate the results. However, an explicit mesh is not used in our method, and we are free to adjust the density of our calculation points in response to the illumination environment. Furthermore, since we compute these view-independent values only as needed, separate form factor and solution stages do not have to complete over the entire scene prior to rendering. This can amount to tremendous savings in large architectural models where only a portion is viewed at any one time.

Figure 1 looks down on a diffuse sphere in a room with indirect lighting only. A blue dot has been placed at the position of each indirect irradiance calculation. Notice that the values are irregularly spaced and denser underneath the sphere, on the sphere and near the walls at the edges of the image. Thus, the spacing of points adapts to changing illumination to maintain constant accuracy with the fewest samples.

To compute the indirect irradiance at a point in our scene, we send a few hundred rays that are uniformly distributed over the projected hemisphere. If any of our rays hits a light source, we disregard it since the direct contribution is computed separately. This sampling process is applied recursively for multiple reflections, and it does not grow exponentially because each level has its own cache of indirect values.

Figure 2

Our hemisphere samples not only tell us the total indirect illumination, they also give us more detailed information about the locations and brightnesses of surfaces visible from the evaluation point. This information may be used to predict how irradiance will change as a function of point location and surface orientation, effectively telling us the first derivative (gradient) of the irradiance function. For example, we may have a bright reflecting surface behind and to the right of a darker surface as shown in Figure 2. Moving our evaluation point to the right would yield an increase in the computed irradiance (i.e. the translational gradient is positive in this direction), and our samples can tell us this. A clockwise rotation of the surface element would also cause an increase in the irradiance value (i.e. the

rotational gradient is positive in this direction), and our hemisphere samples contain this information as well. Formalizing these observations, we have developed a numerical approximation to the irradiance gradient based on hemisphere samples. Unfortunately, its derivation does not fit easily into a general paper, so we refer the reader to the original research [28].

Figure 3

Figure 3a,b. *Plots showing the superiority of gradient interpolation for indirect irradiance values. The reference curve is an exact calculation of the irradiance along the red line in Figure 1. The linear interpolation is equivalent to Gouraud shading between evenly spaced points, as in radiosity rendering. The Hermite cubic interpolation uses the gradient values computed by Radiance, and is not only smoother but demonstrably more accurate than a linear interpolation.*

Knowing the gradient in addition to the value of a function, we can use a higher order interpolation method to get a better irradiance estimate between the calculated points. In effect, we will obtain a smoother and more accurate result without having to do any additional sampling, and with very little overhead. (Evaluating the gradient formulas costs almost nothing compared to computing the hemisphere samples.)

Figure 3a shows the irradiance function across the floor of Figure 1, along the red line. The exact curve is shown overlaid with a linearly interpolated value between regularly spaced calculation points, and a Hermite cubic interpolation using com-

puted gradients. The cubic interpolation is difficult to separate from the exact curve. Figure 3b shows the relative error for these two interpolation methods, clearly demonstrating the advantage of using gradient information.

Caching indirect irradiances has four important advantages over radiosity methods. First, no meshing is required, since a separate octree data structure is used to hold the calculated values. This lifts restrictions on geometric shapes and complexity, and greatly simplifies user input and scene analysis. Second, we only have to compute those irradiances affecting the portion of the scene being viewed. This speeds rendering time under any circumstance, since our *view-independent* values may be reused in subsequent images (unlike values computed with importance-driven radiosity [20]). Third, the density of irradiance calculations is reduced at each level of interreflection, maintaining constant accuracy while reducing the time required to compute additional bounces. Fourth, the technique adapts to illumination by spacing values more closely in regions where there may be large gradients, without actually using the gradient as a criterion. This eliminates errors that result from using initial samples to decide sampling density [12], and improves accuracy overall. The gradient *is* used to improve value interpolation, yielding a smoother and more accurate result without the Mach-bands that can degrade conventional radiosity images.

3.3. Adaptive Sampling of Light Sources

Although sending one sample ray to each light source is quite reasonable for outdoor scenes, such an approach is impractical for indoor scenes that may have over a hundred light sources. Most rays in a typical calculation are in fact shadow rays. It is therefore worth our while to rule out light sources that are unimportant and avoid testing them for visibility.

The method we use in *Radiance* for reducing the number of shadow rays is described in [26]. A prioritized list of potential source contributions is created at each evaluation of Equation (1). The largest potential contributors (contribution being a function of source output, proximity and ρ_{bd}) are tested for shadows first, and we stop testing when the remainder of the source list is below some fraction of the unoccluded contributions. The remaining source contributions are then added based on statistical estimates of how likely each of them is to be visible.

Figure 4 shows a simple example of how this works. The left column represents our sorted list of potential light source contributions for a specific sample point. We proceed down our list, checking the visibility of each source by tracing shadow rays, and summing together the unobstructed contributions. After each test, we check to see if the remainder of our potential contributions has fallen below some specified fraction of our accumulated total. If we set our accuracy goal to 10%, we can stop testing after four light sources because the remainder of the list is less than 10% of our known direct value. We could either add all of the remainder in or throw it away and our value would still be within 10% of the correct answer. But we can do better than that; we can make an educated guess at the visibility of the remaining sources using statistics. Taking the history of obstructed versus unobstructed shadow rays from previous tests of each light source, we multiply this probability of hitting an untested source by the ratio of successful shadow tests at this point over all successful shadow tests ($2/(.9+.55+.65+.95) = 0.65$ in this example), and arrive at a reasonable estimate of the remainder. (If any computed multiplier is greater than 1, 1 is used instead.) Our total estimate of the direct contribution at this point is then the sum of the tested light sources and our statistical estimate of the remainder, or 1616 in this example.

We have found this method to be very successful in reducing the number of shadow test rays required, and it is possible to place absolute bounds on the error of the approximation. Most importantly, this type of adaptive shadow testing emphasizes contrast as the primary decision criterion. Contrast is defined as the difference between the luminance at a point and the background luminance divided by the background luminance. If a shadow boundary is below the visible contrast threshold, then an error in its calculation is undetectable by the viewer. Thus, this method produces no *visible* artifacts in its tradeoff of speed for accuracy. Accuracy is still lost in a controlled way, but the resulting image is subjectively flawless, due to the eye's relative insensitivity to absolute light levels.

Figure 5 shows a theater lighting simulation generated by *Radiance* in 1989. This image contains slightly over a hundred light sources, and originally took about 4 days to render on a Sun-4/260. (The equivalent of about 5 Vax-11/780's.) Using our adaptive shadow testing algorithm reduced the rendering time to 2 days for the same image[†]. The time savings for scenes with more light sources can be better than 70%, especially if the light sources have narrow output distributions, such as the spotlights popular in overlighted retail applications.

A different problem associated with ray-per-source shadow testing is *inadequate sampling* of large or nearby sources, which threatens simulation accuracy. For example, a single ray cannot adequately sample a fluorescent desk lamp for a point directly beneath it. The simplest approach for sources that are large relative to their distance is to send multiple sample rays. Unfortunately, breaking a source into pieces and sending many rays to it is inefficient for distant points in the room. Again, an adaptive sampling technique is the most practical solution.

In our adaptive technique, we send multiple rays to a light source if its extent is large relative to its distance. We recursively divide such sources into smaller pieces until each piece satisfies some size/distance criterion. Figure 6a shows a long, skinny light source that has been broken into halves repeatedly until each source is small enough to keep penumbra and solid angle errors in check. Figure 6b shows a similar subdivision of a

Figure 4

[†]The theater model was later rendered in [2] using automatic meshing and progressive radiosity. Meshing the scene caused it to take up about 100 Mbytes of memory, and rendering took over 18 hours on an SGI R3000 workstation for the direct component alone, compared to 5 hours in 11 Mbytes using *Radiance* on the same computer.

Figure 6

large rectangular source. A point far away from either source will not result in subdivision, sending only a single ray to some (randomly) chosen location on the source to determine visibility.

3.4. Automatic Preprocessing of "Virtual" Light Sources

Thus far we have accounted for direct contributions from known light sources, specular reflections and transmission, and diffuse interreflections. However, there are still transfers from specular surfaces that will not be handled efficiently by our calculation. A mirror surface may reflect sunlight onto a diffuse or semispecular surface, for example. Although the diffuse interreflection calculation could in principle include such an effect, we are returning to the original problem of insufficient sampling of an intense light source. A small source reflected specularly is still too small to find in a practical number of naive Monte Carlo samples. We have to know where to look.

We therefore introduce "virtual" light sources that do not exist in reality, but are used during the calculation to direct shadow rays in the appropriate directions to find reflected or otherwise transferred light sources. This works for any planar surface, and has been implemented for mirrors as well as prismatic glazings (used in daylighting systems [4]). For example, a planar mirror might result in a virtual sun in the mirror direction from the real sun. When a shadow ray is sent towards the virtual sun, it will be reflected off the mirror to intersect the real sun. An example is shown in Figure 7a. This approach is essentially the same as the "virtual worlds" idea put forth by Rushmeier [18] and exploited by Wallace [24], but it is only carried out for light sources and not for all contributing surfaces. Thus, multiple transfers between specular surfaces can be made practical with this method using intelligent optimization techniques.

The first optimization we apply is to limit the scope of a virtual light source to its affected volume. Given a specific source and a specific specular surface, the influence is usually limited to a certain projected volume. Points that fall outside this volume are not affected and thus it is not necessary to consider the source everywhere. Furthermore, multiple reflections of the source are possible only within this volume. We can thus avoid creating virtual-virtual sources in cases where the volume of one virtual source fails to intersect the second reflecting surface, as shown in Figure 7b. The same holds for thrice

Figure 7

redirected sources and so on, and the likelihood that virtual source volumes intersect becomes less likely each time, provided that the reflecting surfaces do not occupy a majority of the space.

To minimize the creation of useless virtual light sources, we check very carefully to confirm that the light in fact has some free path between the source and the reflecting surface before creating the virtual source. For example, we might have an intervening surface that prevents all rays from reaching a reflecting surface from a specific light source, such as the situation shown in Figure 7c. We can test for this condition by sending a number of presampling rays between the light source and the reflecting surface, assuming if none of the rays arrives that the reflecting path must be completely obstructed. Conversely, if none of the rays is obstructed, we can save time during shadow testing later by assuming that any ray arriving at the reflecting surface in fact has a free path to the source, and further ray intersection tests are unnecessary. We have found presampling to be very effective in avoiding wasteful testing of completely blocked or unhindered virtual light source paths.

Figure 8 shows a cross-section of an office space with a light shelf having a mirrored top surface. Exterior to this office is a building with a mirrored glass facade. Figure 9a shows the interior of the office with sunlight reflected by the shelf onto the ceiling. Light has also been reflected by the exterior, glazed building. Light shelf systems utilize daylight very effectively and are finding increasing popularity among designers.

To make our calculation more efficient overall, we have made additional use of "secondary" light sources, described in the next section.

Figure 8

3.5. User-directed Preprocessing of "Secondary" Sources

What happens when daylight enters a space through a skylight or window? If we do not treat such "secondary" emitters specially in our calculation, we will have to rely on the ability of the naive Monte Carlo sampling to find and properly integrate these contributions, which is slow. Especially when a window or skylight is partially obscured by venetian blinds or has a geometrically complex configuration, computing its contribution requires significant effort. Since we know a priori that such openings have an important influence on indoor illumination, we can greatly improve the efficiency of our simulation by removing them from the indirect calculation and treating them instead as part of the direct (i.e. source) component.

Radiance provides a practical means for the user to move such secondary sources into the direct calculation. For example, the user may specify that a certain window is to be treated as a light source, and a separate calculation will collect samples of the transmitted radiation over all points on the window over all directions, a 4-dimensional function. This distribution is then automatically applied to the window, which is treated as a secondary light source in the final calculation. This method was used in Figure 9a not only for the windows, but also for light reflected by the ceiling. Bright solar patches on interior surfaces can make important contributions to interior illumination. Since this was the desired result of our mirrored light shelf design, we knew in advance that treating the ceiling as a secondary light source might improve the efficiency of our calculation. Using secondary light sources in this scene reduced simulation time to approximately one fifth of what it would have been to reach the same accuracy using the default sampling techniques.

Figure 9b shows a Monte Carlo path tracing calculation of the same scene as 9a, and took roughly the same amount of time to compute. The usual optimizations of sending rays to light sources (the sun in this case) and in specular directions were used. Nevertheless, the image is very noisy due to the difficulty of computing interreflection independently at each pixel. Also, locating the sun reflected in the mirrored light shelf is hopeless with naive sampling; thus the ceiling is extremely noisy and the room is not as well lit as it should be.

An important aspect of secondary light sources in *Radiance* is that they have a dual nature. When treated in the direct component calculation, they are merely surfaces with precalculated output distributions. Thus, they can be treated efficiently as light sources and the actual variation that may take place over their extent (e.g. the bright and dark slats of venetian blinds) will

not translate into excessive variance in the calculated illumination. However, when viewed directly, they revert to their original form, showing all the appropriate detail. In our office scene example, we can still see through the window despite its treatment as a secondary light source. This is because we treat a ray coming from the eye differently, allowing it to interact with the actual window rather than seeing only a surface with a smoothed output distribution. In fact, only shadow rays see the simplified representation. Specular rays and other sampling will be carried out as if the window was not a light source at all. As is true with the computation of indirect irradiance described in section 3.2, extreme care must be exercised to avoid double-counting of light sources and other inconsistencies in the calculation.

3.6. Hierarchical Octrees for Spatial Subdivision

One of the goals of our simulation is to model very complicated geometries. Ray-tracing is well-suited to calculations in complicated environments, since spatial subdivision structures reduce the number of ray-surface intersection tests to a tiny fraction of the entire scene. In *Radiance*, we use an octree spatial subdivision scheme similar to that proposed by Glassner [7]. Our octree starts with a cube encompassing the entire scene, and recursively subdivides the cube into eight equal subcubes until each voxel (leaf node) intersects or contains less than a certain number of surfaces, or is smaller than a certain size.

Figure 10. Plot showing sublinear relationship of intersection time to number of surfaces in a scene. The best fit for γ in this test was 0.245, meaning the ray intersection time grew more slowly than the fourth root of N . The spheres were kept small enough so that a random ray sent from the field's interior had about a 50% chance of hitting something. (I.e. the sphere radii were proportional to $N^{1/3}$.) This guarantees that we are really seeing the cost of complicated geometry, since each ray goes by many surfaces.

Although it is difficult to prove in general, our empirical tests show that the average cost of ray intersection using this technique grows as a fractional power of the total number of surfaces, i.e. $O(N^\gamma)$ where $\gamma < 1/2$. The time to create the octree grows linearly with the number of surfaces, but it is usually only a tiny fraction of the time spent rendering. Figure 10 shows the relationship between ray intersection time and number of surfaces for a uniformly distributed random field of spheres.

The basic surface primitives supported in *Radiance* are polygons, spheres and cones. Generator programs provide conversion from arbitrary shape definitions (e.g. surfaces of

revolution, prisms, height fields, parametric patches) to these basic types. Additional scene complexity is modeled using hierarchical *instancing*, similar to the method proposed by Snyder [21]. In our application of instancing, objects to be instanced are stored in a separate octree, then this octree is instanced along with other surfaces to create a second, enclosing octree. This process is repeated as many times and in as many layers as desired to produce the combined scene. It is possible to model scenes with a virtually unlimited number of surfaces using this method.

Figure 11 shows a cabin in a forest. We began with a simple spray of 150 needles, which were put into an octree and instanced many times and combined with twigs to form a branch, which was in turn instanced and combined with larger branches and a trunk to form a pine tree. This pine tree was then put in another octree and instanced in different sizes and orientations to make a small stand of trees, which was combined with a terrain and cabin model to make this scene. Thus, four hierarchical octrees were used together to create this scene, which contains over a million surfaces in all. Despite its complexity, the scene still renders in a couple of hours, and the total data structure takes less than 10 Mbytes of RAM.

3.7. Patterns and Textures

Another practical way to add detail to a scene is through the appropriate infusion of surface detail. In *Radiance*, we call a variation in surface color and/or brightness a *pattern*, and a perturbation of the surface normal a *texture*. This is more in keeping with the English definitions of these words, but sometimes at odds with the computer graphics community, which seems to prefer the term "texture" for a color variation and "bump-map" for a perturbation of the surface normal. In any case, we have extended the notion somewhat by allowing patterns and textures to be functions not only of surface position but also of surface normal and ray direction so that a pattern, for example, may also be used to represent a light source output distribution.

Our treatment of patterns and textures was inspired by Perlin's flexible shading language [17], to which we have added the mapping of lookup functions for multi-dimensional data. Using this technique, it is possible to interpret tabulated or image data in any manner desired through the same functional language used for procedural patterns and textures.

Figure 12 shows a scene with many patterns and textures. The textures on the vases and oranges and lemons are procedural, as is the pattern on the bowl. The pattern on the table is scanned, and the picture on the wall is obviously an earlier rendering. Other patterns which are less obvious in this scene are the ones applied to the three light sources, which define their output distributions. The geometry was created with the generator programs included with *Radiance*, which take functional specifications in the same language as the procedural patterns and textures. The star patterns are generated using a *Radiance* filter option that uses the pixel magnitude in deciding how much to spread the image, showing one advantage of using a floating-point picture format [27]. (The main advantage of this format is the ability to adjust exposure after rendering, taking full advantage of tone mapping operators and display calibration [23,30].)

3.8. Parallel Processing

One of the most practical ways to reduce calculation time is with parallel processing. Ray-tracing is a natural for parallel processing, since the calculation of each pixel is relatively independent. However, the caching of indirect irradiance values in *Radiance* means that we benefit from sharing information between pixels that may or may not be neighbors in one or more images. Sharing this information is critical to the efficiency of a parallel computation, and we want to do this in a system-independent way.

We have implemented a coarse-grained, multiple instruction, shared data (MISD) algorithm for *Radiance* rendering[†]. This technique may be applied to a single image, where multiple processes on one or more machines work on small sections of the image simultaneously, or to a sequence of images, where each process works on a single frame in a long animation. In the latter case, we need only worry about the sharing of indirect irradiance values on multiple active invocations, since dividing the image is not an issue. The method we use is described below.

Indirect irradiance values are written to a shared file whose contents are checked by each process prior to update. If the file has grown, the new values (produced by other processes) are read in before the irradiances computed by this process are written out. File consistency is maintained through the NFS lock manager, thus values may be shared transparently across the network. Irradiance values are written out in blocks big enough to avoid contention problems, but not so big that there is a lot of unnecessary computation caused by insufficient value sharing. We found this method to be much simpler, and about as efficient, as a remote procedure call (RPC) approach.

Since much of the scene information is static throughout the rendering process, it is wasteful to have multiple copies on a multi-processing platform that is capable of sharing memory. As with value sharing, we wished to implement memory sharing in a system-independent fashion. We decided to use the memory sharing properties of the UNIX *fork(2)* call. All systems capable of sharing memory do so during fork on a copy-on-write basis. Thus, a child process need not be concerned that it is sharing its parent's memory, since it will automatically get its own memory the moment it stores something. We can use this feature to our advantage by reading in our entire scene and initializing all the associated data structures before forking a process to run in parallel. So long as we do not alter any of this information during rendering, we will share the associated memory. Duplicate memory may still be required for data that is generated during rendering, but in most cases this represents a minor fraction of our memory requirements.

3.9. Animation

Radiance is often used to create walk-through animations of static environments. Though this is not typically the domain of ray-tracing renderers, we employ some techniques to make the process more efficient. The most important technique is the use of recorded depth information at each pixel to interpolate fully ray-traced frames with a z-buffer algorithm. Our method is similar to the one explained by Chen et al [3], where pixel depths are used to recover an approximate 3-dimensional model of the visible portions of the scene, and a z-buffer is used to make visibility decisions for each intermediate view. This makes it possible to generate 30 very good-looking frames for each second of animation while only having to render about 5 of them. Another technique we use is unique to *Radiance*, which is the sharing of indirect irradiance values. Since these values are view-independent, there is no sense in recomputing them each time, and sharing them during the animation process distributes the cost over so many frames that the incremental cost of simulating diffuse interreflection is negligible.

Finally, it is possible to get interactive frame rates from advanced rendering hardware using illumination maps instead of ray-tracing the frames directly. (An illumination map is a 2-dimensional array of color values that defines the surface shading.) Such maps may be kept separate from the surfaces' own patterns and textures, then combined during rendering. Specular

[†]Data sharing is of course limited in the case of distributed processors, where each node must have its own local copy of scene data structures.

RADIANCE File Types			
Data Type	Format	Created by	Used for
Scene Description	ASCII text	text editor, CAD translator	geometry, materials, patterns, textures
Function File	ASCII text	text editor	surface tessellation, patterns, textures, scattering functions, coordinate mappings, data manipulation
Data File	ASCII integers and floats	luminaire data translator, text editor	N-dimensional patterns, textures, scattering functions
Polygonal Font	ASCII integers	Hershey set, font design system, font translator, text editor	text patterns, label generator
Octree	Binary	scene compiler (oconv)	fast ray intersection, incremental scene compilation, object instancing
Picture	run-length encoded 4-byte/pixel floating-point	renderer, filter, image translator	interactive display, hard copy, lighting analysis, material pattern, rendering recovery
Ambient File	Binary	renderer, point value program	sharing view-independent indirect irradiance values

Table 1. All binary types in *Radiance* are portable between systems, and have a standard information header specifying the format and the originating command(s).

surfaces will not appear correct since they depend on the viewer's perspective, but this may be a necessary sacrifice when user control of the walk-through is desired. Interactive rendering has long been touted as a principal advantage of radiosity, when in fact complete view-independence is primarily a side-effect of assuming diffuse reflection. *Radiance* calculates the same values using a ray-tracing technique, and storage and rendering may even be more efficient since large polygons need not be subdivided into hundreds of little ones -- an illumination map works just as well or better.

3.10. Implementation Issues

Radiance is a collection of C programs designed to work in concert, communicating via the standard data types listed in Table 1. The system may be compiled directly on most UNIX platforms, including SGI, Sun, HP, DEC, Apple (A/UX), and IBM (RS/6000). Portability is maintained over 60,000+ lines of code using the Kernighan and Ritchie standard [11] and conservative programming practices that do not rely on system-specific behaviors or libraries. (In addition to UNIX support, there is a fairly complete Amiga port by Per Bojsen, and a limited MS-DOS port by Karl Grau.)

A typical rendering session might begin with the user creating or modifying a geometric model of the space using a CAD program. (The user spends about 90% of the time on geometric modeling.) The CAD model is then translated into a *Radiance* scene description file, using either a stand-alone program or a function within the CAD system itself. The user might then create or modify the materials, patterns and textures associated with this model, and add some objects from a library of predefined light sources and furnishings. The completed model would then be compiled by **oconv** into an octree file, which would be passed to the interactive renderer, **rview**, to verify the desired view and calculation parameters. Finally, a batch rendering would be started with **rpict**, and after a few minutes or a few hours, the raw picture would be filtered (i.e. anti-aliased via image reduction) by **pfilt** using a suitable exposure level and target resolution. This finished picture may be displayed with **ximage**, translated to another format, printed, or further analyzed using one of the many *Radiance* image utilities. This illustrates the basic sequence of:

model → convert → render → filter → display

all of which may be put in a single pipelined command if desired.

As *Radiance* has evolved over the years, it has become increasingly sophisticated, with nearly 100 programs that do everything from CAD translation to surface tessellation to lighting calculations and rendering to image filtering, composition and conversion. With this sophistication comes great versatility, but learning the ins and outs of the programs, even the few needed for simple rendering, is impractical for most designers.

To overcome system complexity and improve the reliability of rendering results, we have written an executive control program, called **rad**. This program takes as its input a single file that identifies the material and scene description files needed as well as qualitative settings related to this environment and the simulation desired. The control program then calls the other programs with the proper parameters in the proper sequence.

The intricacies of the *Radiance* rendering pipeline are thus replaced by a few intuitive variable settings. For example, there is a variable called "DETAIL", which might be set to "low" for an empty room, "medium" for a room with a few pieces of furniture and "high" for a complicated room with many furnishings and textures. This variable will be used with a few others like it to determine how many rays to send out in the Monte Carlo sampling of indirect lighting, how closely to space these values, how densely to sample the image plane, and so on. One very important variable that affects nearly all rendering parameters is called "QUALITY". Low quality renderings come out quickly but may not look as good as medium quality renderings, and high quality renderings take a while but when they finish, the images can go straight into a magazine article.

This notion of replacing many algorithm-linked rendering parameters with a few qualitative variables has greatly improved the usability of *Radiance* and the reliability of its output. The control program also keeps track of octree creation, secondary source generation, aborted job recovery, image filtering and anti-aliasing, and running the interactive renderer. The encoding of expertise in this program has been so successful, in fact, that we rely on it ourselves almost 100% for setting parameters and controlling the rendering process.

Although the addition of a control program is a big improvement, there are still many aspects of *Radiance* that are not easily accessible to the average user. We have therefore added a number of utility scripts for performing specific tasks from the more general functions that constitute the system. One example of this is the **falsecolor** program, which calls other image filter programs and utilities to generate an image showing luminance contours or other data associated with a scene or

rendering. Figure 9c shows our previous rendering (Figure 9a) superimposed with illuminance contours. These contours tell the lighting designer if there is enough light in the right places or too much light in the wrong places -- information that is difficult to determine from a normal image†.

Even with a competent rendering control program and utility scripts for accomplishing specific tasks, there are still many designers who would not want to touch this system with an extended keyboard. Modern computer users expect a list of pull-down menus with point-and-click options that reduce the problem to a reasonably small and understandable set of alternatives. We are currently working on a graphical user interface (GUI) to the **rad** control program, which would at least put a friendlier face on the standard system. A more effective long-term solution is to customize the rendering interface for each problem domain, e.g. interior lighting design, daylighting, art, etc. Due to our limited resources and expertise, we have left this customization task to third parties who know more about specific applications, and who stand to benefit from putting their GUI on our simulation engine. So far, there are a half dozen or so developers working on interfaces to *Radiance*.

4. Applications and Results

The real proof of a physically-based rendering system is the problems it solves. Here we see how well we have met the challenges and goals we set out. *Radiance* has been used by hundreds of people to solve thousands of problems over the years. In the color pages we have included some of the more recent work of some of the more skilled users. The results have been grouped into two application areas, electric lighting problems and daylighting problems.

4.1. Electric Lighting

Electric lighting was the first domain of *Radiance*, and it continues to be a major strength. A model may contain any number of light sources of all shapes and sizes, and the output distributions may be entered as either near-field or far-field data. The dual nature of light sources (mentioned in section 3.5) also permits detailed modeling of fixture geometry, which is often important in making aesthetic decisions.

There are several application areas where electric lighting is emphasized. The most obvious application is lighting design. Figure 13 shows a comparative study between three possible lighting alternatives in a hotel lobby space. Several other designs were examined in this exploration of design visualization. With such a presentation, the final decision could be safely left to the client.

One design application that requires very careful analysis is indirect lighting. Figure 14 shows a simulation of a new control center for the London Underground. The unusual arrangement of upwardly directed linear fluorescents was designed to provide general lighting without affecting the visibility of the central display panel (image left).

Stage lighting is another good application of physically-based rendering. The designs tend to be complex and changing, and the results must be evaluated aesthetically (i.e. visually). Figure 15 shows a simulation of a scene from the play *Julius Caesar*. Note the complex shadows cast by the many struts in the stage set. Computing these shadows with a radiosity algorithm would be extremely difficult.

†Actually, *Radiance* pictures do contain physical values through a combination of the 4-byte floating-point pixel representation and careful tracking of exposure changes [27], but the fidelity of any physical image presentation is limited by display technology and viewing conditions. We therefore provide the convenience of extracting numerical values with our interactive display program.

4.2. Daylighting

Daylight poses a serious challenge to physically-based rendering. It is brilliant, ever-changing and ever-present. At first, the daylight simulation capabilities in *Radiance* were modest, limited mostly to exteriors and interiors with clear windows or openings. Designers, especially architects, wanted more. They wanted to be able to simulate light through venetian blinds, intricate building facades and skylights. In 1991, the author was hired on sabbatical by EPFL to improve the daylight simulation capabilities of *Radiance*, and developed some of the techniques described earlier in this paper. In particular, the large source adaptive subdivision, virtual source and secondary source calculations proved very important for daylighting problems.

The simplest application of daylight is exterior modeling. Many CAD systems have built-in renderers that will compute the solar position from time of day, year, and location, and generate accurate shadows. In addition to this functionality, we wanted *Radiance* to show the contributions of diffuse skylight and interreflection. Figure 16 shows the exterior of the Mellen-camp Pavillion, an Indiana University project that recently acquired funding (along with its name).

A more difficult daylighting problem is atrium design*. Designing an atrium requires thorough understanding of the daylight availability in a particular region to succeed. Figure 17 shows an atrium space modeled entirely within *Radiance*, without the aid of a CAD program [13]. The hierarchical construction of *Radiance* scene files and the many programmable object generators makes text-editor modeling possible, but most users prefer a "mousier" approach.

