

# Inverse Rendering: From Concept to Applications

Celine Loscos<sup>1</sup> and Katrien Jacobs<sup>1</sup> and Gustavo Patow<sup>2</sup> and Xavier Pueyo<sup>2</sup>

<sup>1</sup>Virtual Environment and Computer Graphics lab, UCL

<sup>2</sup>Girona Graphics Group, Universitat de Girona

---

## Abstract

*Inverse problems are usually of extreme complexity and are an important research topic for the graphics community due to their wide applicability. Those problems are considered for which the input data is captured from reality with a camera, so it is possible to extract information about the scene illumination, reflectance properties, or geometry. Unfortunately, sometimes standard techniques fail due to practical issues: uncertainty in acquisition methods, dynamic behaviors, complexity of the scene and approximate geometrical or lighting model. This course not only aims at the presentation of the fundamental principles behind inverse rendering problems, but also presents some practical considerations that arise in some applications, explains the new problems introduced and shows adequate solutions. In the introduction, a review is given of the course objectives, the context and some of the prerequisites, e.g. a general definition of global illumination. The first part of the course gives a definition of inverse rendering, and presents a selection of inverse rendering solutions present in the literature. The second part of the course, describes practical methods that can be used in specific applications of inverse rendering. Retrieving reflectance and illumination properties of an existing site is of interest in areas as post-production, architectural and urban planning, and interior lighting design. Modeling geometric properties of an object based on illumination output can be used in the design of reflectors in the car industry or urban lighting design. The course is given as a lecture and is supported by slides and videos. Material created by the course organizers and coming from other sources is used to illustrate the ideas presented in the course. At the end of the course, attendees will have a good understanding of inverse rendering problems, and be able to select a method from the literature to apply to a specific application.*


---

## 1. Outline

1. Introduction
2. Inverse Rendering Definition
  - 2.a. Theoretical Background
  - 2.b. Problem Classification
  - 2.c. Characterization of existing techniques
  - 2.d. General solutions/approaches
    - i. State of the art in Inverse Lighting
    - ii. State of the art in Inverse Reflectometry
    - iii. State of the art in Inverse Surface Design
3. Applications to Inverse Lighting
  - 3.a. General problems for complex scenes
  - 3.b. Postproduction: capturing HDRIs for outdoor scenes
  - 3.c. Urban planning: simulating common illumination
  - 3.d. Indoor Lighting Design
4. Applications to Inverse Geometry
  - 4.a. Numerical Aspects of Inverse Surface Design
  - 4.b. General design of reflectors
  - 4.c. User-guided design of reflectors
5. Conclusion and future work

## 2. Necessary background

The attendees of the course need to have a minimum level of understanding in Mathematics (equations, integrals, and derivatives) and a general knowledge of computer graphics and some understanding of the global illumination meth-



ods. In the introduction of the course, these prerequisites are briefly reviewed.

### 3. Resume of the presenters

#### Celine Loscos

*University College London, Virtual Environment and Computer Graphics lab; contact: c.loscos@cs.ucl.ac.uk*

Celine Loscos joined the Department of Computer Science at UCL in 2000 as a post-doc, after completing her PhD on interactive relighting for augmented reality (1999, IMAG-INRIA, France). Since 2001, she has been a lecturer in department and teaches computer graphics to undergrad-

uate and postgraduate students. She is part of Virtual Environments and Computer Graphics (VECG) laboratory. Her research focuses on real-time rendering, animation and interaction in complex environments for mixed reality on which she has co-authored papers published in IEEE, ACM and EG conferences/journals.

#### Katrien Jacobs

*University College London, Virtual Environment and Computer Graphics lab; contact: k.jacobs@cs.ucl.ac.uk*

Katrien Jacobs graduated as a Electrical Engineer at the Katholieke University of Leuven (ESAT-KUL), Belgium, in 2002 and started her PhD in Computer Science at University

College London (UCL) in 2003. Her work includes developing common illumination and relighting solutions for difficult to capture and model environments. Recent published work includes a survey on illumination methods for mixed reality in Computer Graphics Forum. She is supervised by Dr. Celine Loscos.

### **Gustavo Patow**

*Universitat de Girona, Girona Graphics Group; contact: dagush@ima.udg.es*

Gustavo Patow got a degree in physics from the Universidad de La Plata, Argentina, and got his PhD at the Universitat Politècnica de Catalunya at Barcelona, Spain, under the supervision of Xavier Pueyo and Àlvar Vinacua. His thesis topic was the inverse design of reflector surfaces for luminaire design, and his current research continues in the inverse rendering set of problems. He currently holds an associate professor position at the Universitat de Girona, Spain.

### **Xavier Pueyo**

*Universitat de Girona, Girona Graphics Group; contact: Xavier.pueyo@ima.udg.es*

Xavier Pueyo received a PhD in Engineering from Universitat Politècnica de Catalunya (UPC), Spain, in 1986 and the degree of Docteur Ingenieur in Computer Sciences from Université de Rennes I, France, in 1984. He is a professor of Computer Sciences at the Universitat de Girona (UdG), Spain, since 1996 and he has been the dean of its School of Technology. He taught Computer Sciences (Programming and Computer Graphics) at UdG from 1993 to 1996 and at UPC from 1988 to 1993 as an associate professor. His research interest includes computer graphics, realistic computer rendering and visualization. Pueyo has chaired the Eurographics Rendering Working Group and the Spanish Chapter of Eurographics.

# Bibliographic References

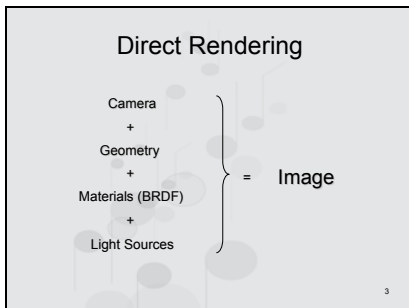
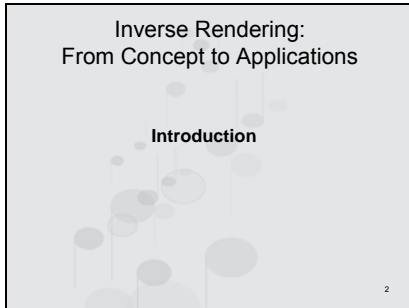
- [Agusanto03] Agusanto, K., Li, L., Chuangui, Z. and Sing, N.W., “Photorealistic rendering for augmented reality using environment illumination”, in proceedings of IEEE/ACM International Symposium on Augmented and Mixed Reality (ISMAR '03), pages 208-216, 2003
- [Boivin01] Boivin, S. and Gagalowicz, A., “Image-based rendering of diffuse, specular and glossy surfaces from a single image”, in proceedings ACM Siggraph '01 (Computer Graphics), ACM Press, pages 107-116, 2001
- [Brickell78] Brickell F. and Westcott B. S., “Phase and Power Distribution on Plane Apertures of Reflector Antennas”, Journal of Physics A: Mathematical and General, vol 11, pages 777-789, 1978
- [CaffarelliOliker99] Caffarelli L. A., Kochengin S. A. and Oliker V. I., “On the Numerical Solution of the Problem of Reflector Design with Given Far-Field Scattering Data”, Contemporary Mathematics, vol 226, 1999
- [Costa99] Cardoso Costa A., Augusto Sousa A. and Nunes Ferreira F., “Lighting Design: A Goal Based Approach Using Optimization”, Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering), Springer-Verlag, New York, NY, pages 317-328, 1999
- [Doyle01] Doyle S., Corcoran D. and Connell J., “A Merit Function for Automated Mirror Design”, Journal of the Illuminating Engineering Society, vol 30 (2), pages 3-11, 2001
- [Doyle99] Doyle S., Corcoran D. and Connell J., “Automated mirror design using an evolution strategy”, Optical Engineering, vol 38 (2), pages 323-333, 1999
- [Debevec97] Debevec P.E. and Malik J., “Recovering high dynamic range radiance maps from photographs”, in proceedings of ACM Siggraph '97 (Computer Graphics), pages 369-378, 1997
- [Debevec03] Debevec P., “Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography”, in proceedings of ACM Siggraph '98 (Computer Graphics), pages 189-198, 1998
- [Drettakis97] Drettakis G., Robert L. and Bougnoux S., “Interactive Common Illumination for Computer Augmented Reality”, Rendering Techniques '97 (Proceedings of the Eighth Eurographics Workshop on Rendering), Springer Wien, New York, NY, pages 45-56, 1997
- [Fournier93] Fournier A., Gunawan S. and Romanzin C., “Common Illumination Between Real and Computer Generated Scenes”, Proceedings of Graphics Interface '93, Morgan Kaufmann, San Francisco, CA, pages 254-262, 1993
- [França98] França F. H. R., Oguma M. and Howell J. R., “Inverse Radiative Heat Transfer with Nongray, Nonisothermal Participating Media”, ASME HTD, R.A. Nelson and T. Chopin and S.T. Thynell eds, vol 361-5, pages 145-151, 1998
- [Gibson01] Gibson, S., Howard, T. and Hubbold, R., “Flexible image-based photometric reconstruction using virtual light sources”, in proceedings of Eurographics 2001, 2001
- [Gibson03] Gibson S., Cook J., Howard T. and Hubbold, R., “Rapid shadow generation in real-world lighting environments”, in proceedings of the 13th Eurographics workshop on Rendering (Rendering Techniques '03), Eurographics Association, pages 219-229, 2003
- [Guillou00] Guillou E., “Simulation d'environnements complexes non lambertiens à partir d'images: Application à la réalité augmentée”, Mathématiques. Informatique, Signal et électronique et Télécommunications. IFSIC/IRISA, 2000 (in french)
- [Haller03] Haller M., Drab S and Hartmann W., “A real-time shadow approach for an augmented reality application using shadow volumes”, in proceedings of the ACM symposium on Virtual reality software and technology (VRST '03), 2003
- [Halstead96] Halstead M. A., “Efficient Techniques for Surface Design Using Constrained

- Optimization”, PhD thesis, University of California, Berkeley, 1996
- [Halstead95] Halstead M. A., Barsky B. A., Klein S. A., and Mandell R. B., “Geometric Modeling of the Cornea Using Videokeratography”, in *Mathematical Methods for Curves and Surfaces*, Daelhen, Morton and Lyche, Tom and Schumaker, Larry L. Eds, Vanderbilt University Press, Nashville, pages 213-223, 1995
- [Halstead96] Halstead M. A., Barsky B. A., Klein S. A. and Mandell R. B., “Reconstructing Curved Surfaces From Specular Reflection Patterns Using Spline Surface Fitting of Normals”, *Computer Graphics, Annual Conference Series*, vol 30, pages 335-342, 1996
- [Harutunian95] Harutunian V., Morales J. C. and Howell J. R., “Radiation Exchange within an Enclosure of Diffuse-Gray Surfaces: The Inverse Problem”, *Inverse Problems in Heat Transfer*, ASME/AICHE National Heat Transfer Conference, Portland, 1995
- [Jacobs04] Jacobs K., Angus C., Loscos C., Nahmias D., Reche A. and Steed A., “Automatic consistent shadow generation for augmented reality”, in *proceedings of Graphics Interface 2005*
- [Jacobs06] Jacobs K. and Loscos C., “A classification of illumination methods for Mixed Reality”, *Computer Graphics Forum*, 2006, vol 26[1]
- [Kochengin03] Kochengin S. A. and Oliker V. I., “Computational algorithms for Constructing Reflectors”, *Computing and Visualization in Science*, vol 6, pages 15-21, 2003
- [Kochengin98] Kochengin S. A. and Oliker V. I., “Determination of Reflector Surfaces from near-field Scattering Data II. Numerical Solution”, *Numer. Math.* vol 79 (4), pages 553-568, 1998
- [Kochengin98] Kochengin S. A., Oliker V. I. and von Tempski O., “On Design of Reflectors with Prespecified Distribution of Virtual Sources and Intensities”, *Inverse Problems*, vol 14 (3), pages 661-678, 1998
- [Loos98] Loos J., Slusallek P. and Seidel H.-P., “Using Wavefront Tracing for the Visualization and Optimization of Progressive Lenses”, *Computer Graphics Forum (Eurographics '98)*, vol 17 (3), pages 255-266, 1998
- [Loscos00] Loscos C., Drettakis G. and Robert L., “Interactive virtual relighting of real scenes”, *IEEE Transactions on Visualization and Computer Graphics*, pages 289-305, 2000
- [Loscos99] Loscos C., Frasson M.-C., Drettakis G., Walter B., Grainer X. and Poulin P., “Interactive Virtual Relighting and Remodeling of Real Scenes”, *Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering)*, Springer Wien, New York, NY, pages 329-340, 1999
- [Marschner98] Marschner S. R., “Inverse Rendering in Computer Graphics”, PhD Thesis, Program of Computer Graphics, Cornell University, Ithaca, NY, 1998.
- [Marschner97] Marschner S. R. and Greenberg D. P., “Inverse Lighting for Photography”, *Proceedings of the IS&T/SID Fifth Color Imaging Conference*, Society for Imaging Science and Technology, Scottsdale, AZ, pages 262-265, 1997
- [Milgram91] Milgram P. and Kishino F., “Taxonomy of mixed reality visual displays”, *IEICE Transactions on Information Systems*, 1994, E77-D, 1321-132
- [Mitsunaga99] Mitsunaga T. and Nayar S.K., “Radiometric self calibration”, in *proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Fort Collins, 1999, 374-380
- [Morales97] Morales J. C., Matsumura M., Oguma M. and Howell J. R., “Computation of Inverse Radiative Heat Transfer Within Enclosures”, *Proc. 1997 ASME National Heat Transfer Conference*, Baltimore, 1997
- [Nakamae86] Nakamae E., Harada K., Ishizaki T. and Nishita T., “A montage method: the overlaying of the computer generated images onto a background photograph”, in *proceedings of ACM Siggraph '86 (Computer Graphics)*, ACM Press, 1986, 207-214
- [Neubauer94] Neubauer A., “The Iterative Solution of a Nonlinear Inverse Problem from Industry: Design of Reflectors”, in *Curves and Surfaces in Geometric Design*, P. J. Laurent ; A. Le Mèhautè and L. L. Schumaker eds, A. K. Peters, Boston, pages 335-342, 1994
- [Neubauer97] Neubauer A., “Design of 3D-Reflectors for Near Field and Far Field Problems”, in *Large Scale Optimization with Applications. Part I: Optimization in Inverse Problems and Design*, IMA Volumes in Mathematics and its Applications, vol 92, Springer, NY, pages 101-118, 1997
- [Oguma95] Oguma M. and Howell J. R., “Solution of the Two-Dimensional Blackbody Inverse

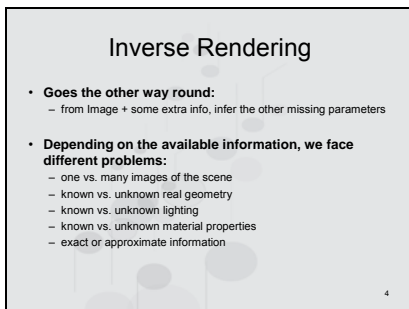
- Radiation Problem by Inverse Monte Carlo Method”, ASME/JSME Thermal Engineering Conference, vol 3, pages 243-250, 1995
- [Oliker02] Oliker V. I., “On the geometry of convex reflectors”, Banach Center Publications, vol 57, pages 155-169, 2002
- [Oliker03] Oliker V. I., “Mathematical aspects of design of beam shaping surfaces in geometrical optics”, Trends in Nonlinear Analysis, S. Kromker and R. Rannacher and F. Tomi eds, pages 193-224, Springer-Verlag, 2003
- [Oliker89] Oliker V. I., “On Reconstructiong a Reflecting Surface from the Scattering data in the Geometric Optics Approximation”, Inverse Problems, vol 5, pages 51-65, 1989
- [Patow06] Patow G., “Image-Based Inverse Rendering”, Research Report IliA 06-09-RR, Institut d'Informàtica i Aplicacions, Universitat de Girona, 2006
- [Patow04] Patow G.; Pueyo X. and Vinacua A., "Reflector Design From Radiance Distributions", International Journal of Shape Modelling, vol 10 (2), pages 211-235, 2004
- [Patow05] Patow G., Pueyo X. and Vinacua A., "User-Guided Inverse Reflector Design", Research Report TR-IliA 04-07-RR, Institut d'Informàtica i Aplicacions, Universitat de Girona, 2004.
- [Poulin92] Poulin P. and Fournier A., “Lights from highlights and shadows”, Computer Graphics, vol 25(2), pages 31-38, 1992
- [Poulin97] Poulin P., Ratib K. and Jacques M., “Sketching Shadows and Highlights to Position Lights”, Proceedings of Computer Graphics International 97, IEEE Computer Society, pages 56-63, 1997.
- [Ramamoorthi02] Ramamoorthi R., "A Signal-Processing Framework for Forward and Inverse Rendering", PhD Thesis, Stanford University,2002.
- [Ramamoorthi01] Ramamoorthi R.and Hanrahan P., "A Signal-Processing Framework for Inverse Rendering", Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001), pages 117-128, August 2001.
- [Sato99a] Sato I., Sato Y. and Ikeuchi K., “Acquiring a Radiance Distribution to Superimpose Virtual Objects onto a Real Scene”, IEEE Transactions on Visualization and Computer Graphics, vol 5 (1), pages 1-12, 99
- [Sato99b] Sato I., Sato Y. and Ikeuchi K., “Illumination distribution from brightness in shadows: adaptive Estimation of illumination distribution with unknown reflectance properties in shadow regions”, Proceedings of IEEE ICCV'99, pages 875-882, 1999
- [Schoeneman93] Schoeneman C., Dorsey J., Smits B., Arvo J. and Greenberg D., “Painting With Light”, Computer Graphics Proceedings, Annual Conference Series, 1993 (ACM SIGGRAPH '93 Proceedings)”, pages, 143-146, 1993
- [Shacked01] Shacked R. and Lischinski D., “Automatic lighting design using a perceptual quality metric”, Computer Graphics Forum (Eurographics 2001), vol 20 (3), pages 215 – 227, 2001
- [Wang96] Wang X.-J., “On the Design of a Reflector Antenna”, Inverse Problems, vol 12, pages 351-375, 1996
- [Ward04] Ward G. , “Fast, robust image registration for compositing high dynamic range photographs from handheld exposures”, Graphics Tools, 2004, 8, 17-30
- [Wescott83] Wescott B. S., “Shaped Reflector Antenna Design”, Letchworth: Research Studies Press, 1983
- [Wescott-Norris75] Westcott B. S. and Norris A. P., “Reflector Synthesis for Generalized far-fields”, Journal of Physics A: Mathematical and General, vol 8 (4), pages 521-532
- [Yu98] Yu Y. and Malik J., “Recovering photometric properties of architectural scenes from photographs”, in proceedings of ACM Siggraph '98 (Computer Graphics), ACM Press, 1998, 207-217
- [Yu99] Yu Y., Debevec P., Malik J. and Hawkins T., “Inverse global illumination: recovering reflectance models of real scenes from photographs”, in proceedings of ACM Siggraph '99 (Computer Graphics), ACM Press/Addison-Wesley Publishing Co., 1999, 215-224



This course is presented by Celine Loscos and Katrien Jacobs from the VECG graphics lab at University College London and Xavier Pueyo and Gustavo Patow from the Graphics Group in Girona. The topic of this course is Inverse Rendering: from concepts to applications.



When the camera position, the scene geometry, materials in the scene and the light source positions and intensities are given, an image can be create of that scene, under the current illumination properties and from the given camera viewpoint using the radiance equation. This is how rendering works.



Inverse rendering goes the other way around, based on an image and some extra information, the missing information in the radiance equation can be calculated. Depending on the amount/type of information available, different strategies apply to calculate the missing information.

### Information availability

- **Geometry information from:**
  - 3D Scanner
  - Photogrammetric techniques
  - Primitives
- **BRDF information from:**
  - Known light probe → pure specular (when building the HDR image)
  - Known materials for some surfaces
- **Lighting information from:**
  - Known, calibrated light sources
  - HDR image from a light probe (computed as a preprocessing step)

5

Various ways exist to calculate/retrieve geometric, BRDF and illumination information.

### Inverse Rendering: Classification

- image + geometry → camera parameters
  - (camera calibration: Well known problem)
- image → geometry
  - (photogrammetry)
- image + geometry + BRDF → lighting
  - (inverse lighting)
- image + geometry + lighting → BRDF/textures
  - (inverse reflectometry)
- image + geometry → BRDF + lighting
  - (inverse combined problem)
- image + lighting + BRDF → geometry
  - (inverse surface design)

6

As mentioned earlier, based on the type/amount of missing information in the radiance equation, a certain strategy applies to retrieve the missing information. On this slide some of these different strategies are listed. For instance when an image and the geometry of a scene are given, the camera properties can be calculated, this type of problem is labelled “Camera calibration”.

### Example Application: Relighting

1. Insert new virtual elements
2. Remove some real elements
3. Change real/virtual object placement
4. Change lighting conditions
5. Change viewpoint
6. ...

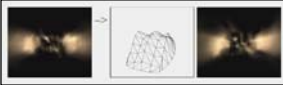
7

Inverse rendering can be used to perform relighting.

Relighting a scene, means that virtually new elements can be added to an existing scene, it can also mean that some existing objects are virtually removed. For each of these cases, relighting will calculate the illumination changes due to the addition or removal of these objects in the scene. An extreme example of relighting is that the scene is virtually ‘relit’ using a completely different illumination distribution than the one present in the real scene.

### Example Application: Reflector Design

- Design of reflector shapes from prescribed optical properties and geometrical constraints.
- Constraints on the shape imposed by industry needs must be taken into account.



8

Another application of Inverse rendering is reflector design.



## Problems with Real Inverse Rendering

Unfortunately, sometimes standard techniques fail:

1. Uncertainty in acquisition methods
2. Dynamic behaviors
3. Approximate geometrical model
  - not completely modeled
  - coarsely modeled
4. Approximate lighting model
  - light sources known only approximately

9

There are some well-known problems with performing inverse rendering. These problems are mainly a result of the failure to capture/model accurately the scene parameters, such as geometry, radiance distribution and light sources.

## Solutions ?

- **Dynamic Scenes:**
  - Automatic HDRI generation of dynamic Environments [Jacobs-Ward-Loscos-05]
- **Approximate geometry and/or approximate light sources**
  - Automatic generation of consistent shadows [Jacobs-et al.-05]
- **Unknown or approximate geometry**
  - Reflector Design from Radiance Distributions [Patow-Pueyo-Vinacua-04]
  - User-guided Inverse Reflector Design [Patow-Pueyo-Vinacua-TR05]
- ...

10

Some solutions are given in the literature to overcome these problems. Some of these solutions are listed on this slide; they will be discussed throughout this course.

## Inverse Rendering: From Concept to Applications

### Outline

11

The following slides detail the outline of this talk.

## Inverse Rendering: From Concept to Applications

### 1. Inverse Rendering Definition

1. General definition of inverse rendering & characterization of existing techniques
2. General solutions/approaches

12

First a definition is given of inverse rendering; this section consists of three subsections.

First a general definition of inverse rendering is given  
 Followed by a characterization of the existing techniques  
 And finally some general solutions and approaches will be described

After this section, the course attendee will have a good understanding of the concept 'inverse rendering' and know which solutions exist and can be applied based on the problem at hand.

**Inverse Rendering:  
From Concept to Applications**

- 1. Inverse Rendering Definition**
- 2. Applications to Inverse Lighting**
  1. General problems for outdoor scenes
  2. Postproduction: capturing HDRI for outdoor scenes
  3. Urban planning: simulating common illumination
  4. Indoor Lighting Design

13

After the break, a first set of applications of inverse rendering are presented, more precisely those related to inverse lighting. This second session consists out of 4 different subsections. Each will be presented and illustrated with demos in the form of videos.

**Inverse Rendering:  
From Concept to Applications**

- 1. Inverse Rendering Definition**
- 2. Applications to Inverse Lighting**
- 3. Applications to Inverse Geometry**
  1. General design of reflectors
  2. User-guided design of reflectors

14

The third section discusses the applications related to inverse geometry. This section consists out of two subsections, and each will be presented in more detail, and illustrated with some demos/videos.

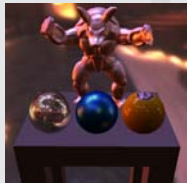
**Inverse Rendering:  
From Concept to Applications**

- 1. Inverse Rendering Definition**
- 2. Applications to Inverse Lighting**
- 3. Applications to Inverse Geometry**
- 4. Conclusion and Future Work**

15



Finally a conclusion and an overview of the future work is given.

**General definition of  
Inverse Rendering Problems**



Gustavo Patow  
Xavier Pueyo

Grup de Gràfics de Girona  
Universitat de Girona

16

In this part of the tutorial we will introduce a general definition and a classification of Inverse Rendering Problems.

### Context: Inverse problems

- In general, they are of an extreme complexity
- Unfortunately, they **are** important
  - e.g.: for lighting engineering, lighting design, animators and lighting for films
- Research in graphics has focused mainly on Direct/Forward Problems
  - Compute the radiance distribution from a known environment

17

Inverse problems are usually of an extreme complexity and are emerging as an important research topic for the graphics community due to their interest in a wide range of fields including lighting engineering and lighting design. Even animators and lighting experts for the film industry would also highly benefit from them. Although progress in rendering to date has mainly focused on improving the accuracy of the physical simulation of light transport and developing algorithms with better performance, some attention has been paid to the problems related to inverse analysis, leading recently to very interesting results. Traditional forward problems in lighting involve the computation of the radiance distribution in an a priori, completely known environment (geometry and materials).


### Inverse Rendering definition

- Inverse Rendering Problems refer to all the problems where
  - Some aspects of the scene are unknown
  - We know in advance the desired illumination at some surfaces of the scene
- The algorithm has to work backwards to establish the missing parameters

18


Inverse Lighting Problems refer to all the problems where, as opposed to what happens with traditional direct lighting problems, some aspects of the scene are unknown. One common characteristic of this kind of problem is that, in general, we know in advance the desired illumination at some surfaces of the scene (their final appearance). Therefore, the algorithm has to work backwards to establish the missing parameters, which would produce the desired illumination.

### Estimating BRDF and Lighting (I)



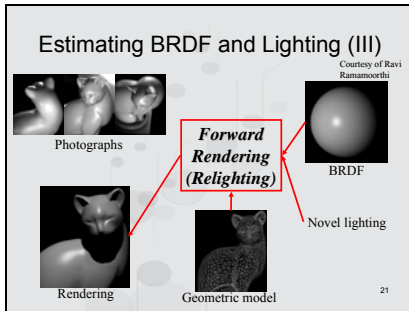
To illustrate this, we can see a simple example. Given a known geometric model like this cat, with completely known material properties, (that is, with known BRDF), under controlled lighting conditions, it can be rendered to a picture like the one shown above (bottom-left side of the slide). Forward or direct rendering aims at making this photograph indistinguishable from the one above, taken with a real camera from a cat figurine.

### Estimating BRDF and Lighting (II)



Instead, if we only know the geometric model and some photographs, we can work the other way round, and try to approximately infer the BRDF and/or the lighting conditions. This process is known as an inverse lighting, inverse reflectometry or a combined problem depending on what we want to find out.

We can say that inverse problems infer parameters of a system from observed or desired data, which define their behavior, in contrast to direct problems which, given all the parameters, simulate the effects.



Now, with the resulting information from the inverse problem, we can generate new images of the cat with different lighting and viewing conditions. Often, this process is called relighting.

- ### From now on
- We have introduced the problem
  - Now, we will
    - Explain the theoretical background
      - Introducing our main classification
      - Presenting the problems behind Inverse Rendering
    - Explain alternative classifications to enrich our view of the problems
- 22

### Theoretical background

- Global Illumination
- Most fundamental magnitude: radiance  $L(\mathbf{r}, \omega)$
- The boundary conditions of the integral form of the transport equation are expressed as

$$L(\mathbf{r}, \omega) = L_e(\mathbf{r}, \omega) + \int_S f_r(\mathbf{r}, \omega_i \rightarrow \omega) L(\mathbf{r}', \omega_i) \cos \theta d\omega_i$$

23

Here we are in the context of the global illumination framework, where the most fundamental magnitude to work with is the radiance  $L$ .

The equation that describes the boundary conditions of the integral form of the transport equation are expressed as that equation in the slide. Here, what we are saying is that the radiance leaving any point in the scene is the sum of the emittance of that point plus all the light that arrives at that point from the other surfaces in the scene, modulated by the material properties of the surface that contains the point, called the Bidirectional Reflectance Distribution Function, or BRDF for short.

This equation is called “the rendering equation”.

Now, we are going to introduce a couple of operators to write it in a clearer and shorter way for our purposes.

### Local reflection operator $\bar{K}$

- Let's define it as [Marschner98]

$$(\bar{K}h)(\mathbf{r}, \omega) \equiv \int_S f_r(\mathbf{r}, \omega_i \rightarrow \omega) h(\mathbf{r}, \omega_i) d\mu(\omega_i)$$

- With  $d\mu(\omega_i) = \cos \theta d\omega_i$
- $\bar{K}$  maps incident light distribution onto the corresponding exiting light distribution resulting from *one* local reflection

24

Firstly, let's introduce the Local Reflection Operator  $K$ , which maps the incident light distribution onto the corresponding exiting light distribution resulting from *one* local reflection. As we can see, this operator takes the integral part of the Rendering Equation in the previous slide. [Marschner98]

### Field radiance operator $\widehat{G}$

- Let's define it as
 
$$(\widehat{G}h)(\mathbf{r}, \omega) \equiv \begin{cases} h(\mathbf{p}(\mathbf{r}, \omega), \omega) & \text{when } v(\mathbf{r}, \omega) < \infty \\ 0 & \text{otherwise} \end{cases}$$
- Where  $v(\mathbf{r}, \omega)$  is the *visible surface function*

$$v(\mathbf{r}, \omega) \equiv \inf\{x > 0: \mathbf{r} + x\omega \in \text{Surfaces in the env.}\}$$
- And the ray casting function
 
$$\mathbf{p}(\mathbf{r}, \omega) = \mathbf{r} + v(\mathbf{r}, \omega)\omega$$

25

Now, let's introduce the Field Radiance Operator  $G$ , which determines what is seen from a point in the scene in a given direction. Basically, it provides information on the nearest point in the scene, given a starting point and a desired direction.

In order to do its job, it needs the definition of the Visible Surface Function, which, by the way, is not a continuous function. This Visible Surface Function, in turn, needs the definition of the ray casting function, which we see at the bottom of the slide

### Final expression

- We can write the rendering equation as:
 
$$L = L_e + \widehat{K}\widehat{G}L$$
- Which leads us to our main classification

26

So, we finally can write the rendering equation as this short expression. With this expression, and following Stephen Marschner's work, we will be able to introduce our main classification.

### Forward rendering

- Solve for  $L$ 

$$L = L_e + \widehat{K}\widehat{G}L$$

27

Direct problems are those which, given known values for  $L_e$ ,  $K$  and  $G$ , solve for  $L$ .

But, if we only have some knowledge of  $L$ , and some of the other parameters, we can pose different kinds of inverse lighting problems.

### Inverse lighting

- Solve for  $L_e$ 

$$L = L_e + \widehat{K}\widehat{G}L$$

28

If  $L_e$  is unknown, and  $K$ ,  $G$  and  $L$  or part of it, are known, we have a problem of *inverse lighting*: given a photograph or any other information that covers part of  $L$ , and a complete model of the scene ( $K$  and  $G$ ), find the emittances ( $L_e$ ) of the luminaries illuminating the scene.

### Inverse reflectometry

- Solve for  $K$  (BRDFs & textures)  $L = L_e + \bar{K}\hat{G}L$

29

If  $K$  is unknown, and  $G$ ,  $L_e$  and part of  $L$  are known, we must solve for information about  $K$ . This problem can, in general, be called *inverse reflectometry*, and a particular case is the one called *image-based reflectometry* in [Marschner98], where images are used as input of the information about  $L$ . As described there, since  $K$  includes information about the variance of the reflectance both spatially and directionally, this can be a very difficult problem since it can be a complex function. Depending on the constraints imposed on the problem, we can subdivide it into the *inverse texture measurement* (constraints on the directional variation), or the *inverse BRDF measurement* (spatial uniformity is assumed).

### Combined problems

- Solve for  $L_e$  and  $K$   $L = L_e + \bar{K}\hat{G}L$

30

If we solve for both  $L_e$  and  $K$ , we are talking about a combined problem, which are more difficult because of the relationships in the simulation process.

### Inverse surfaces

- Solve for  $G$   $L = L_e + \bar{K}\hat{G}L$

31

Finally, if  $G$  is unknown, we have an *inverse geometry* problem, where we want to find a shape from the effect it produces in the scene illumination. Quite often, this surface is called a *reflector*.

### Classification of Inverse Rendering Problems


$L = L_e + \bar{K}\hat{G}L$

	$L$	$L_e$	$K$	$G$
Direct Problems	?	✓	✓	✓
Inverse Lighting	partial	?	✓	✓
Inverse Reflectometry	partial	✓	?	✓
Combined Problems	partial	?	?	✓
Inverse Geometry	✓	✓	✓	?

32

In general, we can build the following table, which shows what we have to know and what we want to find out for each kind of problem.

As we can see, in Inverse Lighting, Inverse Reflectometry and combined problems we can have only partial information about the scene lighting, but Inverse Geometry problems are of such complexity that, we need to know the full illumination information.



Now, let's study the difficulties related to inverse rendering problems. As we can see from the slide, if we start with a highly specular surface and a point light source (bottom left figure), we see that we can perfectly identify the material and the shape of the source. The same happens when we increase the size of the source (bottom row, from left to right). On the other hand, if we start from the first situation, and we change the surface roughness, we get more and more blurred versions of the point light source. The problem is that, in general, we can not distinguish between the situations when the light source is larger with a specular BRDF, and the case when the BRDF is diffuse but the light is a point light. Even worse, as both the BRDF goes diffuse AND the light source gets bigger (softer lighting), we lose the capability of distinguishing which case we are facing. We see that the problem of recovering the information becomes more and more ambiguous, which in mathematical terms means that we have an ill-posed problem.

### Inside the Rendering Equation (I)

- Let's peek at some details inside the equation

$$L(\mathbf{r}, \omega) = L_e(\mathbf{r}, \omega) + \int_{\Omega} f_r(\mathbf{r}, \omega_i \rightarrow \omega) L(\mathbf{r}', \omega_i) \cos \theta d\omega_i$$

- Following [Ramamoorthi-Hanrahan01], we can absorb the  $\cos \theta$  into the BRDF, and if we assume the surfaces do not emit

$$L(\mathbf{r}, \omega) = \int_{\Omega} f_r'(\mathbf{r}, \omega_i \rightarrow \omega) L(\mathbf{r}', \omega_i) d\omega_i$$

34

To have a more precise understanding on the problems we face in Inverse Rendering problems, let's peek into the rendering equation. Following Ramamoorthi and Hanrahan's work, we can include the  $\cos \theta$  term into the BRDF, and if we assume the surfaces do not emit (that is,  $L_e$  is zero for the surfaces we are studying), we get to a slightly simplified expression, the one shown in the bottom of the slide.

### Inside the Rendering Equation (II)

- From now on, let's work on a given vertex
- And assume distant illumination
- Let's write the lighting as

$$L(\omega_i) = \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm} Y_{lm}(\omega_i)$$

with  $Y_{lm}$  a Spherical Harmonic

35


In order to continue our study, we can work on a vertex in isolation, so we don't need to worry about positional problems. Also, we will assume a distant illumination, which will make our formulation independent of the evaluation position on the surface.

We always will be able to express all the quantities as an expansion in terms of basis functions. In this case, we will use spherical harmonics, which form an orthonormal basis in terms of which functions on the sphere can be expanded. The first few functions of this basis is shown in the picture.

We first expand the lighting in global coordinates, arriving at the expression in the slide.

### Inside the Rendering Equation (III)

- Do the same with the **(isotropic)** BRDF:
 
$$f_r'(\omega_i \rightarrow \omega) = \sum_{l,p,q} f_{l,p,q}' Y_l^*(\omega_i) Y_{pq}(\omega)$$
- And with the reflected light field
 
$$B(\mathbf{n}, \omega_o) = \sum_{l,m,p,q} B_{lmpq} \Lambda_l^{-1} R_{lmq}(\mathbf{n}) Y_{pq}(\omega_o)$$
- With  $\Lambda_l = \sqrt{4\pi / (2l + 1)}$
- And  $R_{lmq}(\mathbf{n})$  the rotation coefficients about the surface normal  $\mathbf{n}$ .



36

We now represent the transfer function  $f' = f \cos\theta_i$  in terms of spherical harmonics. We note that  $f'$  is nonzero only over the upper hemisphere, i.e. when  $\cos\theta > 0$  and  $\cos\theta_i > 0$ . We also use a complex conjugate for the first factor, to simplify the final results.

If we assume we are dealing with isotropic BRDFs, which are those where rotating the local tangent frame makes no difference, we get the expression in the slide for  $f'$ . Proceeding similarly with the reflected light field, we arrive at the expression in the slide. We may observe that the  $R_{lmq}(\mathbf{n})$  are the rotation coefficients for the spherical harmonics to accommodate the reference frame to local coordinates given by the local surface normal  $\mathbf{n}$ .

### Inside the Rendering Equation (IV)

- Putting it all together into the equation, we get
 
$$B_{lmpq} = \Lambda_l L_{lm} f'_{lpq}$$
- As we see, we have a "simple" system of algebraic equations
- (which, of course, is not "simple" at all)

37

Putting all together, comparing both expressions and equating coefficients, we arrive at this "simple" expression in the frequency-space domain, as the spherical harmonics we are using are the equivalent, for the sphere, to a Fourier analysis in 2D space. In frequency-space, the reflected light field is obtained simply by multiplying together coefficients of the lighting and BRDF, i.e. by *convolving* the incident illumination with the BRDF.

Looking at this, we could think that this system of equations is easy to solve, but, as we hinted a few slides ago, it is not simple at all!

### Inside the Rendering Equation (V)

- Inverse BRDF:
 
$$f'_{lpq} = \Lambda_l^{-1} B_{lmpq} / L_{lm}$$
  - We can see that BRDF recovery is *well-posed* unless the denominator vanishes
  - But, if  $L_{lm}$  vanishes,  $B_{lmpq}$  has to vanish too to be physically accurate, and the RHS will be indeterminate
  - It's ok, as long as, for all  $l$ , there is at least one value of  $m$  so that  $L_{lm} \neq 0$
- Thus, BRDF recovery is well-conditioned when lighting contains high frequencies (e.g. directional sources)
- And is ill-conditioned for soft lighting

38

Following [Ramamoorthi02] – Observing the expression resulting for  $f'_{lpq}$ , we see that the BRDF estimation process will be well-posed (that is, unambiguous) as long as the denominator on the right-hand side does not vanish. But we must take into account that, in order to be physically accurate, the numerator will also be 0 if the denominator vanishes, as there won't be any reflected light field for that order of the expansion. In this case, the right-hand side will become indeterminate. From the equation in the slide, we see that if, for all  $l$ , there exists at least one value of  $m$  so that  $L_{lm} = 0$ , then the problem of BRDF estimation is well posed. In other words, all orders in the spherical harmonic expansion of the lighting should have at least one coefficient with nonzero amplitude. If any order of the expansion completely vanishes, then we will find we won't be able to estimate the corresponding BRDF coefficients. In signal processing terms, if lighting, which is the input signal, has no amplitude along certain modes of the BRDF (our filter), those modes cannot be estimated. We can say that, when the spherical harmonic expansion of the lighting does not decay rapidly with increasing frequency, the BRDF recovery problem is well conditioned. This is telling us that, when the lighting contains high frequencies like directional sources or sharp edges it is well-conditioned. On the other hand, it is ill-conditioned for soft lighting.

We can see that the equation in the slide gives a precise



mathematical characterization of the BRDF estimation problem, telling us when it is well-posed and when not.

#### Inside the Rendering Equation (VI)

- Inverse Lighting:  
$$L_{in} = A_l^{-1} B_{l,p,q} f_{l,p,q}$$
- We can see lighting recovery is *well-posed* unless the denominator vanishes for all  $p, q$  for some  $l$
- Thus, lighting recovery is well-conditioned when the BRDF contains high frequencies (e.g.: sharp specularities)
- And is ill-conditioned for diffuse surfaces

39

Continuing with the analysis presented in [Ramamoorthi02], now with respect to the Inverse Lighting problem.

We can see that this problem will be well-posed when the denominator does not vanish for all  $p, q$  for some  $l$ . This means that the spherical harmonic expansion of the BRDF transfer function contains at least one non-zero value for all orders.

If we resort again to the signal processing framework, we see that, when the BRDF filter truncates certain frequencies in the input lighting signal, we cannot determine those frequencies from the output signal (our reflected light field). This happens, for instance, if the BRDF acts as a low-pass filter. We can see that Inverse lighting is well-conditioned when the BRDF contains high-frequencies: when its frequency spectrum decays slowly. So, we conclude that inverse lighting is well-conditioned when the BRDF contains sharp specularities, like a mirror surface in the ideal case. On the other hand, inverse lighting from diffuse surfaces is ill-conditioned. We can explain that noting that highly specular surfaces act as high-pass filters, so the resulting images have most of the high frequency content in the lighting. In that case, the lighting can be estimated. On the contrary, diffuse surfaces *blur* the illumination, making it difficult or impossible to recover the high frequencies. We can say that they act as low-pass filters.

#### Inside the Rendering Equation (VII)

- Combined Inverse Problems:
  - The same way, we can see that, up to a global scale, the reflected light field can be factored into the lighting and the BRDF
  - (if the appropriate coefficients do not vanish)
- Inverse Surface Design:
  - Extrapolating the analysis, we can infer that, as the BRDF goes more diffuse, surface recovery becomes harder, reducing convergence of any algorithm we could invent

40

Having analyzed estimation of the BRDF and lighting alone, we can extend the analysis to the problem of *factorizing* the light field, i.e., simultaneously recovering the lighting and BRDF when both are unknown.

The main result is that **the reflected light field can be factored into the lighting and the BRDF**, provided the appropriate coefficients of the reflected light field do not vanish, i.e. the denominators above are nonzero. If the denominators do vanish, the inverse-lighting or inverse-BRDF problems become ill-posed and consequently, the factorization becomes ill-posed.

We can also extrapolate the reasoning to the Inverse Surface Design problem, and we can infer that, as the BRDF goes more and more diffuse, the recovering of a surface becomes more difficult, obviously reducing the convergence of any algorithm we could invent.

## Up to now

- Now, we have presented our main classification based on the compact version of the Rendering Equation

$$L = L_e + \widehat{K}GL$$

- And we have seen the problems involved in Inverse Rendering

41

Up to now we have presented our main classification based on a compact form of the rendering equation, and we have presented the main problems involved in Inverse Rendering problems.

## Alternative classifications

- But the one presented is not the only possible classification we can use, but is the main one
- Alternative classifications can be introduced to enrich our view of the problems and their **solutions**

42

Now, we are going to present some alternative ways of classification of the problems, as the one presented is not the only one we can use. Of course, it will be our main one, but other possible classifications will help us to enrich our view of the problems we are facing and their solutions.

## Alternative Classification: Nature of the solving algorithm

- **Direct-solving**
  - Avoid solving the forward problem at any time
  - Mostly based on the construction of a linear system of equations, and solving it
  - Others use a direct measuring approach
  - A few use Monte Carlo methods
- **Indirect-solving**
  - Based on an optimization procedure
  - Require the solution of a forward problem at least once per iteration
- **Mixed**
  - A combination of both

43

We can present an alternative classification based on the nature of the solving algorithm used. Some of the algorithms are called "direct solving" methods, because they avoid solving the forward rendering problem at all. Most of these approaches build a linear system of equations, trying to solve them, while others take direct measurements of the available information and, finally, there are a few that are based on Monte Carlo methods. Other algorithms are "Indirect-solving" approaches, which need to compute, at least once per iteration, a forward solution to the problem. Those approaches are generally based on an optimization procedure. Finally, we can mention the "mixed" approaches, which are a mixture of the other two, estimating some parameters by direct computation, while leaving others for an optimization process.

## Alternative Classification: Approximations to the equation

- Each approach needs to define how they treat the global illumination equation:

$$L(\mathbf{r}, \omega) = L_e(\mathbf{r}, \omega) + \int_{\Omega} f_r(\mathbf{r}, \omega \rightarrow \omega') L(\mathbf{r}', \omega') \cos \theta d\omega'$$

- **General**
  - General solution independent of any light propagation algorithm
- **Radiosity**
  - Diffuse surfaces, patches with constant radiance
- **Monte Carlo**
  - Firing rays from "known surfaces" towards "unknown surfaces", gathering information to integrate results
- **Local Illumination**
  - Only taking into account paths emitter → surface → registration system (eye)

44

Another factor to take into account is whether the different papers treat the full rendering equation, or a simpler version based on a simplification of the illuminating equation. Basically, we can mention four main approaches: firstly, the general methods, which are independent of the lighting computation procedure chosen, being able to accommodate any rendering algorithm. The second group is the radiosity-based methods, which rely on the computation with diffuse surfaces, and considering patches with constant radiance. The third group is based on Monte Carlo algorithms, which in general are used to compute the unknown information by firing rays from the surfaces with known properties towards the surfaces with unknown properties,

gathering the information finally used. In general, this "gathering" process involves some sort of weighted averaging.

The last group is the one based on local-illumination, considering only point light sources and without considering inter-reflections. Thus, they only consider the paths from the emitter to surface and from this surface to the registration system.

**Alternative Classification:  
BRDF**

- The *Bidirectional Reflectance Distribution Function* (BRDF) plays a key role
- **General BRDF**
  - Uses the BRDF information in a generic way
- **Specific Model**
  - Like the well known Phong, Torrance-Sparrow, ...
- **Diffuse**
  - Just use a real number!
- **Pure Specular**
  - Provides the higher accuracy
  - Probably needs more computations: taking into account many ray bounces

45

It is important to consider the sort of BRDF each approach uses, as the BRDF has a central role in the different inverse problems, as we mentioned before. For example, there are many approaches that try to use a general model for the BRDF, treating the information in a generic way.

Others rely on a specific BRDF model, like a Phong or a Torrance-Sparrow BRDFs. Others, just use a lambertian BRDF, which is characterized just by a real number at each location! Finally, pure specular BRDFs are very commonly used, as are the ones who provide the maximum accuracy for inverse recovery problems, as we saw before. Unfortunately, as the BRDF is more specular, the computation of many light bounces becomes necessary, which could lead to larger computation times.

**Alternative Classification:  
Visibility**

- **Remember:**  
we need to compute the *visible surface function*  
 $v(r, \omega) \equiv \inf\{x > 0: r + x\omega \in \text{Surfaces in the env.}\}$
- That means detecting blockers
  - between the sources and the surfaces
  - between the surfaces and the registration system (eye)
- Not computing this is much faster, but much more inaccurate
  - Could lead to solutions not applicable in real-life

46

It is also important to mention the treatment of visibility in the different approaches: when computing the radiance with the equations we explained before, the visibility problem consists of detecting if there are any blockers between the source and the surface being illuminated, so not adding their contribution in that case. The same is true for the paths from the surface to the eye or the region where the final radiance computations are needed. Most of the reviewed papers omit this treatment, arriving at solutions not applicable in real-life situations.

**Alternative Classification:  
Specific aspects**

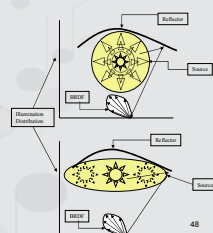
- Each problem also allows to add specific, unique classifications that surely will help us understanding them
- Let's look at each case in particular

47

Finally, we can present a classification with respect to the specific aspects each problem and of their solutions. These are specific, unique classifications that surely will help us understanding them. Let's look at each case in particular.

### Inverse lighting

- Light Simulation
  - General
  - Radiosity
  - Monte Carlo
  - Local Illumination
- Type
  - Emittances
  - Positioning



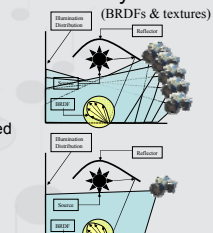
48

In the case of inverse lighting, we can find the four kinds of work: general approaches, radiosity-based, Monte Carlo-based and local illumination-based. But we also can classify them with respect to a specific classification that comes from the two possible sub-problems to solve: inverse emittances, where we only aim at finding the emittances of known light sources, or the problem of where to place a given source, known as "light positioning" problem.

### Inverse Reflectometry

(BRDFs & textures)

- Light Simulation
  - General
  - Radiosity
  - Local Illumination
- Points of view needed
  - One
  - Several.

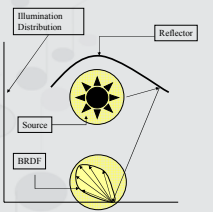


49

Here we find only three kinds of works: to the best of our knowledge, and although is theoretically possible to do it, there are no works that deal with an inverse reflectometry problem, using a Monte Carlo-based approach to solve it. On the other hand, the different approaches could be classified according to the number of images of the scene needed to solve the problem. Normally, this can be wither just one image, or several images, that could come from an animation or several calibrated pictures.

### Combined problems

- Light Simulation
  - Radiosity
  - Local Illumination

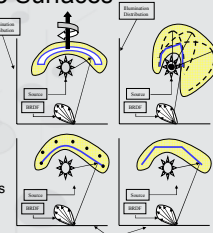


50

For the combined problems we only find radiosity or local illumination-based approaches, although we could think of the possibility of generating solutions with a general approach or with a Monte Carlo method.

### Inverse Surfaces

- Type
  - Analytical
  - Numerical
- Light Simulation
  - "1-to-1"
  - Local illumination
  - Global illumination
- Shape Definition
  - Rotational Symmetric
  - Intersection of quadrics
  - Splines
  - Polygons



51

With respect to inverse surface problems, there are two main groups of works: those that are purely theoretical, showing and demonstrating some properties of the problem and their solutions, and those which try to get a numerical solution that could be implemented in a computer. So, each work treats the light propagation in different ways. Some of the proposed techniques introduce a new sort of lighting method, which we call "1-to-1", which a sever restriction on the local illumination setting. We are going to refer about a little bit later. Finally, it is important to mention that each work uses a different model for the surface, being the rotational symmetric surfaces and the intersection of volumes bounded by quadrics the most common in the theoretical works, while most of the numerical methods either use splines or polygons for the shape to be found.

**Inverse Rendering:  
From Concept to Applications**

- 1. Inverse Rendering Definition**
  1. General definition of inverse rendering & characterization of existing techniques
  2. General solutions/approaches
    1. State of the art in Inverse Lighting
    2. State of the art in Inverse Surface Design
    3. State of the art in Reflectometry
- 2. Applications to Inverse Lighting**
- 3. Applications to Inverse Geometry**
- 4. Conclusion and Future Work**


62

This first section gives an overview of the content of the following surveys and state of the art reports:


- A survey of inverse rendering problems, Gustavo Patow and Xavier Pueyo, *Computer Graphics Forum*, vol 22, num 4, pp. 663-687. 2003.
- A survey of inverse surface design from light transport behaviour specification, Patow, G. and Pueyo, X. *Computer Graphics Forum*. December 2005
- Classification of illumination methods for mixed reality, Katrien Jacobs and Celine Loscos, *Computer Graphics Forum*, vol 25, num 1, 2006.

Instead of a sample of slides, these three papers are included in this document, to give the reader an understanding of what will be presented. The discussion of these three surveys is followed by a break.

**State of the art in Inverse Lighting**



Gustavo Patow  
Xavier Pueyo  
Grup de Gràfics de Girona  
Universitat de Girona

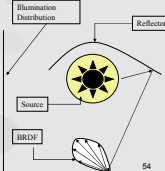


63

Keywords: Inverse Problems, Inverse Rendering, Inverse Geometry, Reflector Design, CAD for luminaries, Optimization

**Inverse lighting**

- Solve for  $L_e$
- We can subdivide the work in this area into
  - **Inverse emittance:** obtaining the emittance
  - **Inverse light positioning:** find the location of the sources
- Different approaches
  - General formulations
  - Monte Carlo
  - Radiosity
  - Local illumination

$$L = L_e + \tilde{K}GL$$


64

**General Formulations (I)**

- Find the luminaries intensities, positions and/or orientations
- Without depending on any particular rendering algorithm

65

### General Formulations (II)

- If we consider  $n$  distant light sources
- Each of them
  - Characterized by a function  $\Phi_i$  (its contr. to the env.)
  - Parameterized by some parameters  $u_i^k$
  - In general, the  $u_i^k$  will be
    - its intensity  $u_i^0$
    - its position  $r_i$  of the light source
    - its orientation  $\omega_i$
- Then, the illumination in a scene would be

$$\Phi = \sum_i^n u_i^0 \Phi_i(u_i^1 \dots u_i^k)$$

56

### General Formulations (III)

- If we measure some intensity values  $\alpha_j$  at some evaluation points, that can be
  - On the surfaces
  - On the screen
- And, if we assume a **linear** relationship with the  $\alpha_j$ , we get

$$\alpha_j = \mathfrak{R}(\Phi) = \sum_i^n u_i^0 \mathfrak{R}(\Phi_i(u_i^1 \dots u_i^k))$$

57


### General Formulations (IV)

- It is very important to notice that
  - As long as the rendering process  $\mathfrak{R}(\Phi)$  conserves this linearity, we can use **any** rendering algorithm
  - The solutions posed this way are *rendering-independent*
  - Just compare patch intensities or images
  - We don't care how the energy got there!
- In general
  - We are facing a non-linear problem
  - Probably with many minima
  - A global optimization procedure is needed

58

### General Formulations (V)

[Patow06]



## General Formulations (VI)

- [Costa-Sousa-Cardoso99]
  - Evaluated objectives defined on the surfaces
    - Intensity over certain surfaces
    - Avoid glares on operator's faces (direct or by specular reflections)
    - Other: the objectives are specified by scripting
  - Based on Simulated Annealing
- [Shackel-Lischinski01]
  - Perception-based image quality function as objective
  - Tries to effectively communicate information about shapes, materials and spatial relationships
  - Steepest descent with "wisely chosen initial values"
- [Patow06]
  - Used images taken with a static camera
  - Objectives defined in image-space
  - Completely rendering-independent
  - Used Simulated Annealing

60

## General Formulations (VII) Inverse Emittance:

- And, if we assume a **linear** relationship with the  $\alpha_j$ , we get
 
$$\alpha_j = \mathfrak{R}(\Phi) = \sum_i u_i \mathfrak{R}(\Phi_i)$$
  - We can try to find the  $u_i$  as a typical least squares problem
- [Schoeneman-Dorsey-Smits-Arvo-Greenberg93]
  - "Painted" lighting requirements on the surfaces of the scene
  - Used a modified Gauss-Seidel iteration plus a regularization process (remember: diffuse BRDF)
- [Marschner-Greenberg97, Marschner98]
  - From a 3D model and a photograph
  - Estimates the directional distribution of the incident light
  - Actually, used a basis  $\phi_j$  for each direction  $j$
  - Least squares, but required regularization

61


## Monte Carlo-Based Formulations (I)

- An Inverse Monte Carlo method is proposed:
  - Fire rays from the "known" surfaces towards the "unknown" surfaces
  - Gather illumination information on some "calculation points" on the "unknowns"
  - Average this info to get the final results
- Has never been used with light positioning problems...

62

## Monte Carlo-Based Formulations (II)

- [Oguma-Howell95]



63

## Radiosity-Based Formulations (I)

- The general problem is reduced to a radiosity one
  - Diffuse BRDF for the surfaces (patches)
  - Constant radiosity over each one

$$B_i = L_{e_i} + \rho_i \sum_j F_{ij} B_j$$

$B_i$  the  $i$ -th patch radiosity  
 $L_{e_i}$  its emittance  
 $\rho_i$  its reflectivity  
 $F_{ij}$  the form factor between elements  $i$  and  $j$

- To the best of our knowledge, never was applied to light positioning problems
  - Diffuse surfaces with sharp lighting...
  - probably very ill-conditioned
  - Unless stated at the full-patch level...

64

## Radiosity-Based Formulations (II)

- [Harutunian-Morales-Howell95]
  - Context: Radiative heat transfer
  - Modified Truncated Singular Value Decomposition
- [França-Oguma-Howell98]
- [Morales-Matsumura-Oguma-Howell97]
  - Same approach, but with participating media
  - Uses a FE approximation for the PDE for the medium
- [Fournier-Gunawan-Romanzin93]
  - Fitted emissions to observed values
- [Drettakis-Robert-Bougnoux97]
  - Improved by using a hierarchical radiosity system
- [Loscos-Frasson-Drettakis-Walter-Granier-Poulin99]
  - Photogrametry: Reconstructed scene geometry from photographs
  - Indirect lighting computed through a hierarchical radiosity system
  - Direct lighting computed separately with ray-casting

65

## Local Illum.-Based Formulations (I)

- Basically, the consider only the direct illumination at a surface
- The rendering equation is actually never used
- Computations are much simpler and faster to do
- Depending on the application, results could be too unrealistic...