Daylighted interiors pose one of the nastiest challenges in rendering. Because sunlight is so intense, it is usually diffused or baffled by louvers or other redirecting systems. Some of these systems can be quite elaborate, emphasizing the need for simulation in their design. Figure 18 shows the interior of the pavillion from Figure 16. Figure 19 shows a library room illuminated by a central skylight. Figure 20a shows a simulation of a daylighted museum interior. Daylight is often preferred in museums as it provides the most natural color balance for viewing paintings, but control is also very important. Figure 20b shows a false color image of the illuminance values on room surfaces; it is critical to keep these values below a certain threshold to minimize damage to the artwork.

5. Conclusion

We have presented a physically-based rendering system that is accurate enough, general enough, and practical enough for the vast majority of lighting design and architectural applications. The simulation uses a light-backwards ray-tracing method with extensions to handle specular, diffuse and directional-diffuse reflection and transmission in any combination to any level in any environment. Not currently included in the calculation are participating media, diffraction and interference, phosphorescence, and polarization effects. There is nothing fundamental preventing us from modeling these processes, but so far there has been little demand for them from our users.

The principle users of *Radiance* are researchers and educators in public and private institutions, and lighting specialists at large architectural, engineering and manufacturing firms. There are between 100 and 200 active users in the U.S. and Canada, and about half as many overseas. This community is continually growing, and as the *Radiance* interface and documentation improves, the growth rate is likely to increase.

*An atrium is an enclosed courtyard with a glazed roof structure for maximizing daylight while controlling the indoor climate.

For the graphics research community, we hope that *Radiance* will provide a basis for evaluating new physically-based rendering techniques. To this end, we provide both the software source code and a set of precomputed test cases on our ftp server. The test suite includes diffuse and specular surfaces configured in a simple rectangular space with and without obstructions. More complicated models are also provided in object libraries and complete scene descriptions.

6. Acknowledgements

Most of the color figures in this paper represent the independent work of *Radiance* users, and are reprinted with permission. Figure 5 was created by Charles Ehrlich from a design by Mark Mack Architects of San Francisco. Figure 11 was created by Cindy Larson. Figures 12 and 13 were created by Martin Moeck of Siemens Lighting, Germany. Figure 14 was created by Steve Walker of Ove Arup and Partners, London. Figure 15 was created by Robert Shakespeare of the Theatre Computer Visualization Center at Indiana University. Figures 16, 18 and 19 were created by Scott Routen and Reuben McFarland of the University Architect's Office at Indiana University. Figures 17 and 20 were created by John Mardaljevic of the ECADAP Group at De Montfort University, Leicester.

The individuals who have contributed to *Radiance* through their support, suggestions, testing and enhancements are too numerous to mention. Nevertheless, I must offer my special thanks to: Peter Apian-Bennowitz, Paul Bourke, Raphael Compagnon, Simon Crone, Charles Ehrlich, Jon Hand, Paul Heckbert, Cindy Larson, Daniel Lucias, John Mardaljevic, Kevin Matthews, Don McLean, Georg Mischler, Holly Rushmeier, Jean-Louis Scartezzini, Jennifer Schuman, Veronika Sumneraur, Philip Thompson, Ken Turkowski, and Florian Wenz.

Work on *Radiance* was sponsored by the Assistant Secretary for Conservation and Renewable Energy, Office of Building Technologies, Buildings Equipment Division of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. Additional funding was provided by the Swiss federal government as part of the LUMEN Project.

7. Software Availability

Radiance is available by anonymous ftp from two official sites:

hobbes.lbl.gov	128.3.12.38	Berkeley, California
nestor.epfl.ch	128.178.139.3	Lausanne, Switzerland

For convenience, *Radiance 2.4* has been included on the CD-ROM version of these proceedings.

From **Mosaic**, try the following URL:

file://hobbes.lbl.gov/www/radiance/radiance.html

8. References

- Baum, Daniel, Holly Rushmeier, James Winget, "Improving Radiosity Solutions Through the Use of Analytically Determined Form-Factors," *Computer Graphics*, Vol. 23, No. 3, July 1989, pp. 325-334.
- Baum, Daniel, Stephen Mann, Kevin Smith, James Winget, "Making Radiosity Usable: Automatic Preprocessing and Meshing Techniques for the Generation of Accurate Radiosity Solutions," *Computer Graphics*, Vol. 25, No. 4, July 1991.
- Chen, Shenchang Eric, Lance Williams, "View Interpolation for Image Synthesis," *Computer Graphics*, August 1993, pp. 279-288.
- Compagnon, Raphael, B. Paule, J.-L. Scartezzini, "Design of New Daylighting Systems Using ADELIN Software," *Solar Energy in Architecture and Urban Planning*, proceedings of the 3rd European Conference on Architecture, Florence, Italy, May 1993.
- Cook, Robert, Thomas Porter, Loren Carpenter, "Distributed Ray Tracing," *Computer Graphics*, Vol. 18, No. 3, July 1984, pp. 137-147.
- Dorsey, Julie O'B., Francois Sillion, Donald Greenberg, "Design and Simulation of Opera Lighting and Projection Effects," *Computer Graphics*, Vol. 25, No. 4, July 1991, pp. 41-50.
- Glassner, Andrew S., "Space subdivision for fast ray tracing" *IEEE Computer Graphics and Applications* Vol. 4, No. 10, October 1984, pp. 15-22.
- Goral, Cindy, Kenneth Torrance, Donald Greenberg, Bennet Bataille, "Modeling the Interaction of Light Between Diffuse Surfaces," *Computer Graphics*, Vol. 18, No. 3, July 1984, pp. 213-222.
- Grynberg, Anat, *Validation of Radiance*, LBID 1575, LBL Technical Information Department, Lawrence Berkeley Laboratory, Berkeley, California, July 1989.
- Kajiya, James T., "The Rendering Equation," *Computer Graphics*, Vol. 20, No. 4, August 1986.
- Kernighan, Brian, Dennis Ritchie, *The C Programming Language*, Prentice-Hall, 1978.
- Kirk, David, James Arvo, "Unbiased Sampling Techniques for Image Synthesis," *Computer Graphics*, Vol 25, No. 4, July 1991, pp. 153-156.
- Mardaljevic, John and Kevin Lomas, "Creating the Right Image," *Building Services / The CIBSE Journal*, Vol 15, No. 7, July 1993, pp. 28-30.
- Mardaljevic, John, K.J. Lomas, D.G. Henderson, "Advanced Daylighting Design for Complex Spaces" *Proceedings of CLIMA 2000*, 1-3 November 1993, London UK.
- Meyer, Gary, Holly Rushmeier, Michael Cohen, Donald Greenberg, Kenneth Torrance, "An Experimental Evaluation of Computer Graphics Imagery," *ACM Transactions on Graphics*, Vol. 5, No. 1, pp. 30-50.
- Nicodemus, F.E., J.C. Richmond, J.J. Hsia, *Geometrical Considerations and Nomenclature for Reflectance*, U.S. Department of Commerce, National Bureau of Standards, October 1977.
- Perlin, Ken, "An Image Synthesizer", *Computer Graphics*, Vol. 19, No. 3, July 1985, pp. 287-296.
- Rushmeier, Holly, *Extending the Radiosity Method to Transmitting and Specularly Reflecting Surfaces*, Master's Thesis, Cornell Univ., Ithaca, NY, 1986.
- Rushmeier, Holly, Charles Patterson, Aravindan Veerasamy, "Geometric Simplification for Indirect Illumination Calculations," *Proceedings of Graphics Interface '93*, May 1993, pp. 227-236.
- Smits, Brian, James Arvo, David Salesin, "An Importance-Driven Radiosity Algorithm," *Computer Graphics*, Vol 26, No. 2, July 1992, pp. 273-282.
- Snyder, John M., Alan H. Barr, "Ray Tracing Complex Models Containing Surface Tessellations," *Computer Graphics* Vol. 21, No. 4, pp. 119-128, July 1987.
- Teller, Seth and Pat Hanrahan, "Global Visibility Algorithms for Illumination Computations," *Computer Graphics*, pp. 239-246, August 1993.
- Tumblin, Jack, Holly Rushmeier, "Tone Reproduction for Realistic Images," *IEEE Computer Graphics and Applications*, Vol. 13, No. 6, November 1993, pp. 42-48.
- Wallace, John, Michael Cohen, Donald Greenberg, "A Two-Pass Solution to the Rendering Equation: A Synthesis of Ray Tracing and Radiosity Methods," *Computer Graphics*, Vol. 21, No. 4, July 1987.
- Ward, Gregory, Francis Rubinstein, Robert Clear, "A Ray Tracing Solution for Diffuse Interreflection," *Computer Graphics*, Vol. 22, No. 4, August 1988.
- Ward, Gregory, "Adaptive Shadow Testing for Ray Tracing," *Second EUROGRAPHICS Workshop on Rendering*, Barcelona, Spain, April 1991.
- Ward, Gregory, "Real Pixels," *Graphics Gems II*, Edited by James Arvo, Academic Press 1991, pp. 80-83.
- Ward, Gregory, Paul Heckbert, "Irradiance Gradients," *Third EUROGRAPHICS Workshop on Rendering*, Bristol, United Kingdom, May 1992.
- Ward, Gregory, "Measuring and Modeling Anisotropic Reflection," *Computer Graphics*, Vol. 26, No. 2, July 1992, pp. 265-272.
- Ward, Gregory, "A Contrast-Based Scalefactor for Luminance Display," *Graphics Gems IV*, Edited by Paul Heckbert, Academic Press, 1994.
- Whitted, Turner, "An Improved Illumination Model for Shaded Display," *Communications of the ACM*, Vol. 23, No. 6, June 1980, pp. 343-349.

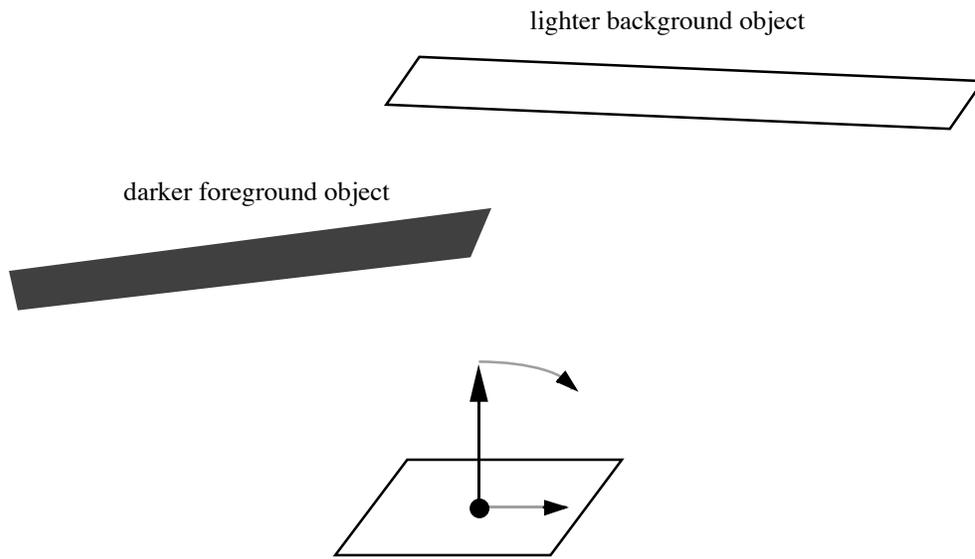


Figure 2: Irradiance gradients due to bright and dark objects in the environment.

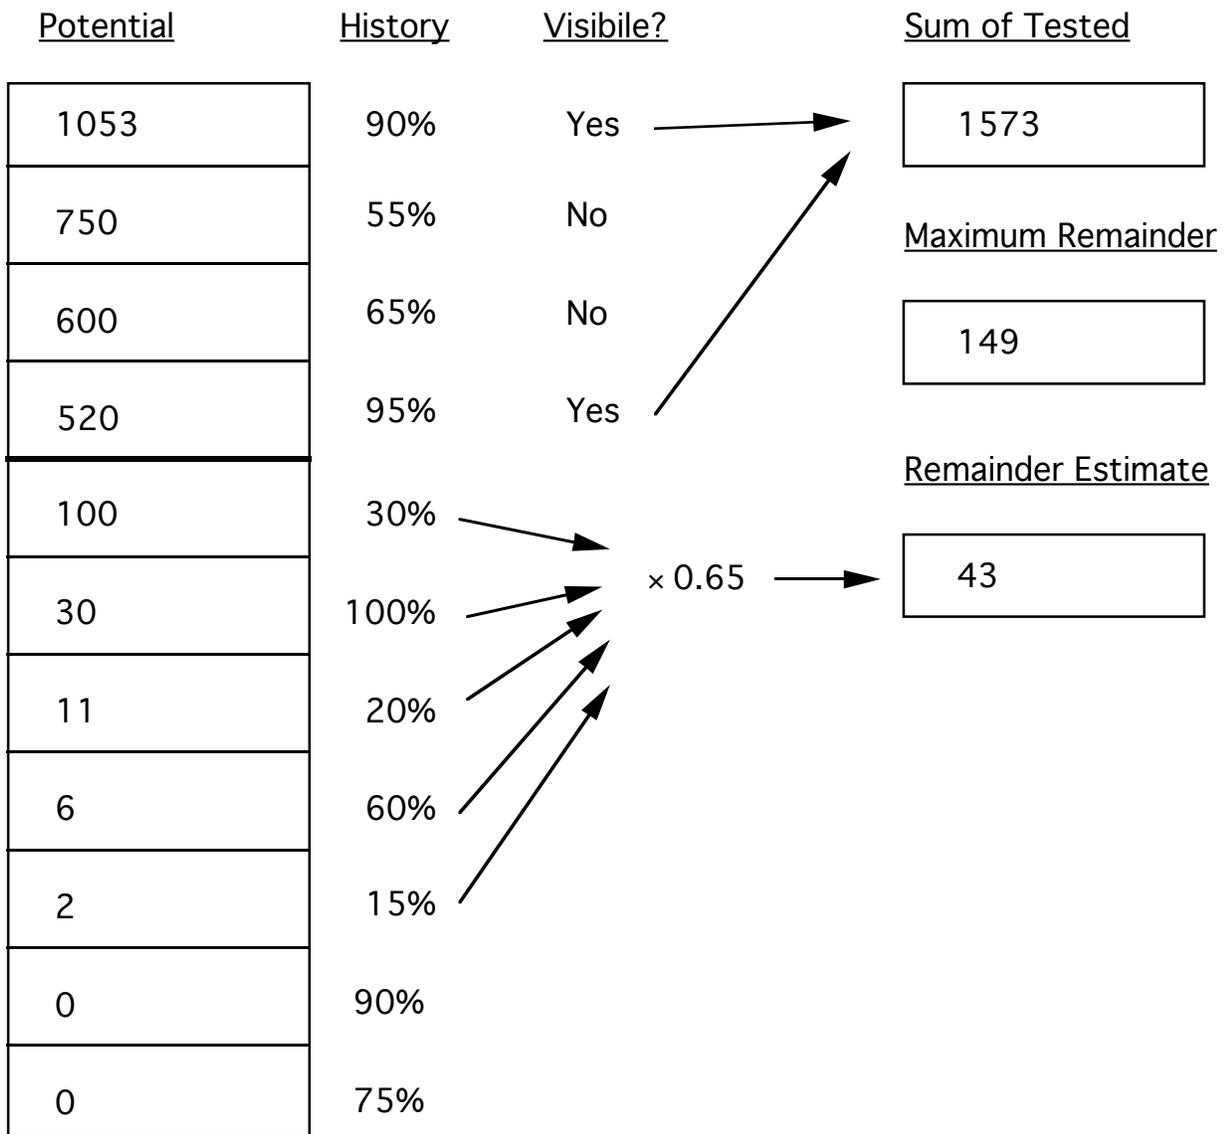


Figure 4: Adaptive shadow testing algorithm, explained in Section 3.4.

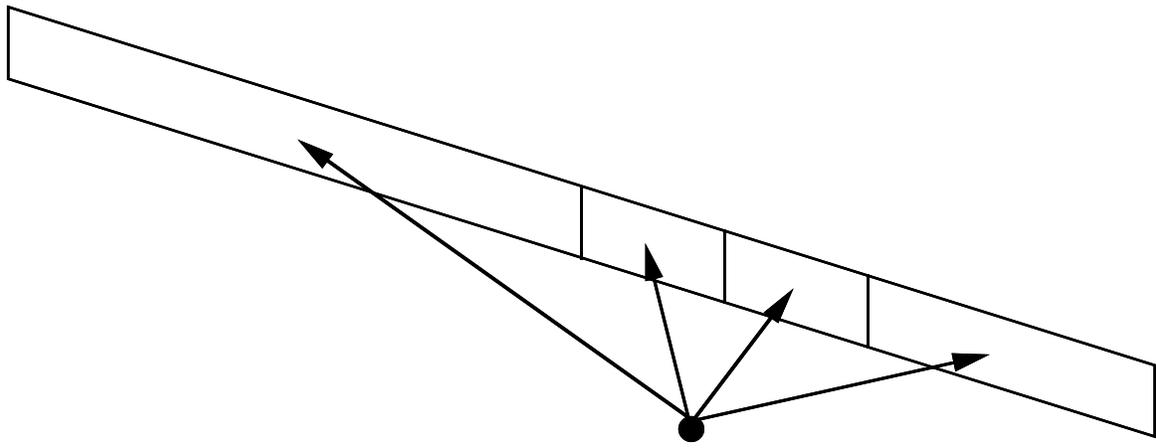


Figure 6a: Linear light source is adaptively split to minimize falloff and visibility errors.

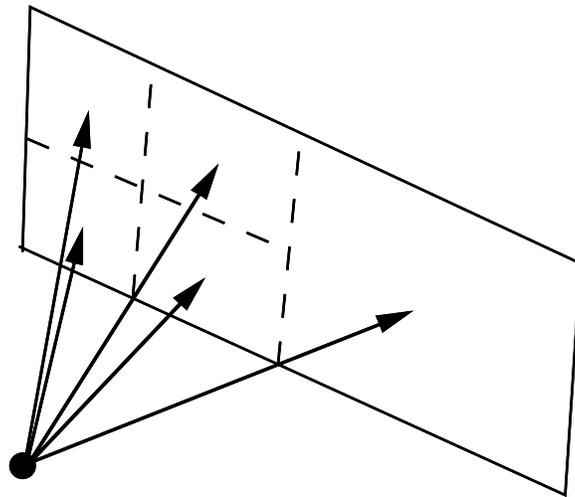


Figure 6b: Area light source is subdivided in two dimensions rather than one.

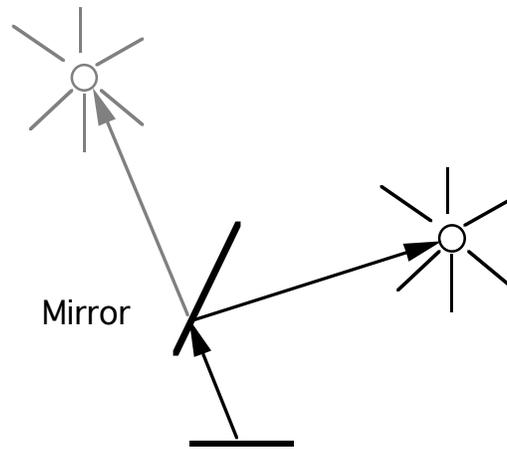


Figure 7a: Virtual source caused by mirror reflection.

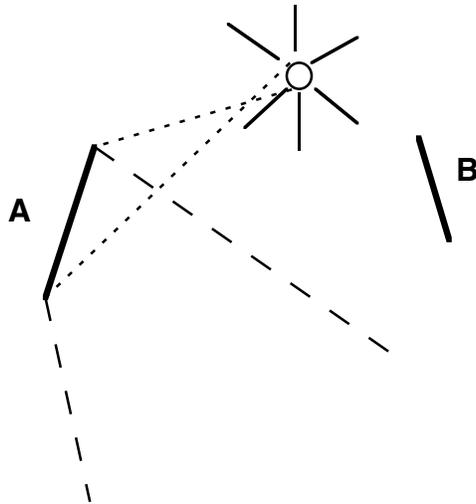


Figure 7b: Source reflection in mirror **A** cannot intersect mirror **B**, so no virtual-virtual source is created.

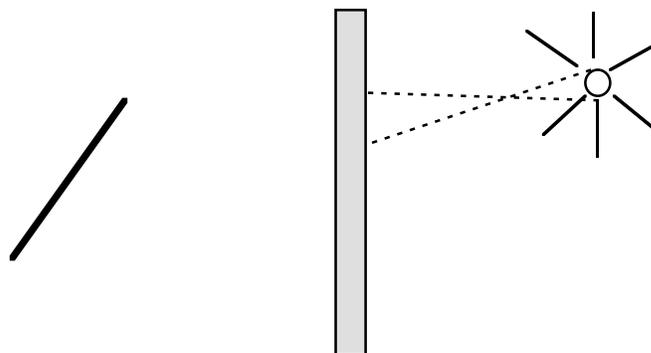


Figure 7c: Source rays cannot reach mirror surface, so no virtual source is created.

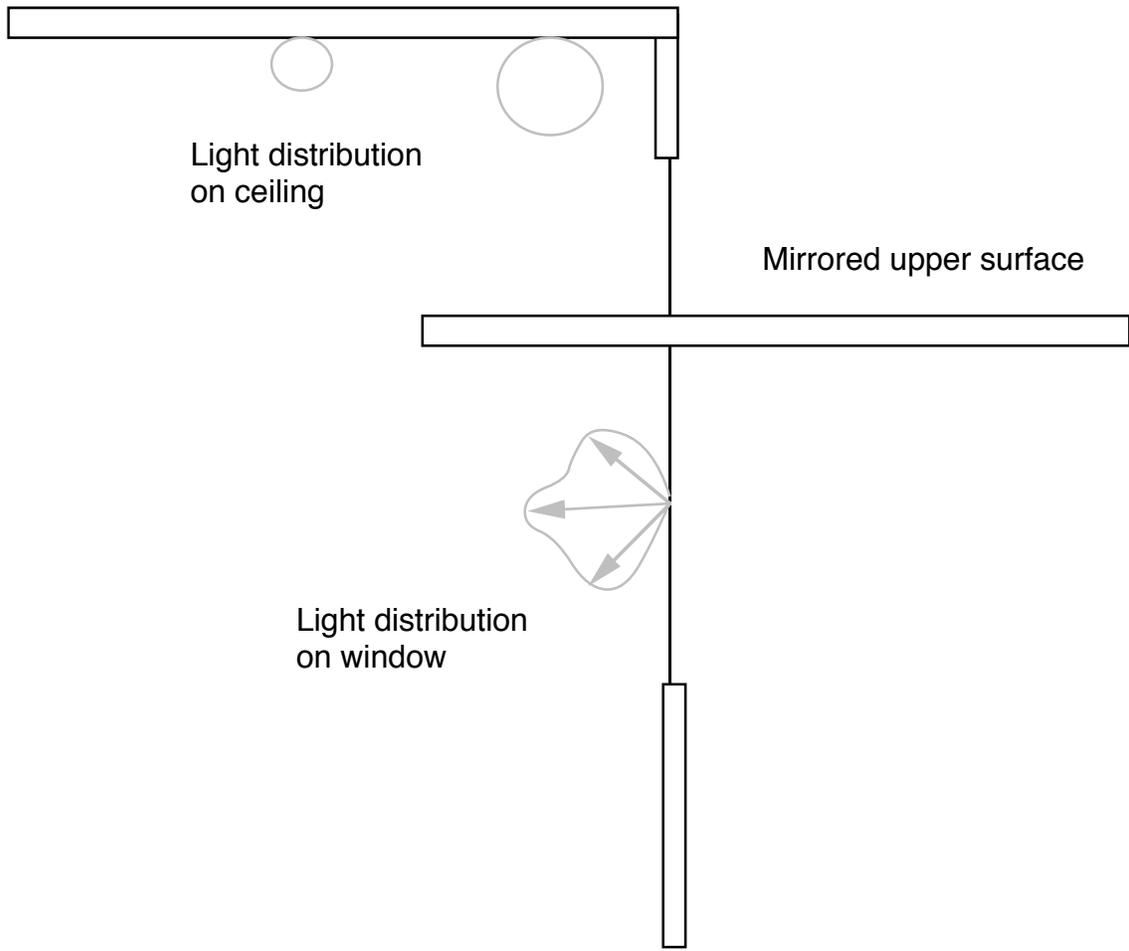
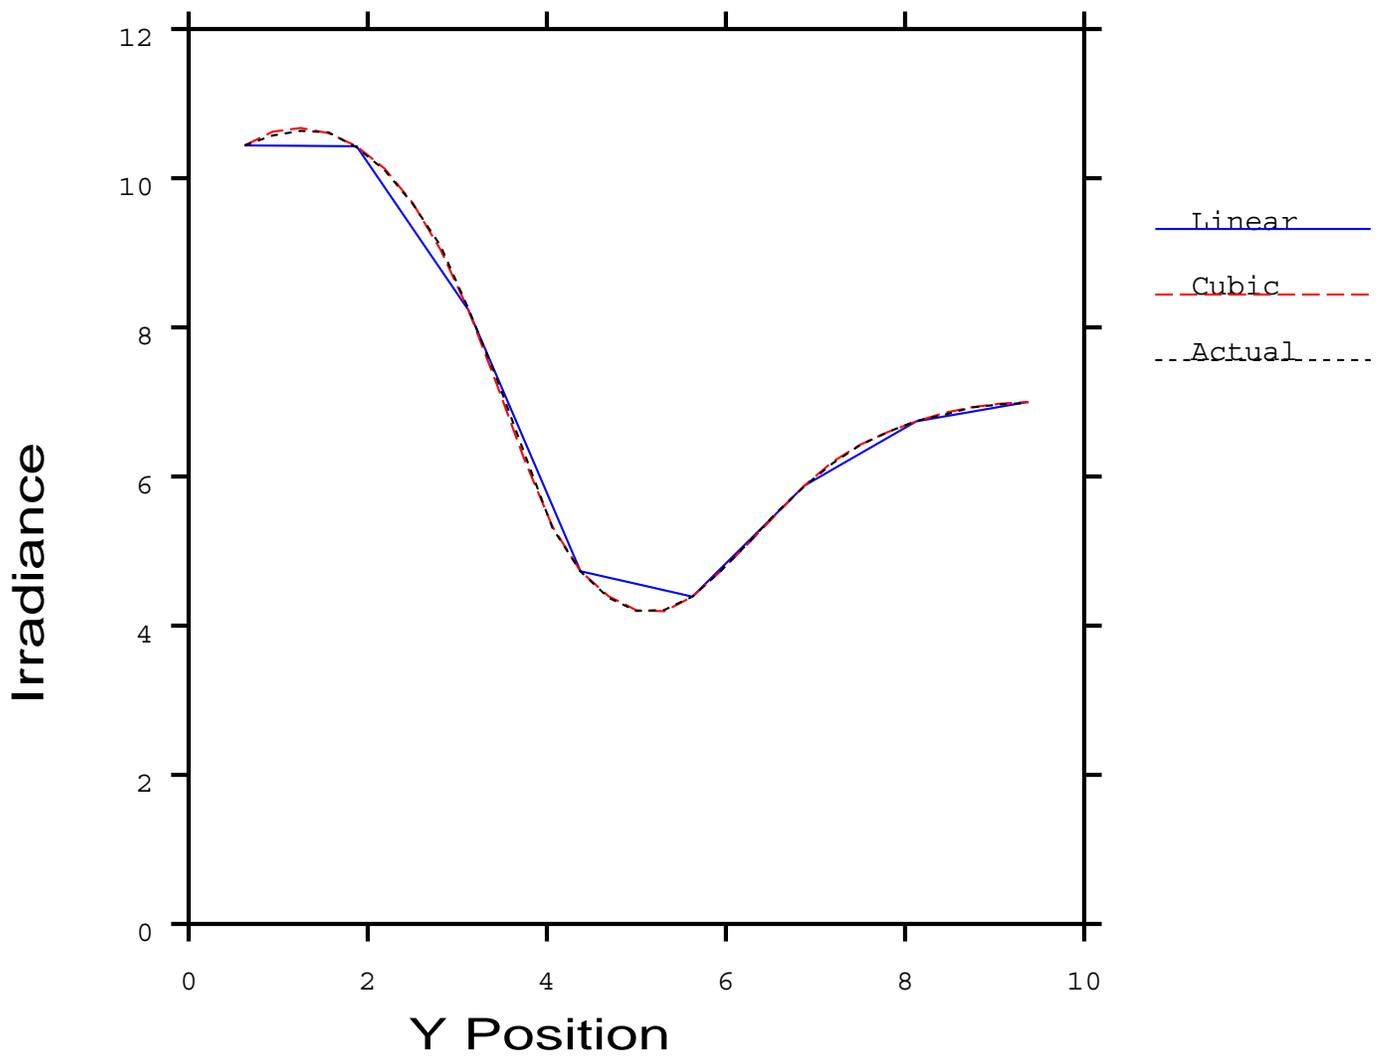


Figure 8: Crosssection of office space with mirrored light shelf.