66

## Local Illum.-Based Formulations (II) Inverse Emittance

- [Ramamoorthi-Hanrahan01]
  - First option: Solve a linear least squares system (like in the general case)
  - Second option
    - Subtract the diffuse component
    - use the resulting mirror-like distribution to recover a high resolution angular space of the illumination
    - Two step process: estimate diffuse from illum, then illum frequency parameters to get sharper results
- [Sato-Sato-Ikeuchi99a,b]
  - Use radiance inside shadows to estimate illum as a collection of point sources
    - uniformly distributed in the scene
    - On a sphere of directions
  - Solved a linear system

67



### Local Illum.-Based Formulations (III) Light Positioning Problems

- [Poulin-Fournier92]
  - Use highlights and shadow for modeling directional and point light sources
- [Poulin-Ratib-Jacques97]
  - Used sketches of desired shadows and highlights
- [Guillou00]
  - Creates many directional light sources from purely diffuse regions in the scene
  - Groups them to estimate a point from the intersections
  - This point is used as starting point to an optimization (Levenberg-Marquardt) of the least squares error



68

### Inverse Rendering: From Concept to Applications

#### 1. Inverse Rendering Definition

1. General definition of inverse rendering & characterization of existing techniques
2. General solutions/approaches
  1. State of the art in Inverse Lighting
  2. State of the art in Inverse Surface Design
  3. State of the art in Reflectometry

#### 2. Applications to Inverse Lighting

#### 3. Applications to Inverse Geometry

#### 4. Conclusion and Future Work

69

### State of the art in Inverse Surfaces



Gustavo Patow  
Xavier Pueyo

Grup de Gràfics de Girona  
Universitat de Girona




70

Keywords: Inverse Problems, Inverse Rendering, Inverse Geometry, Reflector Design, CAD for luminaries, Optimization

### Inverse surfaces

- Type:
  - Analytical
  - Numerical
- Light Simulation:
  - "1-to-1"
  - Local illumination
  - Global illumination
- Shape Definition:
  - Rotational Symmetric
  - Intersection of quadrics
  - Spline
  - Polygons

$$L = L_e + KGL$$



71


Finally, if  $\hat{G}$  is unknown, we have an **inverse geometry** problem. For an in-depth survey on those problems, refer to [cite{PP00b}](#).

## Far- vs. Near- vs. Near-Field

- Far-Field
  - Purely directional distribution
  - Ignores light origin
  - Sources are treated as anisotropic point lights
- Near-Field
  - Directional and spatial distributions
  - Equivalent to a Light-field or a Lumigraph
- Near-Field-RD (for reflector designers)
  - Often, in lighting engineering, they refer a Near-Field as the illumination (irradiance) on plane near the source
  - This way, a Far-Field is a limiting case when the plane is moved infinitely far away

72

## Target Optical Set



73

Firstly, we should define what an *optical set* is.

The reflector shape to be found is just a piece of a set called in lighting engineering an *optical set*, which consists of a light bulb, the reflector and, possibly, the diffusor. The reflector has a border, contained in a plane, that limits its shape. In general, a reflector must fit inside a holding case, so its shape cannot be lower at any point than the plane defined by the border nor higher than a certain threshold defined by the case. We can say that the case defines a bounding box for the reflector.

## Problem Statement

- Given
  - The radiance distribution of the light bulb
  - A reflector border
  - A desired outgoing radiance distribution for the Optical Set
- Find a shape within a user-prescribed tolerance



74

We can state more precisely our problem as:

Given the outgoing radiance distribution of a light bulb and a reflector border, and given a desired optical set-outgoing radiance distribution, find the corresponding shape for the reflector. Do this up to a user-defined tolerance.

## Theoretical Works (I)

- Many of the earlier works were based on the "1-to-1" approximation



- It is a very restrictive approximation...
- [Wescott-Norris75] [Wescott83]


75

## Theoretical Works (II)

- [Brickell-Wescott78]
  - Studied the Near-Field-RD problem
  - Got an equation of Monge-Ampère type
- [Oliker89, 02, 03]
  - Studied the Near-Field-RD problem
  - Worked with rotationally symmetric reflectors
  - Got a more general expression
  - Proved existence and uniqueness for the radially symmetric case
- [Wang96]
  - Existence, uniqueness and smoothness for a far-field problem
  - Based on a differential geometry formulation
  - Showed that the regularity of the solution failed even with simplest cases

76


## Numerical Works General Algorithm



77

## Numerical Methods: Local Illumination (I)

- [Kochengin-Oliker98, 03] [Caffarelli-Kochengin-Oliker99]
- [Kochengin-Oliker-vonTempksi98] (mostly theoretical)  
(refractions → intersecting hyperboloids)



78


## Numerical Methods: Local Illumination (II)

- [Neubauer94, 97]
  - Used a bicubic B-Spline for the surface
  - Minimized the mean square error between
    - Obtained light distribution
    - User prescribed desired irradiances
  - Light propagation by Gaussian quadrature
    - 4 x 4 nodes chosen on the reflector
    - Incoming and outgoing direction calculated
    - Outgoing energy distributed over neighbors in the registration region (far-field)
  - Optimization
    - Powell's method (projected conjugate gradients)
    - Line search via quadratic interpolation
    - Initialization must be quite accurate to avoid wrong minima !

79

## Numerical Methods: Local Illumination (III)

- [Halstead-Barsky-Klein-Mandell95, 96]
- [Halstead96]



## Numerical Methods: Local Illumination (IV)

- [Loos-Slusallek-Seidel98]
  - Designed progressive Lenses for defects in human visual system
  - Lens represented as front surface
    - Back surface is toroidal or spherical
  - Light Propagation
    - Finds point the eye should focus on by tracing a single ray from center of pupil to center of pixel
    - Then, wavefront tracing is used for re-focusing the eye (get a sharp image)
    - Distribution ray-tracing is used to get final pixel color
  - Accurate error functional for
    - Desired properties (effective astigmatism)
    - Optical error across the lens
  - Optimization
    - Variant of Newton iteration
    - Only two iterations needed !

81

## Numerical Methods: Global Illumination

- [Doyle-Corcoran-Connell99, 01]
  - Objective function
    - Element-wise difference between desired and current light distribution
    - Penalty for rays outside this region
    - Difference between reflected and desired light power
  - Light propagation: Flatland (2D)
    - Recursive ray tracing
    - Extended light sources defined as an array of lambertian-like points on regular intervals on a circle
  - Surface
    - Cubic Bézier curve
  - Optimization
    - Evolution strategy (a variant of Genetic Algorithms)
    - Control points from the Bézier curve are the genes of the chromosomes
- [Patow-Pueyo-Vinacua04]
  - Detailed description later

82

## Inverse Rendering: From Concept to Applications

### 1. Inverse Rendering Definition

1. General definition of inverse rendering & characterization of existing techniques
2. General solutions/approaches
  1. State of the art in Inverse Lighting
  2. State of the art in Inverse Surface Design
  3. State of the art in Reflectometry

### 2. Applications to Inverse Lighting

### 3. Applications to Inverse Geometry

### 4. Conclusion and Future Work

83

State of the art in Reflectometry



**Katrien Jacobs**  
**Celine Loscos**


University College London  
Virtual Environments  
and Computer Graphics



84

This talk will be about illumination methods for mixed-reality.

Mixed Reality – definition




[Milgram91]

85

A definition of mixed-reality is given by Milgram et al. In their paper, it is stated that there exist 4 types of worlds.

Mixed Reality – definition




Real Environment (RE)

[Milgram91]

86

A first world consists of real elements only.

Mixed Reality – definition



Augmented Reality (AR)


Real Environment (RE)

CREATE


[Milgram91]

87


In a second world some virtual objects are added to the real world, the resulting environment is called an augmented reality.



In the third world, real elements are added to an entirely virtual environment, the resulting environment is called an augmented virtuality.



Finally the last world consists of virtual elements only, resulting in a virtual environment.




Mixed-reality is the combination of AR and AV.




A mixed-reality can be displayed using see-through glasses (left) or projected on a screen (right).

### Illumination methods


- Illumination consistency between virtual and real object
- 3 groups:



[Jacobs04]  
Common  
Illumination



[Lancos06]  
Relighting



[Boivin01]  
Inverse  
Illumination

92

When an effort is made to match the illumination of the virtual elements and real elements, three types of illumination methods can be distinguished: common illumination, relighting and inverse illumination.



Common illumination: the real illumination is unchanged; the illumination effects induced by the virtual object are simulated (left).

Relighting: the illumination of the real elements is changed and the virtual and real elements are re-lighted using the same illumination (middle)

Inverse illumination: the BRDF of the real elements is calculated in order to achieve a physically correct relighting (right)

### Classification

- Based on input required
  - Geometry
    - One image
    - Several images
    - Many images
  - Scene radiance:
    - One image
    - Several images
    - Many images

93

In our state of the art report presented at EG04 and later published in CGF [Jacobs06], a classification is presented of existing illumination methods for MR. The classification is done based on the type of geometry and radiance information required by the method. The geometry is usually reconstructed using reconstruction software or using scanners. The scene radiance is extracted from input images.

### Classification

1. No geometry, one input image
2. Geometry, one input image
3. Geometry, several input images
4. Geometry, many input images




94

Based on this 4 different classes are identified:  
 No geometry information is required, each scene point needs to be visible in one input image  
 Geometry information is required, each scene point needs to be visible in one input image  
 Geometry information is required, each scene point needs to be visible in several input images  
 Geometry information is required, each scene point needs to be visible in many input images

In what follows a short overview is given of these classes with only a SELECTION (not all) of the methods that exist in the literature that fall under this category.

### 1. No geometry, one input image

- Virtual object inserted inside real scene containing real objects
- Anti-aliasing and distance taken into account
- Shadow projection
- Illumination virtual and real objects consistent
- Fog generation

[Nakamae06] 95

A first method is from Nakamae et al. which uses a photograph and inserts virtual objects without knowing the scene geometry beforehand. They add shadows and fog.

1. No geometry, one input image

4 steps:

1. Acquiring environment illumination (environment map)
2. Pre-filtering environment illumination
3. Render using multiple passes (original texture preserved)
4. shadow generation

[Agusanto03] 96

A second method is from Agusanto et al., this method uses pre-filtered environment maps to simulate the illumination on virtual objects.

1. No geometry, one input image

4 Steps:

1. low-level geometry reconstruction using pair of omni-directional images
2. capturing illumination environment: environment map
3. mapping of environment map onto the geometry
4. rendering VO using Ray Casting method

[Sato-Sato-Beneito99b] 97

Sato et al. present a common illumination method that makes use of environment maps to simulate virtual shadows. The shadow generation is non-real time.

Geometry, one input image

Modified shadow volume algorithm for AR environment:

1. generating shadow volumes
2. rendering algorithm for shadow volumes

[Haller03] 98

Haller et al. present a local common illumination method that generates hard shadows using shadow volumes. The shadows are generated in real time, but problems arise when a virtual shadow overlaps with a real shadow (right). Shadows are generated from the real scene onto the virtual object and from the virtual object onto other virtual objects and onto the real scene.

Geometry, one input image

1. decompose scene into local and distant part
2. calculate diffuse reflectance properties of real objects in the real scene
3. render scene using differential rendering mechanism

[Debevec03] 99

Debevec et al. calculates the diffuse BRDF of the scene materials for the local scene using differential rendering. The result is a global common illumination solution, non-real-time.



### Geometry, one input image

1. Capture scene
2. Split scene into two parts: source patches, receiver patches
3. Construct shaft hierarchy between source and receiver patches
4. Render, assuming diffuse only

[Gibson 03]

rendering      reference

Gibson et al present a non-real time global common illumination method.

### Geometry, one input image

Common illumination between real and computer generated scenes

1. reconstruct the real scene: model the scene with few boxes: low-level geometry reconstruction
2. divide scene in elements and estimate diffuse reflectance properties
3. progressive radiosity is used to calculate global common illumination
4. render scene using ray casting

[Fournier-Gunawan-Romanciu95]

101

Fournier et al. present a non-real time common illumination solution. The diffuse material properties are estimated.

### Geometry, one input image

[Lacour00]

display ratio  $\frac{B_1}{B_0} \cdot \text{Texture}$

relighting allowed

Original photograph      Without shadows      Final rendering 102

This method is discussed in a following section in more detail.

### Geometry, one input image

Image-based rendering of diffuse, specular and glossy surface from one single image

1. Reconstruct scene geometry, light source position
2. BRDF model according to Ward et al.: 3 or 6 parameters: diffuse, specular, roughness
3. Iterative approximation based on one single image:
4. Use scene knowledge to fill in invisible parts of the scene

[Boivin01]

103

In this method from Boivin et al, the specular and diffuse material properties are estimated from one photograph. This method is discussed in more detail in a subsequent section.

### Geometry, several input images

reflectance × confusion → avg. → merged reflectance

- direct lighting: pixel per pixel
- indirect lighting: optimised radiosity solution

[Lacou-Frasson-Drettakis-Walter-Granier-Poulin99] 104

This method from Loscos et al. is explained in more detail in a subsequent section.

### Geometry, several input images

Image based Photometric Reconstruction using Virtual Light Sources

- reconstruct scene
- model direct illumination by a set of virtual light sources
  - spherical illumination surface
  - set of virtual light sources must produce a similar illumination as real illumination
- estimate reflectance of objects in real scene
- render using ray-tracing (not real-time)

original [Gibson01] synthetic light sources augmented 105

In this method from Gibson et al. the scene illumination is reconstructed on a virtual illumination surface. This is used to calculate the BRDF properties. The result is a non-real time global relighting method.

### Geometry, many input images

[Yu99]

Inverse global illumination:

- Reconstruct scene geometry and light source positions
- Diffuse albedo estimated using inverse radiosity
- Roughness and specular albedo estimated using non-linear optimisation
- Rendering with RADIANCE

input relighted scene 106

Yu et al.: calculate the diffuse and specular material properties using a non-linear optimisation. The result is a non-real time global illumination method.

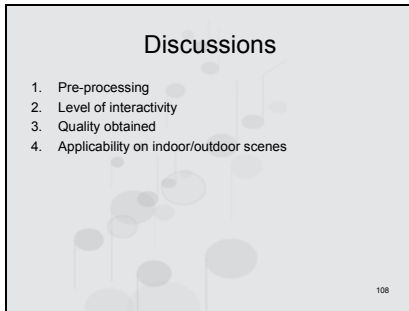
### Geometry, many input images

Recovering photometric properties from photographs

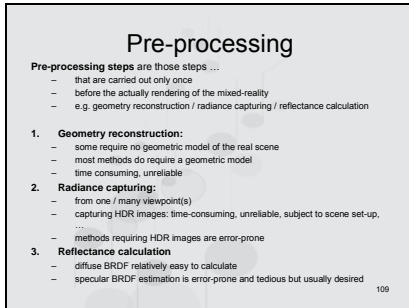
- Separate scene into four parts
- Radiance each point in the scene captured twice
- Two diffuse pseudo BRDF are calculated, specular component estimated (LSE)
- Rendering:
  - sky texture extracted and mapped on hemisphere
  - sky controlled by three parameters: sunset/sunrise, intensity, sun position
  - local scene rendered using pseudo-BRDF's and radiance values

[Yu1998] 107

In this method from Yu et al. diffuse and specular effects are estimated using LSE by separating the scene into four parts (local model, sky, sun and surrounding environment). The radiance of a point needs to be captured once in shadow and once in sunlight, making this method very tedious. Then two BRDFs are calculated (material in sun, material not in sun) and combined.



The presented methods can be discussed and compared based on 4 different parameters:  
 pre-processing time required  
 Level of user interactivity  
 The quality obtained of the simulated illumination  
 The applicability of the method on indoor/outdoor environments

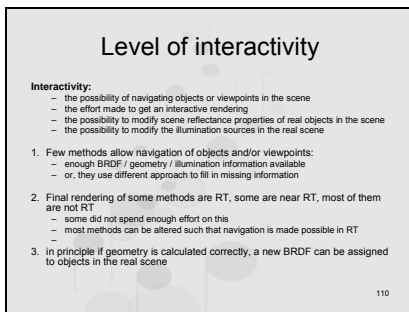


A pre-processing step is a task that is carried out only once, before the actual rendering takes place. For instance geometry reconstruction, radiance capturing and reflectance calculation.

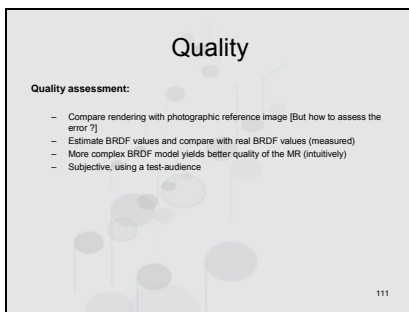
Geometry: the more precise the geometry model needs to be the more time consuming the method will and also the more unreliable as it is not always possible to acquire a decent geometry model

Radiance capture: the more input images needed, the more time-consuming the method and the less reliable, especially in a scene with changing illumination for instance, where it is difficult to capture a scene point from many directions under the same illumination

Reflectance calculation: the more high level the reflectance calculation, the better the quality of the final rendering, but the more tedious the method.



The level of interactivity depends on the possibility to navigate through the scene, to change/move the objects or to change the illumination of the scene.



The quality of a method is very subjective, but can also be measured using more objective measures. For instance it is expected that with a high level BRDF model, the resulting rendering will be better. The quality of a method can also be measured by measuring the resulting rendering against a photographic reference, but in turn this poses the question: what to compare?.

### Applicability on indoor and outdoor scenes

- Outdoor scenes – more complicated than indoor scenes:
  - geometry more difficult to extract (too much detail, moving objects, less controllable than indoor scenes)
  - radiance more difficult to capture (movement, intensity too high, less controllable than indoor scenes)
- Methods using environment maps are usable for outdoor scene
  - if environment map is correct!!!
- Four methods explicitly tested their method on an outdoor scene
  - [Sato-Sato-Ikeuchi99a]
  - [Debevec98]
  - [Gibson03]
  - [Yu98]

112

Most methods are restricted in their applicability on complex scenes, like for instance outdoor scenes, where geometry reconstruction is often difficult, and where radiance capture is often hampered by weather conditions. Several methods presented can be applied on outdoor scenes, though it is a bit questionable if these methods are robust against fast weather changes and uncontrollable object movement.

### Conclusion

- **Mixed-Reality:**
  - Merging of real and virtual elements
  - Illumination consistency: virtual and real elements illuminated in the same manner
  - Three methods to obtain illumination consistency:
    - common illumination
    - relighting
    - inverse illumination methods
- **Illumination methods:**
  - The more information available, the better the quality of the final mixed-reality
  - Therefore classification based on amount of input information available
  - Existing methods are not ideal

113

We've discussed mixed-reality illumination methods. Three types of illumination methods exist. The existing methods can be classified based on their input information required into four different groups.

### Bibliography

- **A survey of inverse rendering problems**, Gustavo Patow and Xavier Pueyo, Computer Graphics Forum, vol 22, num 4, pp. 663-687, 2003.
- **A survey of inverse surface design from light transport behavior specification**, Gustavo Patow and Xavier Pueyo, Computer Graphics Forum, December 2005
- **Classification of illumination methods for mixed reality**, Katrien Jacobs and Celine Loscos, Computer Graphics Forum, vol 25, num 1, January 2006.

114


### Inverse Rendering: From Concept to Applications

1. **Inverse Rendering Definition**
2. **Applications to Inverse Lighting**
  1. General problems with complex & uncontrollable scenes
  2. Postproduction: capturing HDRI for outdoor scenes
  3. Urban planning: simulating common illumination
  4. Indoor Lighting Design
3. **Applications to Inverse Geometry**
4. **Conclusion and Future Work**

115


Before actually discussing some applications of Inverse Rendering, some general problems with rendering methods are discussed first. The focus is on the efficiency of the existing methods in real-life situations. This is to highlight that most existing methods, though they often produce nice results, usually forget the practical applicability of the method.

### General Problems with Complex, Uncontrollable scenes



**Katrien Jacobs  
Celine Loscos**

University College London  
Virtual Environments  
and Computer Graphics




116

### General Problems with Complex, Uncontrollable scenes

**Geometry:**

- Some (outdoor) scenes contain **many (difficult to model) details** (e.g. buildings, trees, fauna, people, cars, ...)
- Some (outdoor) scenes are **less easy to control** (e.g. people passing, cars driving, objects moving because of wind)



117

Geometry extraction of outdoor scenes is complicated. The scene can contain lots of detail, which makes it difficult to model. The scene can be difficult to control, for instance the object might not be static, which makes it again difficult to model.

### General Problems with Complex, Uncontrollable scenes

**Illumination:**


- Complex illumination:
  - Direct: Sun
  - Indirect: Sky & Objects
- Difficult to control:
  - Sun intensity too high to capture with camera
  - Clouds drifting in sky have significant effect on stability outdoor illumination

118

And so is the extraction of the illumination. The illumination might not be static which makes the HDRI capturing error prone. The intensity of the illumination can be too high to capture (e.g. Sun)

### Geometry Loss due to ...

- 1. Reconstruction errors**
  1. Shadows in texture show more detail than the reconstructed object
  2. Outdoor scenes contain more detail than indoor scenes




119

As we can see, the shadow of the statue contains “rounded” edges, this indicates that the representation of the statue by a set of triangles/polygons is incorrect.

### Geometry Loss due to ...

1. Reconstruction errors
2. HDR (high dynamic range) capturing of moving objects
  1. Smeared out in final HDR image
  2. Outdoor scenes are less controllable than indoor scenes

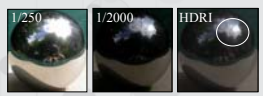


120

Those persons that move during the HDRI capturing appear 'blurred' in the final HDRI. These ghost effects reduce the quality of the HDRI, and the radiance representation of the objects in the scene.

### Incorrect Radiance information due to ...

- Illumination changes during HDR capturing: set of LDR (low dynamic range) images with multiple exposures contain different illumination content, meaningless HDR image
- Illumination changes during radiance capturing entire scene: radiance images are required from different viewpoints.
- Sun intensity too high to capture using conventional ways

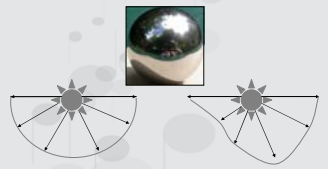


121

The radiance needs to be captured with HDRI images, these HDRI images are often generated from a set of multiple exposures. Problems arise when during the capturing of the multiple exposures, the scene objects move, or when the scene illumination changes. Also, when the intensity in the scene (sun) is too high to be captured with the fastest camera shutter speed.

### Incorrect Radiance information due to ...

- Using one environment map to illuminate the scene: this assumes that the illumination of the sun and the sky is diffuse, while it is not ...



122

When one mirror ball is used to extract the entire illumination, it is implicitly assumed that the scene illumination is diffuse. However, this is not necessarily so.

### Solutions

**Geometry**

- By using image processing techniques as shadow detection/object recognition geometry loss can be compensated for
- Movements during HDRI capturing can be compensated for
- ...

**Illumination**

- Detect illumination changes and compensate
- Use homogenous filter to reduce sun intensity
- Capture light probe at different positions in the scene
- ...

123

## Inverse Rendering: From Concept to Applications

1. Inverse Rendering Definition
2. Applications to Inverse Lighting
  1. General Problems with Complex & Uncontrollable scenes
  2. Postproduction: capturing HDRI for outdoor scenes
  3. Urban planning: simulating common illumination
  4. Indoor Lighting Design
3. Applications to Inverse Geometry
4. Conclusion and Future Work

124

A first application of inverse lighting can be found in post production in which HDRI images of the sky/scene are used to perform inverse illumination. When objects move in the scene, HDRI construction using multiple exposures is error prone. The following talk discusses a solution to minimize the influence of object movement and camera movement on the HDRI generation process.

## Postproduction: Capturing HDRI for outdoor scenes



**Katrien Jacobs**  
**Celine Loscos**  
University College London  
Virtual Environments  
and Computer Graphics


125

## Postproduction: Capturing HDRI for outdoor scenes

1. Introduction
2. What is the problem?
3. Solutions:
  1. Camera Alignment
  2. HCM removal
  3. LCM removal
4. Conclusion
5. Limitations & Future work

126

## HDRI: High Dynamic Range Image?



$dt_1$  >  $dt_2$  >  $dt_3$

- Images captured with # shutter speed (called *exposures* or *LDRs*)
- Camera curve saturation effects: too bright – too dark
- What if dynamic range of the scene is too large?

127

Most cameras allow the user to change the exposure of an image by either changing the shutter speed or the aperture. When the shutter speed is long, the CCD of the camera has enough time to integrate the incoming light and sometimes the CCD might even saturate, like is visible in the image on the left, where the sky is plain white. On the other hand, when the shutter speed is short, the CCD of the camera might not have enough time to integrate the incoming light, especially in dark areas, and as a result some parts in the image can be too dark and not represent the actual light distribution in that area, as is visible in the images on the right.

For a scene with a too high dynamic range, meaning that the scene contains both dark and bright areas, no shutter speed will capture both dark and bright areas without showing some saturation.

This is all due to the characteristics of the CCD integration curve or the camera curve, which relates irradiance in the scene to intensity in the final image in function of exposure. This curve converges to a certain value  $I_{max}$ , and is nearly flat around 0.

### HDR: High Dynamic Range Image?

- HDR from image sequence:
  - Camera curve maps intensities to irradiance
  - Final image is *per-pixel weighted average* of different exposures
  - No saturation effects

A High Dynamic Range Image, is a floating-point image that gives an accurate representation of the irradiance of the points in the scene, rather than the intensity of the pixels in an image.

There are different ways of generating HDRIs, the method presented in this talk makes use of the multiple exposures. This kind of method is the cheapest way of generating HDRIs.

It proceeds in the following manner:

A sequence of exposures is captured of the same scene with different exposures

The reconstructed camera curve is used to transform the intensities to irradiance.

The final HDR is a weighted average of these transformed exposures. The weights are defined per pixels and are based on the level of saturation a certain pixel has in a certain exposure. The task of the weights is basically to discard pixels that show saturation.

The final HDR will be free from saturation, and will actually represent the true irradiance distribution of the scene captured.

### HDR: High Dynamic Range Image?

- What is it used for?
  - High quality images

Now what are these HDRIs used for?

First of all, as they give an accurate representation of the true irradiance distribution in a real scene, they are of more quality than conventional images. After all, using post-processing techniques, any HDR can be transformed into a conventional image, with all necessary saturation if that has to be. Professional software like Photosphere, HDRShop and Photoshop are available to transform exposures into HDRIs.

HDRIs are often used in movie-production and as a result we see companies emerge that design and build HDR cameras, like SpheronVR. And companies that build HDR displays so that the user can actually see the floating point content of an HDR.

### HDR: High Dynamic Range Image?

- What is it used for?
  - High quality images
  - Illumination capturing/generation (environment/cube mapping)


© www.debevec.org

Secondly, since HDRIs represent irradiance, HDRIs are used to capture the illumination in a scene. For instance by using a light probe as can be seen in the image below. This light probe image can then in turn be used to illuminate virtual objects in a real scene.



### Dynamic Scenes ?

- **Camera movement:**
  - Manual capturing without tripod
  - Translation and Rotation effects
- **Object movements:**
  - Controllable: people, cars, ...
  - Uncontrollable: wind

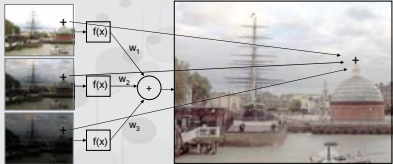


The HDRI generation method presented here, allows the exposures to be captured in a dynamic environment, meaning:

Firstly, that during the exposure capturing process, the camera does not need to be fixed on a tripod, but can be held manually. As a result, the camera position of the different exposures can be slightly different. Secondly, that the objects in the scene can move. Some objects in a scene move in a more or less controllable manner, like cars and people walking by. Other objects move due to uncontrollable powers: like for instance leaves or branches of tree in the wind. [Only visible in the slide when the slide is shown in presentation mode] Since the different exposures need to be captured sequentially, these movements will have an undesired effect on the HDRI generation.

### What is the problem?

- Conventional HDRI generation:
  - Weighted average
  - 1-1 mapping between pixels in different exposure
  - Requirement: *static scene*




So, what is exactly the problem?

Due to the weighted averaging of the pixels to generate the final HDRI the same pixel in the different exposure must correspond exactly to the same physical point in the scene.

Or in other words, the scene has to be static, in order to calculate a correct representation of the irradiance in the scene.

### What is the problem?

- **Dynamic Scene:**
  - Camera misalignment
  - Object movement
- **Consequences:**
  - Ghost effects
  - Blurring



133

When you allow camera misalignments, or object movement, the same pixels in the different exposures will not correspond to the same physical point in the scene. As a result, the final HDRI looks blurred or washed out; a ghost effect becomes visible.

This can be seen in the example below, where the walking people appear multiple times in the final HDRI.

### Postproduction: capturing HDRI for outdoor scenes

1. Introduction
2. What is the problem?
3. **Solutions:**
  1. Camera Alignment
  2. HCM removal
  3. LCM removal
4. Conclusion
5. Limitations & Future work

134

### Solution!

- **Camera misalignment:**
  - Camera alignment
- **Object movement:**
  - High contrast movement removal (HCM)
  - Low contrast movement removal (LCM)

135

Basically there are two sources of error: the camera misalignment and the object movement. As a solution for the camera misalignments a camera alignment algorithm is proposed. The object movement is compensated for by two different algorithms:  
 Low Contrast Movement (LCM) removal  
 And High Contrast Movement (HCM) removal  
 The choice of the algorithm depends on the type of movement the objects in the scene make.