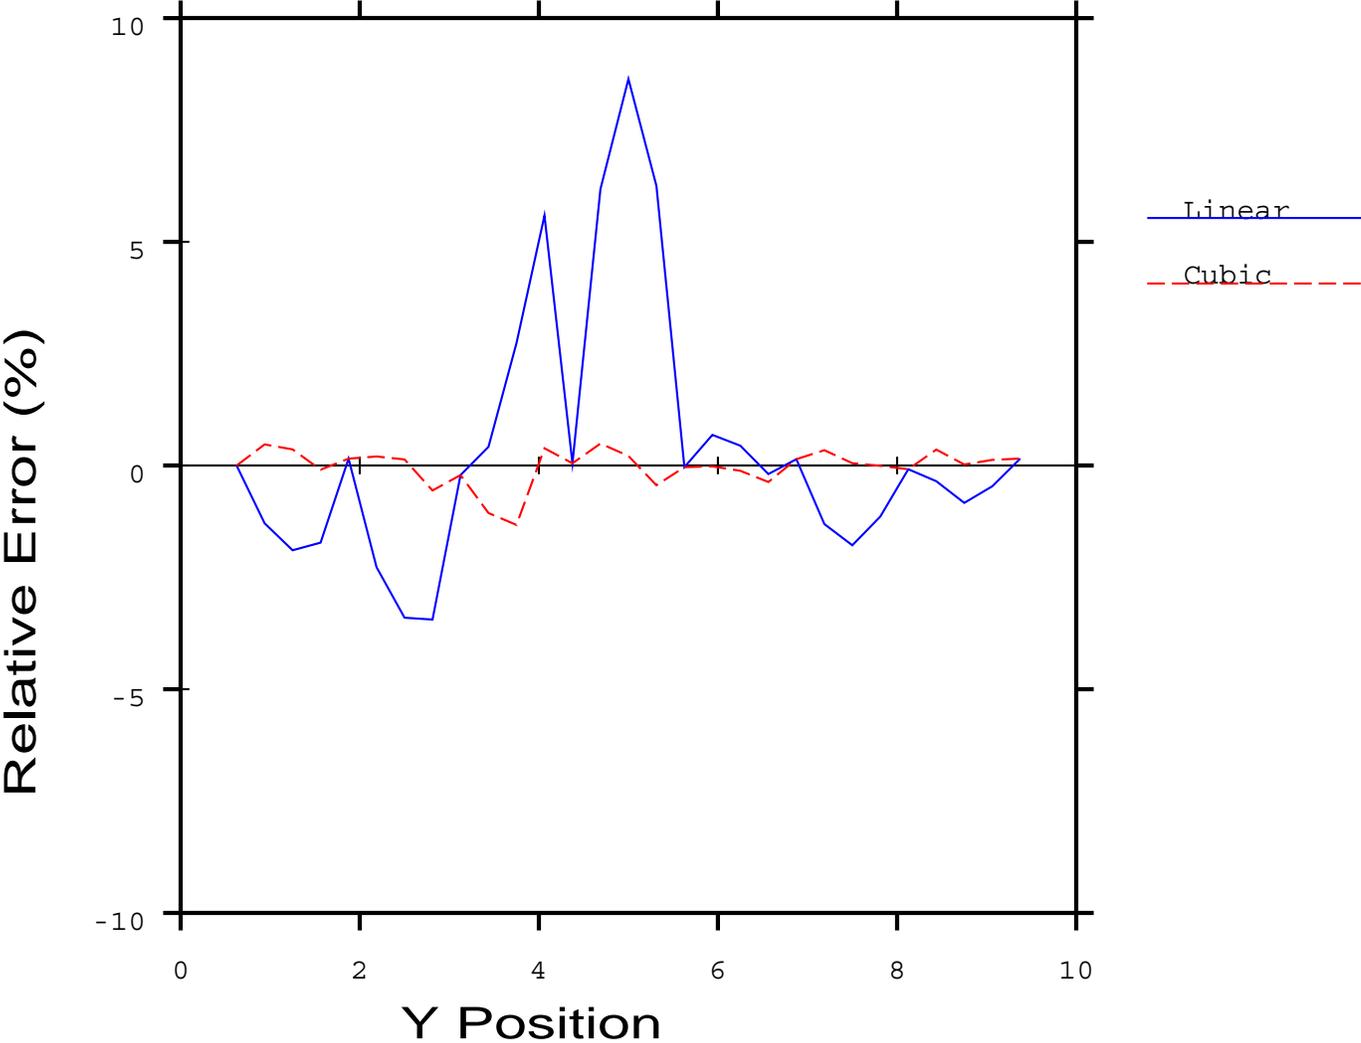
Irradiance Interpolation

$x=6.875$

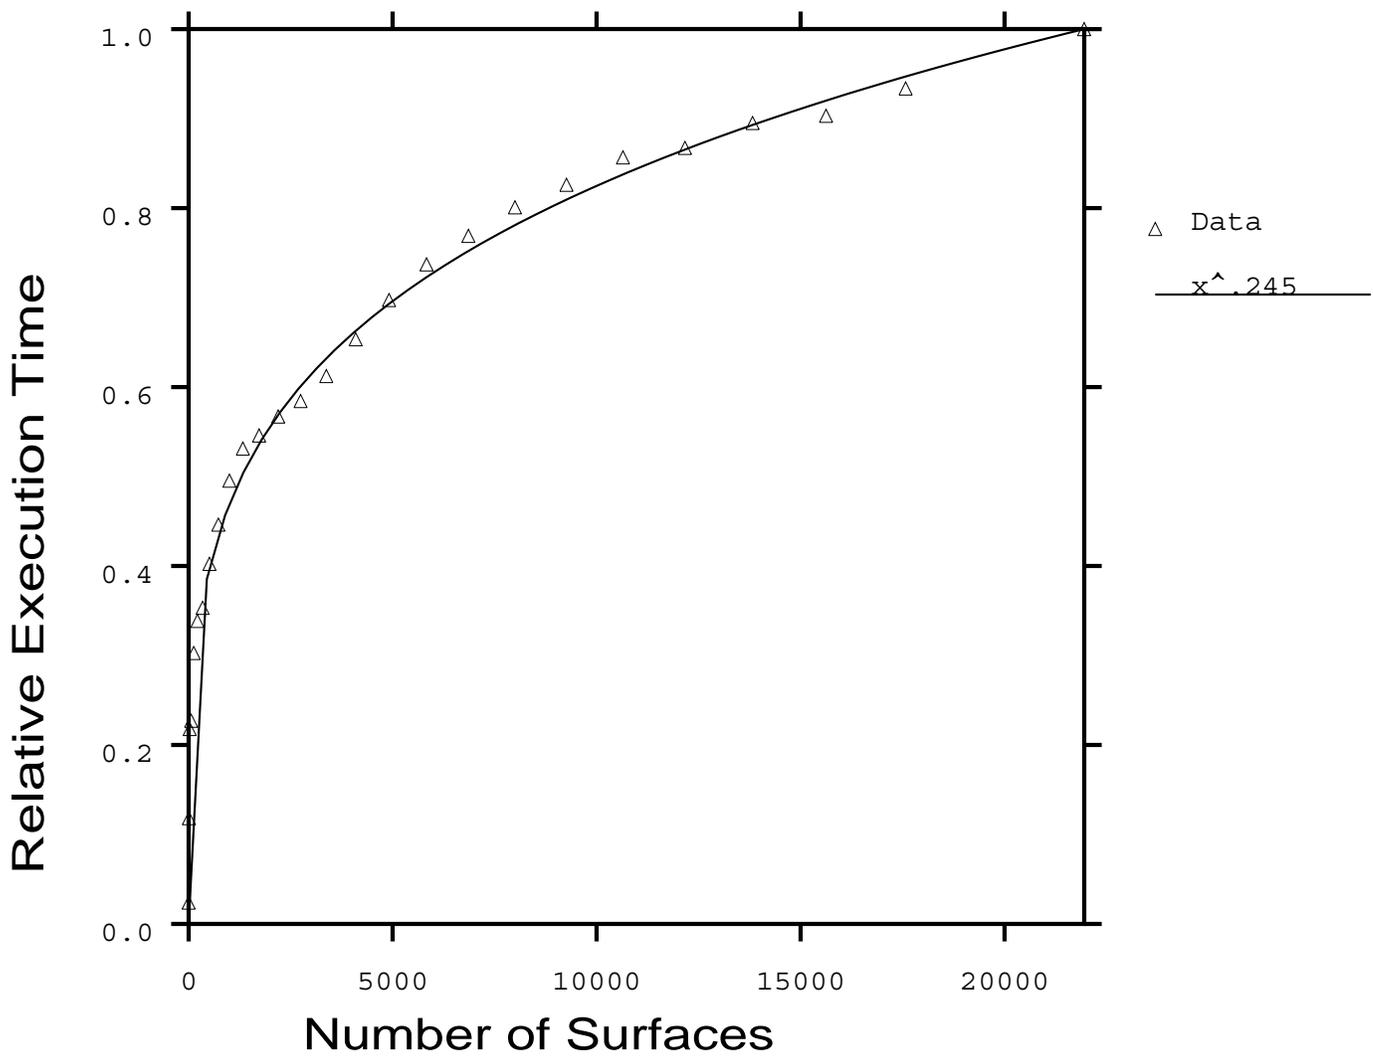


Irradiance Interpolation Error

$x=6.875$



Ray Intersection Time vs. Number of Surfaces



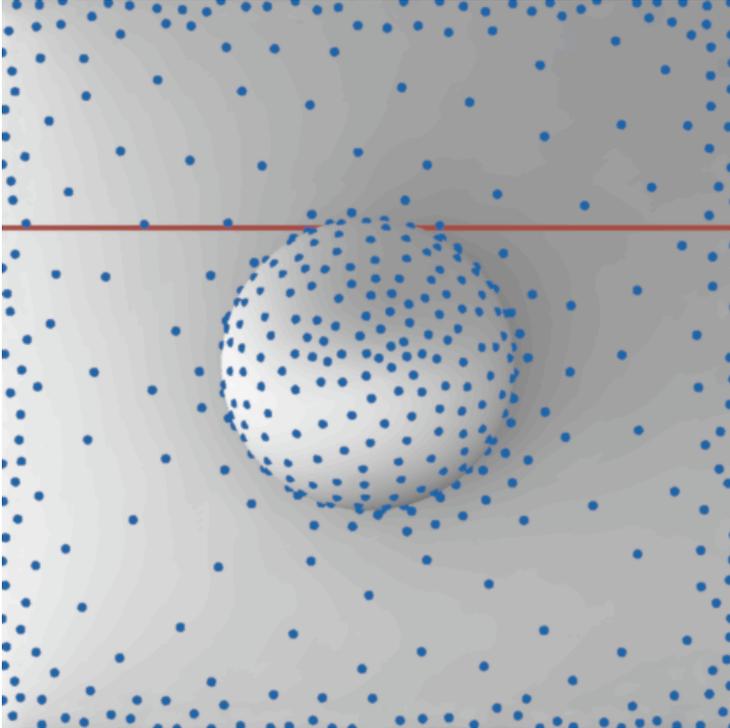


Figure 1. Blue dots show calculation points for diffuse interreflection. (See Figure 3 regarding red line.)

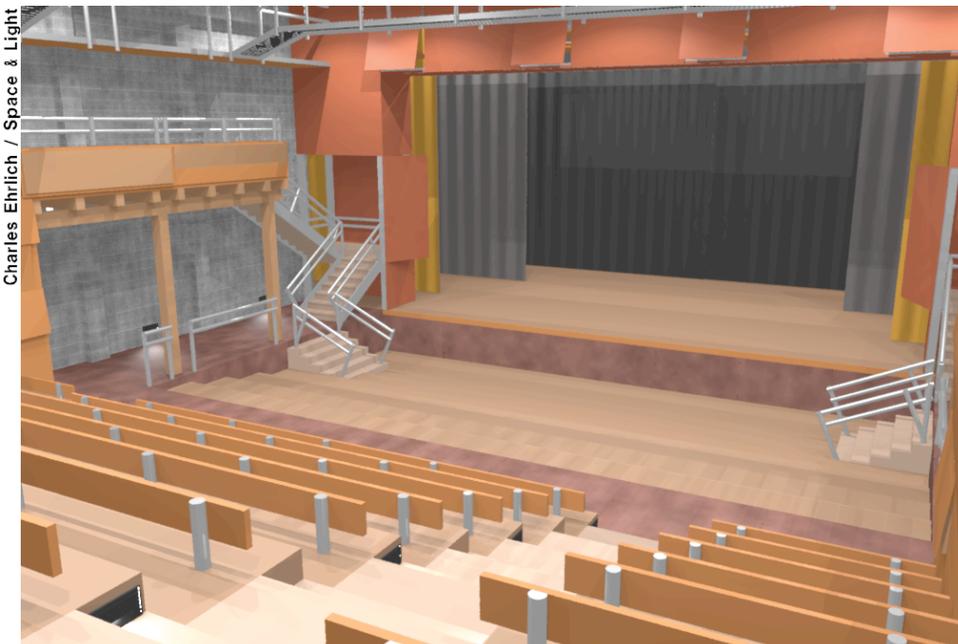


Figure 5. A theater house lighting simulation. The striated shadows on the floor are cast by the catwalks above.

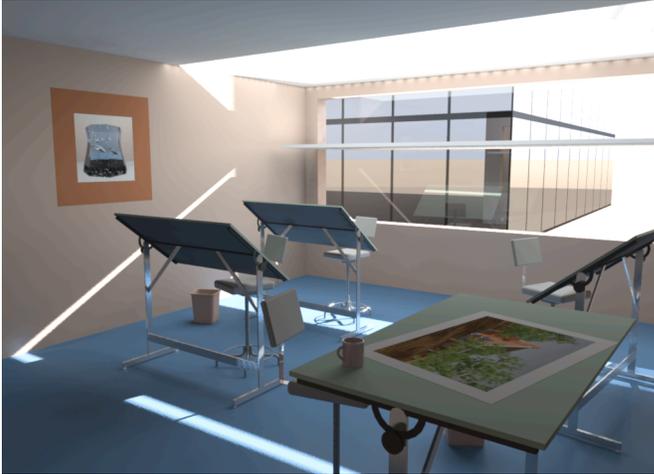


Figure 9a. A drafting office with a mirror light shelf slicing the window and redirecting light upwards.

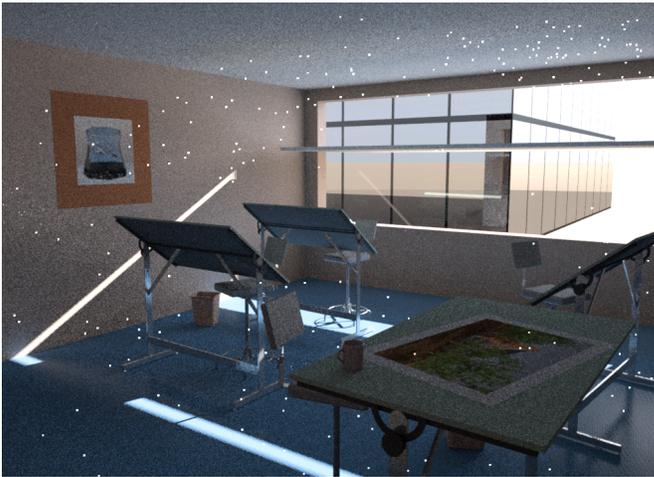


Figure 9b. A Monte Carlo simulation rendered in the same time as 9a, showing the resulting noise.

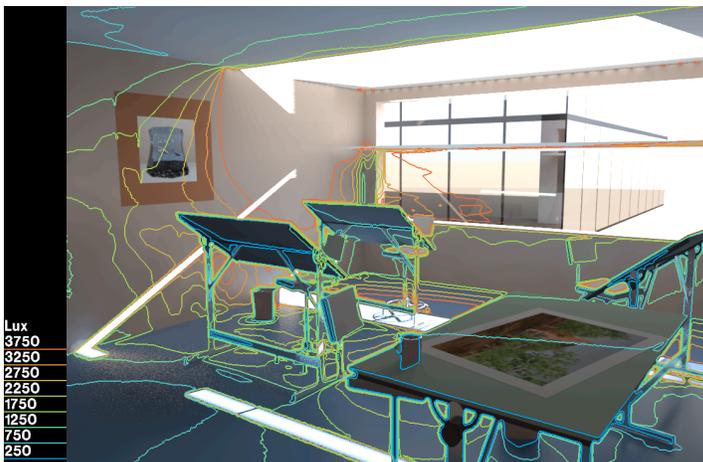


Figure 9c. A visualization of the illuminance levels on the room surfaces in the drafting office.

Cindy Larson



Figure 11. A cabin in a forest using hierarchical octrees to model over a million pine needles.

Martin Moeck



Figure 12. A still life image showing examples of procedural and scanned textures and patterns.

Martin Moeck / Siemens Lighting

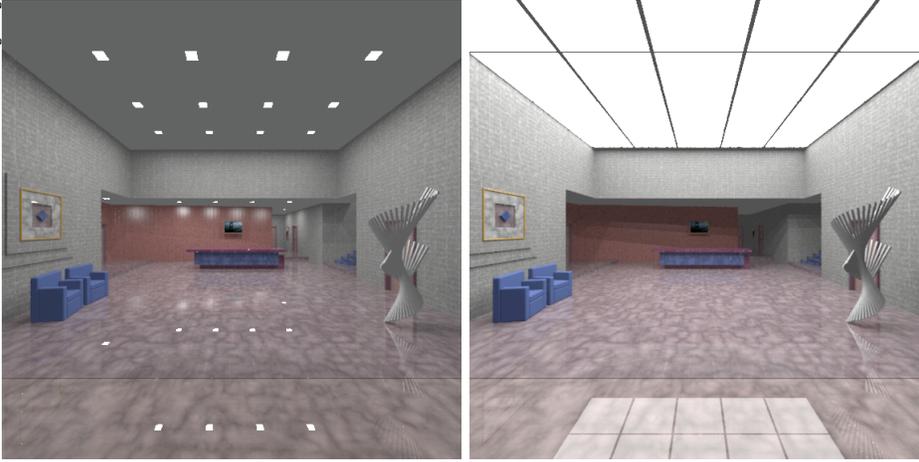
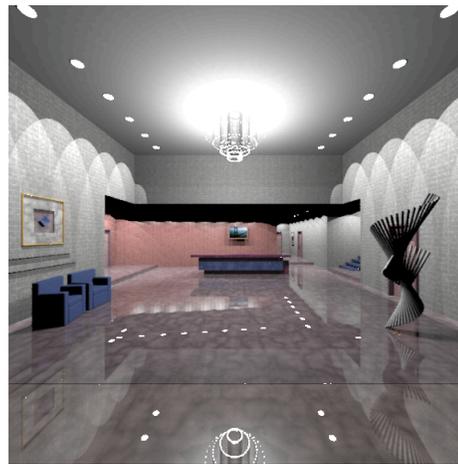


Figure 13. A comparison of three lighting schemes for a hotel lobby space.



Steve Walker / Ove Arup & Partners



Figure 14. An indirect lighting system was designed to reduce glare on monitors in the London Underground control center.

Rob Shakespeare / Indiana University

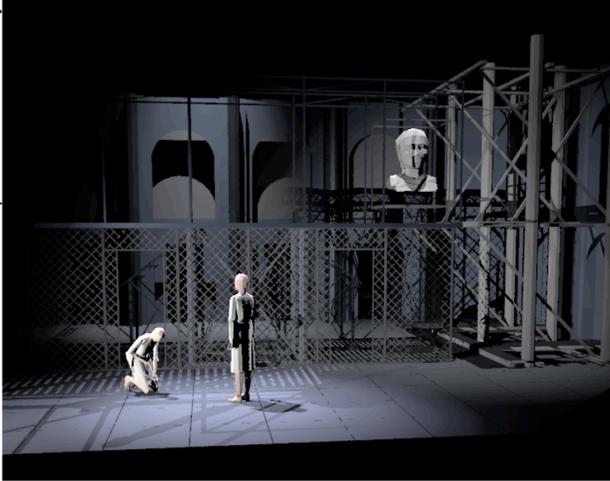


Figure 15. Stage lighting simulation.

Scott Routen & Reuben McFarland / Indiana University

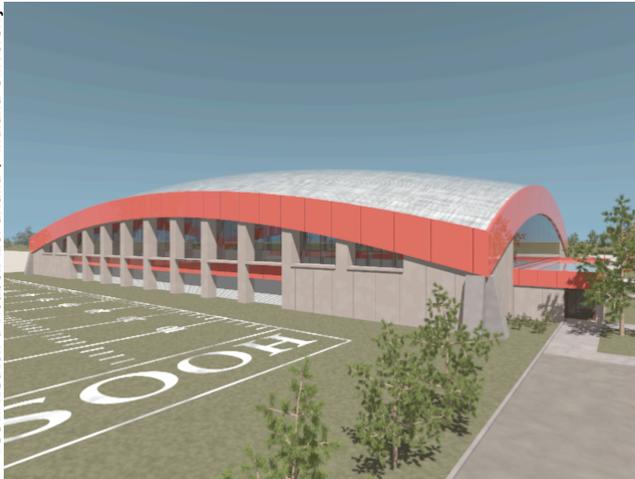


Figure 16. Design of the Mellencamp Pavillion, currently under construction at Indiana University.

John Mardaljevic / De Montfort University



Figure 17. U.K. atrium design modeled without the benefit of a CAD system, using only Radiance scene description utilities.

Scott Routen & Reuben McFarland / Indiana University



Figure 18. Interior view of Mellencamp Pavillion.

Scott Routen & Reuben McFarland / Indiana University



Figure 19. Indiana University library space, illuminated by a central skylight.

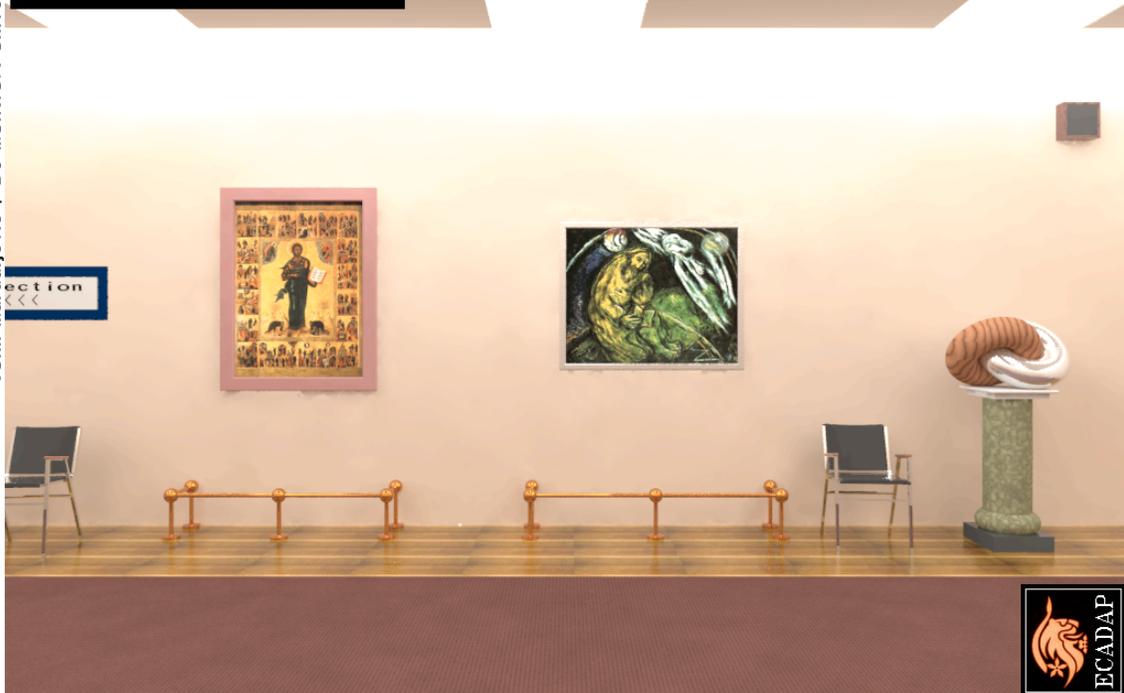


Figure 20a. Art gallery illuminated by skylights.

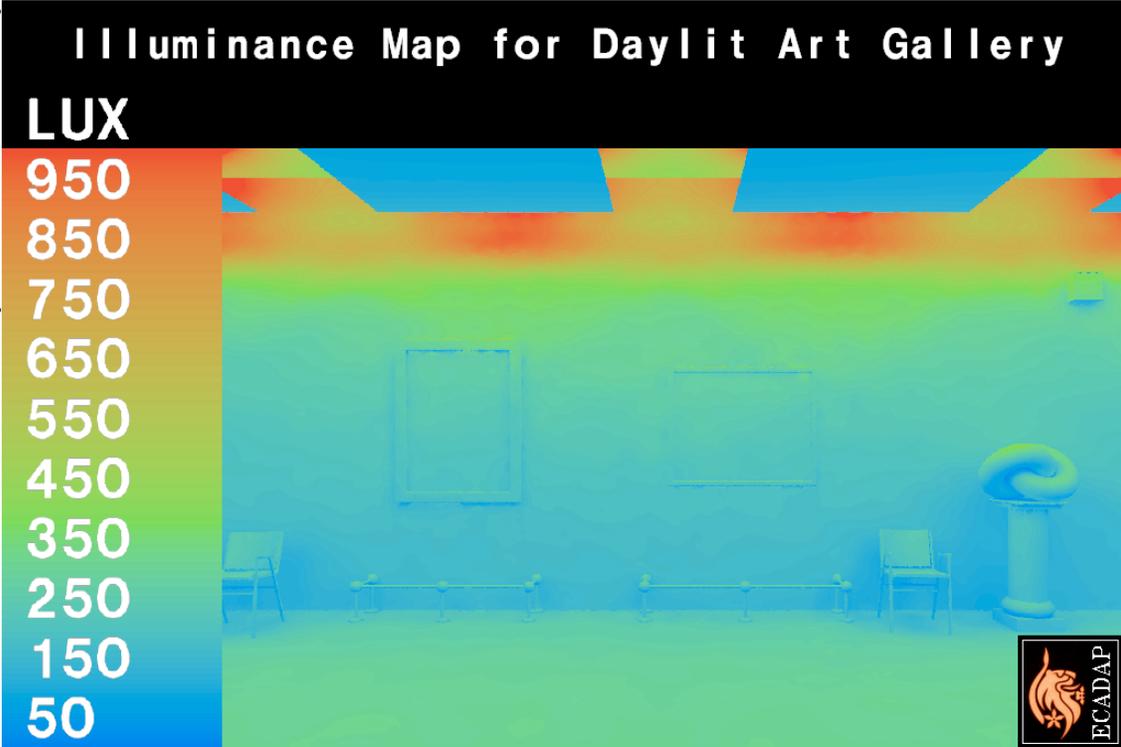


Figure 20b. Corresponding surface illuminances for 20a.

9. Appendix

Table 2 shows some of the more important variables used by the **rad** program, and the effect they have on the rendering process.

Rad Variable Settings			
Variable Name	Interpretation	Legal Values	Affects
DETAIL	geometric detail	High, Med, Low	image sampling, irradiance value density
EXPOSURE	picture exposure	positive real	final picture brightness, ambient approximation
INDIRECT	importance of indirect diffuse contribution	0, 1, 2,...	number of diffuse interreflections
PENUMBRAS	importance of soft shadows	True, False	source subdivision, source sampling, image plane sampling
QUALITY	rendering quality/accuracy	High, Med, Low	nearly everything
VARIABILITY	light distribution in the space	High, Med, Low	indirect irradiance interpolation, hemisphere sampling
ZONE	region of interest	Interior/Exterior keyword plus bounding box	irradiance value density, standard viewpoints

Table 2. Rad rendering variables, their meanings, values and effects.

9.1. Comparison to Other Rendering Environments

Although a comprehensive comparison between rendering systems is beyond the scope of this paper, we can make a few simple observations. Keeping in mind the three challenges of accuracy, generality and practicality, we may judge how well each rendering system fares in light of the goals listed in section 2. Note that these goals are specific to predictive rendering for architectural and lighting design -- a system may fail one or more of these requirements and still be quite useful for other applications.

The most heavily used commercial rendering environments are graphics libraries. These libraries are often developed by computer manufacturers specifically for their graphics hardware. They tend to be very fast and eminently practical, but are not physically accurate or sufficiently general in terms of reflectance, transmittance and geometric modeling to be useful for lighting design. Accuracy and generality have been sacrificed for speed.

So-called "photo-realistic" rendering systems may be general and practical, but they are not accurate. One of the best examples of photo-realistic rendering software is *RenderMan* [Upstill89], which is based on the earlier *REYES* system [Cook87]. Although *Renderman* can generate some beautiful pictures, global illumination is not incorporated, and it is difficult to create accurate light sources. Shadows are usually cast using a z-buffer algorithm that cannot generate penumbra accurately [Reeves87]. However, the system does provide a flexible shading language that can be programmed to simulate some global illumination effects [Cook84b][Hanrahan90].

In recent years, many free rendering packages have been distributed over a network. *RayShade*, one of the best free ray-tracers, does not simulate diffuse interreflection, and uses a non-physical reflection model. As with most photo-realistic rendering software, accuracy is the key missing ingredient. Filling in this gap, some free radiosity programs are starting to show up on the network. Though the author has not had the opportunity to learn about all of them, most appear to use traditional approaches that are limited to diffuse surfaces and simple environments, and therefore are not general or practical enough for lighting design.