### Camera Alignment

- Different exposures → different image features

2 ≠ exposures  
≠ features

136

Most camera alignment methods are feature-based: they align the images based on distinct features in an image, for instance the edges. Unfortunately, feature-based methods are unsuitable when aligning images captured with different exposure. An example is given here:  
 The top two images are images captured at the same time, from approx. the same position, but with a different exposure time.  
 When these images are processed with an edge detector, the two images at the bottom are generated. Different edges are detected and an alignment algorithm based on aligning the edges in the image would most likely fail or at least it will be very unreliable.

### Camera Alignment

- **However:**
  - Between histograms: monotonic transformation
  - Same pixels smaller than median intensity value (approx.)

→ Median Threshold Bitmap Transformation

137

However, when the histograms of the two images are plot, the two histograms are related by a monotonic transformation. Or, the same group of pixels will have intensity values lower than the median intensity value in both exposures.  
 This median intensity value can now be used as a threshold for a Median Threshold Bitmap Transformation.

### Camera Alignment

- Median Threshold Bitmap (MTB) Transformation:


2 ≠ exposures  
comparable features

138

As an example:  
 For the same two images of the same scene captured with a different exposure speed the median intensity value is calculated and used as a threshold. The result is the two binary images at the bottom.  
 This time, these binary images share the same features, and can therefore effectively be used in an alignment algorithm.

### Camera Alignment

- Results:
  - Use MTB to find Translational and Rotational transformation matrix
  - Iterative procedure



Unaligned sequence

139

The final alignment algorithm works in the following manner:


- Take the different exposures
- Calculate the median intensity value for each of these exposures
- Use this value as a threshold to generate the binary images like the once shown in the previous slide

These are then used to calculate the rotational and translational misalignments between the images. The rotation and translation is calculated in an iterative process.

An example of an unaligned sequence of images is shown here, it is (should be) clear that there is a clear rotational and translational misalignment between the images. [Only visible when the slide is shown in presentation mode]

### Camera Alignment

- Results:
  - Use MTB to find Translational and Rotational transformation matrix
  - Iterative procedure




Aligned sequence

140

This sequence shows the same exposures as in the previous slide but this time a rotational and translational transformation is applied to align the sequence. [Only visible when the slide is shown in presentation mode] This method is similar to the method used in Photosphere, but the alignment method is extended to deal with rotational misalignments as well.

### HCM removal

- HCM: High Contrast Movement:
  - High contrast between moving object and background



HDRI

- Can be identified by examining *variance of pixel* over the different exposures

141

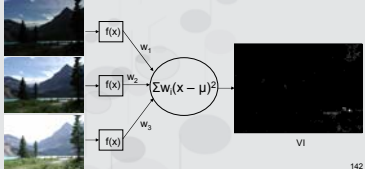
As mentioned earlier, the movement present in the scene is compensated for by two different algorithms. Which algorithm should be used depends on the type of movement present in the scene.

A first type of movement is called High Contrast Movement (HCM) and is displayed here: this type of movement occurs when moving object (walking people) and background are of high contrast.

This type of movement can be identified by calculating the variance of a pixel across the different exposures.

### HCM removal

- Detection: creation of Variance Image (VI)
  - Camera curve maps intensities to irradiance
  - VI is weighted variance image



VI

142

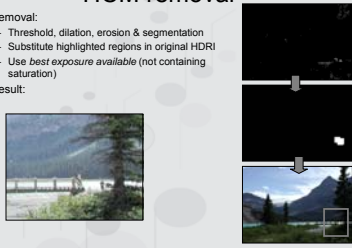
The detection is done by constructing a Variance Image (VI) in the following manner:

First the intensity values in the exposures are mapped to irradiance values using the camera curve. This is necessary in order to being able to actually compare the pixels.

Then VI is constructed by calculating the weighted variance of one pixel across the different exposures.

The weights used here are the same weights used for the HDRI generation: their task is to discard saturated pixels.

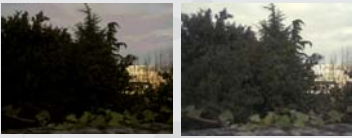
### HCM removal

- Removal:
  - Threshold, dilation, erosion & segmentation
  - Substitute highlighted regions in original HDRi
  - Use best exposure available (not containing saturation)
- Result:
 

The removal of the movement is carried out by the following steps:  
 First a threshold is applied to VI, the resulting binary image is dilated and eroded resulting in the second image on the right: the cluster of white pixels indicates HCM and needs to be removed.  
 The actual removal consists of substituting those identified regions by pixel values by irradiance values coming from one exposure.  
 That exposure is chosen, that shows the least saturation in the identified area.  
 A result is given on the left hand side: the walking people are “de-blurred” and “de-ghosted”

### LCM removal

- LCM: Low Contrast Movement
  - Low contrast between background and moving object



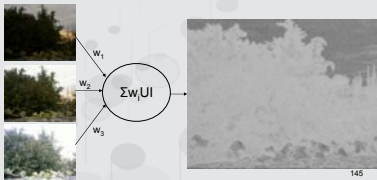
– Movement can be detected with *entropic measurements*

144

The second type of movement is called Low Contrast Movement. It occurs when moving object and background are similar or show low contrast.  
 An example is shown on the right, the leaves and branches move due to the wind. The contrast between moving object and background is low, as they are essentially the same kind of objects. [Only visible when the slide is shown in presentation mode]  
 The resulting HDRI is unclear and fails to represent the moving branches and leaves correctly.  
 This type of movement is identified using entropy.

### LCM removal

- Detection: creation of Uncertainty Image (UI)
  - UI is weighted average of uncertainty measure.

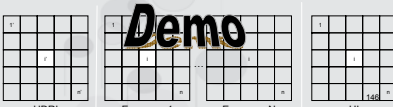


145

The detection of LCM proceeds in a similar fashion as the HCM detection. The detection is done through an Uncertainty Image (UI) and the following slide explains how UI is generated.

### LCM removal

- Uncertainty measure:
  - Generate HDRi & split into cells 1', 2', ..., 1', ..., n'
  - Create empty uncertainty image UI
  - Split each exposure in cells 1, 2, ..., 1, ..., n
  - For each cell in exposure calculate entropy  $H_i$
  - For corresponding cells calculate conditional entropy  $H_{i|j}$
  - Add  $H_i + H_{i|j}$  to value stored in cell I in UI

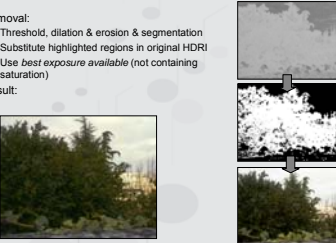


Demo

First an HDRI is generated and split in a set of cells.  
 Then an empty UI image is created, split into the same set of cells.  
 Then for each exposure the following algorithm is applied:

- They are split into the cells
- For each cell the entropy is calculated
- Then the conditional entropy is calculated between the cell in the exposure and in the HDRI
- Finally the sum of the entropy and the conditional entropy is added to the value stored in UI
- This is repeated for all exposures

### LCM removal

- Removal:
  - Threshold, dilation & erosion & segmentation
  - Substitute highlighted regions in original HDRI
  - Use *best exposure available* (not containing saturation)
- Result:
 

The removal of LCM proceeds in a similar fashion as the removal of the HCM:

A threshold is applied to UI, the generated binary image is dilated and eroded. The resulting bright cluster of pixels is identified as pixels affected by LCM

After HDRI generation this LCM region is substituted by pixel values from 1 exposure.

An example is shown in the image on the right: all the leaves are clearly visible, all movement has been de-blurred.

### Conclusion

- Automatic HDRI generation for Dynamic Scenes
- Statistical decision process
- Low cost HDRI generation using multiple exposures
- Camera & object movement removal:
  - Camera alignment
  - HCM & LCM removal
- Without user input

148

In this talk a method is presented that generates HDRIs in a low-cost manner, with low-cost equipment using multiple exposures.

Limited camera and object movement is allowed. These movements are compensated for using a HCM & LCM removal and camera alignment procedures.

All operations are executed automatically, without user input, the decisions made are based on statistical quantities.

### Limitations & Future work

- What if ...
  - too many/large objects would move
  - there is no suitable exposure for the substitution
  - the illumination changes during image capturing (cloud movement)

149

Oblivious there are still limitations to the presented method and there is still some important future work:

What if ...

Too many or large objects would move? Well, most likely the image alignment algorithm will fail. And if that fails, the image calibration, and HDRI generation will be incorrect as well.

There is no suitable exposure, for instance, some light sources (for instance the sun) will always show saturation, and no matter what exposure speed is used? In that case, the HDRI will still show the saturation.

The clouds change that fast, that during the capturing, the illumination has significantly changed as well? Well, in that case no correct HDRI image can be generated. It is essential that the different exposures are captured under the same illumination.

### Inverse Rendering: From Concept to Applications

1. Inverse Rendering Definition
2. Applications to Inverse Lighting
  1. General Problems with Complex & Uncontrollable scenes
  2. Postproduction: capturing HDRI for outdoor scenes
  3. Urban planning: simulating common illumination
  4. Indoor Lighting Design
3. Applications to Inverse Geometry
4. Conclusion and Future Work

150

The next application is that of urban planning. It discusses a method to generate shadows (common illumination solution) in real-time for virtual objects in a scene. This method does not require an accurate scene geometry and allows shadows in the texture model of the scene. Only an approximate light source position needs to be known. The generated virtual shadows are calculated in real-time and are consistent with the shadows already present in the scene.

This method was developed as part of the CREATE project (a European project), and was used to generate consistent shadows for virtual objects inside a model of an existing city.

Urban planning: simulating common illumination



**Katrien Jacobs**  
**Celine Loscos**  
University College London  
Virtual Environments  
and Computer Graphics



151

Urban planning: simulating common illumination

1. Introduction
2. Related work
3. Methodology
4. Results
5. Conclusion

152

Urban planning: simulating common illumination

1. **Introduction**
  1. **Problem description**
  2. Illumination solution
  3. Contribution
2. Related work
3. Methodology
4. Results
5. Conclusion


153

Introduction

**Context: illumination for augmented reality:**

- Augmented reality: real environments augmented by virtual objects/avatars
- Illumination: induced by virtual object

© The CREATE project at FROG, UCL



real scene      geometry extraction      texture mapping & augmentation

154

The context of the presented algorithm is illumination for augmented reality. AR stands for the reality created when you insert virtual objects in a real scene. Illumination stands for the illumination effects that the virtual objects can create. An example is given in the images below. On the left hand side we see a picture of a real scene. Now, in order to insert virtual objects and to generate the illumination effects, the geometry of the scene needs to be reconstructed, the light source positions need to be estimated and the texture of the real scene needs to be extracted as well. A possible illumination effect that a user would like to create is for instance the shadows cast by the avatar on the ground floor as we can see in the example on the right hand side. Now, geometry reconstruction and light source position estimation is error prone and the light source estimation becomes even more complicated when the texture of different parts of the real scene are captured at different times of the day, since then the light sources might actually be different for different parts of the scene.

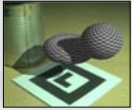
## Introduction

**Context: illumination for augmented reality:**

- Augmented reality: real environments augmented by virtual objects/avatars
- Illumination: induced by virtual object

**Illumination effects:**

- virtual object → real scene
- real scene → virtual object
- virtual objects → virtual objects



155

Now what kind of illumination effects can exist? First of all the virtual object can cast shadows on the real scene, secondly the real scene will illuminate the virtual objects and finally the virtual objects can interfere with each other. An example is given in the image on the right: the torus and the sphere are virtual objects and they generate shadows in the real scene and on top of each other.

## Urban planning: simulating common illumination

1. Introduction
  1. Problem description
  2. **Illumination solution**
  3. Contribution
  2. Related work
  3. Methodology
  4. Results
  5. Conclusion

156

## Introduction

**Common illumination:**

- Virtual shadows consistent with shadows present in real scene
- Illumination in the real scene is unaltered
- No physical understanding of the scene necessary

**Relighting/inverse illumination:**

- Virtual shadows consistent with shadows present in real scene
- Appearance of real objects can be altered, to reflect illumination changes
- Often physical understanding of the scene calculated (BRDF)

• see STAR Jacobs et al. at EG2004: *classification of illumination methods for mixed-reality*

157

There are many different illumination solutions provided in the literature and we can classify them into different classes.

The first class of methods provide common illumination: the virtual object is illuminated conform to the illumination present in the real scene and the virtual shadows cast by the virtual object are generated as well. However, the illumination already present in the real scene can not be changed. Also, usually there is no physical understanding of the scene.

A second class of methods provide relighting. In this case, the virtual object is again illuminated with similar illumination as is present in the real scene and its shadows are generated. This time, a novel illumination can be applied to both the real scene and the virtual object. Usually some form of physical understanding of the scene is at hand, for instance, often the BRDF or reflectance of the materials in the scene are known or at least calculated.

For a complete classification of illumination methods I would like to refer to the following state of the art report presented at Eurographics 2004. Now, the method presented in this talk can be classified under the first group of illumination solutions: it offers common illumination.



**Urban planning: simulating common illumination**

- 1. Introduction**
  1. Problem description
  2. Illumination solution
- 3. Contribution**
  2. Related work
  3. Methodology
  4. Results
  5. Conclusion

158

**Introduction**

**Shadow volumes/maps?**

- Calculation of position of virtual shadow
- Scale pixels underneath (make darker)

- green box is virtual

159

The contribution of the presented method is the generation of consistent shadows for virtual objects in an AR environment using shadow volumes or shadow maps. Shadow maps or shadow volumes generate shadows in real time as the calculation can be done using the graphics hardware. Usually these shadows are hard shadows, but soft shadow generation is possible as well.

**Introduction**

**Shadow volumes/maps?**

- Calculation of position of virtual shadow
- Scale pixels underneath (make darker)

- green box is virtual

160

In short, what shadow volumes and shadow maps does is that it calculates the position of the virtual shadow .....

**Introduction**

**Shadow volumes/maps?**

- Calculation of position of virtual shadow
- Scale pixels underneath (make darker)

- green box is virtual

161

..... and then it scales the pixel intensities of those pixels that fall inside this shadow region as is illustrated here.

**Introduction**

**Shadow volumes/maps?**

**Consistent?**

- Consistent with shadows already present in the scene
- Prevent pixels inside shadow in texture from being scaled
- Texture contains semi-hard shadow

- orange box is real

162

As said earlier, the method provides consistent shadows. Consistent means in this context that the shadows are in the same direction as the shadows already present in the real scene, and with a similar appearance, or colour, which means that we need to find a good scaling factor. In this example we see a real object, the orange box, which casts a clear shadow on the ground.

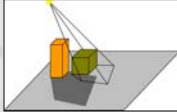


### Introduction

**Shadow volumes/maps?**

**Consistent?**

- Consistent with shadows already present in the scene
- Prevent pixels inside shadow in texture from being scaled
- Texture contains semi-hard shadow



- orange box is real  
- green box is virtual

163

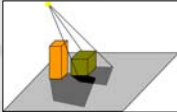
..... when a virtual object is inserted, the position of its shadow is shadow is calculated using an approximate light source estimation. ....

### Introduction

**Shadow volumes/maps?**

**Consistent?**

- Consistent with shadows already present in the scene
- Prevent pixels inside shadow in texture from being scaled
- Texture contains semi-hard shadow



- orange box is real  
- green box is virtual

164

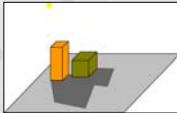
... and then the pixels inside this shadow position are scaled. And as we can see, when virtual and real shadows overlap, using shadow volumes or shadow maps, will generate inconsistent shadows in the overlapping parts. For instance in this example, we see that the colour of the overlapping regions is darker than the rest of the shadow .....

### Introduction

**Shadow volumes/maps?**

**Consistent?**

- Consistent with shadows already present in the scene
- Prevent pixels inside shadow in texture from being scaled
- Texture contains semi-hard shadow



- orange box is real  
- green box is virtual

165

... the presented method will detect the position of the real shadow present in the scene and will not allow scaling pixels in those areas, the result is depicted in the example. The result is that the shadow in the overlapping area is appropriate according to the shadows already present in the scene.


### Urban planning: simulating common illumination

1. Introduction
2. **Related work**
3. Methodology
4. Results
5. Conclusion

166

### Related work

**Sato et al.:** Acquiring a radiance distribution to superimpose virtual objects onto a real scene.




- Common illumination
- Illumination generated using environment maps
- Not real-time
- Inconsistent shadows when real texture contains shadows

[Sato-Sato-Bench99a] 167

In this example from Sato et al. we see a real scene on the left hand side and the real scene augmented with a virtual sphere on the right hand side. The illumination on the sphere is generated using environment maps, the shadow of the virtual object is generated using ray tracing. The method is not real-time. If the virtual shadow would have overlapped with real shadows in the scene, the shadow would have been inconsistent.

### Related work

**Yu et al.:** Inverse global illumination: recovering reflectance models of real scenes from photographs.



- Relighting illumination using inverse illumination
- Reflectance calculation
- Not real-time

[Yu99] 168

This example from Yu et al. illustrates relighting. The top image gives the real scene, the bottom image illustrates the same scene but this time virtually illuminated with a novel illumination pattern. The method does not operate in real time and uses the geometry of the real scene and BRDF information for the materials in the scene to calculate the illumination effects.

### Related work

**Agusanto et al.:** Photorealistic rendering for augmented reality using environment illumination.



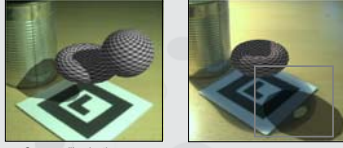
- Common illumination
- Illumination using environment maps
- Real-time using ARToolkit

[Agusanto03] 169

Agusanto et al. provide common illumination using environment maps. A glossy and specular environment map are given on the left hand side. This method operates in real time, the illumination on the virtual objects is generated, no effort is done to include shadows.

### Related work

**Haller et al.:** A real-time shadow approach for an augmented reality application using shadow volumes.



- Common illumination
- Real-time using shadow volumes
- Inconsistent shadows when real texture contains shadows

[Haller03] 170

Finally, an example from Haller et al. which uses shadow volumes to generate the virtual shadows. This method requires a detailed model of the real scene and the light source position. But when we have a look at the image on the right hand side, we see that when virtual and real shadow overlap, the method generates an inconsistent shadow pattern.

Urban planning: simulating common illumination

1. Introduction
2. Related work
3. **Methodology**
4. Results
5. Conclusion

171

Next we will discuss the methodology of the shadow generation algorithm presented in this talk. The methodology consists out of 4 different steps.

Methodology

- geometry
- light source position
- texture
- virtual element

172

The first steps is the input generation. At least four different inputs need to be generated. The first is the geometry of the real scene. Geometry reconstruction, which can be done with methods like image modeler, is very tedious and error prone.

Methodology

- geometry
- light source position
- texture
- virtual element

1. shadow detection

- shadow estimation
- shadow calculation
- scaling factor

- orange box is real

173

Then the shadow detection process proceeds. The following slides give an overview of how this was implemented in the presented method.

Methodology

- geometry
- light source position
- texture
- virtual element

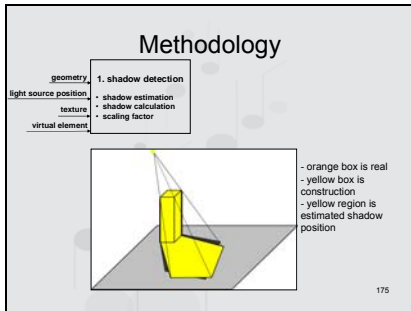
1. shadow detection

- shadow estimation
- shadow calculation
- scaling factor

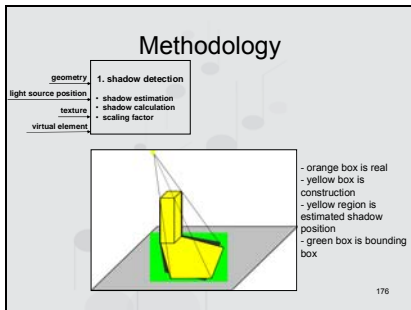
- orange box is real  
- yellow box is construction

174

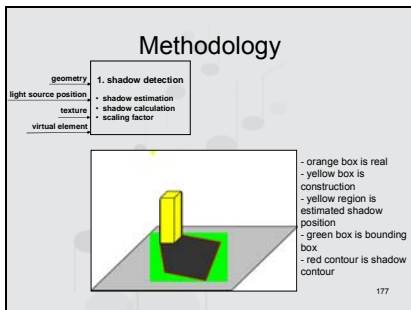
The orange box is a real object, the yellow box is its reconstruction (misaligned with the original object as we can see here, due to numerous errors).



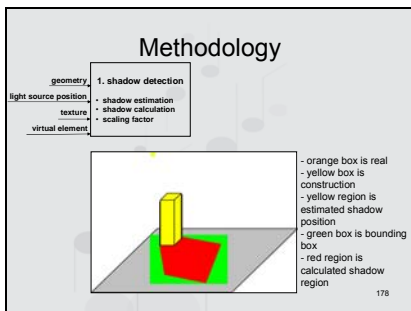
The shadow estimated from the light source estimate and the reconstructed geometry is given in yellow.



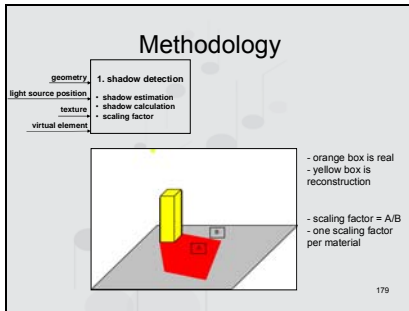
From this estimate a bounding box can be created that is larger than the estimate and will contain the real shadow.



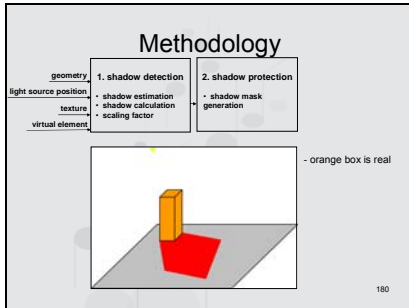
A canny edge detector is used to calculate the position of the shadow contour. This implies that only hard shadows can be detected.



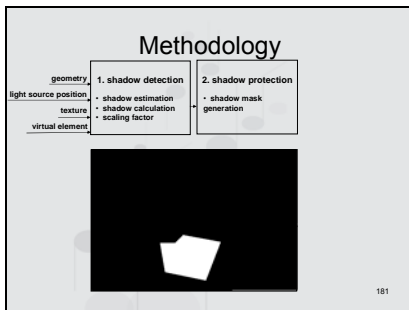
A region filling algorithm is used to find the shadow pixels based on the shadow contours.



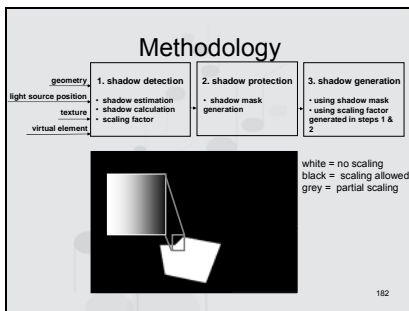
A scaling factor is calculated. This scaling factor relates the intensities of the pixels in shadow with those not in shadow. It is this scaling factor that will be used by the shadow volumes method to generate the virtual shadows. This scaling factor actually varies per material and per pixel, but in the presented method only one scaling factor was used PER MATERIAL. The results, though being an approximation, were satisfying.



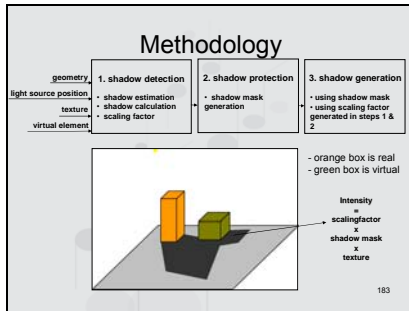
Now, the third step of the algorithm deals with the protection of these real shadow pixels from being scaled during the shadow generation with the shadow volume method.



.. this protection is performed by generating a shadow mask. An example of a binary mask is given here. White indicates that the pixels is inside a real shadow, and hence no scaling is allowed, black indicates that scaling is allowed.



Sometimes the real shadow is not a true hard shadow. To reflect the difference in shadow intensity in penumbra and umbra regions, the shadow mask is in grey scale and a gradient from white to black is introduced along the border of the real shadow. Now, in this shadow mask, white indicates that the pixels cannot be scaled; black indicates the pixels can be scaled and grey indicates they can be scaled, but with a smaller scaling factor than that used for pixels in the black region.



The fourth and final step in the methodology generates the virtual shadows. Now the intensity of a pixel inside the virtual shadow is calculated as the original pixel intensity multiplied by the scaling factor and multiplied by the value in the shadow mask. As we can see, with this method only pixels outside of a real shadow will be scaled.

### Urban planning: simulating common illumination

1. Introduction
2. Related work
3. Methodology
- 4. Results**
  - 1. Demo 1**
  - Demo 2
  - Conclusion

184

### Results

- **Characteristics:**
  - light source static
  - real scene static
  - everything recorded offline
  - RT shadow generation and scene navigation
- **Task flow:**
  - 0. input generation
  - 1. shadow detection
  - 2. shadow protection
  - 3. shadow generation

repeat every frame

Video

In the first demo, the light source and the real scene are static and the texture is captured offline. The shadow generation and scene navigation is calculated online and in RT. Basically this means that the first three steps of the methodology are calculated once, and the final step is repeated every frame. This type of applications performs the shadow detection offline and therefore the necessity of calculating the real shadow positions in real time is less important.

### Results

- **Characteristics:**
  - light source semi-static
  - scene static
  - everything recorded online
  - camera calibration online using ARToolkit
- **Task flow:**
  - 0. input generation
  - 1. shadow detection
  - 2. shadow protection
  - 3. shadow generation

repeat every frame

Video

In the second demo, the light source is semi-static and is allowed to move in a close range around the real light source position. The real scene is static, but the texture is captured at run time. The camera calibration was executed using ARToolkit. In short this means that the input generation can be done offline and only once, but the shadow detection, protection and generation is repeated every frame.

### Conclusion and future work

- **Conclusion:**
  - A real-time shadow detection and generation is presented
  - Common illumination for augmented reality
- **Future work:**
  - Improved shadow detection system (soft shadows)
  - Light source tracking
  - Improved per-pixel scaling factor calculations

187

A real-time shadow detection and generation method is presented, which provided common illumination for augmented reality.

There are a few things that can be improved in the current implementation of the method. For instance, we have only considered semi-hard shadows, but the application field of the method could be broader if a soft shadow detection would have been used in the shadow detection step. Secondly, a light source tracking could be included. And finally the scaling factor calculation could be done per pixel instead of per material, as this would better reflect that pixels inside a shadow cast on one object can have a slightly different appearance due to the orientation of the object towards the light source, or the indirect light the points receive from their neighbouring objects.

### Indoor Lighting Design

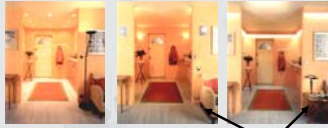


**Celine Loscos**  
**Katrien Jacobs**  
 University College London  
 Virtual Environments  
 and Computer Graphics



188

### Indoor Lighting Design



Lighting modifications      Geometric modifications

Can this be done virtually?

189

When designing interior one would try to imagine different scene set up and lighting set up. It is tedious and often impossible to do it in real life. We could build a virtual model of the room and use traditional computer graphics lighting simulators. But many details and the overall impression would be lost. Could this be done virtually, based on the existing room?

### Different approaches

- **Common goal:**
  - Realistic output
  - Use of global illumination

190

Different approaches exist. All have a common goal to produce a realistic output – meaning that the output can be comparable to what would a photograph of the redesigned room would look like. Global illumination is preferable as it models light transport more realistically.

### Different approaches

- Approach 1: [Loscos00]
  - Input:
    - Partial geometry known
    - One lighting condition known
    - Few input photographs
  - Output:
    - Approximate solution
    - Fixed viewpoint
    - Interactive control

191

We will present 4 different approaches to the problem. The first one was published by Loscos et al. in 2000. Its advantages are that the input can be restrictive. But as a consequence the solution is approximate.

### Different approaches

- Approach 2: [Loscos99]
  - Input:
    - Partial geometry known
    - Several lighting conditions known
    - Few input photographs
  - Output:
    - Approximate solution
    - Fixed viewpoint
    - Interactive control

192

A second approach was published by Loscos et al. in 1999. The input is more restrictive as several lighting conditions need to be known. The results are therefore more reliable than those of the previous approach.

### Different approaches

- Approach 3: [Yu99]
  - Input:
    - Full geometry known
    - One lighting condition known
    - Many input photographs
  - Output:
    - Accurate solution
    - Pre-computed results
    - Possible of changing viewpoint before computation

193

A third approach is of Yu et al. published in 1999. The input needs to be very complete, but the provided solution is very accurate.

### Different approaches

- Approach 4: [Boivin01]
  - Input:
    - Full geometry known
    - One lighting condition known
    - One input photographs
  - Output:
    - Nearly accurate solution
    - Pre-computed results
    - Possible of changing viewpoint before computation

194

The last approach presented is of Boivin et al. in 2001. The output is similar to the one of Yu et al., but the input needs to be less complete.



## Approach 1: Loscos et al. 00

- Input:
  - Partial geometry known
  - One lighting condition known
  - Few input photographs
- Output:
  - Approximate solution
  - Fixed viewpoint
  - Interactive control

195

## Relighting using a single known lighting



Original photograph

Relighting

196

For an original photograph on the left hand side, this method allows to reach results as in the picture shown in the right hand side, where all original light sources were virtually turned off and a new virtual light was inserted. In addition a virtual object was inserted in the scene.

## Relighting using a single known lighting

- Input data
  - 3D-model with textures
  - single known viewpoint
  - single photograph
  - single known lighting



197

A scene is reconstructed from photographs, on which textures taken from one view are mapped.

## Global illumination estimation


- Initialisation of lighting parameters  
*[Fournier-Gunawan-Romanzi93] [Drettakis-Robert-Bougnoux97]*
  - reflectance estimate
    - from pixels of the textures
  - estimation of the light sources exitance
- Initialisation of the lighting system
  - hierarchical radiosity system

198

The technique builds on Fournier et al, and Drettakis et al. to initialise the parameters (initial reflectance, light source exitance). The method also uses a hierarchical radiosity system.

### Global illumination estimation

- Initial radiosity solution (without textures)
  - radiosity for each mesh element




199

From the initial parameters a radiosity solution is computed.

### Adding virtual objects

- Shadow of the virtual object on the table
- Modulate the texture with a display ratio  
[Drettakis et al. 97]

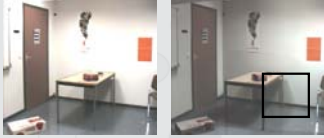
Modified radiosity  
Initial radiosity



To make light changes perceptible after virtual modifications, textures are modulated by a factor of the modified radiosity over the initial radiosity.

### Limitations

- Problem [Drettakis-Robert-Bougnoux97]: shadow already included in the textures



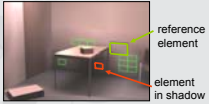
Left real light source virtually switched off

201

But this doesn't allow relighting.

### Creation of new textures without shadows

- For each element in a shadow region
  - add light previously blocked by a real object
  - correction from a reference element chosen automatically

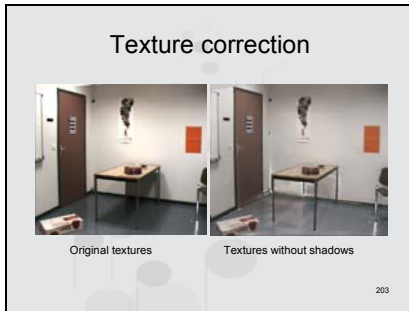


reference element

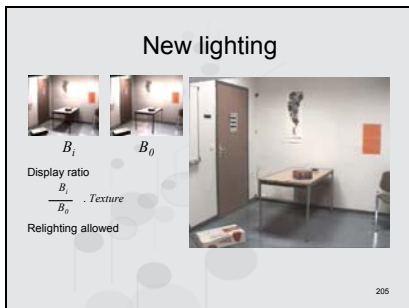

element in shadow

202

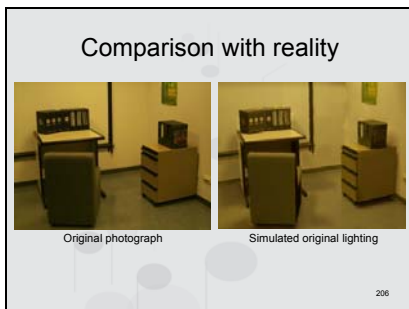
In order to allow relighting new textures are created on which shadow effects are removed.



We can update now a new radiosity  $B_0$  corresponding to the textures without shadows.



The original lighting or any new lighting can then be reproduced by applying a factor to the textures without shadows. The factor is based on the new calculated radiosity over the radiosity without shadows.



## Comparison with reality




Photograph with new lighting

Simulated relighting

207


## HDRIs against RGB format



Textures without shadows



Original photograph



RGB

radiance

208

Using HDRIs helps producing more accurate results.

## Interactive relighting

- Light source exitance modification
  - fast solution
- Add/Move virtual objects and light source
  - incremental algorithm [Drettakis et Sillion 97]

**Videos**

209

## Approach 2: Loscos et al. 99

- Input:
  - Partial geometry known
  - Several lighting conditions known
  - Few input photographs
- Output:
  - Approximate solution (diffuse reflectance only)
  - Fixed viewpoint
  - Interactive control

210

### Input data

- Radiance images from a single viewpoint
  - a single light source per image

Different lighting conditions

211

In a second approach, we will use the fact that we can retrieve reflectance parameters easier when several lighting conditions are known.

### Reflectance estimate pixel per pixel

- For each radiance image
  - reflectance = radiosity / (direct light + indirect light)

Original photograph      Estimated reflectance

- Indirect approximated by an ambient term
  - Leads to errors

212

The reflectance estimate is done pixel per pixel using an approximation of the calculation of the direct lighting and indirect lighting. Because of the approximations the reflectance values are very inaccurate.

### Use of Weighted Average to Merge the Reflectance Values

reflectance      confidence

merged reflectance

213

We use the combination of images to reduce the errors. We assign to each image a confidence value that remove areas in shadows (more prone to errors because of the approximation on the indirect lighting), and for highlights since only diffuse reflectances are treated.

### Lighting simulation

- Direct lighting: Ray casting
- Indirect lighting: optimised radiosity solution

pixel

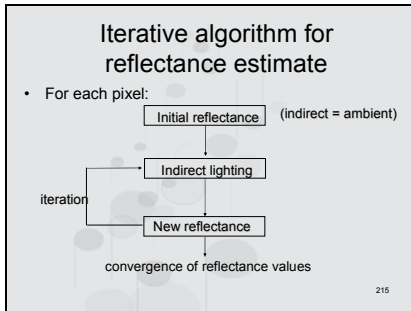
Reflectance

Direct light

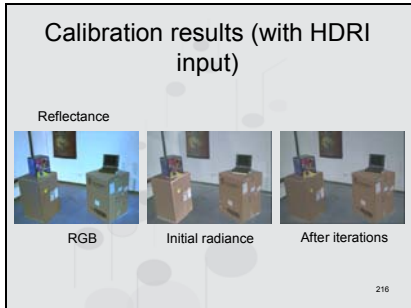
Indirect light

214

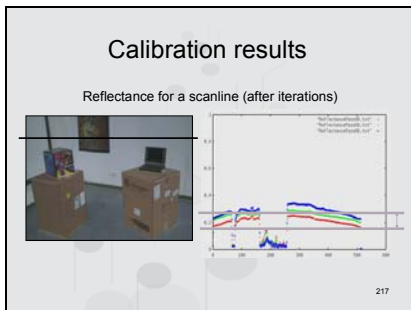
After the reflectance values are calculated, the illumination is calculated per pixel for the direct lighting. The indirect lighting is however calculated using radiosity as it is faster than ray tracing.



An iterative approach is set to update the reflectance values with the indirect lighting calculation improving through the iterations.



The calculated reflectance is shown on those images. HDRIs images help producing results with less visible errors.



The reflectance values are shown in RGB for a scan line. For a same material, values are consistent.



### Approach 3: Yu et al. 99

- Input:
  - Full geometry known
  - One lighting condition known
  - Many input photographs
- Output:
  - Accurate solution
  - Pre-computed results
  - Possible of changing viewpoint before computation

219

### Input

- Reconstruct scene geometry and light source positions



220

In another approach, it is possible to deal with more complex BRDFs and lighting conditions.

### BRDF estimation

- Diffuse albedo estimated using inverse radiosity
- Roughness and specular albedo estimated using non-linear optimisation:
  - based on a large set of images from different viewpoints and specular highlights identified in these images

221

In this method, the diffuse reflectance is computed using a radiosity method.

Whereas the other BRDF parameters are estimated using a non-linear optimisation

### Results

- Rendering using RADIANCE




222

A great advantage of this method is that the computed BRDFs are accurate. The scene can therefore be relighting with very good results.

### Approach 4: Boivin et al. 01


- Input:
  - Full geometry known
  - One lighting condition known
  - One input photograph
- Output:
  - Nearly accurate solution
  - Pre-computed results
  - Possible of changing viewpoint before computation



223

### Method Overview

- Reconstruct scene geometry, light source position
- BRDF model according to Ward et al. : 3 or 5 parameters: diffuse, specular, roughness



224

The method of Boivin et al. is similar to the one of Yu et al.

### Method Overview

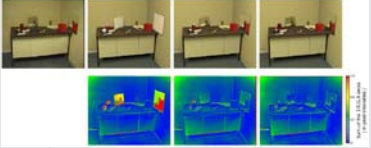
- Iterative approximation based on one single image:
  1. start from initial BRDF estimate
  2. render scene using BRDF, light sources and geometry
  3. calculate difference with input image
  4. if difference too large:
    1. make new BRDF estimate, or
    2. refine BRDF model
    3. go to step 2
  5. else stop!

225

It uses an iterative algorithm but it is based on a single image (rather than a set as in Yu et al.)

### Results


- Output reflectance at each iteration step




226

After each iteration, the errors in computation decrease.

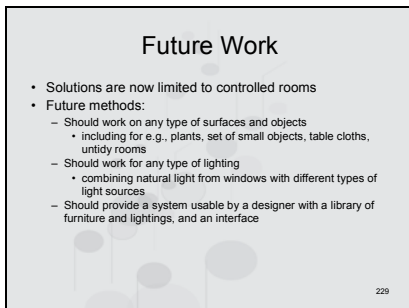




With this method it is possible to modify the view point by completing the unknown regions with the calculated materials.

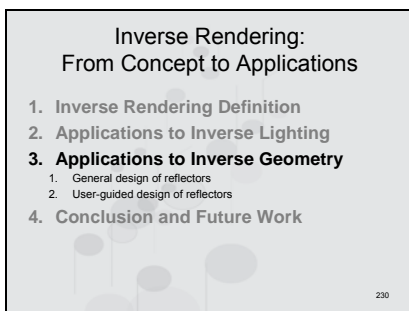


We have presented four different approaches to calculate indoor relighting. The different solutions range from accuracy to interactivity and practicability. All require a fair understanding of the geometry of the scene.




Relighting methods are still at their early stage since for each presented method, the rooms were chosen carefully with controlled geometry and lighting conditions (no windows for example).

It is expected that a new range of methods should come up to allow more complex scenes and conditions to be treated.




In this Section, applications are given of inverse geometry.

### Applications to Inverse Geometry: Surface Design



Gustavo Patow  
Xavier Pueyo

Grup de Gràfics de Girona  
Universitat de Girona



231


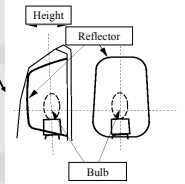
Keywords: Inverse Problems, Inverse Rendering, Inverse Geometry, Reflector Design, CAD for luminaries, Optimization

### Applications to Inverse Geometry: Surface Design

- Introduction
- Previous Work
- Problem Statement
- Foundations
- Numerical Aspects of Inverse Surface Design
- General design of reflectors
- User-guided design of reflectors
- Conclusions
- Future Work

232

### Target Optical Set

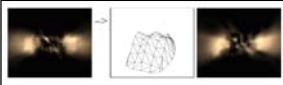



233

Firstly, we should define what an *optical set* is. The reflector shape to be found is just a piece of a set called in lighting engineering an *optical set*, which consists of a light bulb, the reflector and, possibly, the diffusor. The reflector has a border, contained in a plane, that limits its shape. In general, a reflector must fit inside a holding case, so its shape cannot be lower at any point than the plane defined by the border nor higher than a certain threshold defined by the case. We can say that the case defines a bounding box for the reflector.

### Objective

- Design of reflector shapes from prescribed optical properties and geometrical constraints.



234

We will focus our efforts on the following problem: Given a desired optical set-outgoing radiance distribution, find the corresponding shape for the reflector that provides a light distribution similar to the given one up to a user-defined threshold.

### Restrictions (I)

- Surface must be "build-able":  
 $z = f(x,y)$

z  
x/y

Metallic sheet

235

The following constraints are imposed on the surface shape to be built:  
The shape must satisfy certain constructive constraints that amount to requiring that the shape of the reflector be the graph of a function with respect to the plane of the reflector's border.

### Restrictions (II)

- Shape must fit exactly into the border

236

The resulting shape must exactly fit the given border because the case holds the reflector at that curve.

### Restrictions (III)

- Shape can't be lower than border plane (should leave space for light bulb)
- Shape can't be higher than holding case

Holding Case

Border Plane

Light bulb

237

The shape cannot be lower than the border plane ( $z=0$ ), or higher than a certain maximum height (it must fit inside the case).

### Previous Work: Inverse Surfaces

Illumination Distribution

Reflector

Source

BRDF

238

This problem is clearly within the Inverse Geometry problems we mentioned before

### Previous Work: Our case


- To the best of our knowledge, there are no other works dealing with:
  - Inter-reflections inside the reflector
  - General BRDFs
  - Polygon-based surfaces
  - A global strategy for the optimization

239

To the best of our knowledge, there are no other works dealing simultaneously with the treatment of inter-reflections inside the reflector geometry, using general BRDFs for the reflector surface material, a polygonal-based representation of the surface, and finally, a global strategy for the optimization.

### Problem Statement

- Given
  - The radiance distribution of the light bulb
  - A reflector border
  - A desired outgoing radiance distribution for the Optical Set
- Find a shape within a user-prescribed tolerance



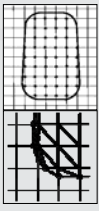
240

We can state more precisely our problem as: Given the outgoing radiance distribution of a light bulb and a reflector border, and given a desired optical set-outgoing radiance distribution, find the corresponding shape for the reflector. Do this up to a user-defined tolerance.

### Surface Representation

Regular Grid:

- + Avoids optical striations
- + Easy to manufacture
- + Allows studying the procedure
- + Border fitting almost trivial
- Polygonal approximation
- $C_0$  Continuity
- Not very flexible
- Industry requires a fine subdivision.



241

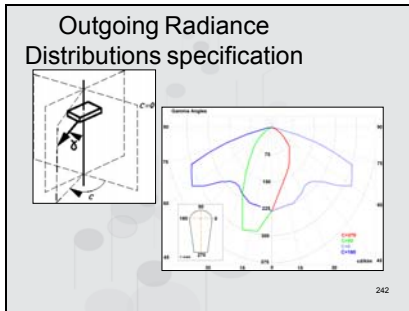
We choose to implement the regular grid representation/scheme because:

Allows to avoid undesirable striations since images (in the optical sense) are formed behind the reflector [CM97]. An accurate template or tool to build a physical reflector from the resulting data can easily be made [CM97]. The overall convergence and stability of the algorithm can be easily tested.

The routines for border matching are very simple: Only the grid points *inside* the border are available for optimization (the ones marked with dots in The Figure above). When assembling the surface, the polygons that join the grid points that define the surface in the interior of the border with the border itself are created, making a polygonal approximation to the originally given border, see Figure. In order to do this, the border is intersected with the grid lines and the intersection points (marked with circles in the lower image) are used to build the polygons. Those polygons are created following a pattern like the one shown in the Figure.

On the other hand, regular grid representations present the following disadvantages  $C^0$  continuity at *all* the triangle lines in the Figure, that are all over the surface. To be flexible to adapt to the illumination requirements imposed by the  $\text{OutRad}_{\text{desired}}$  radiance distribution, the surface built this way needs many vertices, which increases computing times. A fine grid is needed to achieve the smoothness needed to achieve the manufacturing standards of the industry.

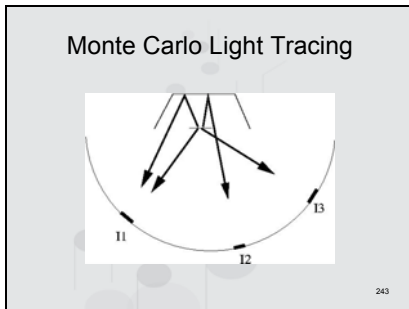
[CM97] J. R. Coaton and A.M. Marsden, “*Lamps and Lighting*”, Ed. Arnold, London, 1997



We choose to use the  $C$ - $\gamma$  coordinate system [CM97] as our discrete representation for the outgoing radiance distributions, because it represents a standard in the lighting engineering industry. This coordinate system is depicted in the Figure. It is worth mentioning that the  $C$ - $\gamma$  representation is equivalent to any other spherical coordinate system (see [Cohen93-RRIS] Section 10.1.4, General Luminaries).

As we can see, the  $C$ - $\gamma$  representation is an angular representation given by two angles,  $c$  and  $\gamma$  ( $c \in [0, 2\pi]$ ,  $\gamma \in [0, \pi]$ ).

[Cohen93-RRIS] Michael F. Cohen and John R. Wallace, “*Radiosity and Realistic Image Synthesis*”, Academic Press Professional, Boston MA, 1993, ISBN 0-12-178270-0



This step has as a main objective the computation of the outgoing  $C$ - $\gamma$  radiance distribution from a given reflector shape.

Light rays are sampled at the light source according to its energy distribution per solid angle, shooting more rays in the direction where more light is emitted. Then, ray paths are followed until they leave the optical set and then their intensities are recorded in the corresponding entry of the  $C$ - $\gamma$  distribution.

- ### Monte Carlo Light Tracing
- Advantages
    - Accurate
    - No Lost Rays
    - General BRDFs
    - Interreflection
    - All kind of optical phenomena
  - Drawbacks
    - Slow
  - Best among ray-based options
- 244

For this, several possibilities were studied, arriving to the conclusion that Monte Carlo Light Tracing is the best option:

This algorithm is, by far, the most efficient and straightforward to code, as it propagates the radiance from the light sources to the rest of the optical set [Lafortune96-MMMCA]. In this method, each fired ray contributes to the final result and there are no lost rays. Thus, we get an optimal usage of the computational effort of tracing rays, since each ray propagates from the light source and each of them contributes significantly to the final result.

On the other hand, is important to notice that this algorithm can easily accommodate general BRDFs (as long as then can be sampled), interreflections and other kind of optical phenomena. For example, if we would like to add the glass diffusor, it would be very easy to include refractions in our current model. Unfortunately, if we want a small variance, we must use many rays for the computations, making this algorithm a slow choice.

It is important to take into account that the *Light Tracing* algorithm was developed with extended, non-isotropic light sources in mind, uniformly sampling locations over the light bulb surface. Thus, although all our experiments were done with point light sources, our algorithms should perform equally well on photometric lamp data provided

by near-field goniophotometers (see [RadiantImaging]).

[Lafortune96-MMMCA] Eric Lafortune, “*Mathematical Models and Monte Carlo Algorithms for Physically Based Rendering*”, Ph.D. thesis, Department of Computer Science, Katholieke Universiteit Leuven, Leuven, Belgium, February, 1996.

[RadiantImaging] Radiant Imaging Photometric, URL: <http://www.rading.com/>

### Error Definition

- Tested several options:
  - Drettakis (% of pixels above threshold)
  - $l^n$  ( $l^2, l^1, l^\infty$ )
- We are going to use the  $l^2$  metric + penalization terms

$$Error_2 = \sqrt{\sum_{i,j} |OutR(\eta)^{ij} - OutR^{ij}_{desired}|^2}$$

Parameters that define a reflector shape

245

The error definition is, by far, one of the most delicate parts of an optimization algorithm. This error is what drives the optimization and is being evaluated every time we need to assess the behavior of any reflector, and so, its definition is crucial for the right behavior of the algorithm. We decided to try several options, among which are those listed in the slide, but with poor success. We finally resorted to the well known  $l^2$  error measure to drive our optimization process.

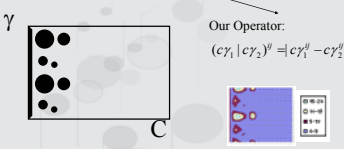
This error is just computed as the square root of the element-wise difference between a light distribution generated with the current reflector (whose shape is given by the parameter vector  $\eta$ ) and the desired one, provided by the user.

To take into account the restrictions imposed by the bounding box, we added penalizing terms to this error, thus changing the problem into an unconstrained one.

### Algorithm Foundations (I)

- Let  $\Delta_i$  be a vertex influence region:
 
$$\Delta_i = OutR(Ref) \setminus OutR(Ref_i)$$

Our Operator:

$$(c\gamma_1 \setminus c\gamma_2)^{ij} = |c\gamma_1^{ij} - c\gamma_2^{ij}|$$



246

In order to be able to explain the behavior of the function and of our algorithm in simple terms, it is convenient to introduce some further definitions.

One of the most important ones is the *differential vertex influence* region,  $\Delta_i$ . This  $\Delta_i$  is computed for the  $i$ -th vertex with our specifically defined operator “ $\setminus$ ”, which is the light distribution that arises as the element-wise absolute value of the difference of two light distributions. Thus,  $\Delta_i$  is defined as the (positive) difference between the light distribution computed for the current reflector, and the light distribution computed for the *same* reflector, but with one of its vertices (the  $i$ -th one) displaced a little bit.

An schematic example of this can be found in the figure at the left.

On the lower right hand side we can find an actual result of a difference between two light distributions, where dark blue means small values (close to zero), brown intermediate values, yellow higher ones and light blue the highest ones.



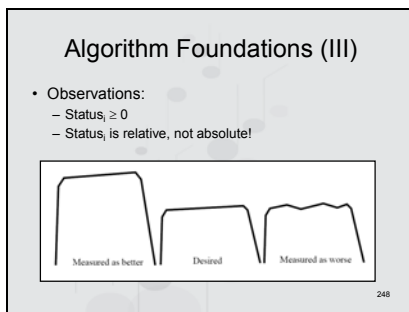
Here we can see how the different operators work together.

Starting from the upper right image, which is the desired light distribution provided by the user, we can compute its difference (with the “|” operator) with the light distribution generated by the current reflector (called “Curr”), getting the light distribution on the upper row of the middle column (Curr | Reference). This distribution shows us the error between the current reflector and the ideal one we are looking for.

We can compute  $\Delta_a$  by making the same operation between  $\square\text{Box}\square$  and the light distribution generated by the same reflector as  $\square\text{Curr}\square$ , but one of its vertices shifted a little bit (Curr | Curr<sub>A</sub><sup>5mm</sup>). This time this was 5 millimeters upwards. This computation shows us the contribution of this particular vertex to the current light distribution.

Now, if we perform an element-wise multiplication of the just computed light distributions ( $\Delta_a$  and Curr | Reference), we get the light distribution on the last column, which shows us the area of contribution of this particular vertex to the overall error.

If we add the values in this light distribution, we get a quantity we called the *status<sub>i</sub>*.




The *status<sub>i</sub>* provides us valuable information, so it is interesting to look at it a little bit more.

The first thing to notice, is that the *status<sub>i</sub>* is a positive defined quantity, that is, is always bigger than zero.

The second one is that it is a *relative* measure, because it measures local shape, not absolute position deviations.

That is, given the reflector shown in the middle, *status<sub>i</sub>* would present the vertices of the reflector on the left hand image as much better than the vertices at the right hand one, as the *shape* is much closer on the left than on the right, although the absolute positions of the vertices are the other way round.


Thus, *status<sub>i</sub>* is **not** a measure we can use to drive our optimization. We only can get hints on where to put more effort, but the global optimization must be still computed by the  $l^2$  error we presented a few slides ago.




We can perform a similar computation with two  $\Delta_a$  and  $\Delta_b$  vertex influence regions (left column), and by element-wise multiplication we can get the product  $\Delta_a * \Delta_b$ , which shows the superposition between both light distributions.

As we can see, there are brown areas in the image, which means a non-zero overlapping area. We can sum those values, and perform that for every pair of vertices. In the lower right image we can see a graph for one particular vertex, and we can observe that we have **no** zero values, meaning that every pair of vertices have overlapping areas. That means that every vertex influences the same regions as every other vertex, thus showing us that we can

not treat a vertex in isolation, as the other vertices would give us a wrong picture during the optimization. We can conclude that our **degrees of freedom**, our vertices, **are strongly coupled**. This is a very bad situation for optimization algorithms which try to optimize each degree of freedom separately, hoping to converge at the end.



Before working on any optimization algorithm, we should make ourselves an idea of the shape of the function we are trying to optimize. In the left picture we have the case when we move the full surface upwards and downwards, with a perfect specular BRDF. In the right picture we move just one vertex, but with a classic Phong BRDF this time. What is important of those pictures is that they both have a clear global minimum, but surrounded by many local minima, which is a very unfortunate situation for local optimization algorithms, as they tend to get trapped by those local minima instead of reaching the deeper global one.

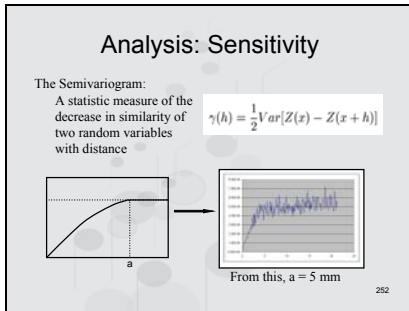


Another point to take into account is the fact that we are using a Monte Carlo (MC) based algorithm for the light propagation evaluation. As is well known, every MC evaluation has an inherent variance associated with it, giving us values that are around the real one but within the associated variance. Of course, the more rays, the more accurate our simulation would be.

For example, if we draw a part of the complete function with  $1e5$  rays (left), we would get a picture which shows a clear MC noise all over its shape. But, if we draw the same function, this time with 100 times more rays, we will get the picture on the right, showing much more clearly the real shape of the function.

So, let's suppose, to start, that we want to perform our optimization only with  $1e5$  rays to fasten calculations. But, we can see from the picture that, if we do movements of about 1 or 2 millimeters (the horizontal axis is millimeters), we won't get any meaningful information, as we are optimizing within MC noise, which hides the real shape. From the left picture we could say that movements of about 5 millimeters are needed to get meaningful information. But we need a well founded way to assess the distances we should move for the small variations in our algorithm...





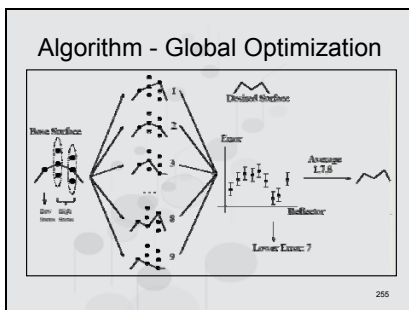
In order to do so, we can compute a *semivariogram*, which is a statistical measure of the way two random variables decrease in similarity as they get apart. The typical shape of a semivariogram is shown in the lower left, where we can define a value called range (“a” in the figure), which is the distance after which we can not get meaningful information about the underlying shape of the function. So, if we want to take a measure about the underlying function, like our differential vertex influence region, we should use values smaller than this range. Performing the measurements for our case, we observe that 5 millimeters is the value of the range. But, as we want to take meaningful measures taking into account the MC noise, we decided to use *exactly* 5 millimeters for our differential measurements. That is why we used that value in the previous slides.

- ### Classic optimization behaviour
- We have tested several alternatives:
    - Conjugate gradients
    - Powell’s method
    - Simplex method
    - Stimulated annealing
  - Vertex coupling + the variance of the MC evaluation made them fail !
- 253

When starting with the optimization stage, we tested several possibilities, those on the slide among others. Unfortunately, the combination of the strong vertex coupling plus the fact that we are performing MC-based computations (every evaluation has a variance associated which creates a noisy landscape) made all of them fail. Thus, we needed to find a different, custom optimization method for our problem.

- ### Summary of problems encountered
- High non-linearity:
    - Many local minimum
    - Local methods tend to fail
  - Strong vertex coupling
    - Can’t optimize one by one
    - Can’t optimize subsets of vertices
  - Noisy Light simulation
    - Values not exact (variance)
- 254

So summarize the problems we found, we can mention A strong non-linearity, with many local minima surrounding a clear global one. As a consequence, unless the starting point is chosen extremely wisely (which is very difficult), local methods tend to fail. Strong vertex coupling that precludes the optimization of vertices one by one, or even in small subsets! We are using a noisy light simulation, measurements have an associated variance: the values we measure are not exact, the real value could lie anywhere within the value plus/minus its variance. Although the probability distribution is a typical Gaussian, we can’t rely only on the central value.