Systems for research in rendering and global illumination algorithms exist at hundreds of universities around the world. Few of these systems ever make it out of the laboratory, so it is particularly difficult to judge them in terms of practicality.

However, from what research has been published, it appears that most of these systems are based on classic or progressive radiosity techniques. As we have noted, radiosity relies on diffuse surfaces and relatively simple geometries, so its generality is limited. Extensions to non-diffuse environments tend to be very expensive in time and memory, since directionality greatly complicates the governing equations of a view-independent solution [Sillion91]. Recent work on extending an adjoint system of equations for a view-dependent solution [Smits92] to non-diffuse environments appears promising, but the techniques are still limited to simple geometries [Christensen93][Aupperle93].

The basic problem with the radiosity method is that it ties illumination information to surfaces, and this approach runs into trouble when millions of surfaces are needed to represent a scene. Rushmeier et al addressed this problem with their "clumping" approach, which partially divorces illumination from geometry [Rushmeier93]. Baum et al [Baum91] developed techniques for automatically meshing large models, which saves on manual labor but does not reduce time and space requirements. The theater model shown in Figure 5 was rendered in [Baum91] using automatic meshing and progressive radiosity. Meshing the scene caused it to take up about 100 Mbytes of memory, and rendering took over 18 hours on an SGI R3000 workstation for the direct component alone, compared to 5 hours in 11 Mbytes using *Radiance* on the same computer.

Some of the larger architecture and engineering firms have the resources to create their own in-house lighting simulation and rendering software. Although it is difficult to speculate as to the nature and capabilities of these systems since they are hidden from public view, the author is aware of at least a half-dozen well-funded projects aimed at putting the state of the art in global illumination into practice. Most of these projects are based on progressive radiosity or something closely related. In at least two cases, Abacus Simulations in Scotland and Siemens Lighting in Germany, in-house software projects have been abandoned after considerable expense in favor of using *Radiance*. At least two other firms, Ove Arup in England and Phillips Lighting in the Netherlands, use *Radiance* side by side with in-house software. Of course, we cannot conclude from this that *Radiance* is the best, but the trend is encouraging.

By far the most relevant class of software to compare is commercial lighting simulation and rendering programs. Most of these systems are practical, or people would not buy them. Most are also accurate, or they would not qualify as lighting

simulations. The problem is lack of generality. *LumenMicro* (Lighting Technologies, Boulder, Colorado) is the biggest-selling program among lighting designers, yet it is limited to modeling environments built from grey, diffuse, axis-aligned rectangles. A more promising product is called *LightScape* (LightScape Graphics Software, Toronto, Canada). This software uses progressive radiosity and graphics rendering hardware to provide real-time update capabilities. *LightScape's* ray tracing module may be used to improve shadow resolution and add specular effects, but this solution is expensive and incomplete. Also, memory costs associated with complicated geometries limit the practicality of this system. To be fair, *LightScape* is in its initial release, and has some very accomplished researchers working on it.

One program that shows great potential has recently been released to the U.S. market, *Arris Integra* (Sigma Design in Baltimore). This program uses a bidirectional ray tracing technique developed by Fujimoto [Fujimoto92], and its capabilities have been demonstrated in [Scully93]. The chief drawback of this system seems to be that it is somewhat difficult and expensive to use, costing roughly 15,000 dollars for the basic software and taking many long hours to perform its calculations.

9.2. Prospects for the Future of Rendering

Today's researchers in global illumination have the opportunity to decide the future direction of rendering for decades to come. Most commercial rendering systems currently pay little attention to the physical behavior of light, providing shortcuts such as Phong shading and lights with linear fall-off that undermine realism and make the results useless for lighting design and other predictive applications. We believe that the golden road to realistic rendering is physical simulation, but it is necessary to decide which phenomena shall be included and which may safely be left out. If we choose a scope that is too broad, it will incur large computational expenses with little payoff for users. If our scope is too narrow, we will limit the application areas and realism and therefore limit our audience. Global illumination researchers must join together to set standards for physically-based rendering; standards that will provide a basis for comparison between techniques, and the stability needed for practical progress.

As part of this larger standardization effort, we would like to see a common scene description format adopted by the rendering community. There are many candidates at this point, but none of them contain physically valid light source and material descriptions. We would welcome the use of the *Radiance* format, but extending a conventional scene description language might work better. We suggest the formation of a small committee of global and local illumination researchers to decide what should be included in such a format. We further suggest that one or two graphics hardware or software vendors could cover expenses for this task. In return, the vendors would get a new, physically valid foundation for building the next generation of rendering solutions.

The future of physically-based rendering depends on cooperation and agreement. We must agree on a starting point and work together towards a goal to bring science to this art.

9.3. Appendix References

- [Aupperle93]
Aupperle, Larry, Pat Hanrahan, "Importance and Discrete Three Point Transport," *Proceedings of the Fourth EUROGRAPHICS Workshop on Rendering*, Paris, France, June 1993, pp. 85-94.
- [Baum91]
Baum, Daniel, Stephen Mann, Kevin Smith, James Winget, "Making Radiosity Usable: Automatic Prepro-

cessing and Meshing Techniques for the Generation of Accurate Radiosity Solutions," *Computer Graphics*, Vol. 25, No. 4, July 1991.

- [Cook84b]
Cook, Robert, "Shade Trees," *Computer Graphics*, Vol. 18, No. 3, July 1984, pp. 223-232.
- [Cook87]
Cook, Robert, Loren Carpenter, Edwin Catmull, "The Reyes Image Rendering Architecture," *Computer Graphics*, Vol. 21, No. 4, July 1987, pp. 95-102.
- [Christensen93]
Christensen, Per, David Salesin, Tony DeRose, "A Continuous Adjoint Formulation for Radiance Transport," *Proceedings of the Fourth EUROGRAPHICS Workshop on Rendering*, Paris, France, June 1993, pp. 95-104.
- [Fujimoto92]
Fujimoto, Akira, Nancy Hays, "Mission Impossible: High Tech Made in Poland," *Computer Graphics and Applications*, Vol. 12, No. 2, March 1992, pp. 8-13.
- [Hanrahan90]
Hanrahan, Pat and Jim Lawson, "A Language for Shading and Lighting Calculations," *Computer Graphics*, Vol. 24, No. 4, August 1990, pp. 289-298.
- [Reeves87]
Reeves, William, David Salesin, Robert Cook, "Rendering Antialiased Shadows with Depth Maps," *Computer Graphics*, Vol. 21, No. 4, July 1987, pp. 283-291.
- [Rushmeier93]
Rushmeier, Holly, Charles Patterson, Aravindan Veerasamy, "Geometric Simplification for Indirect Illumination Calculations," *Proceedings of Graphics Interface '93*, May 1993, pp. 227-236.
- [Scully93]
Scully, Vincent, "A Virtual Landmark," *Progressive Architecture*, September 1993, pp. 80-87.
- [Sillion91]
Sillion, Francois, James Arvo, Stephen Westin, Donald Greenberg, "A Global Illumination Solution for General Reflectance Distributions," *Computer Graphics*, Vol. 25, No. 4, July 1991, pp. 187-196.
- [Upstill89]
Upstill, Steve, *The RenderMan Companion*, Addison-Wesley, 1989.

A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes

Gregory Ward Larson[†]
Building Technologies Program
Environmental Energy Technologies Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
1 Cyclotron Road
Berkeley, California 94720

Holly Rushmeier
IBM T.J. Watson Research Center

Christine Piatko^{††}
National Institute for Standards and Technology

January 15, 1997

This paper is available electronically at:
<http://radsite.lbl.gov/radiance/papers>

Copyright 1997 Regents of the University of California
subject to the approval of the Department of Energy

[†] Author's current address: Silicon Graphics, Inc., Mountain View, CA.

^{††} Author's current address: JHU/APL, Laurel, MD.

A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes

Gregory Ward Larson
Lawrence Berkeley National Laboratory

Holly Rushmeier
IBM T.J. Watson Research Center

Christine Piatko
National Institute for Standards and Technology

ABSTRACT

We present a tone reproduction operator that preserves visibility in high dynamic range scenes. Our method introduces a new histogram adjustment technique, based on the population of local adaptation luminances in a scene. To match subjective viewing experience, the method incorporates models for human contrast sensitivity, glare, spatial acuity and color sensitivity. We compare our results to previous work and present examples of our techniques applied to lighting simulation and electronic photography.

Keywords: Shading, Image Manipulation.

1 Introduction

The real world exhibits a wide range of luminance values. The human visual system is capable of perceiving scenes spanning 5 orders of magnitude, and adapting more gradually to over 9 orders of magnitude. Advanced techniques for producing synthetic images, such as radiosity and Monte Carlo ray tracing, compute the map of luminances that would reach an observer of a real scene. The media used to display these results -- either a video display or a print on paper -- cannot reproduce the computed luminances, or span more than a few orders of magnitude. However, the success of realistic image synthesis has shown that it is possible to produce images that convey the appearance of the simulated scene by mapping to a set of luminances that can be produced by the display medium. This is fundamentally possible because the human eye is sensitive to relative rather than absolute luminance values. However, a robust algorithm for converting real world luminances to display luminances has yet to be developed.

The conversion from real world to display luminances is known as *tone mapping*. Tone mapping ideas were originally developed for photography. In photography or video, chemistry or electronics, together with a human actively controlling the scene lighting and the camera, are used to map real world luminances into an acceptable image on a

display medium. In synthetic image generation, our goal is to avoid active control of lighting and camera settings. Furthermore, we hope to improve tone mapping techniques by having direct numerical control over display values, rather than depending on the physical limitations of chemistry or electronics.

Consider a typical scene that poses a problem for tone reproduction in both photography and computer graphics image synthesis systems. The scene is a room illuminated by a window that looks out on a sunlit landscape. A human observer inside the room can easily see individual objects in the room, as well as features in the outdoor landscape. This is because the eye adapts locally as we scan the different regions of the scene. If we attempt to photograph our view, the result is disappointing. Either the window is over-exposed and we can't see outside, or the interior of the room is under-exposed and looks black. Current computer graphics tone operators either produce the same disappointing result, or introduce artifacts that do not match our perception of the actual scene.

In this paper, we present a new tone reproduction operator that reliably maps real world luminances to display luminances, even in the problematic case just described. We consider the following two criteria most important for reliable tone mapping:

1. Visibility is reproduced. You can see an object in the real scene if and only if you can see it in the display. Objects are not obscured in under- or over-exposed regions, and features are not lost in the middle.
2. Viewing the image produces a subjective experience that corresponds with viewing the real scene. That is, the display should correlate well with memory of the actual scene. The overall impression of brightness, contrast, and color should be reproduced.

Previous tone mapping operators have generally met one of these criteria at the expense of the other. For example, some preserve the visibility of objects while changing the impression of contrast, while others preserve the overall impression of brightness at the expense of visibility.

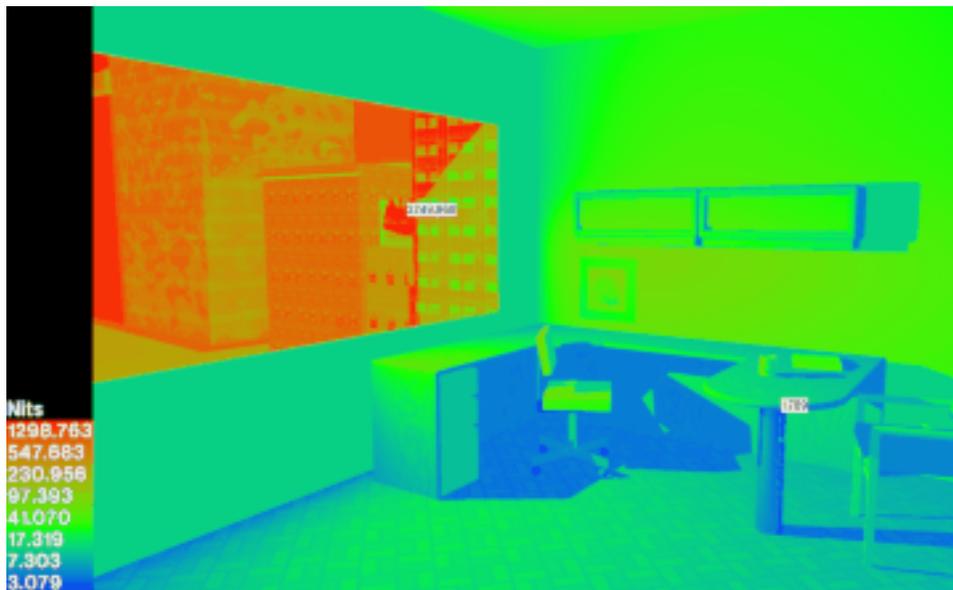


Figure 1. A false color image showing the world luminance values for a window office in candelas per meter squared (cd/m^2 or Nits).

The new tone mapping operator we present addresses our two criteria. We develop a method of modifying a luminance histogram, discovering clusters of adaptation levels and efficiently mapping them to display values to preserve local contrast visibility. We then use models for glare, color sensitivity and visual acuity to reproduce imperfections in human vision that further affect visibility and appearance.



Figure 2. A linear mapping of the luminances in Figure 1 that over-exposes the view through the window.



Figure 3. A linear mapping of the luminances in Figure 1 that under-exposes the view of the interior.



Figure 4. The luminances in Figure 1 mapped to preserve the visibility of both indoor and outdoor features using the new tone mapping techniques described in this paper.

2 Previous Work

The high dynamic range problem was first encountered in computer graphics when physically accurate illumination methods were developed for image synthesis in the 1980's. (See Glassner [Glassner95] for a comprehensive review.) Previous methods for generating images were designed to automatically produce dimensionless values more or less evenly distributed in the range 0 to 1 or 0 to 255, which could be readily mapped to a display device. With the advent of radiosity and Monte Carlo path tracing techniques, we began to compute images in terms of real units with the real dynamic range of physical illumination. Figure 1 is a false color image showing the magnitude and distribution of luminance values in a typical indoor scene containing a window to a sunlit exterior. The goal of image synthesis is to produce results such as Figure 4, which match our impression of what such a scene looks like. Initially though, researchers found that a wide range of displayable images could be obtained from the same input luminances -- such as the unsatisfactory over- and under-exposed linear reproductions of the image in Figures 2 and 3.

Initial attempts to find a consistent mapping from computed to displayable luminances were ad hoc and developed for computational convenience. One approach is to use a function that collapses the high dynamic range of luminance into a small numerical range. By taking the cube root of luminance, for example, the range of values is reduced to something that is easily mapped to the display range. This approach generally preserves visibility of objects, our first criterion for a tone mapping operator. However, condensing the range of values in this way reduces fine detail visibility, and distorts impressions of brightness and contrast, so it does not fully match visibility or reproduce the subjective appearance required by our second criterion.

A more popular approach is to use an arbitrary linear scaling, either mapping the average of luminance in the real world to the average of the display, or the maximum non-light source luminance to the display maximum. For scenes with a dynamic range similar to the display device, this is successful. However, linear scaling methods do not maintain visibility in scenes with high dynamic range, since very bright and very dim values are clipped to fall within the display's limited dynamic range. Furthermore, scenes are mapped the same way regardless of the absolute values of luminance. A scene illuminated by a search light could be mapped to the same image as a scene illuminated by a flashlight, losing the overall impression of brightness and so losing the subjective correspondence between viewing the real and display-mapped scenes.

A tone mapping operator proposed by Tumblin and Rushmeier [Tumblin93] concentrated on the problem of preserving the viewer's overall impression of brightness. As the light level that the eye adapts to in a scene changes, the relationship between brightness (the subjective impression of the viewer) and luminance (the quantity of light in the visible range) also changes. Using a brightness function proposed by Stevens and Stevens [Stevens60], they developed an operator that would preserve the overall impression of brightness in the image, using one adaptation value for real scene, and another adaptation value for the displayed image. Because a single adaptation level is used for the scene, though, preservation of brightness in this case is at the expense of visibility. Areas that are very bright or dim are clipped, and objects in these areas are obscured.

Ward [Ward91] developed a simpler tone mapping method, designed to preserve feature visibility. In this method, a non-arbitrary linear scaling factor is found that preserves the impression of contrast (i.e., the visible changes in luminance) between the real and displayed image at a particular fixation point. While visibility is maintained at this adaptation point, the linear scaling factor still results in the clipping of very high and very low values, and correct visibility is not maintained throughout the image.

Chiu et al. [Chiu93] addressed this problem of global visibility loss by scaling luminance values based on a spatial average of luminances in pixel neighborhoods. Values in bright or dark areas would not be clipped, but scaled according to different values based on their spatial location. Since the human eye is less sensitive to variations at low spatial frequencies than high ones, a variable scaling that changes slowly relative to image features is not immediately visible. However, in a room with a bright source and dark corners, the method inevitably produces display luminance gradients that are the opposite of real world gradients. To make a dark region around a bright source, the transition from a dark area in the room to a bright area shows a decrease in brightness rather than an increase. This is illustrated in Figure 5 which shows a bright source with a dark halo around it. The dark halo that facilitates rendering the visibility of the bulb disrupts what should be a symmetric pattern of light cast by the bulb on the wall behind it. The reverse gradient fails to preserve the subjective correspondence between the real room and the displayed image.

Inspired by the work of Chiu et al., Schlick [Schlick95] developed an alternative method that could compute a spatially varying tone mapping. Schlick's work concentrated on improving computational efficiency and simplifying parameters, rather than improving the subjective correspondence of previous methods.



Figure 5. Dynamic range compression based on a spatially varying scale factor (from [Chiu93]).

Contrast, brightness and visibility are not the only perceptions that should be maintained by a tone mapping operator. Nakamae et al. [Nakamae90] and Spencer et al. [Spencer95] have proposed methods to simulate the effects of glare. These methods simulate the scattering in the eye by spreading the effects of a bright source in an image. Ferwerda et al. [Ferwerda96] proposed a method that accounts for changes in spatial acuity and color sensitivity as a function of light level. Our work is largely inspired by these papers, and we borrow heavily from Ferwerda et al. in particular. Besides maintaining visibility and the overall impression of brightness, the effects of glare, spatial acuity and color sensitivity must be included to fully meet our second criterion for producing a subjective correspondence between the viewer in the real scene and the viewer of the synthetic image.

A related set of methods for adjusting image contrast and visibility have been developed in the field of image processing for image enhancement (e.g., see Chapter 3 in [Green83]). Perhaps the best known image enhancement technique is histogram equalization. In histogram equalization, the grey levels in an image are redistributed more evenly to make better use of the range of the display device. Numerous improvements have been made to simple equalization by incorporating models of perception. Frei [Frei77] introduced histogram hyperbolization that attempts to redistribute perceived brightness, rather than screen grey levels. Frei approximated brightness using the logarithm of luminance. Subsequent researchers such as Mokrane [Mokrane92] have introduced methods that use more sophisticated models of perceived brightness and contrast.

The general idea of altering histogram distributions and using perceptual models to guide these alterations can be applied to tone mapping. However, there are two important

differences between techniques used in image enhancement and techniques for image synthesis and real-world tone mapping:

1. In image enhancement, the problem is to correct an image that has already been distorted by photography or video recording and collapsed into a limited dynamic range. In our problem, we begin with an undistorted array of real world luminances with a potentially high dynamic range.
2. In image enhancement, the goal is to take an imperfect image and *maximize* visibility or contrast. Maintaining subjective correspondence with the original view of the scene is irrelevant. In our problem, we want to maintain subjective correspondence. We want to *simulate* visibility and contrast, not maximize it. We want to produce visually accurate, not enhanced, images.

3 Overview of the New Method

In constructing a new method for tone mapping, we wish to keep the elements of previous methods that have been successful, and overcome the problems.

Consider again the room with a window looking out on a sunlit landscape. Like any high dynamic range scene, luminance levels occur in clusters, as shown in the histogram in Figure 6, rather than being uniformly distributed throughout the dynamic range. The failure of any method that uses a single adaptation level is that it maps a large range of sparsely populated real world luminance levels to a large range of display values. If the eye were sensitive to absolute values of luminance difference, this would be necessary. However, the eye is only sensitive to the fact that there are bright areas and dim areas. As long as the bright areas are displayed by higher luminances than the dim areas in the final image, the absolute value of the difference in luminance is not important. Exploiting this aspect of vision, we can close the gap between the display values for high and low luminance regions, and we have more display luminances to work with to render feature visibility.

Another failure of using a uniform adaptation level is that the eye rapidly adapts to the level of a relatively small angle in the visual field (i.e., about 1°) around the current fixation point [Moon&Spencer45]. When we look out the window, the eye adapts to the high exterior level, and when we look inside, it adapts to the low interior level. Chiu et al. [Chiu93] attempted to account for this using spatially varying scaling factors, but this method produces noticeable gradient reversals, as shown in Figure 5.

Rather than adjusting the adaptation level based on spatial location in the image, we will base our mapping on the population of the luminance adaptation levels in the image. To identify clusters of luminance levels and *initially* map them to display values, we will use the cumulative distribution of the luminance histogram. More specifically, we will start with a cumulative distribution based on a logarithmic approximation of brightness from luminance values.

Histogram of Brightness

Window Office

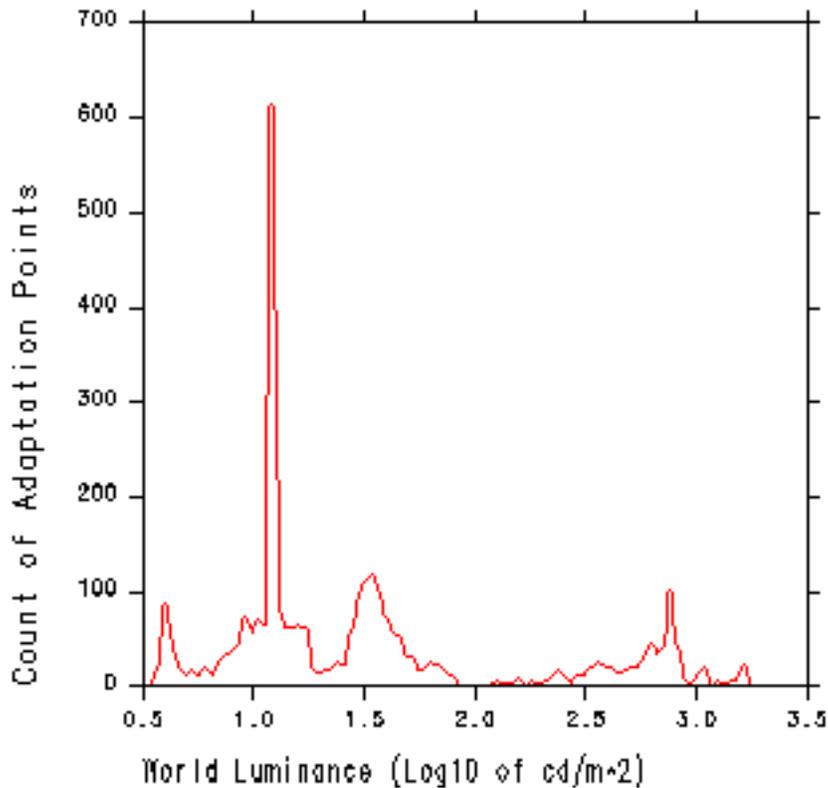


Figure 6. A histogram of adaptation values from Figure 1 (1° spot luminance averages).

First, we calculate the population of levels from a luminance image of the scene in which each pixel represents 1° in the visual field. We make a crude approximation of the brightness values (i.e., the subjective response) associated with these luminances by taking the logarithm of luminance. (Note that we will not display logarithmic values, we will merely use them to obtain a distribution.) We then build a histogram and cumulative distribution function from these values. Since the brightness values are integrated over a small solid angle, they are in some sense based on a spatial average, and the resulting mapping will be local to a particular adaptation level. Unlike Chiu's method however, the mapping for a particular luminance level will be consistent throughout the image, and will be order preserving. Specifically, an increase in real scene luminance level will always be represented by an increase in display luminance. The histogram and cumulative distribution function will allow us to close the gaps of sparsely populated luminance values and avoid the clipping problems of single adaptation level methods. By deriving a single, global tone mapping operator from locally averaged adaptation levels, we avoid the reverse gradient artifacts associated with a spatially varying multiplier.

We will use this histogram only as a starting point, and impose restrictions to preserve (rather than maximize) contrast based on models of human perception using our knowledge of the true luminance values in the scene. Simulations of glare and variations in spatial acuity and color sensitivity will be added into the model to maintain subjective correspondence and visibility. In the end, we obtain a mapping of real world to display luminance similar to the one shown in Figure 7.

For our target display, all mapped brightness values below 1 cd/m^2 (0 on the vertical axis) or above 100 (2 on the vertical axis) are lost because they are outside the displayable range. Here we see that the dynamic range between 1.75 and 2.5 has been compressed, yet we don't notice it in the displayed result (Figure 4). Compared to the two linear operators, our new tone mapping is the only one that can represent the entire scene without losing object or detail visibility.

World to Display Luminance Mapping Window Office

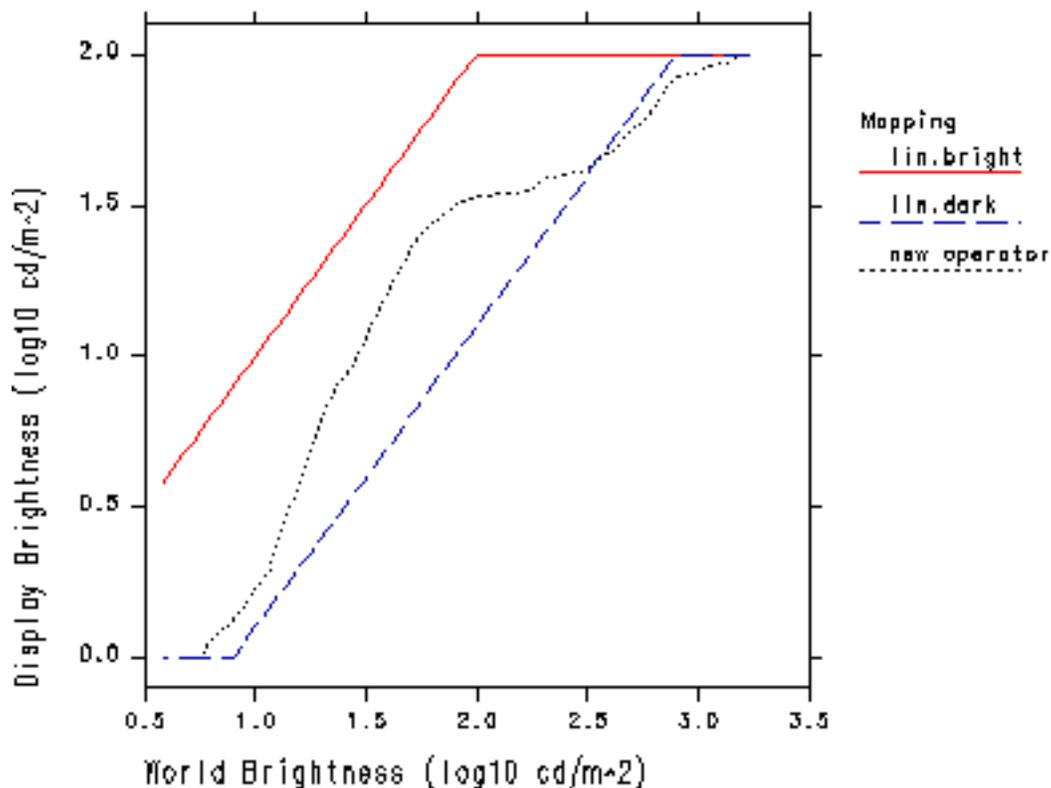


Figure 7. A plot comparing the global brightness mapping functions for Figures 1, 2, and 3, respectively.

In the following section, we illustrate this technique for histogram adjustment based on contrast sensitivity. After this, we describe models of glare, color sensitivity and visual

acuity that complete our simulation of the measurable and subjective responses of human vision. Finally, we complete the methods presentation with a summary describing how all the pieces fit together.

4 Histogram Adjustment

In this section, we present a detailed description of our basic tone mapping operator. We begin with the introduction of symbols and definitions, and a description of the histogram calculation. We then describe a naive equalization step that partially accomplishes our goals, but results in undesirable artifacts. This method is then refined with a linear contrast ceiling, which is further refined using human contrast sensitivity data.

4.1 Symbols and Definitions

L_w	= world luminance (in candelas/meter ²)
B_w	= world brightness, $\log(L_w)$
L_{wmin}	= minimum world luminance for scene
L_{wmax}	= maximum world luminance for scene
L_d	= display luminance (in candelas/meter ²)
L_{dmin}	= minimum display luminance (black level)
L_{dmax}	= maximum display luminance (white level)
B_{de}	= computed display brightness, $\log(L_d)$ [Equation (4)]
N	= the number of histogram bins
T	= the total number of adaptation samples
$f(b_i)$	= frequency count for the histogram bin at b_i
b	= the bin step size in $\log(cd/m^2)$
$P(b)$	= the cumulative distribution function [Equation (2)]
$\log(x)$	= natural logarithm of x
$\log_{10}(x)$	= decimal logarithm of x

4.2 Histogram Calculation

Since we are interested in optimizing the mapping between world adaptation and display adaptation, we start with a histogram of world adaptation luminances. The eye adapts for the best view in the fovea, so we compute each luminance over a 1° diameter solid angle corresponding to a potential foveal fixation point in the scene. We use a logarithmic scale for the histogram to best capture luminance population and subjective response over a wide dynamic range. This requires setting a minimum value as well as a maximum, since the logarithm of zero is -∞. For the minimum value, we use either the minimum 1° spot average, or 10⁻⁴ cd/m² (the lower threshold of human vision), whichever is larger. The maximum value is just the maximum spot average.