So, the core of our optimization algorithm works as follows:  
 Sorts the available vertices by their respective status.  
 Then takes the N with highest values.  
 Assigns to each chosen vertex a span of possible positions (based on their respective status).  
 Then, we compute tests/evaluations of the performance of each member of a family of reflectors obtained by iteratively combining the addition of an increment to each one of the vertices.  
 Then, we kept the best one: the one with smallest error.  
 Unfortunately, just keeping the one with lower central

value is not right, as each value has an associated variance bar. So, we must keep the best one as well as all the other reflectors that gave an error whose variance bar overlaps with the one of the best one.

What we can do is to consider those values as being measurements of an ideal reflector we don't know, so we can average them to get the final reflector for this iteration.

```

Algorithm

Reflector := create a low-res reflector
while (not converged) and (not userDefinedStop)
  FreeVertexList := all vertices in Reflector
  WrappedRef := wrap (Reflector, FreeVertexList)
  while (not converged) and (FreeVertexList is not empty)
    addVertices (WrappedRef, FreeVertexList)
    optimize (WrappedRef)
  if (not converged)
    increaseResolution (Reflector)
  
```

Unfortunately, this algorithm has an exponential behavior each time we increase the resolution of the reflector: we increase resolution by doubling the number of vertices in each direction. This is done whenever we finish an iteration and do not reach the user defined threshold and the number of iterations do not exceed a certain defined number. So, each time we have 4 times more vertices, thus getting an exponential behavior.

In order to avoid such behavior, a strategy called wrapping is introduced. The basic idea is to expose the optimization stage only a carefully chosen subset of the available vertices, interpolating the others.

- ### Algorithm - Wrapping
- Avoids an exponential behavior
  - Consists of two parts
    - Wrapper
    - New vertex addition strategy
  - Surface Wrapper:
    - Initialized w/ vertices at previous resolution
    - Interpolates wrapped vertices from the non-wrapped ones (Akima)

Thus, the wrapping strategy consists of two parts: the wrapper itself and the strategy for the addition of new vertices every time more flexibility is needed.


The surface wrapper is initialized at each iteration with the vertices at the previous resolution/iteration.

The non-exposed vertices are interpolated by a polynomial scheme.

- ### Algorithm - Wrapping
- Addition of new vertices:
    - Only when more flexibility is needed
    - Sort unwrapped vertices by Status<sub>i</sub>
    - Choose worst N
    - For each, select wrapped neighbors
    - Add *m* with maximum free-area-coverage.
    - Remove from Free-Vertex-List

With respect to the addition (exposure) of new vertices, the basic idea is to add new vertices only when more flexibility is need. The vertices are added by sorting the current ones by their *status* and by choosing the N with worst values, and the areas surrounding them are refined (remember that *status* is a local measure of the wrongness in shape of a vertex neighborhood). Thus, we add some of those vertices, in a way controlled by the user.

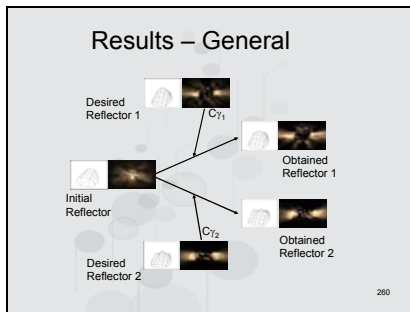
In our case, we decided to add a fixed number of vertices each iteration, having a linear behavior in the number of vertices to optimize.



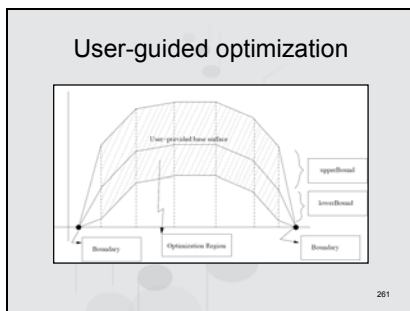
Before showing the results for this algorithm, let's study its accuracy. We can see that, as expected, the variance diminishes as the number of rays increases, a classical behavior for MC algorithms. What is interesting is that the variance also increases as the BRDF goes more diffuse. This result can be explained as diffuse BRDFs tend to "blur" the information, mixing the rays because of the high number of interreflections.

It is important to mention that this variance also


represents a bound for the performance of our algorithm. For example, if we are using the Phong BRDF defined by the diffuse to specular ratio 0.05, and we decide to use 10000 rays, we have a variance of about 0.14. This means that we will not be able to get closer to zero than 0.14, because if we get closer we won't be able to distinguish the real function of the MC noise. Thus, the variance gives us a bound on the convergence of our algorithm.



Here we can see a couple of examples: starting from an initial reflector with a corresponding outgoing light distribution (left of the slide), and given a desired light distribution like the one in the middle-lower part of the slide, we get a reflector like the one on the lower right part of the slide, with a distribution that is very similar to the desired one. If we compare the obtained reflector with the one we used to generate the desired light distribution (the reflector shown in the middle-lower image), we can see that they are quite similar. The same can be seen in the upper part of the image for a different situation. Unfortunately, this algorithm is slow: it took about 10 days to get from the left part of the slide to the right one!



So, a logical possibility is to incorporate knowledge provided by the user to the algorithm: We expect the natural user of this system to be a lighting engineer, who generally has, by experience, a reasonable idea of what the resulting reflector should look like. So, we request him to provide a starting reflector plus two confidence values that tell us his own confidence on his own guess. We are going to use this information to define bounds on the optimization region for the possible reflectors during the optimization stage.



In order to take maximum advantage of that information the user provided, we added an extra step before the global optimization described a few slides ago. Basically, what we intend to do is to modify the user-provided surface without changing its overall shape, just trying to adjust it to be as close as possible to the desired one before starting the global optimization.

So, this new step consists of two parts: Firstly, we move and stretch the surface upwards and downwards. Then, we shift the vertex positions horizontally, in the plane of the reflector border, trying to get the maximum flexibility of the vertices we have. Observe that we continue having a regular grid, we just compressed/expanded and shifted the vertices as a whole.

### User-guided: Vertex Selection

- Problem:
  - We don't have a "previous" surface to choose vertices from.
  - We can't use the full user-provided surface.
- Solution:
  - an heuristic that iteratively chooses vertices
    - With maximum coverage area
    - As far as possible to the already placed ones

$$\sqrt{(r_{x_i}^2 - r_{x_j}^2)^2 + (r_{y_i}^2 - r_{y_j}^2)^2}^{1/4}$$

263

Now, before we can start with the global optimization, we must face another small problem: In the previous version of the algorithm, we initialized the wrapper with the vertices used at the previous iteration. But now, we *don't have* a previous iteration to start with. So, we must find a way to choose the initial vertices for the first iteration. We decided to use a simple heuristic that iteratively chooses vertices with maximum free area (without other vertices) around them, being at the same time as far as possible to the already placed ones. Distances between vertices are measured with the standard 2D distance.

### User-guided: bounded optimization

264

The new algorithm works as follows: We start from a user-provided initial surface, plus two confidence values. We perform firstly a vertical accommodation step and then a horizontal accommodation one. Finally, we start a global optimization as before, but taking into account the original bounds we were provided, not generating any vertex outside this allowed region.

### User-guided : Test Cases

265

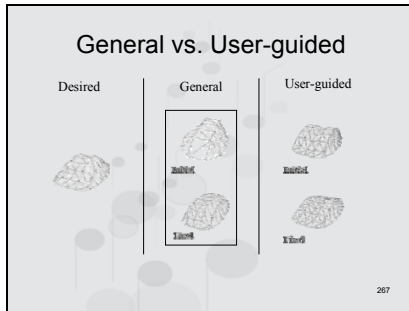
As test reflectors, we built for each desired one (see upper row), a set of user-provided starting reflectors with different geometric properties (other rows).

### General vs. User-guided

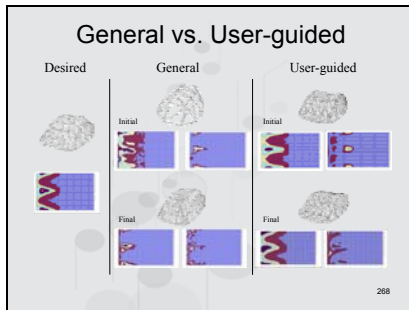
DIFFUSE	rays	Initial Error	Final Error	Time	
Diffuse to specular ratio	B	1000	0.000 ± 0.000	0.000 ± 0.000	100.000
		100	0.000 ± 0.000	0.000 ± 0.000	2.000
		10	0.000 ± 0.000	0.000 ± 0.000	0.100
	BOLD	1000	0.000 ± 0.000	0.000 ± 0.000	100.000
		100	0.000 ± 0.000	0.000 ± 0.000	2.000
		10	0.000 ± 0.000	0.000 ± 0.000	0.100
	BOLD	1000	0.000 ± 0.000	0.000 ± 0.000	100.000
		100	0.000 ± 0.000	0.000 ± 0.000	2.000
		10	0.000 ± 0.000	0.000 ± 0.000	0.100
	I	1000	0.000 ± 0.000	0.000 ± 0.000	100.000
		100	0.000 ± 0.000	0.000 ± 0.000	2.000
		10	0.000 ± 0.000	0.000 ± 0.000	0.100

266

As we can see from this table, times were reduced from the order of 10 days to the order of only one. Also, by observing the fourth column, we can see the error is much smaller than the previous version of the algorithm (called "old"). But this affirmation is not completely fair, as we were *forced* to use many more rays for the new optimization as the user provided a surface close to the desired one, with an associated error that quickly collided with the variance for that number of rays. So, we must use more rays to be able to do any meaningful optimization. So, the advantages are twofold: on the one side, we reduced the computing time from 10 days to about one. On the other hand, we even were able to increase the accuracy of our algorithm!



To visually compare both algorithms, we can provide the algorithms with the desired light distribution resulting from a reflector we know (but the algorithm doesn't), shown in the right column. The general algorithm started from the one shown in the middle column, upper reflector, and got to the lower one. Instead, the second algorithm started from the reflector in the right column, upper row, and got to the lower one, which is much more similar to the desired one than the one we got from a generic starting point.



If we want to observe the outgoing light distributions and their respective errors, we can compare the respective light distributions shown below each reflector (left side when there are two distributions). By visual comparison, we can see that the resulting light distribution from the final reflector in the user-guided algorithm is much closer to the desired light distribution than the final distribution from the general algorithm. The respective error graphs are shown on the right of these distributions, and we can see that the resulting distributions from the user-guided case present a much smaller error than their general counterpart (remember: light blue mean high values, so the errors for the user-guided case are much smaller than the general ones as they do not present colors for higher values in their distribution).

- ### Conclusions
- A good solution for design of reflectors with:
    - Interreflections
    - General BRDFs
    - A polygonal surface
    - A global strategy for the optimization
  - Robust solutions:
    - From a generic start
    - With a user-guided setting
- 269

We have presented a new solution for design of reflectors taking into account multiple interreflections, a general BRDF model, a polygonal surface and a global strategy for the optimization. Also, we have presented robust solutions, both starting from a generic starting point and in a user-guided setting.

- ### Future Work
- Different reflector shape definitions:
    - Non-regular distribution of vertices + triangularization
    - Splines
  - Other error metrics (entropy)
  - Recent theoretical developments for optimization in restricted cases
  - Generalization to the near-field (in the lighting engineering sense) problem
  - Parallelism
    - Distribution of the optimization work among several CPUs
    - Change the **Light Propagation Stage** to a GPU-based solution
- 270

## bibliography

- **A Survey on Inverse Rendering Problems**, Gustavo Patow and Xavier Pueyo, Computer Graphics Forum 22(4), pp. 663-687.
- **Reflector Design From Radiance Distributions**, Gustavo Patow, Xavier Pueyo and Alvar Vinacua, International Journal of Shape Modeling, vol 10 (2004), number 2, pp. 211-235.
- **A Survey of Inverse Surface Design From Light Transport Behavior Specification**, Gustavo Patow and Xavier Pueyo, Computer Graphics Forum, December 2005.
- **User-Guided Inverse Reflector Design**, Gustavo Patow, Xavier Pueyo and Alvar Vinacua, Research Report TR-IIA-04-07-RR, Institut d'Informàtica i Aplicacions, Universitat de Girona, 2004.

271

## Inverse Rendering: From Concept to Applications

### Conclusions

272

### Conclusions (I)

- A definition of Inverse Rendering: Goes the opposite way of classical rendering
  - Based on an image (or a lighting description) and some extra information, the missing information in the radiance equation can be calculated.
  - Depending on the amount/type of information available, different strategies apply to calculate the missing information.

273

### Conclusions (II)

- Classification of existing problems
  - Inverse lighting problems
  - Inverse reflectometry problems
  - Combined lighting-reflectometry problems
  - Inverse surface design problems
- Classification of existing solution methods
  - For general applications
  - For mixed reality

274

### Conclusions (III)

- From theory to practice is not straightforward due to:
  - Capturing failure
  - Uncontrollability of scene
  - Complexity of the scene

275

### Conclusions (IV)

- Some solutions to overcome these problems were presented
  - Capturing illumination in complex (e.g.: outdoor) scenes
  - HDR/ construction with moving objects
  - Common illumination with inaccurate geometry and lighting
  - Indoor lighting design (accuracy vs. interactivity)
  - Inverse reflector design
- But further research is required!

276

# A Survey of Inverse Rendering Problems\*

Gustavo Patow and Xavier Pueyo

Grup de Gràfics de Girona  
Institut d'Informàtica i Aplicacions,  
Universitat de Girona,  
Campus de Montilivi,  
E-17003 Girona, Spain.  
E-mail: {dagush,xavier}@iia.udg.es

---

## Abstract

*Inverse rendering problems usually represent extremely complex and costly processes, but their importance in many research areas is well known. In particular, they are of extreme importance in lighting engineering, where potentially costly mistakes usually make it unfeasible to test design decisions on a model. In this survey we present the main ideas behind these kinds of problems, characterize them, and summarize work developed in the area, revealing problems that remain unsolved and possible areas of further research.*

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture I.4.7 [Image Processing and Computer Vision]: Feature Measurement I.4.8 [Image Processing and Computer Vision]: Scene Analysis

---

## 1. Introduction

Inverse problems are usually of an extreme complexity and are emerging as an important research topic for the graphics community due to their interest in a wide range of fields including lighting engineering and lighting design. Although progress in rendering to date has mainly focused on improving the accuracy of the physical simulation of light transport and developing algorithms with better performance, some attention has been paid to the problems related to inverse analysis, leading recently to interesting results.

In computer graphics this sort of problem is not completely new: the problem of inverse kinematics has been widely applied for animation <sup>WW92</sup>, and an excellent survey on inverse placement of cameras, curves and objects for animation can be found in <sup>Kas92</sup> with references therein.

We can say that inverse problems infer parameters of a system from observed or desired data which define their behavior, in contrast to direct problems which, given all the parameters, simulate the effects. Tradi-

tional direct problems in lighting involve the computation of the radiance distribution in an a priori, completely known environment (geometry and materials). These problems can be proven to be well-posed <sup>HM95</sup>. The inverse rendering problems lack at least one of the Hadamard <sup>Had02</sup> criteria for being well-posed: the solution does not depend continuously on the data, which means that small errors in measurements may cause large errors in the solution (see <sup>HM95</sup>).

Inverse Lighting Problems refer to all the problems where, as opposed to what happens with traditional direct lighting problems, several aspects of the scene are unknown. One common characteristic of this kind of problem is that, in general, we know in advance the desired illumination at some surfaces of the scene (their final appearance). Therefore, the algorithm has to work backwards to establish the missing parameters. Such a tool is of extreme importance in lighting engineering, and animators and lighting experts for the film industry would also benefit highly from it.

Inverse illumination problems are intimately related



to the computer vision field. The specific problems investigated include shape from shading<sup>HB89</sup>, direction of the luminary from images and identifying surface characteristics from an image or a sequence of images. Fortunately, in a computer graphics system the viewing parameters and the exact scene geometry (or at least part of it) are known and therefore many problems become easier to solve. Our situation also differs from the one in computer vision in that we want the user to control the illumination on a surface. As such, we expect the user to provide options and feedback into the system and thus solve ambiguities when they arise. Thus, the aim of this paper is leaning towards the Computer Graphics field, but the interested reader is referred elsewhere<sup>HB89</sup> for related developments in the Computer Vision field.

The paper is organized as follows: in Section 2 the theoretical background is given and a classification based on the rendering equation is presented. Next, in Section 3 the approaches developed for each possible problem, according to the previous classification, are explained, characterized and analyzed. In Section 6 the conclusions and open lines of research are presented.

## 2. Theoretical Background

Global illumination is related to transport theory and can be viewed as a special case of it<sup>CW93, SP94</sup>. The behavior of transported light is characterized by the properties of the particles (photons) when traversing the environment. Global illumination's most fundamental magnitude is radiance  $L(\mathbf{r}, \omega)$  which is defined as the power radiated at a given point  $\mathbf{r}$  in a given direction  $\omega$  per unit of projected area perpendicular to that direction per unit solid angle for a given frequency ( $Watt\ m^{-2}\ sr^{-1}$ ).

The boundary conditions of the integral form of the transport equation are expressed as

$$L(\mathbf{r}, \omega) = L_e(\mathbf{r}, \omega) + \int_{\mathcal{S}_i} f_r(\mathbf{r}, \omega, \omega_i) L(\mathbf{r}, \omega_i) \cos \theta d\omega_i \quad (1)$$

for points  $\mathbf{r}$  in surfaces, being  $f_r$  the bidirectional reflection (and/or transmission) distribution function (BRDF),  $\theta$  the angle between the surface normal at  $\mathbf{r}$  and  $\omega$ ,  $\mathcal{S}_i$  the hemisphere of incoming directions with respect to  $\mathbf{r}$  and  $\omega_i$  an incoming direction.

This classical governing equation can be concisely expressed as a linear operator equation<sup>Arv95a, Arv95b</sup>. First, define the *local reflection operator*  $\hat{K}$  by

$$(\hat{K}h)(\mathbf{r}, \omega) \equiv \int_{\mathcal{S}_i} k(\mathbf{r}; \omega' \rightarrow \omega) h(\mathbf{r}, \omega') d\mu(\omega')$$

which accounts for the scattering of incident radiant

energy. Here  $h$  is a field radiance function, corresponding to all incident light. The  $\hat{K}$  operator maps the incident light distribution onto the corresponding exiting light distribution that results from *one local reflection*.

Next, we can define the *field radiance operator*  $\hat{G}$ , that transforms an exiting light distribution into the incident light distribution that results from surfaces illuminating one another:

$$(\hat{G}h)(\mathbf{r}, \omega) \equiv \begin{cases} h(\mathbf{p}(\mathbf{r}; -\omega), \omega) & \text{when } \nu(\mathbf{r}, \omega) < \infty \\ 0 & \text{otherwise} \end{cases}$$

where  $\nu(\mathbf{r}, \omega)$  is the *visible surface function* and is defined<sup>Arv95b</sup> as  $\nu(\mathbf{r}, \omega) \equiv \inf\{x > 0 : \mathbf{r} + x\omega \in \text{Surfaces in the environment}\}$ .

Defining these operators we can factor out the implicit function  $\mathbf{r}(\mathbf{r}, \omega)$  from the integral equation 1 and we may write:

$$L = L_e + \hat{K}\hat{G}L \quad (2)$$

Following the outlines in Stephen Marschner's Ph.D. thesis introduction<sup>Mar98</sup>, we can classify the different papers on inverse lighting problems according to which of the quantities of the above equation is unknown:

Direct problems are those which, given known values for  $L_e$ ,  $\hat{K}$  and  $\hat{G}$ , solve for  $L$ . But, if we have some knowledge of  $L$ , we can pose different kinds of inverse lighting problems.

If  $L_e$  is unknown, and  $\hat{K}$ ,  $\hat{G}$  and  $L$  or part of it, are known, we have a problem of *inverse lighting*: given a photograph or any other information that covers part of  $L$ , and a complete model of the scene ( $\hat{K}$  and  $\hat{G}$ ), find the emittances ( $L_e$ ) of the luminaries illuminating the scene.

If  $\hat{K}$  is unknown, and  $\hat{G}$ ,  $L_e$  and part of  $L$  are known, we must solve for information about  $\hat{K}$ . This problem can, in general, be called *inverse reflectometry*, and a particular case is the one called *image-based reflectometry* in<sup>Mar98</sup>, where images are used as input to the information about  $L$ . As described there, since  $\hat{K}$  includes information about the variance of the reflectance both spatially and directionally, this can be a very difficult problem since it can be a very complex function. Depending on the constraints imposed on the problem, we can subdivide it into the *inverse texture measurement* (constraints on the directional variation), or the *inverse BRDF measurement* (spatial uniformity is assumed). In Table 1 we present the different papers surveyed in our work classified according to this scheme.

Finally, if  $\hat{G}$  is unknown, we have an *inverse geometry* problem. For an in-depth survey on those problems, refer to<sup>PP00</sup>.

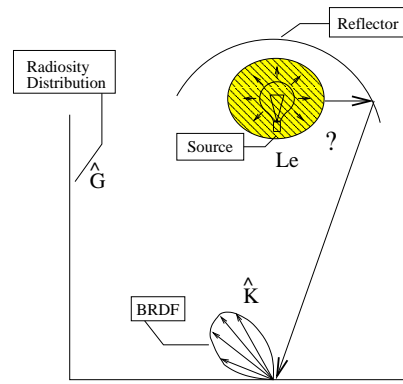
Inverse Lighting ( $L_e$ )	SDS <sup>+</sup> 93, MG97, Mar98, CSF99, SL01, MAB <sup>+</sup> 97, HMH95, FGR93, DRB97, LFD <sup>+</sup> 99, FOH98, MMOH97, OH95, RH01b, SSI99c, SSI99b, SSI99a, PF92, PRJ97, Gui00
Inverse Reflectometry ( $\hat{K}$ )	Deb98, BG01, FGR93, DRB97, YDMH99, LFD <sup>+</sup> 99, LDR00, IS91, KC94, POF98, OSRW97, DvGNK99, LKG <sup>+</sup> 01, Pou93, PF95, Mar98, MWL <sup>+</sup> 99, SI96, SWI97, DHT <sup>+</sup> 00, YM98, Mar98, RH01b
Combined Problems ( $L_e$ and $\hat{K}$ )	KPC93, RH01b, SSI99b, SSI99c

**Table 1:** Papers reviewed in this work, classified according to equation 2.

Equation 1 can be regarded in a signal processing framework<sup>RH01b</sup> under the restrictions of distant illumination, no inter-reflections, isotropic BRDFs and known geometry and camera parameters. So, the reflected light field integral is regarded as the convolution of two signals: the bidirectional reflectance function and the incident lighting; i.e. by filtering the illumination using the BRDF. Inverse rendering can simply be viewed as a deconvolution of the two signals. This framework<sup>RH01b</sup> let the authors conclude that BRDF recovery is well-conditioned (in a mathematical sense) when lighting contains high frequencies (e.g., directional sources) and is ill-conditioned for soft lighting. Alternatively, inverse lighting is well-conditioned for BRDFs with high-frequency components (specular peaks) and ill-conditioned for diffuse surfaces.

Another factor to take into account is whether the different papers treat the full global illumination equation, Equation 2, or a simpler local-illumination version based on a simplification of the illuminating equation, considering only point light sources and without considering inter-reflections. It is also important to mention the treatment of visibility in the different approaches reviewed: when computing the radiance with the above equations, the visibility problem consists of detecting if there are any blockers between the source and the surface being illuminated, and not adding their contribution in that case. The same is true for the paths from the surface to the eye or the region where the final radiance computations are needed. Most of the reviewed papers omit this treatment, arriving at solutions not applicable in real-life situations.

Other kinds of methods are based on a treatment of the problem in its global form, considering the inter-reflections of light on the whole scene, that is, in the



**Figure 1:** Diagram of Inverse Lighting Problems, where the problem is to characterize the illumination ( $L_e$ ) on a scene, either by finding the emissivities of already positioned sources, or by finding their locations in the scene

context of global illumination. Most of them work with some kind of projection space where they project the solution, transforming the integro-differential problem of computing the illumination into a matricial one.

### 3. Inverse Lighting Problems (ILP)

The inverse lighting problems are those problems where the unknown is the lighting of the scene (see figure 1 where a very simple scene is depicted with an unknown light source). These problems can be further classified into problems of inverse emittances and inverse light positioning. In the former the unknowns are the emittances of a given subset of surfaces of the scene. In the second the problem is to find the locations of the light sources (or luminaries) in order to achieve a desired illumination. Another possible classification of ILP problems naturally arises when considering the treatment each one gives to Equation 1, where the different approaches could be posed in a way such that any lighting algorithm would fit it. These are the classical F.E. radiosity setting (with BRDFs and radiance constant all over each patch surface), an inverse Monte Carlo framework or even by posing them as local illumination problems. Table 2 presents a classification of Inverse Lighting Problems papers according to these two criteria.

#### 3.1. Emittance Problem (EP)

As stated above, these problems deal with obtaining the emittance of a subset of the patches, the light sources. The different works dealing with this type of problem can be grouped according to the restrictions

	emittances	positioning
general	SDS <sup>+</sup> 93, MG97, Mar98	CSF99, SL01, MAB <sup>+</sup> 97
radiosity	HMH95, FGR93, FOH98, DRB97, LFD <sup>+</sup> 99, MMOH97	
Monte Carlo	OH95	
Local	RH01b, SSI99c, SSI99b, SSI99a	PF92, PRJ97, Gui00

**Table 2:** Classification of Inverse Lighting Problems according to whether they compute surface emittances or light source positioning. Also, the different works can be organized with respect to the treatment they give to Equation 1: General, Radiosity, Monte Carlo and Local.

they introduce to solve to Equation 1 in the following classes:

- **General formulations.** We consider  $n$  distinct light sources illuminating a scene, each of them being characterized by a function  $\Phi_i$ , which represents the isolated contribution of the  $i$ -th light source, with unit intensity, on the environment. Thus, we could describe the illumination in a scene by a linear combination of the form:  $\Phi = \sum_i^n u_i \Phi_i$ , where  $u_i$  is the non-negative weight of the  $i$ -th light on the environment. By assuming a linear relationship  $\mathfrak{R}$  with some measured intensity values  $\alpha_j$  in the scene or in a screen, we arrive at an expression of the form

$$\alpha_j = \mathfrak{R}(\Phi) = \sum_i^n u_i \mathfrak{R}(\Phi_i) \quad (3)$$

which can be regarded as a typical least squares problem. As we can see, the problem formulated this way is rendering-independent, since any approximation for Equation 1 can be used without changing the formulation.

- **Radiosity-based formulations.** Here the general problem is reduced to a radiosity setting by making the following approximation to Equation 1: a purely diffuse BRDF for the surfaces (patches), that is constant all over each one. In this case, Equation 1 is reduced to the form:

$$B_i = Le_i + \rho_i \sum_j F_{ij} B_j \quad (4)$$

where  $B_i$  is the  $i$ -th patch radiosity,  $Le_i$  its emittance,  $\rho_i$  its reflectivity (diffuse BRDF) and  $F_{ij}$  is the form factor from element  $i$  to element  $j$ . From there, re-writing it in the form of Equation 3 is triv-

ial, as shown in Section 3.1.2. In general, in those problems the patches can be grouped according to whether their radiosities  $B_i$  and their emittances  $E_i$  are known or not, resulting in a system of equations with some  $E_i$  and some  $B_i$  as unknowns.

- **Monte Carlo formulations.** An Inverse Monte Carlo method is proposed and proceeds by firing a set of rays from the surfaces with known properties towards points on surfaces with unknown properties, and gathering illumination information from surfaces of the first type to the surfaces of the second type.
- **Local Illumination formulations.** A simple local illumination model is used instead of a global one. So, Equation 1 is actually not used.

Marschner and Greenberg<sup>MG97, Mar98</sup> studied the ill-conditioning of the Inverse Lighting Problem in the case of diffuse surfaces, and later Ramamoorthi and Hanrahan<sup>RH01b</sup> presented a signal-processing framework for inverse rendering, showing that inverse lighting is well-conditioned only when the BFDR has high-frequency components (sharp specularities), and is ill-conditioned for diffuse surfaces. They show that, for the special case of a mirror BRDF, the lighting coefficients of a Spherical Harmonics decomposition of the Lighting correspond in a very direct way to the reflected light field, thus being a well-conditioned inverse problem. Instead, for Lambertian objects the lighting recovery is ill-conditioned for frequencies above the second order in a Spherical Harmonics decomposition (low frequencies). For Phong BRDFs, it is shown that inverse lighting calculations are well-conditioned only up to order of the square root of the shininess, while for the Torrance-Sparrow micro-facet model it is well-conditioned only for frequencies up to order of the inverse of the roughness.