We start by filtering our original floating-point image down to a resolution that roughly corresponds to 1° square pixels. If we are using a linear perspective projection, the pixels on the perimeter will have slightly smaller diameter than the center pixels, but they will still be within the correct range. The following formula yields the correct resolution for

1° diameter pixels near the center of a linear perspective image:

$$S = 2 \tan(\theta/2) / 0.01745 \quad (1)$$

where:

$$\begin{aligned} S &= \text{width or height in pixels} \\ \theta &= \text{horizontal or vertical full view angle} \\ 0.01745 &= \text{number of radians in } 1^\circ \end{aligned}$$

For example, the view width and height for Figure 4 are 63° and 45° respectively, which yield a sample image resolution of 70 by 47 pixels. Near the center, the pixels will be 1° square exactly, but near the corners, they will be closer to 0.85° for this wide-angle view. The filter kernel used for averaging will have little influence on our result, so long as every pixel in the original image is weighted similarly. We employ a simple box filter.

From our reduced image, we compute the logarithms of the floating-point luminance values. Here, we assume there is some method for obtaining the absolute luminances at each spot sample. If the image is uncalibrated, then the corrections for human vision will not work, although the method may still be used to optimize the visible dynamic range. (We will return to this in the summary.)

The histogram is taken between the minimum and maximum values mentioned earlier in equal-sized bins on a log(luminance) scale. The algorithm is not sensitive to the number of bins, so long as there are enough to obtain adequate resolution. We use 100 bins in all of our examples. The resulting histogram for Figure 1 is shown in Figure 6.

4.2.1 Cumulative Distribution

The cumulative frequency distribution is defined as:

$$P(b) = \frac{\sum_{b_i < b} f(b_i)}{T} \quad (2)$$

where:

$$T = \sum_{b_i} f(b_i) \text{ (i.e., the total number of samples)}$$

Later on, we will also need the derivative of this function. Since the cumulative distribution is a numerical integration of the histogram, the derivative is simply the histogram with an appropriate normalization factor. In our method, we approximate a continuous distribution and derivative by interpolating adjacent values linearly. The derivative of our function is:

$$\frac{dP(b)}{db} = \frac{f(b)}{T \cdot b} \quad (3)$$

where:

$$b = \frac{[\log(L_{wmax}) - \log(L_{wmin})]}{N} \text{ (i.e., the size of each bin)}$$



Figure 8. Rendering of a bathroom model mapped with a linear operator.

4.3 Naive Histogram Equalization

If we wanted all the brightness values to have equal probability in our final displayed image, we could now perform a straightforward histogram equalization. Although this is not our goal, it is a good starting point for us. Based on the cumulative frequency distribution just described, the equalization formula can be stated in terms of brightness as follows:

$$B_{de} = \log(L_{dmin}) + [\log(L_{dmax}) - \log(L_{dmin})] P(B_w) \quad (4)$$

The problem with naive histogram equalization is that it not only compresses dynamic range (contrast) in regions where there are few samples, it also *expands* contrast in highly populated regions of the histogram. The net effect is to exaggerate contrast in large areas of the displayed image. Take as an example the scene shown in Figure 8. Although we cannot see the region surrounding the lamps due to the clamped linear tone mapping operator, the image appears to us as more or less normal. Applying the naive histogram equalization, Figure 9 is produced. The tiles in the shower now have a mottled appearance. Because this region of world luminance values is so well represented, naive

histogram equalization spreads it out over a relatively larger portion of the display's dynamic range, generating superlinear contrast in this region.



Figure 9. Naive histogram equalization allows us to see the area around the light sources but contrast is exaggerated in other areas such as the shower tiles.

4.4 Histogram Adjustment with a Linear Ceiling

If the contrast being produced is too high, then what is an appropriate contrast for representing image features? The crude answer is that the contrast in any given region should not exceed that produced by a linear tone mapping operator, since linear operators produce satisfactory results for scenes with limited dynamic range. We will take this simple approach first, and later refine our answer based on human contrast sensitivity.

A linear ceiling on the contrast produced by our tone mapping operator can be written thus:

$$\frac{dL_d}{dL_w} \leq \frac{L_d}{L_w} \quad (5a)$$

That is, the derivative of the display luminance with respect to the world luminance must not exceed the display luminance divided by the world luminance. Since we have an expression for the display luminance as a function of world luminance for our naive histogram equalization, we can differentiate the exponentiation of Equation (4) using the chain rule and the derivative from Equation (3) to get the following inequality:

$$\exp(B_{de}) \frac{f(B_w)}{T} \frac{b}{\log(L_{dmax}) - \log(L_{dmin})} \frac{L_d}{L_w} \quad (5b)$$

Since L_d is equal to $\exp(B_{de})$, this reduces to a constant ceiling on $f(b)$:

$$f(b) \leq \frac{T}{\log(L_{dmax}) - \log(L_{dmin})} \quad (5c)$$

In other words, so long as we make sure no frequency count exceeds this ceiling, our resulting histogram will not exaggerate contrast. How can we create this modified histogram? We considered both truncating larger counts to this ceiling and redistributing counts that exceeded the ceiling to other histogram bins. After trying both methods, we found truncation to be the simplest and most reliable approach. The only complication introduced by this technique is that once frequency counts are truncated, T changes, which changes the ceiling. We therefore apply iteration until a tolerance criterion is met, which says that fewer than 2.5% of the original samples exceed the ceiling.¹ Our pseudocode for `histogram_ceiling` is given below:

```

boolean function histogram_ceiling()
tolerance := 2.5% of histogram total
repeat {
    trimmings := 0
    compute the new histogram total T
    if T < tolerance then
        return FALSE
    foreach histogram bin i do
        compute the ceiling
        if  $f(b_i) > \text{ceiling}$  then {
            trimmings +=  $f(b_i) - \text{ceiling}$ 
             $f(b_i) := \text{ceiling}$ 
        }
} until trimmings <= tolerance
return TRUE

```

This iteration will fail to converge (and the function will return `FALSE`) if and only if the dynamic range of the output device is already ample for representing the sample luminances in the original histogram. This is evident from Equation (5c), since b is the world brightness range over the number of bins:

$$f(b_i) = \frac{T}{N} \frac{[\log(L_{wmax}) - \log(L_{wmin})]}{[\log(L_{dmax}) - \log(L_{dmin})]} \quad (5d)$$

¹The tolerance of 2.5% was chosen as an arbitrary small value, and it seems to make little difference either to the convergence time or the results.

If the ratio of the world brightness range over the display brightness range is less than one (i.e., our world range fits in our display range), then our frequency ceiling is less than the total count over the number of bins. Such a condition will never be met, since a uniform distribution of samples would still be over the ceiling in every bin. It is easiest to detect this case at the outset by checking the respective brightness ranges, and applying a simple linear operator if compression is unnecessary.

We call this method *histogram adjustment* rather than histogram equalization because the final brightness distribution is *not equalized*. The net result is a mapping of the scene's high dynamic range to the display's smaller dynamic range that minimizes visible contrast distortions, by compressing under-represented regions without expanding over-represented ones.

Figure 10 shows the results of our histogram adjustment algorithm with a linear ceiling. The problems of exaggerated contrast are resolved, and we can still see the full range of brightness. A comparison of these tone mapping operators is shown in Figure 11. The naive operator is superlinear over a large range, seen as a very steep slope near world luminances around $10^{0.8}$.



Figure 10. Histogram adjustment with a linear ceiling on contrast preserves both lamp visibility and tile appearance.

Brightness Mapping Function

Bathroom

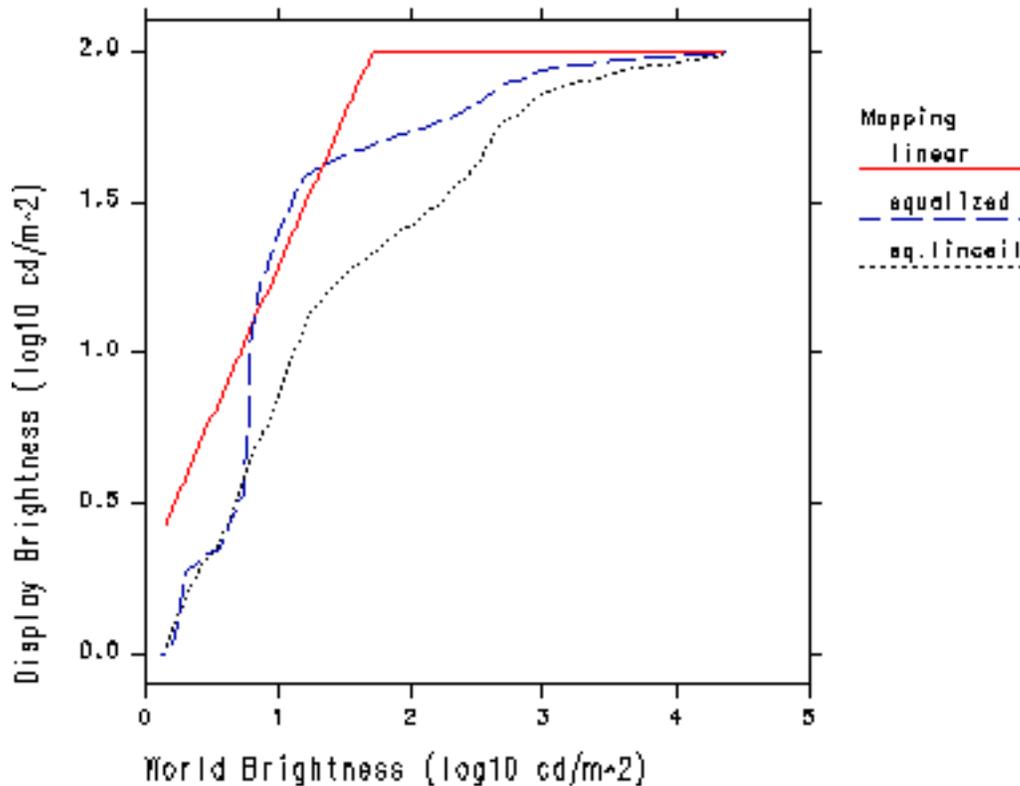


Figure 11. A comparison of naive histogram equalization with histogram adjustment using a linear contrast ceiling.

The method we have just presented is itself quite useful. We have managed to overcome limitations in the dynamic range of typical displays without introducing objectionable contrast compression artifacts in our image. In situations where we want to get a good, natural-looking image without regard to how well a human observer would be able to see in a real environment, this may be an optimal solution. However, if we are concerned with reproducing both visibility and subjective experience in our displayed image, then we must take it a step further and consider the *limitations* of human vision.

4.5 Histogram Adjustment Based on Human Contrast Sensitivity

Although the human eye is capable of adapting over a very wide dynamic range (on the order of 10^9), we do not see equally well at all light levels. As the light grows dim, we have more and more trouble detecting contrast. The relationship between adaptation luminance and the minimum detectable luminance change is well studied [CIE81]. For consistency with earlier work, we use the same detection threshold function used by Ferwerda et al. [Ferwerda96]. This function covers sensitivity from the lower limit of

human vision to daylight levels, and accounts for both rod and cone response functions. The piecewise fit is reprinted in Table 1.

log10 of just noticeable difference	applicable luminance range
-2.86	$\log_{10}(L_a) < -3.94$
$(0.405 \log_{10}(L_a) + 1.6)^{2.18} - 2.86$	$-3.94 \leq \log_{10}(L_a) < -1.44$
$\log_{10}(L_a) - 0.395$	$-1.44 \leq \log_{10}(L_a) < -0.0184$
$(0.249 \log(L_a) + 0.65)^{2.7} - 0.72$	$-0.0184 \leq \log_{10}(L_a) < 1.9$
$\log_{10}(L_a) - 1.255$	$\log_{10}(L_a) \geq 1.9$

Table 1. Piecewise approximation for $L_t(L_a)$.

We name this combined sensitivity function:

$$L_t(L_a) = \text{"just noticeable difference" for adaptation level } L_a \quad (6)$$

Ferwerda et al. did not combine the rod and cone sensitivity functions in this manner, since they used the two ranges for different tone mapping operators. Since we are using this function to control the maximum reproduced contrast, we combine them at their crossover point of $10^{-0.0184}$ cd/m².

To guarantee that our display representation does not exhibit contrast that is more noticeable than it would be in the actual scene, we constrain the slope of our operator to the ratio of the two adaptation thresholds for the display and world, respectively. This is the same technique introduced by Ward [Ward91] and used by Ferwerda et al. [Ferwerda96] to derive a global scale factor. In our case, however, the overall tone mapping operator will not be linear, since the constraint will be met at all potential adaptation levels, not just a single selected one. The new ceiling can be written as:

$$\frac{dL_d}{dL_w} = \frac{L_t(L_d)}{L_t(L_w)} \quad (7a)$$

As before, we compute the derivative of the histogram equalization function (Equation (4)) to get:

$$\exp(B_{de}) \frac{f(B_w)}{T b} \frac{\log(L_{dmax}) - \log(L_{dmin})}{L_w} = \frac{L_t(L_d)}{L_t(L_w)} \quad (7b)$$

However, this time the constraint does not reduce to a constant ceiling for $f(b)$. We notice that since L_d equals $\exp(B_{de})$ and B_{de} is a function of L_w from Equation (4), our

ceiling is completely defined for a given $P(b)$ and world luminance, L_w :

$$f(B_w) = \frac{L_t(L_d)}{L_t(L_w)} \frac{T b L_w}{[\log(L_{dmax}) - \log(L_{dmin})] L_d} \quad (7c)$$

where:

$$L_d = \exp(B_{de}), B_{de} \text{ given in Equation (4)}$$

Once again, we must iterate to a solution, since truncating bin counts will affect T and $P(b)$. We reuse the `histogram_ceiling` procedure given earlier, replacing the linear contrast ceiling computation with the above formula.

Brightness Mapping Function Bathroom

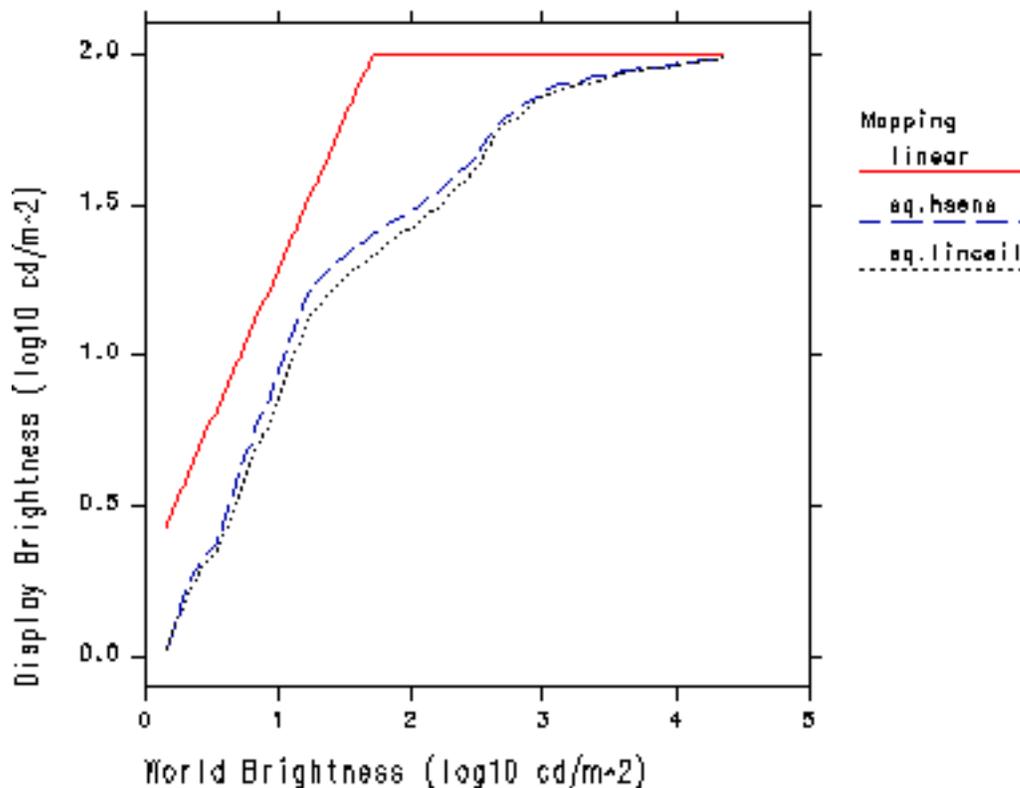


Figure 12. Our tone mapping operator based on human contrast sensitivity compared to the histogram adjustment with linear ceiling used in Figure 10. Human contrast sensitivity makes little difference at these light levels.

Figure 12 shows the same curves for the linear tone mapping and histogram adjustment with linear clamping shown before in Figure 11, but with the curve for naive histogram

equalization replaced by our human visibility matching algorithm. We see the two histogram adjustment curves are very close. In fact, we would have some difficulty differentiating images mapped with our latest method and histogram adjustment with a linear ceiling. This is because the scene we have chosen has most of its luminance levels in the same range as our display luminances. Therefore, the ratio between display and world luminance detection thresholds is close to the ratio of the display and world adaptation luminances. This is known as Weber's law [Riggs71], and it holds true over a wide range of luminances where the eye sees equally well. This correspondence makes the right-hand side of Equations (5b) and (7b) equivalent, and so we should expect the same result as a linear ceiling.

Brightness Mapping Function

Dim Bathroom

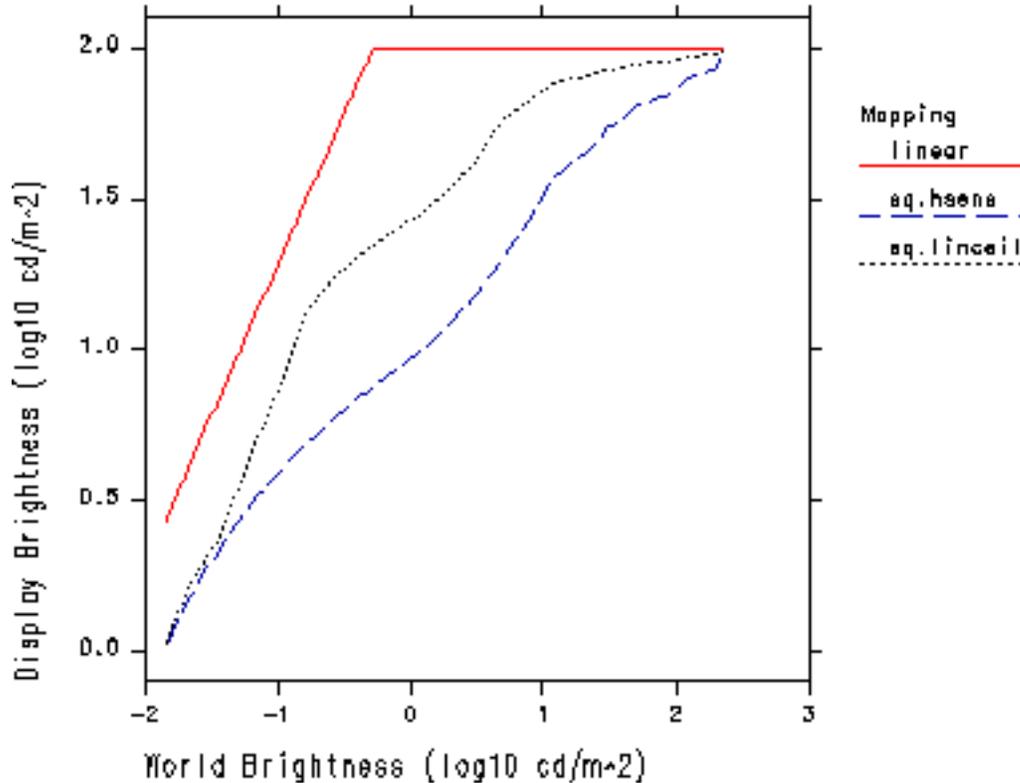


Figure 13. The brightness map for the bathroom scene with lights dimmed to 1/100th of their original intensity, where human contrast sensitivity makes a difference.

To see a contrast sensitivity effect, our world adaptation would have to be very different from our display adaptation. If we reduce the light level in the bathroom by a factor of 100, our ability to detect contrast is diminished. This shows up in a relatively larger detection threshold in the denominator of Equation (7c), which reduces the ceiling for the

frequency counts. The change in the tone mapping operator is plotted in Figure 13 and the resulting image is shown in Figure 14.

Figure 13 shows that the linear mapping is unaffected, since we just raise the scale factor to achieve an average exposure. Likewise, the histogram adjustment with a linear ceiling maps the image to the same display range, since its goal is to reproduce linear contrast. However, the ceiling based on human threshold visibility limits contrast over much of the scene, and the resulting image is darker and less visible everywhere except the top of the range, which is actually shown with higher contrast since we now have display range to spare.

Figure 14 is darker and the display contrast is reduced compared to Figure 10. Because the tone mapping is based on *local adaptation* rather than a single global or spot average, threshold visibility is reproduced *everywhere* in the image, not just around a certain set of values. This criterion is met within the limitations of the display's dynamic range.



Figure 14. The dimmed bathroom scene mapped with the function shown in Figure 13.

5 Human Visual Limitations

We have seen how histogram adjustment matches display contrast visibility to world visibility, but we have ignored three important limitations in human vision: glare, color sensitivity and visual acuity. Glare is caused by bright sources in the visual periphery, which scatter light in the lens of the eye, obscuring foveal vision. Color sensitivity is lost in dark environments, as the light-sensitive rods take over for the color-sensitive cone system. Visual acuity is also impaired in dark environments, due to the complete loss of cone response and the quantum nature of light sensation.

In our treatment, we will rely heavily on previous work performed by Moon and Spencer [Moon&Spencer45] and Ferwerda et al. [Ferwerda96], applying it in the context of a locally adapted visibility-matching model.

5.1 Veiling Luminance

Bright *glare sources* in the periphery reduce contrast visibility because light scattered in the lens and aqueous humor obscures the fovea; this effect is less noticeable when looking directly at a source, since the eye adapts to the high light level. The influence of glare sources on contrast sensitivity is well studied and documented. We apply the work of Holladay [Holladay26] and Moon and Spencer [Moon&Spencer45], which relates the effective adaptation luminance to the foveal average and glare source position and illuminance.

In our presentation, we will first compute a low resolution veil image from our foveal sample values. We will then interpolate this veil image to add glare effects to the original rendering. Finally, we will apply this veil as a correction to the adaptation luminances used for our contrast, color sensitivity and acuity models.

Moon and Spencer base their formula for adaptation luminance on the effect of individual glare sources measured by Holladay, which they converted to an integral over the entire visual periphery. The resulting glare formula gives the effective adaptation luminance at a particular fixation for an arbitrary visual field:

$$L_a = 0.913L_f + \frac{K}{\pi} \int_{\theta > \theta_f} \frac{L(\theta, \phi)}{\theta^2} \cos(\theta) \sin(\theta) d\theta d\phi \quad (8)$$

where:

- L_a = corrected adaptation luminance (in cd/m²)
- L_f = the average foveal luminance (in cd/m²)
- $L(\theta, \phi)$ = the luminance in the direction (θ, ϕ)
- θ_f = foveal half angle, approx. 0.00873 radians (0.5°)
- K = constant measured by Holladay, 0.0096

The constant 0.913 in this formula is the remainder from integrating the second part assuming one luminance everywhere. In other words, the periphery contributes less than 9% to the average adaptation luminance, due to the small value Holladay determined for K . If there are no bright sources, this influence can be safely neglected. However, bright sources will significantly affect the adaptation luminance, and should be considered in our model of contrast sensitivity.

To compute the veiling luminance corresponding to a given foveal sample (i.e., fixation point), we can convert the integral in Equation (8) to an average over peripheral sample values:

$$L_{vi} = 0.087 \frac{\sum_j L_j \cos(\theta_{i,j})}{\sum_j \theta_{i,j}^2 \cos(\theta_{i,j})} \quad (9)$$

where:

- L_{vi} = veiling luminance for fixation point i
- L_j = foveal luminance for fixation point j
- $\theta_{i,j}$ = angle between sample i and j (in radians)

Since we must compute this sum over all foveal samples j for each fixation point i, the calculation can be very time consuming. We therefore reduce our costs by approximating the weight expression as:

$$\frac{\cos \theta}{\theta^2} \approx \frac{\cos \theta}{2 - 2\cos \theta} \quad (10)$$

Since the angles between our samples are most conveniently available as vector dot products, which is the cosine, the above weight computation is quite fast. However, for large images (in terms of angular size), the L_{vi} calculation is still the most computationally expensive step in our method due to the double iteration over i and j.

To simulate the effect of glare on visibility, we simply add the computed veil map to our original image. Just as it occurs in the eye, the veiling luminance will obscure the visible contrast on the display by adding to both the background and the foreground luminance.² This was the original suggestion made by Holladay, who noted that the effect glare has on luminance threshold visibility is equivalent to what one would get by adding the veiling luminance function to the original image [Holladay26]. This is quite straightforward once we have computed our foveal-sampled veiling image given in Equation (9). At each image pixel, we perform the following calculation:

$$L_{pvk} = 0.913L_{pk} + L_v(k) \quad (11)$$

where:

- L_{pvk} = veiled pixel at image position k
- L_{pk} = original pixel at image position k
- $L_v(k)$ = interpolated veiling luminance at k

The $L_v(k)$ function is a simple bilinear interpolation on the four closest samples in our veil image computed in Equation (9). The final image will be lighter around glare sources, and just slightly darker *on* glare sources (since the veil is effectively being spread away from bright points). Although we have shown this as a luminance calculation, we retain color information so that our veil has the same color cast as the responsible glare source(s).

²The contrast is defined as the ratio of the foreground minus the background over the background, so adding luminance to both foreground and background reduces contrast.

Figure 15 shows our original, fully lit bathroom scene again, this time adding in the computed veiling luminance. Contrast visibility is reduced around the lamps, but the veil falls off rapidly (as $1/r^2$) over other parts of the image. If we were to measure the luminance detection threshold at any given image point, the result should correspond closely to the threshold we would measure at that point in the actual scene.



Figure 15. Our tone reproduction operator for the original bathroom scene with veiling luminance added.

Since glare sources scatter light onto the fovea, they also affect the local adaptation level, and we should consider this in the other parts of our calculation. We therefore apply the computed veiling luminances to our foveal samples as a correction *before* the histogram generation and adjustment described in Section 4. We deferred the introduction of this correction factor to simplify our presentation, since in most cases it only weakly affects the brightness mapping function.

The correction to local adaptation is the same as Equation (11), but without interpolation, since our veil samples correspond one-to-one:

$$L_{ai} = 0.913L_i + L_{vi} \quad (12)$$

where:

$$\begin{aligned} L_{ai} &= \text{adjusted adaptation luminance at fixation point } i \\ L_i &= \text{foveal luminance for fixation point } i \end{aligned}$$

We will also employ these L_{ai} adaptation samples for the models of color sensitivity and visual acuity that follow.

5.2 Color Sensitivity

To simulate the loss of color vision in dark environments, we use the technique presented by Ferwerda et al. [Ferwerda96] and ramp between a scotopic (grey) response function and a photopic (color) response function as we move through the mesopic range. The lower limit of the mesopic range, where cones are just starting to get enough light, is approximately 0.0056 cd/m². Below this value, we use the straight scotopic luminance. The upper limit of the mesopic range, where rods are no longer contributing significantly to vision, is approximately 5.6 cd/m². Above this value, we use the straight photopic luminance plus color. In between these two world luminances (i.e., within the mesopic range), our adjusted pixel is a simple interpolation of the two computed output colors, using a linear ramp based on luminance.

Since we do not have a value available for the scotopic luminance at each pixel, we use the following approximation based on a least squares fit to the colors on the Macbeth ColorChecker Chart™:

$$Y_{scot} = Y \cdot 1.33 \cdot \left(1 + \frac{Y + Z}{X} \right)^{-1.68} \quad (13)$$

where:

$$\begin{aligned} Y_{scot} &= \text{scotopic luminance} \\ X, Y, Z &= \text{photopic color, CIE } 2^\circ \text{ observer (Y is luminance)} \end{aligned}$$

This is a very good approximation to scotopic luminance for most natural colors, and it avoids the need to render another channel. We also have an approximation based on RGB values, but since there is no accepted standard for RGB primaries in computer graphics, this is much less reliable.

Figure 16 shows our dimmed bathroom scene with the human color sensitivity function in place. Notice there is still some veiling, even with the lights reduced to 1/100th their normal level. This is because the relative luminances are still the same, and they scatter in the eye as before. The only difference here is that the eye cannot adapt as well when there is so little light, so everything appears dimmer, including the lamps. The colors are clearly visible near the light sources, but gradually less visible in the darker regions.



Figure 16. Our dimmed bathroom scene with tone mapping using human contrast sensitivity, veiling luminance and mesopic color response.

5.3 Visual Acuity

Besides losing the ability to see contrast and color, the human eye loses its ability to resolve fine detail in dark environments. The relationship between adaptation level and foveal acuity has been measured in subject studies reported by Shaler [Shaler37]. At daylight levels, human visual acuity is very high, about 50 cycles/degree. In the mesopic range, acuity falls off rapidly from 42 cycles/degree at the top down to 4 cycles/degree near the bottom. Near the limits of vision, the visual acuity is only about 2 cycles/degree. Shaler's original data is shown in Figure 17 along with the following functional fit:

$$R(L_a) = 17.25 \arctan(1.4 \log_{10}(L_a) + 0.35) + 25.72 \quad (15)$$

where:

$$\begin{aligned} R(L_a) &= \text{visual acuity in cycles/degree} \\ L_a &= \text{local adaptation luminance (in cd/m}^2\text{)} \end{aligned}$$

Human Visual Acuity Function (Foveal) due to Shaler

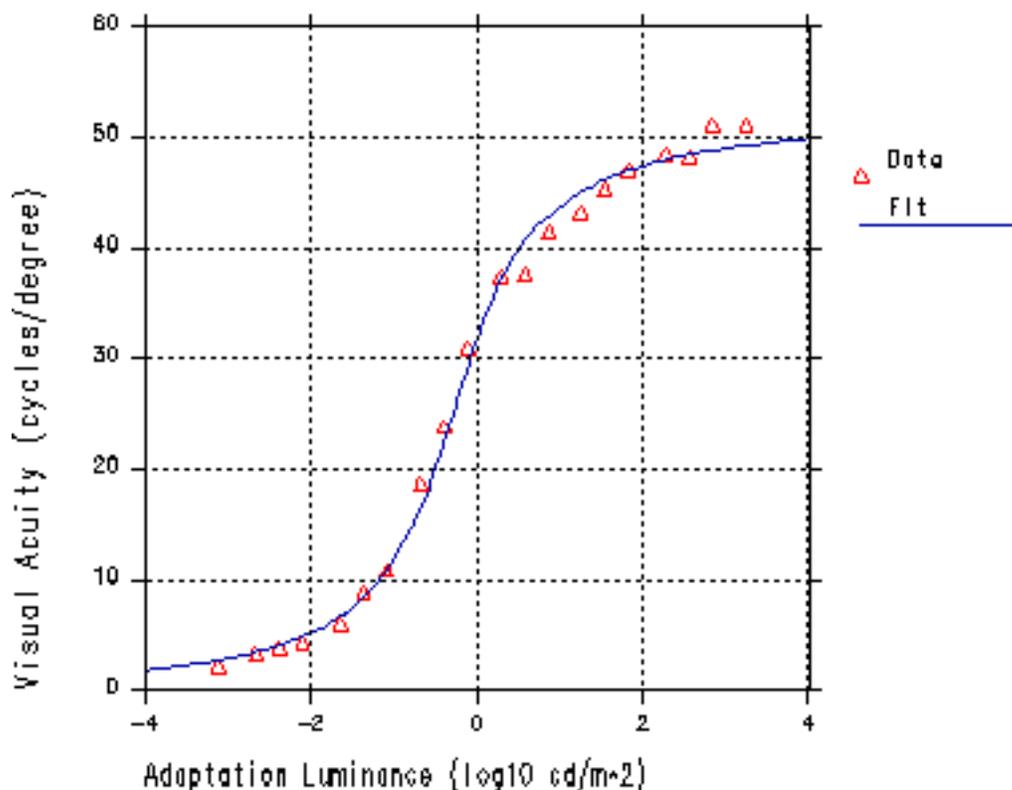


Figure 17. Shaler's visual acuity data and our functional fit to it.

In their tone mapping paper, Ferwerda et al. applied a global blurring function based on a single adaptation level [Ferwerda96]. Since we wish to adjust for acuity changes over a wide dynamic range, we must apply our blurring function locally according to the foveal adaptation computed in Equation (12). To do this, we implement a variable-resolution filter using an image pyramid and interpolation, which is the *mip map* introduced by Williams [Williams83] for texture mapping. The only difference here is that we are working with real values rather than integer pixels.

At each point in the image, we interpolate the local acuity based on the four closest (veiled) foveal samples and Shaler's data. It is very important to use the foveal data (L_{ai}) and not the original pixel value, since it is the fovea's adaptation that determines acuity. The resulting image will show higher resolution in brighter areas, and lower resolution in darker areas.

Figure 18 shows our dim bathroom scene again, this time applying the variable acuity operator applied together with all the rest. Since the resolution of the printed image is low, we enlarged two areas for a closer look. The bright area has an average level around

25 cd/m², corresponding to a visual acuity of about 45 cycles/degree. The dark area has an average level of around 0.05 cd/m², corresponding to a visual acuity of about 9 cycles/degree.

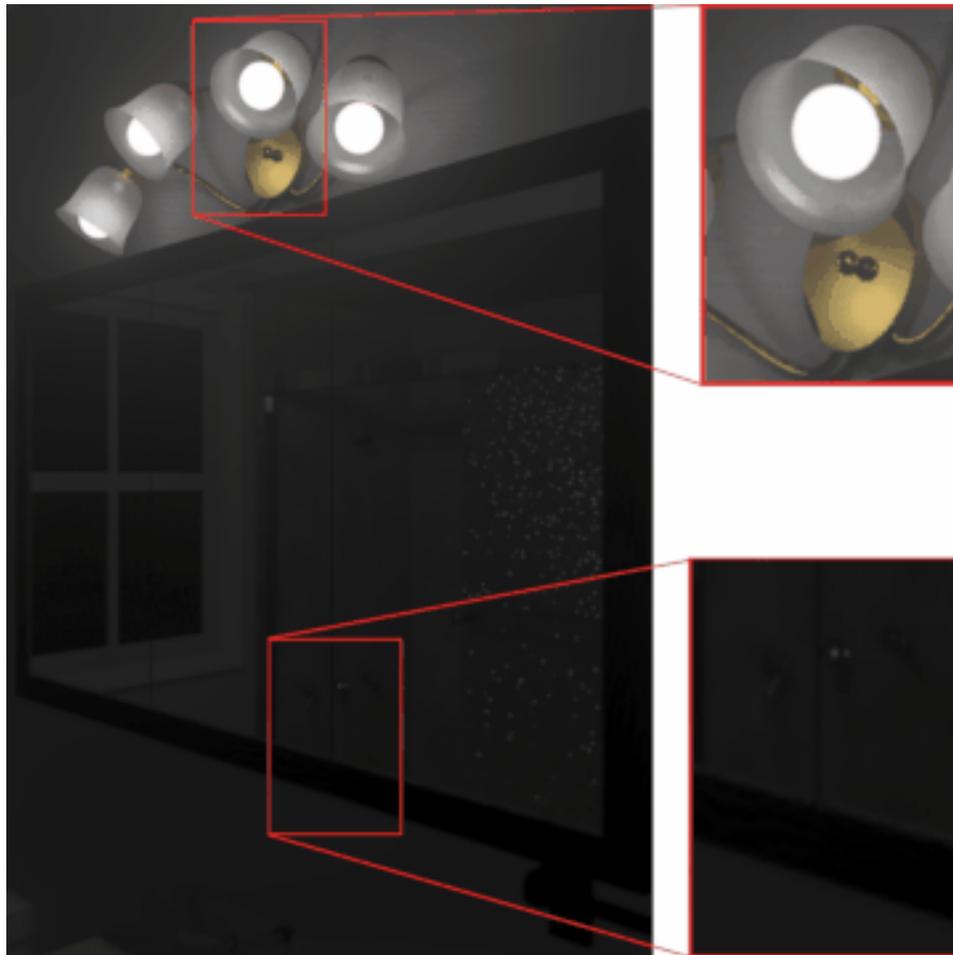


Figure 18. The dim bathroom scene with variable acuity adjustment. The insets show two areas, one light and one dark, and the relative blurring of the two.

6 Method Summary

We have presented a method for matching the visibility of high dynamic range scenes on conventional displays, accounting for human contrast sensitivity, veiling luminance, color sensitivity and visual acuity, all in the context of a local adaptation model. However, in presenting this method in parts, we have not given a clear idea of how the parts are integrated together into a working program.

The order in which the different processes are executed to produce the final image is of critical importance. These are the steps in the order they must be performed:

```
procedure match_visibility()  
  compute 1° foveal sample image  
  compute veil image  
  add veil to foveal adaptation image  
  add veil to image  
  blur image locally based on visual acuity function  
  apply color sensitivity function to image  
  generate histogram of effective adaptation image  
  adjust histogram to contrast sensitivity function  
  apply histogram adjustment to image  
  translate CIE results to display RGB values  
end
```

We have not discussed the final step, mapping the computed display luminances and chrominances to appropriate values for the display device (e.g., monitor RGB settings). This is a well studied problem, and we refer the reader to the literature (e.g., [Hall89]) for details. Bear in mind that the mapped image accounts for the black level of the display, which must be subtracted out before applying the appropriate gamma and color corrections.

Although we state that the above steps must be carried out in this order, a few of the steps may be moved around, or removed entirely for a different effect. Specifically, it makes little difference whether the luminance veil is added before or after the blurring function, since the veil varies slowly over the image. Also, the color sensitivity function may be applied anywhere after the veil is added so long as it is before histogram adjustment.

If the goal is to optimize visibility and appearance without regard to the limitations of human vision, then all the steps between computing the foveal average and generating the histogram may be skipped, and a linear ceiling may be applied during histogram adjustment instead of the human contrast sensitivity function. The result will be an image with all parts visible on the display, regardless of the world luminance level or the presence of glare sources. This may be preferable when the only goal is to produce a nice-looking image, or when the absolute luminance levels in the original scene are unknown.

7 Results

In our dynamic range compression algorithm, we have exploited the fact that humans are insensitive to relative and absolute differences in luminance. For example, we can see that it is brighter outside than inside on a sunny day, but we cannot tell how much brighter (3 times or 100) or what the actual luminances are (10 cd/m² or 1000). With the additional display range made available by adjusting the histogram to close the gaps between luminance levels, visibility (i.e., contrast) within each level can be properly preserved. Furthermore, this is done in a way that is compatible with subjective aspects of vision.

In the development sections, two synthetic scenes have served as examples. In this section, we show results from two different application areas -- lighting simulation and electronic photography.



Figure 19. A simulation of a shipboard control panel under emergency lighting.



Figure 20. A simulation of an air traffic control console.



Figure 21. A Christmas tree with very small light sources.

7.1 Lighting Simulation

In lighting design, it is important to simulate what it is like to *be* in an environment, not what a photograph of the environment looks like. Figures 19 and 20 show examples of real lighting design applications.

In Figure 19, the emergency lighting of a control panel is shown. It is critical that the lighting provide adequate visibility of signage and levers. An image synthesis method that cannot predict human visibility is useless for making lighting or system design judgments.

Figure 20 shows a flight controller's console. Being able to switch back and forth between the console and the outdoor view is an essential part of the controller's job. Again, judgments on the design of the console cannot be made on the basis of ill-exposed or arbitrarily mapped images.

Figure 21 is not a real lighting application, but represents another type of interesting lighting. In this case, the high dynamic range is not represented by large areas of either high or low luminance. Very high, almost point, luminances are scattered in the scene. The new tone mapping works equally well on this type of lighting, preserving visibility

while keeping the impression of the brightness of the point sources. The color sensitivity and variable acuity mapping also correctly represent the sharp color view of areas surrounding the lights, and the greyed blurring of more dimly lit areas.



Figure 22. A scanned photograph of Memorial Church.

7.2 Electronic Photography

Finally, we present an example from electronic photography. In traditional photography, it is impossible to set the exposure so all areas of a scene are visible as they would be to a human observer. New techniques of digital compositing are now capable of creating images with much higher dynamic ranges. Our tone reproduction operator can be applied to appropriately map these images into the range of a display device.

Figure 22 shows the interior of a church, taken on print film by a 35mm SLR camera with a 15mm fisheye lens. The stained glass windows are not completely visible because the recording film has been saturated, even though the rafters on the right are too dark to see. Figure 23 shows our tone reproduction operator applied to a high dynamic range version of this image, called a *radiance map*. The radiance map was generated from 16 separate exposures, each separated by one stop. These images were scanned, registered, and the full dynamic range was recovered using an algorithm developed by Debevec and Malik

[Debevec97]. Our tone mapping operator makes it possible to retain the image features shown in Figure 23, whose world luminances span over 6 orders of magnitude.

The field of electronic photography is still in its infancy. Manufacturers are rapidly improving the dynamic range of sensors and other electronics that are available at a reasonable cost. Visibility preserving tone reproduction operators will be essential in accurately displaying the output of such sensors in print and on common video devices.



Figure 23. Histogram adjusted radiance map of Memorial Church.

8 Conclusions and Future Work

There are still several degrees of freedom possible in this tone mapping operator. For example, the method of computing the foveal samples corresponding to viewer fixation points could be altered. This would depend on factors such as whether an interactive system or a preplanned animation is being designed. Even in a still image, a theory of probable gaze could be applied to improve the initial adaptation histogram. Additional modifications could easily be made to the threshold sensitivity, veil and acuity models to simulate the effects of aging and certain types of visual impairment.

This method could also be extended to other application areas. The tone mapping could be incorporated into global illumination calculations to make them more efficient by

relating error to visibility. The mapping could also become part of a metric to compare images and validate simulations, since the results correspond roughly to human perception [Rushmeier95].

Some of the approximations in our operator merit further study, such as color sensitivity changes in the mesopic range. A simple choice was made to interpolate linearly between scotopic and photopic response functions, which follows Ferwerda et al. [Ferwerda96] but should be examined more closely. The effect of the luminous surround on adaptation should also be considered, especially for projection systems in darkened rooms. Finally, the current method pays little attention to absolute color perception, which is strongly affected by global adaptation and source color (i.e., white balance).

The examples and results we have shown match well with the subjective impression of viewing the actual environments being simulated or recorded. On this informal level, our tone mapping operator has been validated experimentally. To improve upon this, more rigorous validations are needed. While validations of image synthesis techniques have been performed before (e.g., Meyer et al. [Meyer86]), they have not dealt with the level of detail required for validating an accurate tone operator. Validation experiments will require building a stable, non-trivial, high dynamic range environment and introducing observers to the environment in a controlled way. Reliable, calibrated methods are needed to capture the actual radiances in the scene and reproduce them on a display following the tone mapping process. Finally, a series of unbiased questions must be formulated to evaluate the subjective correspondence between observation of the physical scene and observation of images of the scene in various media. While such experiments will be a significant undertaking, the level of sophistication in image synthesis and electronic photography requires such detailed validation work.

The dynamic range of an interactive display system is limited by the technology required to control continual, intense, focused energy over millisecond time frames, and by the uncontrollable elements in the ambient viewing environment. The technological, economic and practical barriers to display improvement are formidable. Meanwhile, luminance simulation and acquisition systems continue to improve, providing images with higher dynamic range and greater content, and we need to communicate this content on conventional displays and hard copy. This is what tone mapping is all about.

9 Acknowledgments

The authors wish to thank Robert Clear and Samuel Berman for their helpful discussions and comments. This work was supported by the Laboratory Directed Research and Development Funds of Lawrence Berkeley National Laboratory under the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

10 References

[Chiu93]

K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang and K. Zimmerman
"Spatially nonuniform scaling functions for high contrast images,"
Proceedings of Graphics Interface '93, Toronto, Canada, May 1993, pp. 245-253.

- [CIE81] CIE (1981) *An analytical model for describing the influence of lighting parameters upon visual performance*, vol 1. Technical foundations. CIE 19/2.1, Technical committee 3.1
- [Debevec97] Debevec, Paul and Jitendra Malik, "Recovering High Dynamic Range Radiance Maps from Photographs," *Proceedings of ACM SIGGRAPH '97*.
- [Ferwerda96] J. Ferwerda, S. Pattanaik, P. Shirley and D.P. Greenberg. "A Model of Visual Adaptation for Realistic Image Synthesis," *Proceedings of ACM SIGGRAPH '96*, p. 249-258.
- [Frei77] W. Frei, "Image Enhancement by Histogram Hyperbolization," *Computer Graphics and Image Processing*, Vol 6, 1977 286-294.
- [Glassner95] A. Glassner, *Principles of Digital Image Synthesis*, Morgan Kaufman, San Francisco, 1995.
- [Green83] W. Green *Digital Image Processing: A Systems Approach*, Van Nostrand Reinhold Company, NY, 1983.
- [Hall89] R. Hall, *Illumination and Color in Computer Generated Imagery*, Springer-Verlag, New York, 1989.
- [Holladay26] Holladay, L.L., *Journal of the Optical Society of America*, 12, 271 (1926).
- [Meyer86] G. Meyer, H. Rushmeier, M. Cohen, D. Greenberg and K. Torrance. "An Experimental Evaluation of Computer Graphics Imagery," *ACM Transactions on Graphics*, January 1986, Vol. 5, No. 1, pp. 30-50.
- [Mokrane92] A. Mokrane, "A New Image Contrast Enhancement Technique Based on a Contrast Discrimination Model," *CVGIP: Graphical Models and Image Processing*, 54(2) March 1992, pp. 171-180.
- [Moon&Spencer45] P. Moon and D. Spencer, "The Visual Effect of Non-Uniform Surrounds", *Journal of the Optical Society of America*, vol. 35, No. 3, pp. 233-248 (1945)
- [Nakamae90] E. Nakamae, K. Kaneda, T. Okamoto, and T. Nishita. "A lighting model aiming at drive simulators," *Proceedings of ACM SIGGRAPH 90*, 24(3):395-404, June, 1990.
- [Rushmeier95] H. Rushmeier, G. Ward, C. Piatko, P. Sanders, B. Rust, "Comparing Real and Synthetic Images: Some Ideas about Metrics," *Sixth Eurographics Workshop on Rendering*, proceedings published by Springer-Verlag. Dublin, Ireland, June 1995.

- [Schlick95]
C. Schlick, "Quantization Techniques for Visualization of High Dynamic Range Pictures," *Photorealistic Rendering Techniques* (G. Sakas, P. Shirley and S. Mueller, Eds.), Springer, Berlin, 1995, pp.7-20.
- [Spencer95]
G. Spencer, P. Shirley, K. Zimmerman, and D. Greenberg, "Physically-based glare effects for computer generated images," *Proceedings ACM SIGGRAPH '95*, pp. 325-334.
- [Stevens60]
S. S. Stevens and J.C. Stevens, "Brightness Function: Parametric Effects of adaptation and contrast," *Journal of the Optical Society of America*, 53, 1139. 1960.
- [Tumblin93]
J. Tumblin and H. Rushmeier. "Tone Reproduction for Realistic Images," *IEEE Computer Graphics and Applications*, November 1993, 13(6), 42-48.
- [Ward91]
G. Ward, "A contrast-based scalefactor for luminance display," In P.S. Heckbert (Ed.) *Graphics Gems IV*, Boston, Academic Press Professional.
- [Williams83]
L. Williams, "Pyramidal Parametrics," *Computer Graphics*, v.17,n.3, July 1983.

High Dynamic Range Imaging

*Greg Ward
Exponent – Failure Analysis Assoc.
Menlo Park, California*

Abstract

The ultimate in color reproduction is a display that can produce arbitrary spectral content over a 300-800 nm range with 1 arc-minute resolution in a full spherical hologram. Although such displays will not be available until next year, we already have the means to calculate this information using physically-based rendering. We would therefore like to know: how may we represent the results of our calculation in a device-independent way, and how do we map this information onto the displays we currently own? In this paper, we give an example of how to calculate full spectral radiance at a point and convert it to a reasonably correct display color. We contrast this with the way computer graphics is usually done, and show where reproduction errors creep in. We then go on to explain reasonable short-cuts that save time and storage space without sacrificing accuracy, such as illuminant discounting and human gamut color encodings. Finally, we demonstrate a simple and efficient tone-mapping technique that matches display visibility to the original scene.

Introduction

Most computer graphics software works in a 24-bit RGB space, with 8-bits allotted to each of the three primaries in a power-law encoding. The advantage of this representation is that no tone-mapping is required to obtain a reasonable reproduction on most commercial CRT display monitors, especially if both the monitor and the software adhere to the sRGB standard, i.e., CCIR-709 primaries and a 2.2 gamma [1]. The disadvantage of this practice is that colors outside the sRGB gamut cannot be represented, particularly values that are either too dark or too bright, since the useful dynamic range is only about 90:1, less than 2 orders of magnitude. By contrast, human observers can readily perceive detail in scenes that span 4-5 orders of magnitude in luminance through local adaptation, and can adapt in minutes to over 9 orders of magnitude. Furthermore, the sRGB gamut only covers about half the perceivable colors, missing large regions of blue-greens and violets, among others. Therefore, although 24-bit RGB does a reasonable job of representing what a CRT monitor can display, it does a poor job representing what a human observer can see.

Display technology is evolving rapidly. Flat-screen LCD displays are starting to replace CRT monitors in many offices, and LED displays are just a few years off. Micromirror projection systems with their superior dynamic

range and color gamut are already widespread, and laser raster projectors are on the horizon. It is an important question whether we will be able to take full advantage and adapt our color models to these new devices, or will we be limited as we are now to remapping sRGB to the new gamuts we have available -- or worse, getting the colors wrong? Unless we introduce new color models to our image *sources* and do it soon, we will never get out of the CRT color cube.

The simplest solution to the gamut problem is to adhere to a floating-point color space. As long as we permit values greater than one and less than zero, any set of color primaries may be linearly transformed into any other set of color primaries without loss. The principal disadvantage of most floating-point representations is that they take up too much space (96-bits/pixel as opposed to 24). Although this may be the best representation for color computations, storing this information to disk or transferring it over the internet is a problem. Fortunately, there are representations based on human perception that are compact and sufficiently accurate to reproduce any visible color in 32-bits/pixel or less, and we will discuss some of these in this paper.

There are two principal methods for generating high dynamic-range source imagery: physically-based rendering (e.g., [2]), and multiple-exposure image capture (e.g., [3]). In this paper, we will focus on the first method, since it is most familiar to the author. It is our hope that in the future, camera manufacturers will build HDR imaging principles and techniques into their cameras, but for now, the easiest path to full gamut imagery seems to be computer graphics rendering.

Computer graphics lifts the usual constraints associated with physical measurements, making floating-point color the most natural medium in which to work. If a renderer is physically-based, it will compute color values that correspond to spectral radiance at each point in the rendered image. These values may later be converted to displayable colors, and the how and wherefore of this tone-mapping operation is the main topic of this paper. Before we get to tone-mapping, however, we must go over some of the details of physically-based rendering, and what qualifies a renderer in this category. Specifically, we will detail the basic lighting calculation, and compare this to common practice in computer graphics rendering. We highlight some common assumptions and approximations, and describe alternatives when these assumptions fail. Finally, we demonstrate color and tone mapping methods for converting the computed spectral radiance value to a displayable color at each pixel.

The Spectral Rendering Equation

$$R_o(\omega_o, \lambda) = \int f_r(\omega_o; \omega_i, \lambda) R_i(\omega_i, \lambda) \cos \theta_i d\omega_i \quad (1)$$

The spectral rendering Eq. (1) expresses outgoing spectral radiance R_o at a point on a surface in the direction ω_o (θ_o, ϕ_o) as a convolution of the bidirectional reflectance distribution function (BRDF) with the incoming spectral radiance over the projected hemisphere. This equation is the basis of many physically-based rendering programs, and it already contains a number of assumptions:

1. Light is reflected at the same wavelength at which it is received; i.e., the surface is not fluorescent.
2. Light is reflected at the same position at which it is received; i.e., there is no subsurface scattering.
3. Surface transmission is zero.
4. There are no polarization effects.
5. There is no diffraction.
6. The surface does not spontaneously emit light.

In general, these assumptions are often wrong. Starting with the first assumption, many modern materials such as fabrics, paints, and even detergents, contain “whitening agents” which are essentially phosphors added to absorb ultraviolet rays and re-emit them at visible wavelengths. The second assumption is violated by many natural and man-made surfaces, such as marble, skin, and vinyl. The third assumption works for opaque surfaces, but fails for transparent and thin, translucent objects. The fourth assumption fails for any surface with a specular (shiny) component, and becomes particularly troublesome when skylight (which is strongly polarized) or multiple reflections are involved. The fifth assumption fails when surface features are on the order of the wavelength of visible light, and the sixth assumption is violated for light sources.

Each of these assumptions may be addressed and remedied as necessary. Since a more general rendering equation would require a long and tedious explanation, we merely describe what to add to account for the effects listed. To handle fluorescence, the outgoing radiance at wavelength ω_o may be computed from an integral of incoming radiance over all wavelengths ω_i , which may be discretized in a matrix form [4]. To handle subsurface scattering, we can integrate over the surface as well as incoming directions, or use an approximation [5]. To handle transmission, we simply integrate over the sphere instead of the hemisphere, and take the absolute value of the cosine for the projected area [2]. To account for polarization, we add two terms for the transverse and parallel polarizations in each specular direction [4] [6]. To handle diffraction, we fold interactions between wavelength, polarization, amplitude and direction into the BRDF and the aforementioned extensions [7]. Light sources are the simplest exception to handle – we simply add in the appropriate amount of spontaneous radiance output as a function of direction and wavelength.

Participating Media

Implicitly missing from Eq. (1) is the interaction of light with the atmosphere, or participating media. If the space between surfaces contains significant amounts of dust, smoke, or condensation, a photon leaving one surface may be scattered or absorbed along the way. An additional equation is therefore needed to describe this volumetric effect, since the rendering equation only addresses interactions at surfaces.

$$\frac{dR(s)}{ds} = -\sigma_a R(s) - \sigma_s R(s) + \int \frac{\sigma_s}{4\pi} R_i(\theta_i) P(\theta_i) d\omega \quad (2)$$

Eq. (2) gives the differential change in radiance as a function of distance along a path. The coefficients σ_a and σ_s give the absorption and scattering densities respectively at position s , which correspond to the probabilities that light will be absorbed or scattered per unit of distance traveled. The scattering phase function, $P(\theta_i)$, gives the relative probability that a ray will be scattered in from direction θ_i at this position. All of these functions and coefficients are also a function of wavelength.

The above differential-integral equation is usually solved numerically by stepping through each position along the path, starting with the radiance leaving a surface given by Eq. (1). Recursive iteration from a sphere of scattered directions can quickly overwhelm such a calculation, especially if it is extended to multiple scattering events. Without going into details, Rushmeier et al. approached the problem of globally participating media using a zonal approach akin to radiosity that divides the scene into a finite set of voxels whose interactions are characterized in a form-factor matrix [8]. More recently, a modified ray-tracing method called the photon map has been applied successfully to this problem by Wann Jensen et al. [9]. In this method, photons are tracked as they scatter and are stored in the environment for later resampling during rendering.

Solving the Rendering Equation

Eq. (1) is a Fredholm integral equation of the second kind, which comes close to the appropriate level of intimidation but fails to explain why it is so difficult to solve in general [10]. Essentially, the equation defines outgoing radiance as an integral of incoming radiance at a surface point, and that incoming radiance is in turn defined by the same integral with different parameters evaluated at another surface point. Thus, the surface geometry and material functions comprise the boundary conditions of an infinitely recursive system of integral equations. In some sense, it is remarkable that researchers have made any progress in this area at all, but in fact, there are many people in computer graphics who believe that rendering is a solved problem.

For over fifteen years, three approaches have dominated research and practice in rendering. The first approach is usually referred to as the *local illumination* approximation, and is the basis for most graphics rendering hardware, and

much of what you see in movies and games. In this approximation, the integral equation is converted into a simple sum over light sources (i.e., concentrated emitters) and a general ambient term. The second approach is called *ray tracing*, and as its name implies, this method traces additional rays to determine specular reflection and transmission, and may be used to account for more general interreflections as well [11] [12]. The third approach is called *radiosity* after the identical method used in radiative transfer, where reflectances are approximated as Lambertian and the surfaces are divided into patches to convert the integral equation into a large linear system that may be solved iteratively [13]. Comparing these three approaches, local illumination is the cheapest and least accurate. Ray tracing has the advantage of coping well with complex geometry and materials, and radiosity does the best job of computing global interactions in simpler, diffuse environments.

In truth, none of the methods currently in use provides a complete and accurate solution to the rendering equation for general environments, though some come closer than others. The first thing to recognize in computer graphics, and computer simulation in general, is that the key to getting a reasonable answer is finding the right approximation. The reason that local illumination is so widely employed when there are better techniques available is not simply that it's cheaper; it provides a reasonable approximation to much of what we see. With a few added tricks, such as shadow maps, reflection maps and ambient lights, local illumination in the hands of an expert does a very credible job. However, this is not to say that the results are correct or accurate. Even in perceptual terms, the colors produced at each pixel are usually quite different from those one would observe in a real environment. In the entertainment industry, this may not be a concern, but if the application is prediction or virtual reenactment, better accuracy is necessary.

For the remainder of this paper, we assume that accuracy is an important goal, particularly color accuracy. We therefore restrict our discussion of rendering and display to physically-based global illumination methods, such as ray-tracing and radiosity.