### 3.1.1. General EP

One of the first approaches to the emittance problem was by Chris Schoeneman et al<sup>SDS<sup>+</sup>93</sup>, where the user defines the light features by “spraying” color onto surfaces. Later, Stephen Marschner and Donald Greenberg<sup>MG97, Mar98</sup> presented their *re-lighting system*, which, from a photograph and a 3D surface model of the object pictured (and a model of the camera used to take the picture), estimates the directional distribution of the incident light. In the first case, a modified Gauss-Seidel iteration was implemented to solve Equation 3, in a way such that, at each iteration, the negative values for the weights (light intensities) are clipped to zero. In this technique, the “sprayed” colors are the  $\alpha_i$ . As mentioned above, the system is ill-conditioned if the BRDF used is too diffuse, so the authors decided to add a first order linear *regularization*, and to use a generalized Singular Value Decom-

position (SVD) to allow the regularization parameter to be adjusted interactively. Here, the  $\alpha_i$  are the observed pixel values on the photographs.

In the case of Schoeneman et al., interactivity was achieved by accounting only for direct illumination in their final implementation, but the method is independent of the illumination algorithm used. Marschner and Greenberg's approach has the advantage of using a photograph as objective, rather than manually user-defined objectives, and of using a generic, scene-independent set of basis light sources. Unfortunately, this leads to a system which is more ill-conditioned than the system that comes from a set of focused light sources, thus requiring a regularization procedure.

### 3.1.2. Radiosity-based EP

The formulations presented in this subsection are closely based on the Radiosity approximations of Equation 1: assuming a purely diffuse BRDF for the surfaces (patches), and assuming it is constant all over them. In this case, Equation 1 reduces to the known:

$$B_i = Le_i + \rho_i \sum_j F_{ij} B_j$$

where  $B_i$  is the  $i$ -th patch radiosity,  $E_i$  its emittance,  $\rho_i$  its reflectivity (diffuse BRDF) and  $F_{ij}$  is the form factor from element  $i$  to element  $j$ . Re-writing this equation in the form of Equation 3 can be done trivially by defining a matrix  $M_{ij} = \delta_{ij} - \rho_i F_{ij}$  which allows us to write  $B_i = \sum_j Le_j M_{ij}^{-1}$ .

A remarkable work with these sorts of inverse problems was developed by Harutunian et al.<sup>HMH95</sup>, done in the context of radiative heat transfer. The authors observed that the resulting set of equations for the inverse problem is ill-conditioned, and thus the need to resort to the Modified Truncated SVD<sup>HSS92</sup> (MTSVD) matrix inversion method to compute  $Le_i$ . Working with the same approach (linearizing the system to solve it and MTSVD or TSVD to invert it due to its ill-conditioning), França et al.<sup>FOH98</sup> and Morales et al.<sup>MMOH97</sup> solved the inverse problem of source emissivities, this time in the presence of participating media. In the latter two cases, the problem requires the introduction of a system of equations resulting from the discretization of the medium into volume elements in order to solve the corresponding partial differential equations (PDE).

Computer Augmented Reality enables users to mix real and virtual worlds. As such, it requires the precise characterization of the geometry, source-emittances and surface reflectances of the real scene through a given set of photographs. To compute these emittances, Fournier, Gunawan and Romanzin<sup>FGR93</sup> also base their formulation on the radiosity approximation,

directly fitting the element emissions to the observed values ( $\alpha_i$ ). Drettakis, Robert and Bougnoux<sup>DRB97</sup> improved the work of Fournier et al. by setting up a hierarchical radiosity system. Loscos et al.<sup>LFD+99</sup> approximately reconstructed real scene geometry from photographs taken from several different viewpoints. There, indirect illumination is computed with a hierarchical radiosity system as before, while the direct component is calculated separately using ray-casting on a per pixel basis.

### 3.1.3. Inverse Monte Carlo EP

As is well known, radiative heat transfer problems are equivalent to lighting problems. We summarize here an interesting work on inverse radiative heat transfer, finding the temperatures of emitting surfaces, carried out by Masahito Oguma and John Howell<sup>OH95</sup>. The authors developed an Inverse Monte Carlo method. In this paper, surfaces are perfect lambertian reflectors, although the generalization to more general BRDFs seems straightforward. The method starts by letting users choose a set of points (calculation points) on the surfaces where they are interested in finding the temperature (light) distribution. Then, it randomly chooses a set of points at the surfaces with completely known properties, and for each of those points casts a given number  $N_s$  of rays. Although never stated explicitly in the paper, the points seem to be chosen following a uniform distribution, and the rays are fired by subdividing the hemisphere of directions above each surface point in  $N_s$  equal intervals, and firing one ray for each interval. These rays fly to points at the surfaces with unknown temperature (emittance). For each ray, the algorithm finds the expected energy it should carry. When enough rays have arrived, each calculation point on an "unknown" surface is assigned a temperature that is a weighted-sum of the temperatures of the rays. The weights are a set of position-related coefficients, because the exiting point of the incident intensity is often not a calculation point. Finally, the whole process is iterated until the radiance distribution on "known" surfaces led by the calculated temperature distribution of "unknown" surfaces satisfies the required heat flux distribution of "known" surfaces within a user-provided threshold.

### 3.1.4. Local Illumination EP

Based on the features presented in the introduction of Section 3.1, Ramamoorthi and Hanrahan<sup>RH01b</sup> proposed an algorithm that only recovers frequencies below a cutoff of the order of the inverse of the roughness. Thus, two possible ways are shown: solving a linear least-squares system for the lighting coefficients like Equation 3, or subtracting the diffuse component and

using the resulting mirror-like object to recover a high-resolution angular-space version of the illumination. In the second case, a two-step process is presented, where the first phase estimates the diffuse components of the reflected field from the estimated illumination frequency parameters, and the second phase does it the other way round to achieve sharper results.

Sato et al.<sup>SSI99c</sup> use the radiance information inside shadows to estimate the illumination distribution of a scene as a collection of imaginary point light sources uniformly distributed over the scene. To do this, it is necessary for the BRDF to be Lambertian, and to solve a system like Equation 3. This work was later improved by the authors<sup>SSI99b</sup>, see section 5, but the introduction of regularization using user-weighted penalty terms was required in both works, and the computational complexity limited the formulations to a coarse discretization of the sphere. Instead, the method proposed by Ramamoorthi and Hanrahan<sup>RH01b</sup> required no explicit regularization and yielded results with better sharpness and overall quality than the approaches of Sato et al. On the other hand, the methods proposed by Sato et al. are easier to extend to concave surfaces. Finally, Sato et al.<sup>SSI99a</sup> proposed the use of two omni-directional stereo images to construct a geometric model of the scene: extracting common feature points, generating a triangular mesh and finally mapping the radiance over the mesh. The radiance of the whole scene was constructed from a sequence of omni-directional high dynamic range radiance images<sup>DM97</sup> and mapped onto the constructed geometric model.

### 3.2. Light Source Positioning Problem (LSPP)

In this subsection we review the papers which perform computations on the position and orientation of light sources. As mentioned above, these works can also be classified according to the treatment they give to Equation 1.

- **General formulations.** These methods try to find luminary locations and/or orientations without relying on any particular illumination algorithm. In general, the luminary position/orientation is chosen as an optimizable variable of a certain objective function, which, in turn, is optimized with a problem-independent optimization algorithm like Stimulated Annealing<sup>PTVF92, PT91</sup>.
- **Radiosity-based formulations.** Problems which try to locate light sources in the context of the classic radiosity approximation have not been presented yet, but we believe this is possible. Unfortunately, this case corresponds to purely diffuse BRDFs, which show themselves as producers of a

severe ill-conditioning in lighting-characterization problems<sup>RH01a</sup>. The overall method would assume small polygonal emitters and use the knowledge of the illumination they produce to compute their Form Factors ( $F_{ij}$  in Equation 4). From this knowledge we should find the desired vertices as degrees of freedom of an optimization problem.

- **Monte Carlo formulations.** Although, to the best of our knowledge there are no works using this kind of formulation, we believe that this is a feasible approach, too. By firing rays from the illumination-constrained areas of the scene, it is possible to gather information and qualify areas in the scene space which may contain the sources, and by further refinement, either automatic or interactive, finding their exact location within a certain threshold.
- **Local illumination formulations.** These formulations are based on the simplification of Equation 1 to take into account only local illumination, and using the observation of this local illumination (highlights and shadows) to position the corresponding light sources. Thus, instead of Equation 1, we could introduce the expression for a point light source<sup>CW93</sup>,

$$L(\mathbf{r}, \omega) = L_e(\mathbf{r}, \omega) + \frac{\Phi}{4\pi|r - r_s|^2} f_r(\mathbf{r}, \omega, \omega_s) \cos \theta_s d\omega_s$$

where  $\theta_s$  is the angle between the normal and the light source direction. When using shadow information for positioning the light sources, points on the shadow boundary must be determined and joined to their corresponding blocking silhouette in order to get a reliable direction. When enough pairs are defined, a least squares procedure can be performed. On the other hand, when using highlights for source locations, points on the desired highlight maximum and its boundary (defined to be the line where the highlight falls below a given threshold) must be given. This procedure strongly depends on the BRDF chosen, in general a simple Phong<sup>BT75</sup> formula.

#### 3.2.1. General LSPP Formulations

Costa et al.<sup>CSF99</sup> implemented an automatic method to search for the best placement of luminaries, as well as their relative intensities. The method implements a preprocessing step where the user-defined requirements (called Inverse Luminaries, IL) are considered as sources of unit importance, which propagate through the environment as the dual of radiance<sup>Gla95</sup>. This allows using any global illumination engine to run the simulation backwards, from the Inverse Luminaries to a user-defined set of surfaces where the importance distribution is computed. Basically, what is presented is a validating preprocess step which tries

to find incompatibilities between design goals and already placed design elements (light sources), followed by a calculation step which attempts to find the best placement and orientation for the light sources by optimizing a user-defined objective function. The minimization of the objective function is performed with the Simulated Annealing algorithm <sup>PTVF92, PT91</sup>. The objective function is given to the system by means of a script language specially developed, which allows the specialized user to define the lighting goals to achieve (positive ILs), the illumination constraints to avoid (negative ILs, like having too much glare into a virtual character's face), and the geometric constraints the user might impose on the location or direction of the sources. Although this method of scripting the design goals seems very promising, a higher abstraction level should be achieved in order to allow the non-programming-skilled designers to be able to use the presented tools.

Instead, Shacked and Lichinski<sup>SL01</sup> presented an approach to lighting design based on the optimization of an objective function which is a *perception-based* image quality function. This function was designed to yield compressible images of 3D scenes, trying to effectively communicate information about shapes, materials and spatial relationships. Their current implementation was based in an OpenGL rendering engine and a local steepest descent optimization scheme, although the authors stated that local minima were found to be quite satisfactory if initial values were chosen wisely<sup>Sha01</sup>. The main difference with <sup>CSF99</sup> is the choice of the perception-based optimization function, as well as the use of a local optimization method vs. the global algorithm used before.


An entirely different approach for exploring the space of lighting designs was presented by Marks et al.<sup>MAB+97</sup>. In a framework named Design Galleries, they try to optimally disperse the space of solution images in terms of perceptual quality, and allow the user to browse and combine them to achieve a desired solution. This is clearly not an automatic process, since user input is required.

### 3.2.2. Local Illumination LSPP Formulations

One of the first works on inverse problems in the context of local illumination was done by Poulin and Fournier, <sup>PF92</sup>. They proposed using the highlights and shadows on the scene's objects in the modeling of the light sources. In the case of highlights, the authors considered the specular term of Phong <sup>BT75</sup> shading as expressed by Blinn <sup>Bli77</sup>. By letting the user manually point at the desired maximum intensity of this highlight, they were able to analytically optimize this expression and determine the light direction. By de-

termining another point on the surface, the user specifies where the specular term reaches a fixed threshold, and thus the roughness exponent can be computed. The use of further restrictions, like the light being on a given plane, is suggested for other cases. Observe that this method gives only directional light sources, that is, point light sources at infinity. For a general point light source, the shadow volume <sup>FvDFH90</sup> generated by it must be used <sup>Pou93</sup>. In order to specify the direction of a directional light one simply chooses two arbitrary distinct points in the scene, the second being along the shadow cast by the first one. For extended linear or polygonal (planar) light sources, new point light sources that define the vertices of the light source are needed. For general extended light sources a divide and conquer strategy was presented: if both the light and the object being shaded are divided into convex elements, the whole shadow is the *union* in 3D of all the shadow convex hulls.

Poulin, Ratib, and Jacques <sup>PRJ97</sup> find the position of point light sources by sketches of shadows or highlights, and extended light sources are positioned by sketches of umbra or penumbra. The user introduces the sketches as continuous strokes of points that are immediately transformed to 3D <sup>HH90</sup>, where the sketched points are considered to be all enclosed by the real shadow of the object, and similarly for a highlight. The method starts by considering each point forming the sketch of the shadow and defining for it the cone of possible positions for the light (Figure 2). The volume where the light can be is the intersection of all these cones: if the volume is infinite, a directional light is computed; otherwise a point light is used. The problem is presented as a constrained optimization problem by defining as an objective function the distance between the sketched points and the light source, and maximizing it. The constraints defined are that the point light source must lie inside all cones, that the light must stay on the same side as the normal vector at the sketched point, and that the light position is on the right side of the half-cone for this sketched point, oriented along the axis that joins this point and the center of the occluder. The initial position for the solver is chosen as a small distance above the occluder surface aligned with the center of mass of the sketched points. For extended light sources, the user sketches the umbra or penumbra, relaxing the inclusion condition to force all points to lie within all cones simultaneously in the case of umbras, and for penumbras the condition is that at least one point of the light must belong to each cone (a test that the intersection of the cone and the light is not null must be made). When sketching highlights, a point on a surface is considered within a highlight if the evaluation of Phong specular function at this point is higher than a certain



**Figure 2:** Elements for the computation of the position of a light source from a sketch of a shadow.

threshold  $\tau$ . It could be that the cones do not intersect at all, due to the curvature of the surface. Then, the roughness coefficient  $n$  (the exponent in Phong specular term) is lowered, broadening the cone, until an intersection is possible. The main difference with the previously mentioned research<sup>PF92</sup> is that the input method, sketching shadows and highlights, is far superior in terms of user interaction and greatly improves upon the design of simple illumination in a computer graphics scene.

A different approach was taken by Guillou<sup>Gui00</sup>, who developed a local illumination-based method to determine the position of  $n$  light sources. His method starts by computing  $m$  directional light sources from user-defined regions of a known, purely diffuse scene. Then, these  $m$  directional sources are grouped into sets and each set is used to *estimate* the position of a point light source by trying to optimize an intersection point from the directions of the distant light sources. Finally, the light position and photometric parameters are found by minimizing (Levenberg-Marquardt numerical minimization method) a least squares error, using the previous estimations as a starting point. As can be seen, this method is strongly local-illumination-based, since it relies on an estimation step that closely follows this assumption. Also, this distant light estimation step can only serve as a reference as long as the user-defined regions are cleverly chosen, since it does not take into account the superposition of illumination from the different sources at the points where the estimations are computed. With respect to the previous works mentioned, it has the clear advantage of working simultaneously with  $n$  point light sources in a diffuse environment.

### 3.3. Conclusions on ILP

Analyzing the problems concerning EP (Subsection 3.1), we see that:

- Most of the papers deal with perfectly diffuse

BRDFs. The more relevant exceptions to this rule are <sup>SDS+93</sup>, <sup>MG97, Mar98</sup> and <sup>RH01b</sup>, that deal with linear combinations of photographs and present algorithms that are independent of the BRDF used because they only depend on the obtained image, not on the method to compute it. (<sup>RH01b</sup> assumes a local illumination model to deconvolute illumination from the known BRDF.)

- The assumption that the BRDF is constant on the patch surface is found in most of the papers, too. The exceptions to this rule are <sup>OH95</sup> and <sup>SDS+93, MG97, Mar98</sup>, mainly because the first paper presents an algorithm that samples on the surfaces pointwise and the others are BRDF independent.
- Only França et al.<sup>FOH98</sup> and Morales et al.<sup>MMOH97</sup> solved the inverse problem of source emissivities in the presence of participating media. The rest of the reviewed works only deal with non-participating media. It is important to note the added cost: in this case, the problem requires the introduction of a system of equations resulting from the discretization into volume elements of the medium in order to solve the corresponding partial differential equations (PDE).
- Among the reviewed papers, only <sup>SDS+93, MG97, Mar98</sup> and <sup>KPC93</sup> present viewing-dependent goals, especially the first three because they deal with input given by photographs, while the last one can also be used with view-independent goals.
- It is also important to point out that, although the method proposed by Schoeneman et al.<sup>SDS+93</sup> achieved interactivity by accounting only for direct illumination in their final implementation, it is independent of the illumination algorithm used. The same happens with Marschner and Greenberg's approach<sup>MG97, Mar98</sup>, which has the advantage of using a photograph as objective, rather than manually user-defined objectives as before, and of using a generic scene-independent set of basis light sources. Unfortunately, this leads to a more ill-conditioned system than the system that comes from a set of focused light sources, thus requiring a regularization procedure.
- Sato et al.<sup>SSI99c, SSI99b</sup> required the introduction of regularization using user-weighted penalty terms, and the computational complexity limited their formulations to a coarse discretization of the sphere. Instead, the method proposed by Ramamoorthi and Hanrahan<sup>RH01b</sup> required no explicit regularization and yielded results with better sharpness and overall quality than the approaches of Sato et al. On the other hand, the methods proposed by Sato et al. are easier to extend to concave surfaces.

With respect to LSPP (Subsection 3.2), it is very evident that two ways of attacking this problem, that

is Radiosity-based formulations and Monte Carlo formulations, have not been studied in the literature. Nevertheless, in our opinion, these ways are feasible in spite of the possible ill-conditioning they might present (which would surely be alleviated by the use of some sort of regularization procedure).

In the reviewed articles we can also notice that:

- It is remarkable that <sup>CSF99, SL01</sup> and <sup>MAB+97</sup> are the only global illumination-based approaches among all the reviewed papers for this kind of problem. Contrary to what happens in Inverse Reflectometry problems, all studied approaches rely on some sort of optimization procedure to achieve their results. This is so because of the high complexity of the problem faced, since finding absolute locations of light sources or types (and other characteristics) of luminaries involves using indirectly measured information. The choice of the optimization method is, to our knowledge, quite arbitrary and the papers present several different approaches to this point. It is clear that the best optimization method to use is still an open research area.
- Although the method by Costa et al.<sup>CSF99</sup> for scripting the design of goals seems very promising, a higher abstraction level should be achieved in order to allow the non-programming-skilled designers to be able to use the presented tools.
- On the other hand, the main difference between Shacked and Lichinski<sup>SL01</sup> and Costa et al.<sup>CSF99</sup> is the choice in the first case of the perception-based optimization function, as well as the usage of a local optimization method vs. the global algorithm used before.
- With respect to the work done by Marks et al.<sup>MAB+97</sup>, it can be clearly seen that this process does not involve an automatic optimization at all, the user being responsible for all the decisions towards the final result.
- The main difference of the work done by Poulin, Ratib and Jacques<sup>PRJ97</sup> with the work by Pouin and Fournier<sup>PF92</sup> is that the input method, sketching shadows and highlights, is far superior in terms of user interaction and greatly improves upon the design of simple illumination in a computer graphics scene.
- Finally, for Guillou<sup>Gui00</sup> we can say that it has the clear advantage of working simultaneously with  $n$  point light sources in a diffuse environment, whilst the other methods presented in the same subsection do not.

#### 4. Inverse Reflectometry Problems (IRP)

Inverse reflectometry problems are those where  $\hat{K}$  in equation 2 is unknown, and  $\hat{C}$ ,  $L_e$  and part of  $L$  are

General	Deb98, BG01
Radiosity	FGR93, DRB97, YDMH99, LFD+99, LDR00
Texture-based	IS91, KC94, POF98, OSRW97, DvGNK99, LKG+01
Angular-based	Pou93, PF95, Mar98, MWL+99, SI96, SWI97, DHT+00
Local	General YM98, Mar98, RH01b

**Table 3:** Classification of Inverse Reflectometry Problems organized with respect the treatment they give to Equation 1: General, Radiosity, Monte Carlo and Local.

known (Figure 3). Thus, we must solve for information about  $\hat{K}$ . The methods studied in this section can be classified according to the illumination approach used. As above, this can be either local-based, general global illumination, Monte Carlo-based or radiosity-based, see Table 3.

- **General formulations.** These methods try to find reflectance properties without relying on any particular illumination algorithm.
- **Radiosity-based formulations.** As stated in section 3.1, here the general problem is reduced to a radiosity setting by making the following approximation to Equation 1: a purely diffuse BRDF for the surfaces (patches), that is constant all over each one. In this case, Equation 1 is reduced to the form:

$$B_i = Le_i + \rho_i \sum_j F_{ij} B_j$$


where  $B_i$  is the  $i$ -th patch radiosity,  $Le_i$  its emittance,  $\rho_i$  its reflectivity (diffuse BRDF) and  $F_{ij}$  is the form factor from element  $i$  to element  $j$ . If we know  $B_j$  and  $Le_i$  for every surface, and with form factors  $F_{ij}$  known if the geometry is known, finding the reflectances is reduced to:

$$\rho_i = (B_i - Le_i) / \sum_j F_{ij} B_j$$

The most common way of knowing  $B_i$  and  $Le_i$  is by using an image of each surface taken by a camera, and retrieving the information from there. But if we work without an image for each surface, some heuristics must be used. The most common approach is to start from an initial estimate of the average reflectivity  $\rho^*$  and estimating the surface reflectance as

$$\rho_i^* = \frac{B_i^*}{B_A^*} \times \rho^*$$





**Figure 3:** Diagram for Inverse Reflectometry Problems.

where  $B_i^*$  is the average intensity of the pixels recovered from the camera observation and  $B_A^*$  is the ambient radiosity defined as

$$B_A^* = \frac{\sum_{xy} p_{xy}}{N\rho^*}$$

where the sum is performed over the  $x \times y = N$  pixels of the image with intensities  $p_{xy}$ . The procedure is iterated until some satisfactory threshold is achieved.

- **Monte Carlo formulations.** To the best of our knowledge, there is no research using this kind of formulation, but we strongly believe that this is a feasible and sensible approach: Simply fire rays from the camera or the surfaces with known properties and continue its path until hitting a surface with unknown BRDF. Then, using the gathered information we could estimate the reflectance parameters. Also, a bidirectional approach could be used, given known lighting conditions. In any case, it is clear that BRDF recovery is feasible, but certainly will have to deal with the inherent variance problems Monte Carlo methods present, which can only get worse for an inverse problem of this kind.
- **Local Illumination formulations.** These formulations are based on the simplification of Equation 1 to take into account only local illumination, and using the observation of this local illumination to obtain values for the  $f_r$  (BRDF) coefficients.

Ramamoorthi and Hanrahan<sup>RH01b</sup> have studied the Inverse Reflectometry Problem under a signal processing framework and arrived at the conclusion that BRDF recovery is feasible (well-conditioned in the mathematical sense) when the known lighting contains high frequencies like directional sources, and is ill-conditioned for soft lighting.

#### 4.1. General Global Illumination based IRP

Paul Debevec, in <sup>Deb98</sup>, introduces the concept of a *light-based model*, a representation of a scene that consists of radiance information, possibly with specific reference to light leaving the surfaces, but not necessarily containing BRDF information. He presents a method that uses the measured scene's radiances and global illumination in order to add new objects to light-based models with correct lighting. The light-based model is constructed from an approximate geometric model of the scene and by using a light probe to measure the incident illumination at the location of the synthetic objects. To do that he divides the scene into three main regions: the distant scene, represented with an environment map; the local (or near) scene which is going to photometrically interact with the synthetic objects and whose geometry must be well known; and synthetic objects. To estimate the local scene BRDF, he assumes a reflectance model (e.g. diffuse, specular, ...) with approximate initial values, and iteratively computes the global illumination solution for the local scene with the current parameters with respect to the observed lighting configurations. By comparing the appearance of the rendered local scene to the actual appearance, he decides whether to continue iterating with adjusted parameters, or not. In the case of purely diffuse reflectors, the next estimate of the reflectances is the ratio from the resulting radiance to the observed value. Any other case is left as future work, manually estimating the specular coefficients for the non-diffuse objects in his test scenes.

Boivin and Gagalowicz developed<sup>BG01</sup> a reflectance recovery algorithm that starts with a pure Lambertian model and successively tries more complex BRDF models until a fit between the original image and its synthetic reproduction is achieved. For the simpler models (diffuse, perfect and almost-perfect specular), an iterative correction is applied based on the object image to synthetic image ratio of the previous iteration, while in more sophisticated models a Simplex method is used. If the whole hierarchy of models fails to provide a good fit, the method proceeds to a plain texture extraction, using methods from any of the examples in section 4.2. This method has the clear advantage over the previous one of being able to work with non-diffuse BRDFs without requiring a manual user intervention.

#### 4.2. Radiosity-based IRP

Computer augmented reality also requires the computation of reflectances from images. All research works have a preprocess stage where the scene geometry is approximately reconstructed with photogrammetric techniques. The work from Fournier et al. <sup>FGR93</sup>, as-

sumes reflectance is constant across each patch and uses a heuristic method that assigns each patch a reflectance that is an average reflectivity multiplied by the ratio between the element radiosity (computed as the average of all the visible pixels it contains) and an average radiosity. Average values are computed directly from the images, weighting the obtained values with the respective areas of the patches. In <sup>DRB97</sup>, textures of arbitrary resolution are extracted by de-warping (extracting the textures from the images by reversing the perspective-introduced distortion) the original image and bringing it back to the plane of the previously built polygon. Unfortunately, their method sacrifices the usage of dynamic cameras and real scenes to gain speed, and the quality of the obtained images is slightly degraded due to the use of a polygon-projection method instead of ray tracing.

Yu et al. <sup>YDMH99</sup> have presented an inverse radiosity method to account for mutual illumination in estimating spatially varying diffuse and piecewise constant specular properties within a room from a sparse set of photographs. Their technique is based on the usage of a low-parameter reflectance model (metals and plastics treated differently, see below), allowing the diffuse component to vary freely over surfaces while assuming non-diffuse characteristics remain constant across particular regions. As input the method receives a geometric model of the scene and a set of calibrated, high dynamic range photographs <sup>DM97</sup> taken with known direct illumination. The algorithm proceeds by hierarchically partitioning the scene into a polygonal mesh and, by using image-based rendering techniques, it computes an estimate of both the radiance and irradiance of each patch. Using the known geometry and light source positions, it computes the estimate placement of the specular highlights falling inside the radiance images, and runs an iterative optimization procedure to recover the diffuse and specular reflectance parameters of each region. These results are used to update the hierarchical system. Then, the estimation-update procedure for the BRDF parameters is repeated. The iteration is performed several times to obtain the final solution of the BRDFs for all surfaces. Two different formulas are used for isotropic or anisotropic BRDFs (Ward's model), and it is suggested that metals be treated differently from plastics: for plastics, they consider the specular coefficient constant and the diffuse one variable, while for metals they do the opposite. The selection between both models was performed through a simple test: a metallic surface has its specular reflectance larger than the estimated diffuse component.

In <sup>LFD+99</sup> Loscos et al. a different approach for reflectance recovery is presented: a set of images from a fixed viewpoint but with controlled, varying illu-

mination (no shadows) is combined with confidence weight factors that represent the visibility with respect to the light source for each pixel in the images. The resulting textures are de-warped as before. Later, this was extended <sup>LDR00</sup> and an algorithm for interactive re-lighting was introduced, based on a preprocessing step that reconstructs geometry and creates unoccluded illumination textures (thus taking into account the effect of shadows). Unfortunately, only diffuse surfaces can be considered. The creation of these unoccluded textures has two steps: firstly they add the light that was blocked in a hierarchical radiosity solution, and then a heuristic correction is applied. This correction is computed by finding an appropriate unoccluded reference patch which will give an indication of the desired color, and computing a modulation factor consisting of the ratio of form-factors of each patch to the light source. In <sup>LD00</sup> this method was improved by adding a low-cost photometric calibration method which improves the reflectance estimate of real scenes. This was achieved by adapting a high-dynamic range image creation to a low-cost camera, and an iterative approach to correct reflectance estimation using a radiosity algorithm for indirect light calculation. Unlike previous work, it allows for a restricted set of BRDFs (purely diffuse) to be recovered, but works with the simplest capture process since it doesn't need user-controlled specific lighting. Most important, this last piece of research does not attempt to perform a reflectance estimation, since it uses a simple texture modulation for display. Also, when compared with the work by Yu et al. <sup>YDMH99</sup>, we see that the last one is far from interactive, due basically to the generality of the algorithm used for the light propagation (despite the fact that their method is specifically tailored to the radiosity setting), but has the advantage of handling any viewpoint in the environment.