Tone Mapping

By computing an approximate solution to Eq. (1) for a given planar projection, we obtain a spectral rendering that represents each image point in physical units of radiance per wavelength (e.g., SI units of watts/steradian/meter²/nm). Whether we arrive at this result by ray-tracing, radiosity, or some combination, the next important task is to convert the spectral radiances to pixel color values for display. If we fail to take this step seriously, it almost doesn't matter how much effort we put into the rendering calculation – the displayed image will look wrong.

Converting a spectral image to a display image is usually accomplished in two stages. The first stage is to convert the spectral radiances to a tristimulus space, such as CIE XYZ. This is done by convolving each radiance

spectrum with the three standard CIE observer functions. The second stage is to map each tristimulus value into our target display's color space. This process is called *tone-mapping*, and depending on our goals and requirements, we may take different approaches to arrive at different results. Here are a few possible *rendering intents*:

1. Colorimetric intent: Attempt to reproduce the exact color on the display, ignoring viewer adaptation.¹
2. Saturation intent: Maintain color saturation as far as possible, allowing hue to drift.
3. Perceptual intent: Attempt to match perception of color by remapping to display gamut and viewer adaptation.

The rendering intents listed above have been put forth by the ICC profile committee, and their exact meaning is somewhat open to interpretation, especially for out-of-gamut colors. Even for in-gamut colors, the perceptual intent, which interests us most, may be approached in several different ways. Here are a few possible techniques:

- A. Shrink the source (visible) gamut to fit within the display gamut, scaling uniformly about the neutral line.
- B. Same as A, except apply relative scaling so less saturated colors are affected less than more saturated ones. The extreme form of this is gamut-clipping.
- C. Scale colors on a curve determined by image content, as in a global histogram adjustment.
- D. Scale colors locally based on image spatial content, as in Land's retinex theory.

To any of the above, we may also add a white point transformation and/or contrast adjustment to compensate for a darker or brighter surround. In general, it is impossible to reproduce exactly the desired observer stimulus unless the source image contains no bright or saturated colors or the display has an unusually wide gamut and dynamic range.²

Before we can explore any gamut-mapping techniques, we need to know how to get from a spectral radiance value to a tristimulus color such as XYZ or RGB. The calculation is actually straightforward, but the literature on this topic is vast and confusing, so we give an explicit example to make sure we get it right.

Correct Color Rendering

Looking at the simplest case, spectral reflection of a small light source from a diffuse surface in Eq. (1) reduces to the following formula for outgoing radiance:

$$R_o(\lambda) = \frac{\rho_d(\lambda)}{\pi} E_i(\lambda) \quad (3)$$

¹ The ICC Colorimetric intent is actually divided into *relative* and *absolute* intents, but this distinction is irrelevant to our discussion.

² See www.hitl.washington.edu/research/vrd/ for information on Virtual Retinal Display technology.

where $\rho_d(\lambda)$ is the diffuse reflectance as a function of wavelength, and $E_i(\lambda)$ is the spectral irradiance computed by integrating radiance over the projected source. To convert this to an absolute XYZ color, we apply the standard CIE conversion, given below for SI units [16]:

$$\begin{aligned} X &= 683 \int \bar{x}(\lambda) R(\lambda) d\lambda \\ Y &= 683 \int \bar{y}(\lambda) R(\lambda) d\lambda \\ Z &= 683 \int \bar{z}(\lambda) R(\lambda) d\lambda \end{aligned} \quad (4)$$

At this point, we may wish to convert to an opponent color space for gamut-mapping, or we may wait until we are in the device color space. If our tone-mapping is a simple scale factor as described in technique A above, we may apply it in any linear color space and the results will be the same. If we convert first to a nonlinear device color space, we need to be aware of the meaning of out-of-gamut colors in that space before we map them back into the legal range of display values. We demonstrate a consistent and reasonable method, then compare to what is usually done in computer graphics.

BlueFlower Example

To compute the absolute CIE color for a surface point, we need to know the spectra of the source and the material. Fig. 1 shows the source spectra for standard illuminant A (2856K tungsten), illuminant B (simulated sunlight), and illuminant D65 (6500K daylight). Fig. 2 shows the reflected spectral radiance of the BlueFlower patch from the MacBeth chart under each of these illuminants. To these curves, we apply the CIE standard observer functions using Eq. (4).

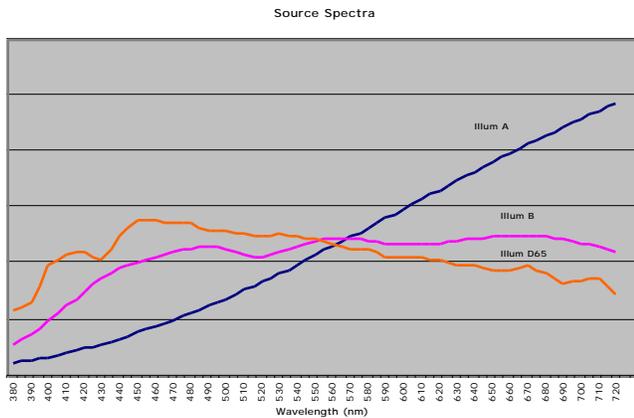


Figure 1. Spectral power of three standard illuminants.

The resulting XYZ values for the three source conditions is given in the first row of Table 1. Not surprisingly, there is a large deviation in color under different illuminants, especially tungsten. We can convert these colors to their RGB equivalents using Eq. (5), as given in the second row of Table 1. If we were to directly display the colors from the

illuminant A and B conditions on the screen, they would likely appear incorrect because the viewer would be adapted to the white point of the monitor rather than the white point of the original scenes being rendered. If we assume the scene white point is the same color as the illuminant and the display white point is D65, then a white point adjustment is necessary for the other illuminants (A and B), as given in the third row of Table 1.

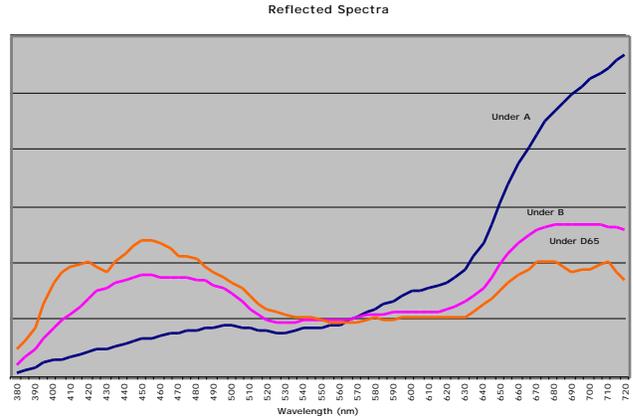


Figure 2. Spectral radiance of MacBeth BlueFlower patch under three standard illuminants.

Source	Illum D65 (.3127,.3290)	Illum B (.3484,.3516)	Illum A (.4475,.4075)
BlueFlower CIE (x,y)	0.274	0.280	0.302
CIE XYZ	0.248 0.456	0.248 0.356	0.248 0.145
709 RGB (absolute)	0.279 0.219 0.447	0.349 0.209 0.341	0.525 0.179 0.119
709 RGB (adjusted)	0.279 0.219 0.447	0.285 0.218 0.444	0.306 0.215 0.426

Table 1. Computed color values for BlueFlower under three standard illuminants.

$$\begin{aligned} R &= X \\ G &= C Y \\ B &= Z \end{aligned} \quad \begin{matrix} \\ \\ \\ \end{matrix} \quad \begin{matrix} \\ \\ \\ \end{matrix} \quad \begin{matrix} \\ \\ \\ \end{matrix} \quad (5)$$

$$\mathbf{C}_{709} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix}$$

We use a linear transform to adjust the white point from that of the illuminant to that of the display, which we assume to be D65 in this example. Eq. (5) gives the absolute transformation from XYZ to CCIR-709 linear RGB, and

this is all we need for the D65 illuminant condition. For the others, we apply the transformation shown in Eq. (6).

Eq. (6) is the linear von Kries adaptation model with the CMCCAT2000 primary matrix [14], which does a reasonable job of accounting for chromatic adaptation when shifting from one dominant illuminant to another [15]. The original white point primaries (R_w, G_w, B_w) are computed from the illuminant XYZ using the \mathbf{M}_{CMCCAT} matrix, and the destination primaries (R_w', G_w', B_w') for D65 are computed using the same transform to be (0.9478, 1.0334, 1.0850).

$$\begin{array}{rcccl}
 X & & R_w/R_w & 0 & 0 & X \\
 Y & = \mathbf{M}^{-1} & 0 & G_w/G_w & 0 & \mathbf{M} Y \\
 Z & & 0 & 0 & B_w/B_w & Z \\
 \\
 R_w & & X_w & & & \\
 G_w & = \mathbf{M} & Y_w & & & \\
 B_w & & Z_w & & & \\
 \\
 \mathbf{M}_{CMCCAT} & = & 0.7982 & 0.3389 & -0.1371 \\
 & & -0.5918 & 1.5512 & 0.0406 \\
 & & 0.0008 & 0.0239 & 0.9753
 \end{array} \quad (6)$$

The combined matrices for a white shift from standard illuminants B and A to D65 (whose chromaticities are given at the top of Table 1) and subsequent conversion from CIE XYZ to CCIR-709 RGB color space, are given in Eq.(7) as \mathbf{C}_B and \mathbf{C}_A . Matrix \mathbf{C}_{709} from Eq. (5) was concatenated with the matrix terms in Eq. (6) to arrive at these results, which may be substituted for \mathbf{C}_{709} in Eq. (5) to get the adjusted RGB colors in the third row of Table 1 from the absolute XYZ values in the first row.

$$\begin{array}{rcccl}
 & 3.1273 & -1.6836 & -0.4867 \\
 \mathbf{C}_B & = & -0.9806 & 1.9476 & 0.0282 \\
 & 0.0605 & -0.2036 & 1.3404 \\
 \\
 & 2.9355 & -2.0416 & -0.5116 \\
 \mathbf{C}_A & = & -1.0247 & 2.1431 & -0.0500 \\
 & 0.0732 & -0.1798 & 3.0895
 \end{array} \quad (7)$$

Conventional CG Calculation

The standard approach in computer graphics color calculations is to assume all light sources are perfectly white and perform calculations in RGB color space. To display the results, a linear scale factor may be applied to bring the results into some reasonable range, and any values outside the sRGB gamut will be clamped.

We obtain an RGB value for the BlueFlower material from its published (x,y) chromaticity of (0.265,0.240) and reflectance of 24.3%. These published values correspond to viewing under standard illuminant C (simulated overcast),

which is slightly bluer than D65. The linear RGB color for the flower material using the matrix \mathbf{C}_{709} from Eq. (5) is (0.246,0.217,0.495), which differs from the D65 results in Table 1 by 10 E* units using the CIE L*uv perceptual metric [16]. Most of this difference is due to the incorrect scene illuminant assumption, since the E* between illuminant C and D65 is also around 10. This demonstrates the inherent sensitivity of color calculations to source color. Using the color corresponding to the correct illuminant is therefore very important.

The reason CG lighters usually treat sources as white is to avoid the whole white balancing issue. As evident from the third row in Table 1, careful accounting of the light source and chromatic adaptation is almost a no-op in the end. For white points close to the viewing condition of D65, the difference is small: a difference of just 1 E* for illuminant B. However, tungsten is very far from daylight, and the E* for illuminant A is more than 5, which is definitely visible. Clearly, if we include the source spectrum, we need to include chromatic adaptation in our tone-mapping. Otherwise, the differences will be very visible indeed -- a E* of 22 for illuminant B and nearly 80 for illuminant A!

What if we include the source color, but use an RGB approximation instead of the full spectral rendering? Errors will creep in from the reduced spectral resolution, and their significance will depend on the source and reflectance spectra. Computing everything in CCIR-709 RGB for our BlueFlower example, the E* from the correct result is 1 for illuminant B and nearly 8 for illuminant A. These errors are at least as large as ignoring the source color entirely, so there seems to be little benefit in this approach.

Relative Color Approximation

An improved method that works well for scenes with a single dominant illuminant is to compute the absolute RGB color of each material under the illuminant using a spectral precalculation from Eqs. (3) and (4). The source itself is modeled as pure white (Y,Y,Y) in the scene, and sources with a different color are modeled relative to this illuminant as $(R_s/R_w, G_s/G_w, B_s/B_w)$, where (R_w, G_w, B_w) is the RGB value of the dominant illuminant and (R_s, G_s, B_s) is the color of the other source. In our example, the RGB color of the BlueFlower material under the three standard illuminants are those given in the second row of Table 1.

Prior to display, the von Kries chromatic adaptation in Eq. (5) is applied to the image pixels using the dominant source and display illuminants. The incremental cost of our approximation is therefore a single transform on top of the conventional CG rendering, and the error is zero by construction for direct reflection from a single source type. There may be errors associated with sources having different colors and multiple reflections, but these will be negligible in most scenes. Best of all, no software change is required -- we need only precalculate the correct RGB values for our sources and surfaces, and the rest comes for free.

It is even possible to save the cost of the final von Kries transform by incorporating it into the precalculation,

computing adjusted rather than absolute RGB values for the materials, as in Eq. (7). We would prefer to keep this transform separate to preserve the colorimetric nature of the rendered image, but as a practical matter, it is often necessary to record a white-balanced image, anyway. As long as we record the scene white point in an image format that preserves the full gamut and dynamic range of our tristimulus pixels, we insure our ability to correctly display the rendering in any device's color space, now and in the future.

High Dynamic Range Images

Real scenes and physically-based renderings of real scenes do not generally fit within a conventional display's gamut using any reasonable exposure value (i.e., scale factor). If we compress or remap the colors to fit an sRGB or similar gamut, we lose the ability to later adjust the tone-scale or show off the image on a device with a larger gamut or wider dynamic range. What we need is a truly device-independent image representation, which doesn't take up too much space, and delivers superior image quality whatever the destination. Fortunately, such formats exist.

Since its inception in 1985, the *Radiance* physically-based renderer has employed a 32-bit/pixel RGBE (Red-Green-Blue-Exponent) format to store its high dynamic range output [17]. Predating *Radiance*, Bill Reeves of **Pixar** created a 33-bit log RGB format for the REYES rendering system, and this format has a public version contributed by Dan McCoy in 1996 to Sam Leffler's free TIFF library (www.libtiff.org). While working at **SGI**, the author added to the same TIFF library a LogLuv format that captures 5 orders of magnitude and the full visible gamut in 24 bits using a perceptual color encoding [18]. The 32-bit version of this format holds up to 38 orders of magnitude, and often results in smaller files due to run-length encoding [19]. Both LogLuv formats combine a logarithmic encoding of luminance with a linear encoding of CIE (u',v') chromaticity to cover the full visible gamut as opposed to the gamut of a specific device or medium.

Of the formats mentioned, only SGI's LogLuv TIFF encoding covers the full gamut and dynamic range of perceivable colors. The *Radiance* RGBE format spans a large dynamic range but is restricted to positive RGB values, so there are visible chromaticities it cannot represent. There is an XYZE version of the same format, but the associated quantization errors make it a poor choice. The Pixar 33-bit log format also has a restricted RGB gamut and only covers 3.8 orders of magnitude, which is marginal for human perception. Since the TIFF library is well tested and free, there is really no reason not to use LogLuv, and many rendering packages now output in this format. Even shareware browsers such as *ACDSee* are able to read and display LogLuv TIFF's.

Gamut Mapping

In order to fit a high dynamic range image into the limited color space of a conventional display, we need to apply one

of the gamut compression techniques mentioned at the beginning of this section.

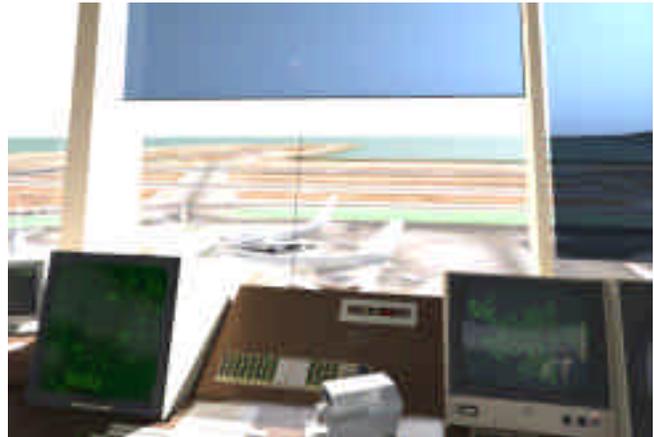


Figure 3. Radiance rendering of control tower clamped to limited display gamut and dynamic range.



Figure 4. The same rendering displayed using a visibility-preserving tone operator including glare effects.



Figure 5. A tone operator designed to optimize print contrast.

Specifically, we show how one might apply the third approach to display an image:

C. Scale colors on a curve determined by image content, as in a global histogram adjustment.

We assume that the rendering system has calculated the correct color at each pixel and stored the result in a high dynamic-range image format. Our task is then to examine this image and choose an appropriate mapping to our display. This is a difficult process to automate, and there is no guarantee we will achieve a satisfactory result in all cases. The best we can do is codify a specific set of goals and requirements and optimize our tone-mapping accordingly.

One possible goal of physically-based rendering is to assess visibility in some hypothetical environment or situation, or to recreate a situation that is no longer readily available (e.g., a plane crash). In such cases, we want to say that anything visible to an observer in the actual scene will be visible on the tone-mapped display. Conversely, if something is not visible on the display, we want to say that it would not be visible to an observer in the actual scene. This kind of visibility-matching operator was described in [20], and we show the result in Fig. 4. Fig. 3 shows the image mapped to an sRGB gamut using technique B to desaturate out-of-gamut colors. As we can see, some of the detail in the planes outside the window was lost to clamping, where it is preserved in the visibility-matching histogram-adjustment procedure in Fig. 4. An optional feature of our tone operator is the ability to simulate disability glare, which reduces visible contrast due to the harsh backlighting in the tower environment. This is visible as a slight haze in front of the monitors in Fig. 4.

Fig. 5 demonstrates another type of tone operator. This is also a histogram adjustment method, but instead of attempting to reproduce visibility, this operator seeks to optimize contrast over the entire image while keeping colors within the printable gamut. Especially in digital photo printers, saturated colors may be difficult to reproduce, so it may be desirable to darken an image to avoid desaturating some regions. We see that this method produces good contrast over most of the image.

Fig. 6 shows the global mapping of these three operators from world (rendered) luminance to display value (fraction of maximum). Where the naive linear operator clamps a lot of information off the top end, the two histogram adjustment operators present this information at a reduced contrast. This compression is necessary in order to bring out detail in the darker regions. We can see that the slopes match the linear operator near black in Fig. 7, deviating from the linear clamping operator above a certain level, where compression begins.

Fig. 8 plots the contrast optimizing tone operator against the world luminance distribution. Peaks in the luminance histogram correspond to increases in contrast, visible in the tone-mapping as a slight increase in slope. Since this is a log-log luminance plot, a small change in slope corresponds to a large change in contrast. The dip between 1.5 and 2.0 corresponds to a more gradual slope in

the tone-mapping and lower contrast. In the low end, we see that this operator tends to provide more contrast to compensate for veiling reflection typical of glossy prints.

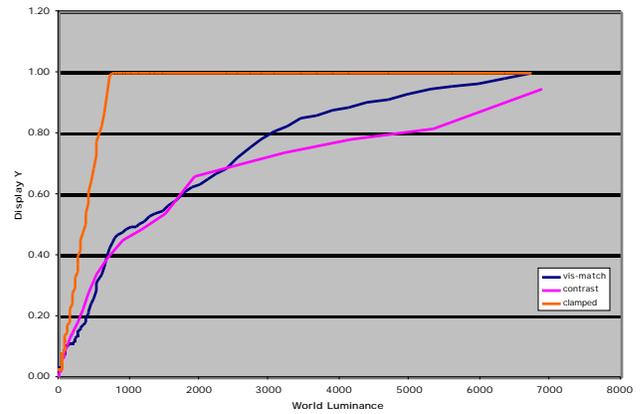


Figure 6. Comparison between three tone-mapping operators.

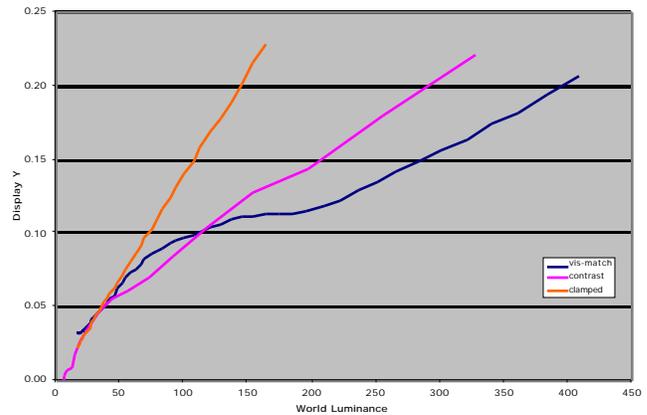


Figure 7. Close-up on darker region of tone-mappings.

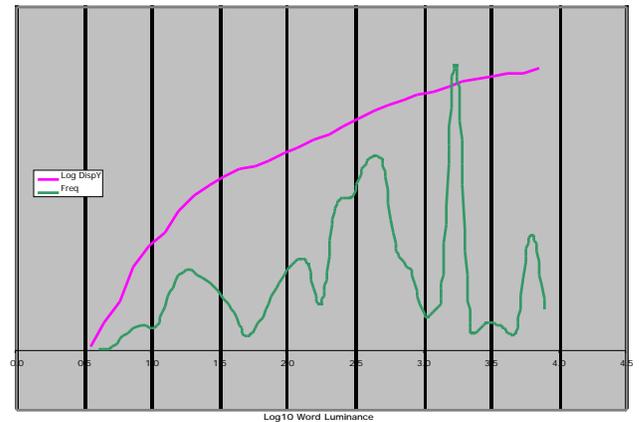


Figure 8. Good global tone operators produce greater contrast at peaks in the input histogram.

Conclusion

The recommendations we make in this paper for accurate color rendering may be summarized as follows:

1. Use a global illumination method with appropriate solutions for all of the phenomena being simulated.
2. Follow accurate spectral calculations with a good chromatic adaptation model to avoid color casts in the displayed image.
3. Substitute full spectral rendering with a relative color approximation for scenes with a single dominant illuminant.
4. Record images in a high dynamic range format to preserve display options (i.e., SGI LogLuv TIFF).
5. Base tone-mapping and gamut-mapping operators on specific goals, such as matching visibility or optimizing color or contrast.

Floating-point spectral calculations and high dynamic-range image manipulation are critical to accurate color rendering. The original approach of rendering directly in 24-bit RGB was recognized as hopeless and abandoned decades ago, but much of the mentality behind it remains with us today.

The methods outlined in this paper are not particularly expensive, neither in terms of implementation effort nor rendering cost. It's simply a matter of applying the right approximation. The author is not aware of any commercial software package that follows more than one or two of these principles, and it seems like a question of priorities.

Most of the money in rendering is spent by the entertainment industry, either in movies or in games. Little emphasis has been placed on accurate color rendering, but with the recent increase in mixed-reality rendering, this is beginning to change. Mixed-reality special effects and games require rendered imagery to blend seamlessly with film or live footage. Since reality follows physics and color science, rendering software will have to do likewise. Those of us whose livelihood depends on predictive rendering and accurate color stand to benefit from this shift.

References

1. Michael Stokes, Matthew Anderson, Srinivasan Chandrasekar, Ricardo Motta, A Standard Default Color Space for the Internet, www.w3.org/Graphics/Color/sRGB
2. Greg Ward, The RADIANCE Lighting Simulation and Rendering System, *Computer Graphics (Proceedings of SIGGRAPH 94)*, ACM, 1994.
3. Paul Debevec, Jitendra Malik, Recovering High Dynamic Range Radiance Maps from Photographs, *Computer Graphics (Proceedings of SIGGRAPH 97)*, ACM, 1997.
4. Alexander Wilkie, Robert Tobler, Werner Purgathofer, Combined Rendering of Polarization and Fluorescence Effects, *Proceedings of 12th Eurographics Workshop on Rendering*, June 2001.
5. Henrik Wann Jensen, Stephen Marschner, Marc Levoy, Pat Hanrahan, A Practical Model for Subsurface Light Transport, *Computer Graphics (Proceedings of SIGGRAPH 01)*, ACM, 2001.
6. Xiaodong He, Ken Torrance, François Sillion, Don Greenberg, A Comprehensive Physical Model for Light Reflection, *Computer Graphics (Proceedings of SIGGRAPH 91)*, ACM, 1991.
7. Jay Gondek, Gary Meyer, Jon Newman, Wavelength Dependent Reflectance Functions, *Computer Graphics (Proceedings of SIGGRAPH 94)*, ACM, 1994.
8. Holly Rushmeier, Ken Torrance, The Zonal Method for Calculating Light Intensities in the Presence of a Participating Medium, *Computer Graphics (Proceedings of SIGGRAPH 87)*, ACM, 1987.
9. Henrik Wann Jensen, Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps, *Computer Graphics (Proceedings of SIGGRAPH 98)*, ACM, 1998.
10. Jim Kajiya, The Rendering Equation, *Computer Graphics (Proceedings of SIGGRAPH 86)*, ACM, 1986.
11. Greg Ward Larson, Rob Shakespeare, *Rendering with Radiance*, Morgan Kaufmann Publishers, 1997.
12. Henrik Wann Jensen, *Realistic Image Synthesis Using Photon Mapping*, A.K. Peters Ltd., 2001.
13. Francois Sillion, Claude Puech, *Radiosity and Global Illumination*, Morgan Kaufmann Publishers, 1994.
14. C. Li, M.R. Luo, B. Rigg, Simplification of the CMCCAT97, *Proc. IS&T/SID 8th Color Imaging Conference*, November 2000.
15. Sabine Süsstrunk, Jack Holm, Graham Finlayson, Chromatic Adaptation Performance of Different RGB Sensors, *IS&T/SPIE Electronic Imaging*, SPIE Vol. 4300, January 2001.
16. Günter Wyszecki, W.S. Stiles, *Color Science*, J. Wiley, 1982.
17. Greg Ward, Real Pixels, *Graphics Gems II*, edited by James Arvo, Academic Press, 1992.
18. Greg Ward Larson, Overcoming Gamut and Dynamic Range Limitations in Digital Images, *IS&T/SID 6th Color Imaging Conference*, November 1998.
19. Greg Ward Larson, The LogLuv Encoding for Full Gamut, High Dynamic Range Images, *Journal of Graphics Tools*, 3(1):15-31 1998.
20. Greg Ward Larson, Holly Rushmeier, Christine Piatko, A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 3, No. 4, December 1997.

Biography

Greg Ward (a.k.a. Greg Ward Larson) graduated in Physics from UC Berkeley in 1983 and earned a Master's in Computer Science from SF State University in 1985. Since 1985, he has worked in the field of light measurement, simulation, and rendering variously at the Berkeley National Lab, EPFL Switzerland, Silicon Graphics Inc., Shutterfly, and Exponent. He is author of the widely used *Radiance* package for lighting simulation and rendering.

Picture Perfect *RGB* Rendering Using Spectral Prefiltering and Sharp Color Primaries

Greg Ward[†]

Elena Eydelberg-Vileshin[‡]

Abstract

Accurate color rendering requires the consideration of many samples over the visible spectrum, and advanced rendering tools developed by the research community offer multispectral sampling towards this goal. However, for practical reasons including efficiency, white balance, and data demands, most commercial rendering packages still employ a naive RGB model in their lighting calculations. This often results in colors that are qualitatively different from the correct ones. In this paper, we demonstrate two independent and complementary techniques for improving RGB rendering accuracy without impacting calculation time: spectral prefiltering and color space selection. Spectral prefiltering is an obvious but overlooked method of preparing input colors for a conventional RGB rendering calculation, which achieves exact results for the direct component, and very accurate results for the interreflected component when compared with full-spectral rendering. In an empirical error analysis of our method, we show how the choice of rendering color space affects final image accuracy, independent of prefiltering. Specifically, we demonstrate the merits of a particular transform that has emerged from the color research community as the best performer in computing white point adaptation under changing illuminants: the Sharp RGB space.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Color, shading, shadowing, and texture

1. Introduction

It is well-known that the human eye perceives color in a three-dimensional space, owing to the presence of three types of color receptors. Early psychophysical research demonstrated conclusively that three component values are sufficient to represent any perceived color, and these values may be quantified using the CIE *XYZ* tristimulus space¹⁹. However, because the spectrum of light is continuous, the interaction between illumination and materials cannot be accurately simulated with only three samples. In fact, no finite number of fixed spectral samples is guaranteed to be sufficient — one can easily find pathological cases, for example, a pure spectral source mixed with a narrow band absorber, that require either component analysis or a ludicrous number of fixed samples to resolve. If the rendered spectrum is

inaccurate, reducing it to a tristimulus value will usually not hide the problem.

Besides the open question of how many spectral samples to use, there are other practical barriers to applying full spectral rendering in commercial software. First, there is the general dearth of spectral reflectance data on which to base a spectral simulation. This is consistent with the lack of any kind of reflectance data for rendering. We are grateful to the researchers who are hard at work making spectral data available^{3, 18}, but the ultimate solution is to put the necessary measurement tools in the hands of people who care about accurate color rendering. Hand-held spectrophotometers exist and may be purchased for the cost of a good laser printer, but few people apply them in a rendering context, and to our knowledge, no rendering package takes spectrophotometer data as direct input.

The second practical barrier to spectral rendering is white balance. This is actually a minor issue once you know how to address it, but the first time you render with the correct

[†] Exponent – Failure Analysis Associates, Menlo Park, California

[‡] Department of Computer Science, Stanford University, Palo Alto, California

source and reflectance spectra, you are likely to be disappointed by the strong color cast in your output. This is due to the change in illuminant from the simulated scene to the viewing condition, and there is a well-known method to correct for this, which we will cover in Section 2.

The third practical barrier to the widespread acceptance of spectral rendering is what we call the “data mixing problem.” What if the user goes to the trouble of acquiring spectral reflectances for a set of surfaces, but they also want to include materials that are characterized in terms of *RGB* color, or light sources that are specified to a different spectral resolution? One may interpolate and extrapolate to some extent, but in the end, it may be necessary to either synthesize a spectrum from *RGB* triples a la Smits’ method¹³, or reduce all the spectral data to *RGB* values and fall back on three component rendering again.

The fourth practical barrier to full spectral rendering is cost. In many renderings, shading calculations dominate the computation, even in *RGB*. If all of these calculations must be carried out at the maximum spectral resolution of the input, the added cost may not be worth the added benefit.

Many researchers in computer graphics and color science have addressed the problem of efficient spectral sampling^{8,7}. Meyer suggested a point-sampling method based on Gaussian quadrature and a preferred color space, which requires only 4 spectral samples and is thus very efficient¹⁰. Like other point sampling techniques, however, Meyer’s method is prone to problems when the source spectrum has significant spikes in it, as in the case of common fluorescent lighting. A more sophisticated approach employing orthonormal basis functions was presented by Peercy, who uses characteristic vector analysis on combinations of light source and reflectance spectra to find an optimal, orthonormal basis set². Peercy’s method has the advantage of handling spiked and smooth spectra with equal efficiency, and he demonstrated accurate results with as few as three orthonormal bases. The additional cost is comparable to spectral sampling, replacing N multiplies in an N -sample spectral model with $M \times M$ multiplies in an M -basis vector model. Examples in his paper showed the method significantly out-performing uniform spectral sampling for the same number of operations. The cost for a 3-basis simulation, the minimum for acceptable accuracy in Peercy’s technique, is roughly three times that of a standard *RGB* shading calculation.

In this paper, we present a method that has the same overall accuracy as Peercy’s technique, but without the computational overhead. In fact, no modification at all is required to a conventional *RGB* rendering engine, which multiplies and sums its three color components separately throughout the calculation. Our method is not subject to point sampling problems in spiked source or absorption spectra, and the use of an *RGB* rendering space all but eliminates the data mixing problem mentioned earlier. White adaptation is also accounted for by our technique, since we ask the user to iden-

tify a dominant source spectrum for their scene. This avoids the dreaded color cast in the final image.

We start with a few simple observations:

1. The direct lighting component is the first order in any rendering calculation, and its accuracy determines the accuracy of what follows.
2. Most scenes contain a single dominant illuminant; there may be many light sources, but they tend to all have the same spectral power distribution, and spectrally differentiated sources make a negligible contribution to illumination.
3. Exceptional scenes, where spectrally distinct sources make roughly equal contributions, cannot be “white balanced,” and will look wrong no matter how accurately the colors are simulated. We can be satisfied if our color accuracy is no worse on average than standard methods in the mixed illuminant case.

The spectral prefiltering method we propose is quite simple. We apply a standard CIE formula to compute the reflected *XYZ* color of each surface under the dominant illuminant, then transform this to a white-balanced *RGB* color space for rendering and display. The dominant sources are then replaced by white sources of equal intensity, and other source colors are modified to account for this adaptation. By construction, the renderer gets the exact answer for the dominant direct component, and a reasonably close approximation for other sources and higher order components.

The accuracy of indirect contributions and spectrally distinct illumination will depend on the sources, materials, and geometry in the scene, as well as the color space chosen for rendering. We show by empirical example how a sharpened *RGB* color space seems to perform particularly well in simulation, and offer some speculation as to why this might be the case.

Section 2 details the equations and steps needed for spectral filtering and white point adjustment. Section 3 shows an example scene with three combinations of two spectrally distinct light sources, and we compare the color accuracy of naive *RGB* rendering to our prefiltering approach, each measured against a full spectral reference solution. We also look at three different color spaces for rendering: CIE *XYZ*, linear *sRGB*, and the Sharp *RGB* space. Finally, we conclude with a summary discussion and suggestions for future work.

2. Method

Spectral prefiltering is just a straightforward transformation from measured source and reflectance spectra to three separate color channels for rendering. These input colors are then used in a conventional rendering process, followed by a final transformation into the display *RGB* space. Chromatic adaptation (i.e., white balancing) may take place either before or after rendering, as a matter of convenience and efficiency.

2.1. Color Transformation

Given a source $I(\lambda)$ and a material $\rho_m(\lambda)$ with arbitrary spectral distributions, the CIE describes a standard method for deriving a tristimulus color value that quantifies what the average human observer sees. The XYZ tristimulus color space is computed from the CIE “standard observer” response functions, \bar{x} , \bar{y} , and \bar{z} , which are integrated with an arbitrary source illuminant spectrum and surface reflectance spectrum as shown in Eq. (1), below:

$$\begin{aligned} X_m &= \int I(\lambda) \rho_m(\lambda) \bar{x}(\lambda) d\lambda \\ Y_m &= \int I(\lambda) \rho_m(\lambda) \bar{y}(\lambda) d\lambda \\ Z_m &= \int I(\lambda) \rho_m(\lambda) \bar{z}(\lambda) d\lambda \end{aligned} \quad (1)$$

For most applications, the 1971 2° standard observer curves are used, and these may be found in Wyszecki and Stiles⁹.

Eq. (1) claims to quantify the exact color an observer would experience if she were to look at a diffuse color patch with the given reflectance spectrum under the given illuminant, but does it really work? In reality, there is a strong tendency for viewers to discount the illuminant in their observations, and the color one sees depends strongly on the ambient lighting in the environment. For example, Eq. (1) might predict a yellow-orange color for a white patch under a tungsten illuminant, while a human observer would still call it “white” if they were in a room lit by the same tungsten source. In fact, a standard photograph of the patch would show its true yellow-orange color, and most novice photographers have the experience of being startled when the colors they get back from their indoor snapshots are not as they remembered them.

To provide for the viewer’s chromatic adaptation and thus avoid a color cast in our image after all our hard work, we apply a von Kries style linear transform to our values prior to display¹⁶. This transform takes an XYZ material color computed under our scene illuminant, and shifts it to the equivalent, apparent color XYZ' under a different illuminant that corresponds to our display viewing condition. All we need are the XYZ colors for white under the two illuminants as computed by Eq. (1) with $\rho_m(\lambda) = 1$, and a 3×3 transformation matrix, M_C , that takes us from XYZ to an appropriate color space for chromatic adaptation. (We will discuss the choice of M_C shortly.) The combined adaptation and display transform is given in Eq. (2), below:

$$\begin{bmatrix} R'_m \\ G'_m \\ B'_m \end{bmatrix} = M_D M_C^{-1} \begin{bmatrix} R'_w & 0 & 0 \\ 0 & G'_w & 0 \\ 0 & 0 & B'_w \end{bmatrix} M_C \begin{bmatrix} X_m \\ Y_m \\ Z_m \end{bmatrix}, \quad (2)$$

where

$$\begin{bmatrix} R_w \\ G_w \\ B_w \end{bmatrix} = M_C \begin{bmatrix} X_w \\ Y_w \\ Z_w \end{bmatrix}$$

for the scene illuminant, and similarly for the display white point, (X'_w, Y'_w, Z'_w) .

The display matrix, M_D , that we added to the standard von Kries transform, takes us from CIE XYZ coordinates to our display color space. For an $sRGB$ image or monitor with D65 white point¹⁴, one would use the following matrix, followed by a gamma correction of 1/2.2:

$$M_{sRGB} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix}$$

If we are rendering a high dynamic-range scene, we may need to apply a tone-mapping operator such as Larson et al⁵ to compress our values into a displayable range. The tone operator of Pattanaik et al even incorporates a partial chromatic adaptation model¹¹.

The choice of which matrix to use for chromatic adaptation, M_C , is an interesting one. Much debate has gone on in the color science community over the past few years as to which space is most appropriate, and several contenders seem to perform equally well in side-by-side experiments². However, it seems clear that RGB primary sets that are “sharper” (more saturated) tend to be more plausible than primaries that are inward of the spectral locus⁴. In this paper, we have selected the *Sharp* adaptation matrix for M_C , which was proposed based on spectral sharpening of color-matching data¹⁶:

$$M_{Sharp} = \begin{bmatrix} 1.2694 & -0.0988 & -0.1706 \\ -0.8364 & 1.8006 & 0.0357 \\ 0.0297 & -0.0315 & 1.0018 \end{bmatrix}$$

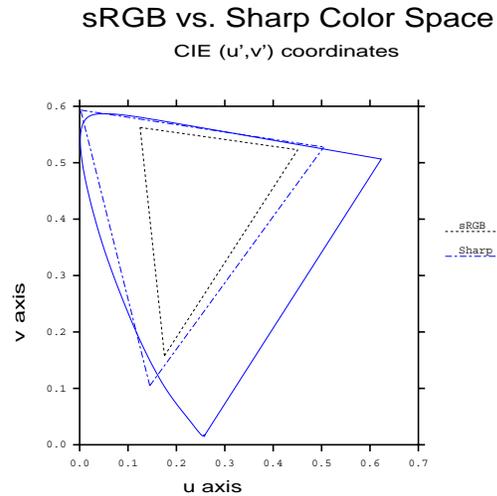


Figure 1: A plot showing the relative gamuts of the $sRGB$ and Sharp color spaces.

Figure 1 shows a CIE (u', v') plot with the locations of the *sRGB* and Sharp color primaries relative to the visible gamut. Clearly, one could not manufacture a color monitor with Sharp primaries, as they lie just outside the spectral locus. However, this poses no problem for a color transform or a rendering calculation, since we can always transform back to a displayable color space.

In fact, the Sharp primaries may be preferred for rendering and *RGB* image representation simply because they include a larger gamut than the standard *sRGB* primaries. This is not an issue if one can represent color values less than zero and greater than one, but most image formats and some rendering frameworks do not permit this. As we will see in Section 3, the choice of color space plays a significant role in the final image accuracy, even when gamut is not an issue.

2.2. Application to Rendering

We begin with the assumption that the direct-diffuse component is most important to color and overall rendering accuracy. Inside the shader of a conventional *RGB* rendering system, the direct-diffuse component is computed by multiplying the light source color by the diffuse material color, where color multiplication happens separately for each of the three *RGB* values. If this calculation is accurate, it must give the same result one would get using Eq. (1) followed by conversion to the rendering color space. In general, this will not be the case, because the diffuse *RGB* for the surface will be based on some other illuminant whose spectrum does not match the one in the model.

For example, the CIE (x, y) chromaticities and *Y*-reflectances published on the back of the Macbeth ColorChecker chart⁹ are measured under standard illuminant *C*, which is a simulated overcast sky. If a user wants to use the color Purple in his *RGB* rendering of an interior space with an incandescent (tungsten) light source, he might convert the published (Y, x, y) reflectances directly to *RGB* values using the inverse of M_{sRGB} given earlier. Unfortunately, he makes at least three errors in doing so. First, he is forgetting to perform a white point transform, so there is a slight red shift as he converts from (Y, x, y) under the bluish illuminant *C* to the more neutral D65 white point of *sRGB*. Second, the tungsten source in his model has a slight orange hue he forgets to account for, and there should be a general darkening of the surface under this illuminant, which he fails to simulate. Finally, the weak output at the blue end of a tungsten spectrum makes purple very difficult to distinguish from blue, and he has failed to simulate this metameric effect in his rendering. In the end, the rendering shows something more like violet than the dark blue one would actually witness for this color in such a scene.

If the spectra of all the light sources are equivalent, we can precompute the correct result for the direct-diffuse component and replace the light sources with neutral (white)

emitters, inserting our spectrally prefiltered *RGB* values as the diffuse reflectances in each material. We need not worry about how many spectral samples we can afford, since we only have to perform the calculation once for each material in a preprocess. If we intend to render in our display color space, we may even perform the white balance transform ahead of time, saving ourselves the final 3×3 matrix transform at each pixel.

In Section 3, we analyze the error associated with three different color spaces using our spectral prefiltering method, and compare it statistically to the error from naive rendering. The first color space we apply is CIE *XYZ* space, as recommended by Borges¹. The second color space we use is linear *sRGB*, which has the CCIR-709 *RGB* color primaries that correspond to nominal CRT display phosphors⁴. The third color space is the same one we apply in our white point transformation, the Sharp *RGB* space. We look at cases of direct lighting under a single illuminant, where we expect our technique to perform well, and mixed illuminants with indirect diffuse and specular reflections, where we expect prefiltering to work less effectively.

When we render in CIE *XYZ* space, it makes the most sense to go directly from the prefiltered result of Eq. (1) to *XYZ* colors divided by white under the same illuminant:

$$X'_m = \frac{X_m}{X_w} \quad Y'_m = \frac{Y_m}{Y_w} \quad Z'_m = \frac{Z_m}{Z_w}$$

We may then render with light sources using their absolute *XYZ* emissions, and the resulting *XYZ* direct diffuse component will be correct in absolute terms, since they will be remultiplied by the source colors. The final white point adjustment may then be combined with the display color transform exactly as shown in Eq. (2).

When we render in *sRGB* space, it is more convenient to perform white balancing ahead of time, applying both Eq. (1) and Eq. (2) prior to rendering. All light sources that match the spectrum of the dominant illuminant will be modeled as neutral, and spectrally distinct light sources will be modeled as having their *sRGB* color divided by that of the dominant illuminant.

When we render in the Sharp *RGB* space, we can eliminate the transformation into another color space by applying just the right half of Eq. (2) to the surface colors calculated by Eq. (1):

$$\begin{bmatrix} R'_m \\ G'_m \\ B'_m \end{bmatrix} = \begin{bmatrix} \frac{1}{R_w} & 0 & 0 \\ 0 & \frac{1}{G_w} & 0 \\ 0 & 0 & \frac{1}{B_w} \end{bmatrix} M_{\text{Sharp}} \begin{bmatrix} X_m \\ Y_m \\ Z_m \end{bmatrix},$$

Dominant illuminants will again be modeled as neutral, and spectrally distinct illuminants will use:

$$R'_s = \frac{R_s}{R_w} \quad G'_s = \frac{G_s}{G_w} \quad B'_s = \frac{B_s}{B_w}$$

The final transformation to the display space will apply the

remaining part of Eq. (2):

$$\begin{bmatrix} R_d \\ G_d \\ B_d \end{bmatrix} = M_D M_{\text{Sharp}}^{-1} \begin{bmatrix} R'_w & 0 & 0 \\ 0 & G'_w & 0 \\ 0 & 0 & B'_w \end{bmatrix} \begin{bmatrix} R'_m \\ G'_m \\ B'_m \end{bmatrix}.$$

3. Results

Our test scene was constructed using published spectral data and simple geometry. It consists of a square room with two light sources and two spheres. One sphere is made of a smooth plastic with a 5% specular component, and the other sphere is made of pure, polished gold (24 carat). The diffuse color of the plastic ball is Macbeth Green⁹. The color of elemental gold is computed from its complex index of refraction as a function of wavelength. The ceiling, floor, and far wall are made of the Macbeth Neutral.8 material. The left wall is Macbeth Red, and the right wall is Macbeth Blue. The near wall, seen in the reflection of the spheres, is the Macbeth BlueFlower color. The left light source is a 2856°K tungsten source (i.e., Standard Illuminant A). The right light source is a cool white fluorescent.

All spectral data for our scene were taken from the material tables in Appendix G of Glassner’s *Principles of Digital Image Synthesis*⁵, and these are also available in the Materials and Geometry Format (MGF)¹⁷. For convenience, the model used in this paper have been prepared as MGF files and included with our image comparisons in the supplemental materials.

Figure 2 shows a Monte Carlo path tracing of this environment with fluorescent lighting using 69 evenly spaced spectral samples from 380 to 720 nm, which is the resolution of our input data. Using our spectral prefiltering method with the cool white illuminant, we recomputed the image using only three *sRGB* components, taking care to retrace exactly the same ray paths. The result shown in Figure 3 is nearly indistinguishable from the original, with the possible exception of the reflection of the blue wall in the gold sphere. This can be seen graphically in Figure 5, which plots the CIE 1994 Lab ΔE^* color difference in false color. A ΔE^* value of one is just noticeable if the colors are adjacent, and we have found values above five or so to be visible in side-by-side image comparisons.

Using a naive assumption of an equal-energy illuminant, we recomputed the *sRGB* material colors from their reflectance spectra and rendered the scene again, arriving at Figure 4. The rendering took the same time to finish, about a third as long as the full-spectral rendering, and the results are quite different. Both the red wall and the green sphere have changed lightness and saturation from the reference image, the blue wall is reflected as purple in the gold sphere, and the ΔE^* errors shown in Figure 6 are over 20 in large regions. Clearly, this level of accuracy is unacceptable for critical color evaluations, such as selecting a color to repaint the living room.

Illum	Method	XYZ		<i>sRGB</i>		Sharp	
		50%	98%	50%	98%	50%	98%
tung	naive	8.4	39.2	3.9	16.3	0.8	4.4
	prefilt	1.3	6.6	0.3	2.4	0.2	0.9
fluor	naive	5.3	25.5	5.8	29.3	0.9	4.6
	prefilt	0.8	6.3	0.2	1.3	0.1	0.9
both	naive	5.0	27.1	4.2	14.0	0.6	2.5
	prefilt tung	3.5	14.9	0.6	2.3	0.7	2.2
	prefilt fluor	4.6	46.1	0.6	6.8	0.7	8.1
Average		4.1	23.7	2.2	10.4	0.6	3.4

Table 1: ΔE^* percentiles for our example scene.

We repeated the same comparisons in CIE *XYZ* and Sharp *RGB* color spaces, then changed the lighting configuration and ran them again. Besides the fluorescent-only lighting condition, we looked at tungsten-only and both sources together. Since the lumen output of the two sources is equal, it was not clear which one to choose as the dominant illuminant, so we applied our prefiltering technique first to one source then to the other. Altogether, we compared 21 combinations of light sources, color spaces, and rendering methods to our multispectral reference solution. The false color images showing the ΔE^* for each comparison are included in the supplemental materials, and we summarize the results statistically in Table 1 and Figure 7.

Table 1 gives the 50th percentile (median) and 98th percentile ΔE^* statistics for each combination of method, lighting, and color space. These columns are averaged to show the relative performance of the three rendering color spaces at the bottom. Figure 7 plots the errors in Table 1 as a bar chart. The 50th percentile errors are coupled with the 98th percentile errors in each bar. In every simulation but one, the Sharp *RGB* rendering space keeps 98% of the pixels below a ΔE^* of five relative to the reference solution, a level at which it is difficult to tell the images apart in side-by-side comparisons. The smallest errors are associated with the Sharp color space and spectral prefiltering with a single illuminant, where 98% of the pixel differences are below the detectable threshold. In the mixed illuminant condition, spectral prefiltering using tungsten as the dominant illuminant performs slightly better than a naive assumption, and prefiltering using cool white as the dominant illuminant performs slightly worse. The worst performance by far is seen when we use CIE *XYZ* as the rendering space, which produces noticeable differences above five for over 2% of the pixels in every simulation, and a *median* ΔE^* over five for all the naive simulations.

4. Conclusions

In our experiments, we found spectral prefiltering to minimize color errors in scenes with a single dominant illuminant spectrum, regardless of the rendering color space. The

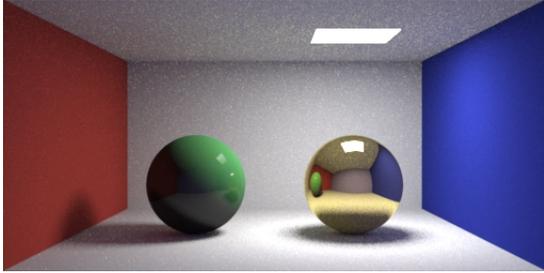


Figure 2: Our reference multi-spectral solution for the fluorescent-only scene.

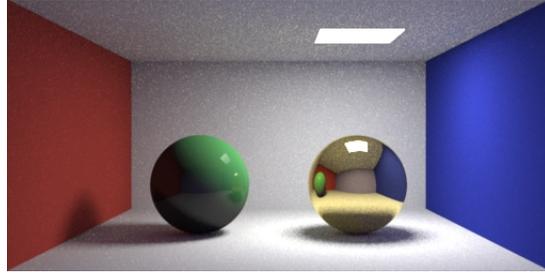


Figure 3: Our prefiltered sRGB solution for the fluorescent-only scene.

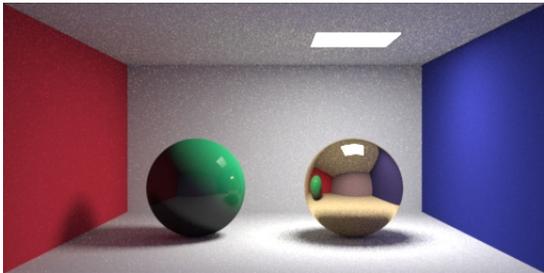


Figure 4: Our naive sRGB solution for the fluorescent-only scene.

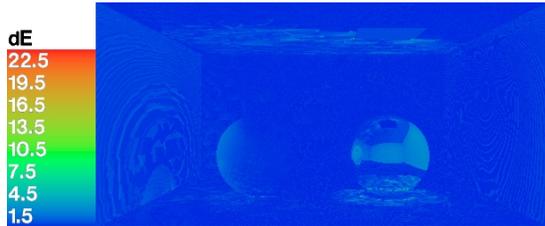


Figure 5: The ΔE^* error for the prefiltered sRGB solution.

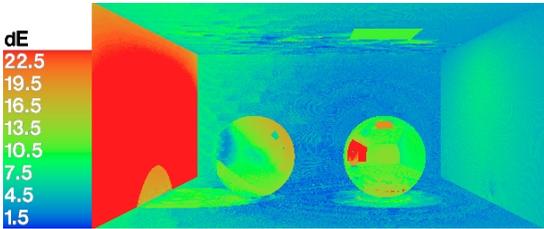


Figure 6: The ΔE^* error for the naive sRGB solution.

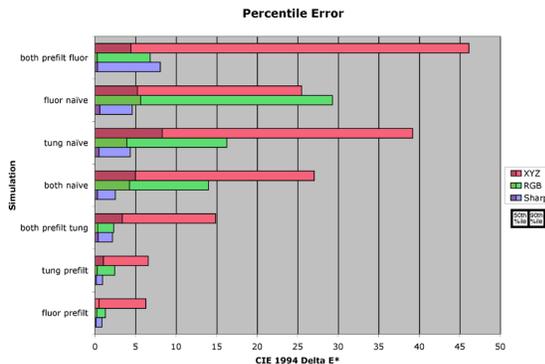


Figure 7: Error statistics for all solutions and color spaces.

median CIE 1994 Lab ΔE^* values were reduced by a factor of eight, to levels that were below the detectable threshold when using the sRGB and Sharp color spaces. Of the three color spaces we used for rendering, the CIE XYZ performed the worst, generating median errors that were just above the detectable threshold even with prefiltering, and five times the threshold without prefiltering, meaning the difference was clearly visible over most of the image in side-by-side comparisons to the reference solution. In contrast, the Sharp RGB color space, favored by the color science community for chromatic adaptation transforms, performed exceptionally well in a rendering context, producing median error levels that were below the detectable threshold both with and without prefiltering.

We believe the Sharp RGB space works especially well for rendering because it minimizes the representation error for tristimulus values by aligning its axes along the densest regions of XYZ space. This property is held in common with the AC_1C_2 color space recommended by Meyer for rendering for this very reason¹⁰. In fact, the AC_1C_2 space has also been favored for chromatic adaptation, indicating the strong connection between rendering calculations and von Kries style transforms. This is apparent when we notice how the white points in the diagonal matrix of Eq. (2), are multiplied in separate channels, analogous to the color calculations inside a three-component shader.

The combination of spectral prefiltering and the Sharp RGB space is particularly effective. With prefiltering under

a single illuminant, 98% of the pixels were below the detectable error threshold using the Sharp *RGB* space, and only certain reflections in the gold sphere were visibly different in a side-by-side comparison. We included a polished gold sphere because we knew its strong spectral selectivity and specularly violated one of our key assumptions, which is that the direct-diffuse component dominates the rendering. We saw in our results that the errors using prefiltering for the gold sphere are no worse than without, and it probably does not matter whether we apply our prefiltering method to specular colors or not. However, rendering in a sharpened *RGB* space always seemed to help.

We also tested the performance of prefiltering when we violated our second assumption of a single, dominant illuminant spectrum. When both sources were present and equally bright, the median error was still below the visible threshold using prefiltering in either the *sRGB* or Sharp color space. Without prefiltering, the median jumped significantly for the *sRGB* space, but was still below threshold for Sharp *RGB* rendering. Thus, prefiltering performed no worse on average than the naive approach for mixed illuminants, which was our goal as stated in the introduction.

In conclusion, we have presented an approach to *RGB* rendering that works within any standard framework, adding virtually nothing to the computation time while reducing color difference errors to below the detectable threshold in typical environments. The spectral prefiltering technique accommodates sharp peaks and valleys in the source and reflectance spectra, and user-selection of a dominant illuminant avoids most white balance problems in the output. Rendering in a sharpened *RGB* space also greatly improves color accuracy independent of prefiltering. Work still needs to be done in the areas of mixed illuminants and colored specular reflections, and we would like to test our method to a greater variety of example scenes.

Acknowledgments

The authors would like to thank Maryann Simmons for providing timely reviews of the paper in progress, and Albert Meltzer for critical editing and LaTeX formatting assistance.

References

1. C. Borges. Trichromatic Approximation for Computer Graphics Illumination Models. *Proc. Siggraph '91*.
2. Anthony J. Calabria and Mark D. Fairchild. Herding CATs: A Comparison of Linear Chromatic-Adaptation Transforms for CIECAM97s. *Proc. 9th Color Imaging Conf.*, pp. 174–178, 2001.
3. Kristin J. Dana, Bram van Ginneken, Shree K. Nayar and Jan J. Koenderink. Reflectance and Texture of Real World Surfaces. *ACM TOG*, **15**(1):1–34, 1999.
4. G. D. Finlayson and P. Morovic. Is the Sharp Adaptation Transform more plausible than CMCCAT2000? *Proc. 9th Color Imaging Conf.*, pp. 310–315, 2001.
5. Andrew S. Glassner. Principles of Digital Image Synthesis. Morgan Kaufmann, 1995.
6. G. W. Larson, H. Rushmeier and C. Piatko. A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes. *IEEE Transactions on Visualization and Computer Graphics*, **3**(4) (December 1997).
7. Laurence T. Maloney. Evaluation of Linear Models of Surface Spectral Reflectance with Small Numbers of Parameters. *J. Optical Society of America A*, **3**(10):1673–1683 (October 1986).
8. David Marimont and Brian Wandell. Linear Models of Surface and Illuminant Spectra. *J. Optical Society of America A*, **9**(11):1905–1913 (November 1992).
9. C. S. McCamy, H. Marcus and J. G. Davidson. A color-rendition chart. *J. Applied Photographic Engineering*, **2**(3):95–99 (summer 1976).
10. Gary Meyer. Wavelength Selection for Synthetic Image Generation. *Computer Vision, Graphics and Image Processing*, **41**:57–79, 1988.
11. Sumanta N. Pattanaik, James A. Ferwerda, Mark D. Fairchild and Donald P. Greenberg. A multiscale model of adaptation and spatial vision for realistic image display. *Proc. Siggraph '98*.
12. Mark S. Peercy. Linear color representations for full speed spectral rendering. *Proc. Siggraph '93*.
13. Brian Smits. An RGB to Spectrum Conversion for Reflectances. *J. Graphics Tools*, **4**(4):11–22, 1999.
14. Michael Stokes et al. A Standard Default Color Space for the Internet – sRGB. Ver. 1.10, November 1996. <http://www.w3.org/Graphics/Color/sRGB>.
15. S. Sueeprasan and R. Luo. Incomplete Chromatic Adaptation under Mixed Illuminations. *Proc. 9th Color Imaging Conf.*, pp. 316–320, 2001.
16. S. Süsstrunk, J. Holm and G. D. Finlayson. Chromatic Adaptation Performance of Different *RGB* Sensors. *IS&T/SPIE Electronic Imaging*, SPIE **4300**, Jan. 2001.
17. Greg Ward et al. Materials and Geometry Format. <http://radsite.lbl.gov/mgf>.
18. Harold B. Westlund and Gary W. Meyer. A BRDF Database Employing the Beard-Maxwell Reflection Model. *Graphics Interface 2002*.
19. Günter Wyszecki and W. S. Stiles. Color Science: Concepts and Methods, Quantitative Data and Formulae. John Wiley & Sons, New York, 2nd ed., 1982.