### 4.3. Local Illumination Based IRP

Although a BRDF is a function  $f_r(\mathbf{r}, \omega, \omega_i)$  where  $\mathbf{r}$  is a point on a surface and  $\omega$  and  $\omega_i$  the outgoing and incoming directions, for classification purposes it is convenient to subdivide the different methods that attempt to recover  $f_r$  as texture-based, angular-based and general BRDF-based. Texture-based are those that only consider spatial variations in  $\mathbf{r}$  disregarding angular variations and assuming a constant angular behavior, generally purely Lambertian. Instead, Angular-based methods consider the BRDF as a function of  $\omega$  and  $\omega_i$  only, disregarding spatial variations.

#### 4.3.1. Texture-based BRDF Recovery

In 1991, Ikeuchi and Sato<sup>IS91</sup>, using one intensity and one range map (z-buffer), and assuming constant

material regions over an object, were able to obtain estimates of diffuse and specular parameters for a specifically tailored version of the Torrance-Sparrow BRDF (by means of an iterative least-squares fitting algorithm). Later, Kay and Caelli<sup>KC94</sup> extended this approach without requiring BRDF to be constant over object regions, estimating the parameters at each point on the object. They applied a photometric stereo method to a range map and a number of intensity maps, and inverted the illumination model at each point on the object. (For non-highlighted regions, they used a linear least-squares method while for highlight regions a nonlinear separable least-squares method with regularization was used.) They were able to recover highly textured surfaces by using enough intensity maps to recover diffuse and specular parameters at each point. If enough intensity maps were not available, they showed how to recover the BRDF when the specular component is assumed to vary smoothly over the material.

Poulin et al.<sup>POF98</sup> describe an interactive system to reconstruct 3D geometry and extract textures from a set of photographs. The authors describe a three step process: A least-squares problem is first solved for the camera parameters, and then for the 3D geometry. Once a satisfying 3D model is recovered, its color textures are extracted by sampling the re-projected texels in the corresponding images. All the textures associated with a polygon are fitted to each other, and the corresponding colors are combined according to a set of custom criteria in order to form a unique texture. For each texel, the size in pixels of its projection in the images is used as an indication of the quality of the extracted color.

Ofek et al.<sup>OSRW97</sup> present a method which deals with the problem of recovering and blending textures from different images, but also discusses the problem of removing highlights and reflections. Multi-resolution textures are stored in a quad-tree data structure, which is filled by a recursive level-of-detail projection algorithm of the image to the texture space, followed by a push-pull procedure to propagate information all along the tree. The algorithm thus calculates an approximation of the view-independent color for each texel, and calculates the average of the projected texture area for every texture pixel that is near the mentioned estimation. The main advantage of this algorithm is its ability to account for the different sampling rates that result from different views of the surface. Unfortunately, they do not attempt to model the surface reflectance, but use texture maps to record the diffuse component of the radiance reflected under the lighting conditions at the time the photographs were taken.

Dana et al.<sup>DvGNK99</sup> applied reflectometry techniques to the domain of textured objects by using a spectrophotometer to carefully measure spectral BRDFs without separating diffuse or specular components. Of course, this becomes impractical for complex BRDFs due to the high storage costs, and requires the lighting to be totally known for the images.

Lensch et al.<sup>LKG+01</sup> presented an image-based measuring method that robustly detects the different materials of objects, organizes them in clusters of similar materials and fits an average BRDF to each of them (Levenberg-Marquardt method initialized by an average BRDF). Although their method is BRDF-independent, they used a Lafortune<sup>LFTG97</sup> model as target BRDF. In order to model local changes, they projected the measured data for each surface point into a basis formed by the recovered BRDFs, leading to a spatially varying representation. To do this they introduced the concept of a *lumitexel*, a data structure that stores all geometric and photometric information for a surface point. The general idea behind the BRDF classification algorithm is to start with a cluster containing all the samples and recursively subdivide it, classifying the lumitexels according to the error they have with respect to each cluster-computed BRDF. This method has the advantage of not requiring a specific BRDF or a homogeneous material, as most of the previously mentioned approaches, but requires full knowledge of the object geometry, lights and camera.

#### 4.3.2. Angular-based BRDF Recovery

Poulin and Forunier<sup>Pou93, PF95</sup> deal with the problem of determining the characteristics of surface materials (some parameters of their BRDF) by using a painting paradigm where the user simply paints color points on a surface. The system attempts to find the best values for the surface characteristics such that the points will retain their assigned color in the final rendering. Depending on the number of constraints (color points) given by the user, the problem can either be a non-linear constrained optimization one (when there are less color points than variables to find out) or a weighted least-squares fitting problem with penalty functions to constrain the values of surface parameters (otherwise). For the former, each color point is considered as a volume in the 3D-color space of acceptable colors, introducing two inequality constraints in each direction (and no value can be negative). For the non-linear, least-squares fitting problem, the authors use penalty functions to introduce the constraints. Also, the system is modified to assign different weights depending on the location of the color points, e.g. at the dark side of an object the ambient term dominates.

Another inverse reflectometry problem posed by

Marschner<sup>Mar98, MWL<sup>+</sup>99</sup> was called *image-based BRDF measurement*, which presents a system that measures reflectance quickly without special equipment. The method works by taking a series of photographs of a curved object, each image capturing light reflected by differently oriented parts of the surface. The photographs are analyzed to determine the BRDF by using a curved test sample with known shape, an imaging detector and automated photogrammetry to measure the camera position, light source location and sample placement. At first, the *geometric calibration* stage uses machine-readable targets with embedded identification codes placed near the sample to allow photogrammetric techniques<sup>Mar98</sup> to be able to locate those samples. The information derived in this stage is the position of the light source, the camera location for each measurement and the location of the sample. The next step, *radiometric calibration*, obtains the relationship between the radiance reflected to the camera and the irradiance due to the source. An important assumption done in those measurements is the approximation of the source as a single point, which is correct when the source is small compared to the distance to the sample. To get the absolute magnitude of the BRDF correctly, they measured the intensity of the light source relative to the camera sensitivity by photographing a diffuse white reference sample in a known position. The last step, *data processing*, is performed by the *de-renderer*, which uses standard rendering techniques<sup>MWL<sup>+</sup>99</sup> to find the intersection point of each pixel's viewing ray with the sample surface and to compute the radiance from the source. To obtain the desired BRDF value, the de-renderer divides the pixel's measured radiance by the irradiance. The de-renderer's output is a list of BRDF samples, each including the incident direction, the exitant direction and the value for that configuration.

Debevec et al.<sup>DHT<sup>+</sup>00</sup> recovered a two-parameter BRDF model for the human skin with color space analysis techniques from a set of photographs with varying illumination. To do that, they assumed that the specular component was the same color as the incident light, while the diffuse one was obtained in a two-step manner: first, by fitting a Lambertian lobe to obtain the surface normal, and then finding the parameters in their model by fitting them to the observed chromaticities in the original un-separated reflectance function. Previously, Sato et al.<sup>SWI97</sup> had presented a similar algorithm to retrieve the shape and the BRDF of convex objects by using a turntable and a single point source. The main difference between these algorithms for BRDF recovery is that, in the second step in Sato et al., the diffuse lobe is fitted to the diffuse term in a modified Torrance-Sparrow model. Their method required 120 color images and 12 range maps to com-

pute the BRDF parameters. This work is an extension of a method<sup>SI96</sup> that recovered a simplified Torrance-Sparrow reflection model for an isolated object from a sequence of range images and a reconstructed 3D model, constraining the camera parameters and light source position. (They used the Levenberg-Marquardt numerical minimization method.) This way, they were able to separate the diffuse and specular components and recover the uniform reflectance of the surface. Unlike Marschner et al.<sup>Mar98, MWL<sup>+</sup>99</sup>, Sato et al. sacrificed the generality of measuring a full BRDF at each surface point and instead used a single-formula model of specular and diffuse reflectance to extrapolate the appearance at novel viewpoints.

#### 4.3.3. General BRDF Recovery

Yu and Malik<sup>YM98</sup> presented an approach to produce photo-realistic computer renderings of real outdoor architectural scenes (building facades) under novel lighting conditions. Their system uses a small set of photographs as input, along with a geometric model of the scene generated with photogrammetric techniques. The input photographs are taken with a hand-held CCD camera and converted into radiance images<sup>DM97</sup>. They defined two pseudo-BRDFs, one corresponding to the spectral distribution of the sun (modeled as a parallel light source) and one corresponding to the integrated light from sky (they fit a sky model to a set of calibrated photographs) and the environment (modeled through a low-resolution Spherical Environment Map). The incident radiance is obtained from the sun, sky and environment, while the outgoing diffuse radiance is taken from the photographs in directions away from specular reflection. Each face of a building must appear in at least two photographs, one with direct illumination from the sun and the other without it. Each polygon in the original geometric model is first triangulated and a dense grid is set up on each triangle in order to capture the spatial variations in the pseudo-BRDF. The specular lobes are recovered with an empirical model<sup>LF<sup>+</sup>TG97</sup>. The sky and environment are divided into small pieces and the vector flux is plugged from each piece into the specular model, thus resulting in a least-square minimization problem. The authors assume that each visibility-blocking surface in the model has the same specular lobe (except for windows which are left for further investigation).

Together with the *re-lighting* system<sup>MG97</sup> described in section 3, Marschner's thesis presents two inverse reflectometry problems, the first one being *Photographic Texture Measurement*<sup>Mar98</sup>. Its purpose is to construct a representation of the spatially-varying parameters of the BRDF in equation 2, based on samples provided by a set of photographs of the object (each with known lighting and camera position). For

that reason, a BRDF with a small number of parameters is chosen. In particular, for many of the examples, they choose a pure lambertian BRDF. The algorithm first gathers all the observations of radiance reflected from a particular surface point by sampling all the user-provided photographs in which the surface point is visible and illuminated. From these measurements, and using the known geometry of the surface, the incident and exitant directions and the reflectance value are computed for each observation. These samples are used at that point to estimate parameters of the BRDF model. To obtain an estimate of the Lambertian component of surface reflection, for example, it requires a least-squares fitting of the values, with weights that depend on the incident and exitant directions, i.e. values with view or illumination directions near the surface normal are more reliable than samples that are near-grazing, and samples that include significant contribution from specular reflection are less reliable than those that do not. The specular part of the BRDF is handled by combining measurements from different points, based on the supposition that some parameters of the BRDF are spatially constant while others may vary. Sample points are chosen at the vertices of an optimized triangulation of the object's surface, performing a linear interpolation for the points in between.

Ramamoorthi and Hanrahan<sup>RH01b</sup> estimate low-parameter BRDFs using a three-component model of the reflected light field: a diffuse component, specularities from the slowly-varying lighting and specular highlights from the fast-varying lighting component (obtained from a Spherical Harmonics decomposition). It is shown that two loops estimate the parameters of a simplified Torrance-Sparrow BRDF: the outer one, through a simplex algorithm adjusts the non-linear parameters, while the inner loop performs a linear optimization of the diffuse and specular weights of the formula. For spatially varying BRDFs, a loop over the different surface points is added and the mentioned algorithm is repeated for each location. The main difference with previous work, and especially the ones able to work with outdoor scenes, is that it does not assume a simple parametric model for skylight like Yu and Malik<sup>YM98</sup>, nor does it requires highly controlled lighting conditions (generally by careful active positioning of a single source) like Marschner et al.<sup>MG97</sup>. Also, Ramamoorthi and Hanrahan were the first to solve the IRP for general illumination (irradiance), without requiring simple BRDFs or low resolution textures, as in the previously mentioned work.

#### 4.4. Conclusions on IRP

The first important thing to notice is that, to the best of our knowledge, there are no efforts to solve these problems with Monte Carlo-based methods, although they could be well suited for this kind of problem.

In general, we can make a few observations:

- In the category of General Inverse Reflectometry Problems, we can observe that the method used by Boivin and Galgalowicz<sup>BG01</sup> presents the clear advantage over other image-based methods (like Debevec et al.<sup>Deb98</sup>, Yu et al.<sup>YDMH99</sup>, Loscos et al.<sup>LFD+99</sup>, Fournier et al.<sup>FGR93</sup> and Drettakis et al.<sup>DRB97</sup>) in that it uses the area covered by the projection of an object in the real image or images to determine its reflectance, and thus avoids producing large errors for small objects, since this method uses a feedback through the comparison between real and synthetic images (significantly reducing bias). Also, it does not need a specular highlight for each surface to appear in at least one image, as Yu et al.<sup>YDMH99</sup> do. On the other hand, observe that the method by Yu et al.<sup>YDMH99</sup> and the one by Boivin and Galgalowicz<sup>BG01</sup> work on full scenes, while Debevec<sup>Deb98</sup> only estimated material properties on a part of the scene.
- With respect to radiosity-based approaches, it is important to note that in <sup>DRB97</sup> textures of arbitrary resolution are extracted, but unfortunately, their method sacrifices the usage of dynamic cameras and real scenes to gain speed, and the quality of the obtained images is slightly de-gradated due to the use of a polygon-projection method instead of ray tracing. Unlike other previous research, Loscos et al.<sup>LFD+99</sup> allow for a restricted set of BRDFs (purely diffuse) to be recovered, but work with the simplest capture process since user-controlled specific lighting is not needed. Most important, this last work does not attempt to perform a reflectance estimation, since it uses a simple texture modulation for display. Also, in comparison, the work by Yu et al.<sup>YDMH99</sup> is far from interactive, basically because of the generality of the algorithm used for the light propagation (RADIANCE) despite the fact that the method is specifically tailored to the radiosity setting. However, it has instead the clear advantage of handling any viewpoint in the environment.
- In the Local-based approaches subsection, we observe that the method by Poulin et al.<sup>POF98</sup> and the one by Poulin and Fournier<sup>PF92</sup> are the only interactive methods that allow some sort of texture recovery. Instead, Ofek et al.<sup>OSRW97</sup> do not attempt to do real-time processing, but to compute high quality multi-resolution textures from image sequences.

The work by Dana et al.<sup>DvGNK99</sup> has the problem of becoming impractical for complex BRDFs due to its high storage costs, and requires the lighting to be totally known for the images. On the other hand, the method by Lensch et al.<sup>LKG+01</sup> has the advantage of not requiring a particular BRDF or a homogeneous material, like most of the previously mentioned approaches, but requires full knowledge of the object geometry, lights and camera.

The main difference between the works by Debevec et al.<sup>DHT+00</sup> and Sato et al.<sup>SWI97</sup> for BRDF recovery is that, in the second step in the Sato et al.<sup>SWI97</sup> work, the diffuse part is fitted to the diffuse term in a modified Torrance-Sparrow model. Unlike Marschner et al.<sup>Mar98, MWL+99</sup>, Sato et al. sacrificed the generality of measuring a full BRDF at each surface point and used a model of specular and diffuse reflectance to extrapolate the appearance of novel viewpoints.

Finally, the main advantage over previous methods of the work by Ramamoorthi and Hanrahan<sup>RH01b</sup> is that it does not assume a simple parametric model for skylight, nor does it require highly controlled lighting conditions. Also, Ramamoorthi and Hanrahan were the first to solve the IRP for general illumination (irradiance), without requiring simple BRDFs or low resolution textures, as in the previously mentioned work. Thus, we see that they presented a general solution based on a convolution principle that enables a reasonably easy method for BRDF, texture and/or lighting recovery.

Also, we can make some further comparisons on methods on different sections by noticing that:

- We can see that <sup>Deb98</sup> BG01 <sup>YDMH99</sup> <sup>FGR93</sup> <sup>DRB97</sup> <sup>BG01</sup> <sup>LFD+99</sup> <sup>LD00</sup> <sup>LDR00</sup> are papers that use the full global illumination equations, while the rest of the reviewed papers in this section use only the simplified local versions.
- We also see that <sup>POF98</sup>, <sup>DHT+00</sup>, <sup>SWI97</sup> and the works on Computer Augmented Reality present an approach that does not rely on any optimization procedure, taking the measurements directly from the user's input.
- The same happens with <sup>Mar98</sup>, <sup>MWL+99</sup>, which take the information directly from the user-provided photographs, without needing an intermediate optimization process to get the results.
- Instead, <sup>YDMH99</sup> uses a hybrid approach, measuring the diffuse component of the BRDF and resorting to another approach for the other parameters in the BRDF used.
- One of the most important things to notice is that most of the papers are based on small-parameter BRDFs, in general of two main types:

	Several PoV	Single PoV
Local	SI96, OSRW97, YM98, DHT+00, LKG+01	SWI97, KC94, DvGNK99, LKG+01
General	Deb98	BG01
Radiosity	YDMH99, LDR00	LFD+99, DRB97, FGR93

**Table 4:** Inverse reflectometry papers that use images from different Points of View (PoV) as input.

- The scene can be decomposed in regions with arbitrary variation of the diffuse reflectance.
- The scene must be decomposed in regions without spatial variation of the reflectance, but with the possibility of using a BRDF function with higher dimensionality (more complex glossi behaviour).
- The papers that solve this sort of problems and use images as input can be further classified<sup>BG01</sup> according to whether they used multiple points of view or a single point of view for the images in their computations. This classification can be found in Table 4.

## 5. Combined Inverse Lighting and Inverse Reflectometry Problems (CILRP)

Combined problems are those where  $\hat{K}$  and  $L_e$  are unknown in equation 2, but  $\hat{G}$  and part of  $L$  are known. The methods studied in this section also can be further classified according to the illumination approach used: As in Sections 3 and 4, this can be either local-based, General Global Illumination, Monte Carlo-based or radiosity-based.

- **General formulations.** These methods try to find reflectance properties and lighting conditions without relying on any particular illumination algorithm. There are no works using this sort of approach, probably due to the high complexity of the optimization of both problems at the same time.
- **Radiosity-based formulations.** Here the general problem is reduced to a radiosity setting by making the same approximations to Equation 1 as the ones described in Section 3: a purely diffuse BRDF for the surfaces (patches), that is constant all over each one.

- **Monte Carlo formulations.** Once again and to the best of our knowledge, there are no works using this kind of formulation, but the solution would be to simply fire rays from the camera or the surfaces with known properties and continue its path until hitting a surface with unknown BRDF or emittance, and using the gathered information to estimate the parameters. Also, a bidirectional approach could be thought of, too, where rays are fired as before and from the light sources with unknown emittance, and using the gathered information to recover the missing information. In any case, it is clear that simultaneous recovery of both BRDF and lighting conditions is feasible, but certainly the inherent variance problems presented by the Monte Carlo methods will have to be dealt with.
- **Local Illumination formulations.** These formulations are based on a local illumination formulation to obtain values for the lighting parameters and the BRDF coefficients.

In the work by Ramamoorthi and Hanrahan<sup>RH01b</sup>, the combined problem was studied: Inverse Lighting and Inverse Reflectometry Problem in a signal processing framework, and it was concluded that, up to a global scale, the reflected light field can be separated into the lighting and the BRDF, provided that the appropriate coefficients of the reflected light field, in a Spherical Harmonics representation, do not vanish. It is important to notice that this factorization can be done up to a global scaling factor.

### 5.1. Radiosity-based CILRP

John Kawai, James Painter and Michael Cohen<sup>KPC93</sup> use the radiosity to minimize the global energy of the scene, instead of the mean-squared difference between the desired radiosity values and current values at the patches as before. This global energy is given by the area-weighted sum of the element radiosities, plus a user-defined weighted sum of *Physical Terms* and *Human Perception Based Terms*. The *Physical Terms* include:

- radiosities
- emissions
- directionality and distribution of the light sources
- patch reflectances

while the *Human Perception Based Terms* are a quantification of the subjective impression of clearness, pleasantness or privacy based on the scene's brightness. Constraints are imposed by the user to the objective function as explicit weighted penalty terms. The resulting unconstrained problem is solved by the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method. In order to speed up computations, a Hierarchical

Radiosity (HR) solution was used<sup>HSA91</sup> to compute the initial baseline rendering and they reuse the computed links to propagate the increments used by a finite difference scheme to find an approximation to the derivatives of the radiosities with respect to the unknowns. This is possible since these derivatives obey the same equation as the ordinary radiosity problem. In the mentioned user-defined weighted sum of *Physical Terms*, they were able to include a variable element reflectivity term (constant over the patch surface), thus allowing the combined optimization of both emissivities and reflectivities in the same process. The partial derivative of the radiosities with respect to the reflectivities was computed by "shooting" the un-shot radiosity  $\Delta\rho$  due to the change in reflectivity,  $B_i\Delta\rho$ .

### 5.2. Local Illumination-based CILRP

Sato et al.<sup>SSI99c</sup> use the radiance information inside shadows to derive the illumination distribution of a real scene when the BRDF is Lambertian, recovering the diffuse coefficient up to a scaling factor, which is obtained from the camera calibration (see Section 4). This work was later improved<sup>SSI99b</sup> by using a non-uniform, adaptive discretization of the directions of illumination. Also, the need to know the reflectance properties of the shaded surface was not longer required, but the limitations of using only distant lighting, no interreflections, uniform reflectance and known camera and object shape remain unchanged. The algorithm proceeds by two nested loops, the outer one estimating radiance values of imaginary directional light sources (a linearized set of equations of the influences of the light sources vs. the pixels of the shadow surface) and the inner one estimating the reflectance parameters of the surface in shadows (Powell's method on the RMS difference of the actual pixel value and the estimated, local-illumination only value).

Ramamoorthi and Hanrahan<sup>RH01b</sup> presented a remarkable work from the perspective of signal processing, requiring a single manually specified directional source to recover the roughness of the surface. Absolute reflectance cannot be found from this method, so an ad-hoc relationship between the diffuse and specular weights is established. The algorithm consists of two nested loops, the outer one being an inverse BRDF-problem (see above, section 4) while the inner one is an estimation of the lighting with known BRDF parameters (see section 3). This work can be considered an extension of the previous one<sup>SSI99b</sup>, since it does not require shadow information and presents improved methods for estimating the illumination. It also addresses a more general setting, being able to work with spatially varying materials.

### 5.3. Conclusions on CILRP

One of the most important things to notice in this section is that there are only a few approaches that deal with both problems, illumination and BRDF recovery at the same time. In particular, Radiosity-based solutions only include the work by John Kawai, James Painter and Michael Cohen <sup>KPC93</sup>, but it is important to notice that their work presents a fairly efficient solution given the diffuse-surface and constant-patch-radiosity approximations involved.

On the other hand, Local Illumination-based CILRP is included in two works, with the one by Ramamoorthi and Hanrahan <sup>RH01b</sup> being able to be considered as an extension of Sato et al. <sup>SSI99c, SSI99b</sup> because it does not require a classification of the scene into shadow regions and it works in a more general setting, and is even able to deal with spatially varying materials.

Finally, it is important to note that no General CILRP and no Monte Carlo-based CILRP were presented, mostly because of the inherent high complexity involved, of special importance in the former. However, we think that Monte Carlo-based solutions are a sensible way of trying to solve these combined problems, despite the high variance inherent in the application of those methods to inverse problems.

## 6. Conclusions and Open Problems

Here we will summarize the conclusions of the surveyed work, to show common problems, characterize approaches to solutions and suggest open issues. We can organize the studied methods depending on the type of illumination model (general, radiosity, Monte Carlo and local) they use and according to the sort of problem they solve (Tables 2 and 2).

In Tables 5 and 6 we summarize the papers on Inverse Emittance and Inverse Reflectance respectively. We then show the main features of each method: type of approach used to solve it (direct or indirect) and numerical methods to implement it. These mathematical aspects are summarized in the Appendix.

As we can see in the tables, there are three kinds of approaches used to solve the different problems faced:

- *Indirect-solving* approaches, or *optimization-based* approaches, where the solution is obtained by finding the minimum (or maximum) of an adequately defined objective function. These methods generally require solving the direct problem at least once per iteration <sup>SDS+93, KPC93, MG97, Mar98, PF95, Mar98, YDMH99</sup>.
- *Direct-solving* approaches, where the goal is to

find methods to invert the equation without solving the direct problem at any time. Among those, we can mention the matrix-based approaches, where the goal is to invert a highly ill-conditioned algebraic system of equations resulting from a finite element approximation of the underlying equations <sup>HMH95, FOH98, MMOH97</sup>. Another approach <sup>OH95</sup> implements an Inverse Monte Carlo method to find the emissivities for a set of surfaces.

- *Mixed*, that use a combination of the two above-mentioned approaches. The only examples reviewed with this kind of approach are <sup>SI96, SWI97</sup> and <sup>YM98</sup>.

Analyzing the Inverse Emittance problems, with the aid of Table 5, we see that most of the papers that work on global illumination formulations deal with purely lambertian BRDFs, with only a few exceptions. (<sup>MG97, Mar98</sup> and <sup>RH01b</sup>, that deal with combinations of photographs and present algorithms that are independent of the BRDF used). Only two works pay any attention to the problem of Participating Media, <sup>HMH95</sup> and <sup>MMOH97</sup>.

With respect to the Inverse Reflectometry problems, we see in Table 3 that <sup>Deb98</sup> and <sup>BG01</sup> are papers that present an illumination-independent formulation, while the rest of the reviewed papers in this section use only radiosity or the simplified local version. We also see that <sup>POF98</sup> presents a direct approach, and together with <sup>Mar98, MWL+99, FGR93, DRB97, LFD+99, OSRW97</sup> and <sup>DvGNK99</sup> are the purely direct approaches among those surveyed. On the other hand, <sup>YDMH99, SI96</sup> and <sup>SWI97</sup> use hybrid approaches, using a direct measurement for one of the BRDF components (generally the diffuse component of the BRDF) and resorting to an indirect approach for the other parameters in the BRDF used. This assumes that the BRDF is separable into different parts. The most important thing to notice is that all papers are based on BRDFs of two main types:

- The scene can be decomposed into regions with arbitrary variation of the diffuse reflectance (directional reflectance properties remain constant on each area).
- The scene must be decomposed into regions without spatial variation of the reflectance while using a BRDF function with higher dimensionality.

In general, we see that *indirect* methods use some sort of least squares for the optimization process, with the exception being <sup>YDMH99</sup> that uses the Nelder-Mead with Simulated Annealing for this part. From the reviewed results, it seems that the Levenberg-Marquardt numerical minimization method is one of the best suited for non-linear BRDF fitting.

Of the reviewed papers that deal with Inverse Emittance or combined problems, only <sup>MG97, Mar98</sup>



and <sup>KPC93</sup> present view-dependent goals, especially the first one because it deals with input given by photographs, while the second can also be used with view-independent goals.

From our survey we may derive that several inverse problems remain open:

- The application of Monte Carlo methods to Inverse Rendering problems has remained a vastly untouched topic with only one exception, and is an option that should be thoroughly researched and exploited in the future.
- The problem of Inverse Emittance in the presence of participating media has only been superficially touched. More research and development is needed.
- General, illumination-independent, Inverse Reflectometry approaches were also barely touched, both in the single-image and in the multiple-image variants.
- Combined problems, especially those with non-local illumination models, remain vastly untouched beyond the initial efforts analyzed in section 5. This is especially true for non-lambertian environments, since the seminal work by Kawai, Painter and Cohen<sup>KPC93</sup> covers the radiosity problem in a quite satisfactory way. We can also see that studying combined inverse reflectometry and lighting problems in the context of the full-radiance equation is a very promising and open line of research, since local illumination entails a rough approximation.
- Combining lighting reconstruction for different points<sup>RH01b</sup> across a scene, and performing correlations of the obtained information (a generalization of the work by Sato et al.<sup>SSI99c</sup>) seems an efficient and unexplored way to obtain the complete lighting of the scene, which could be very efficiently stored in some sort of lighting textures<sup>Hec90</sup>. This would allow the absolute positioning of light sources in the scene to be obtained.
- Choosing the luminaries' shapes among a small set of possible shapes is another option that is suggested by the results of Costa et al.<sup>CSF98</sup>.
- Letting an algorithm automatically choose among several possibilities the more suited material for a reflector to get a given illumination at a given position in the scene.

## 7. Acknowledgments

We would like to thank Pierre Poulin for his insights and Jack Howell for his reprints and for proofreading an earlier version of this paper, as well to the reviewers for their useful comments. This work was partially done under grant TIC2001-2392-C03-01 of CI-CYT and 2001SGR-0296 of Generalitat de Catalunya.

## Appendix A: Appendix

In this appendix we analyze, for each surveyed paper the mathematical aspects of its solution method. Tables 7 and 8 explain, for Inverse Emittance and Reflectometry respectively, the constraints imposed either on the solution method (for optimization approaches) or on the system of equations for direct inversion based approaches. In the first case the restrictions are incorporated into the optimization algorithm, while in the second the constraints are used in the system's building process, generally in a manual way. The third column presents the starting point used by the different methods when the used technique requires it. The fourth column presents either the objective function to optimize in the case of optimization approaches, or the method used to solve the direct problem presented.

Work	Type of approach	Approximations	Method used
SDS <sup>+</sup> 93	Indirect	Lambertian surfaces. Final impl: local illum.	CLS
OH95	Direct	Lambertian surfaces. Uniform spatial and angular subdiv.	Inverse Monte Carlo
HMH95	Direct	Lambertian surfaces + Radiosity constant on patches. Truncation of singular values.	MTSVD
FOH98	Direct + Finite Volume method for media + Radiosity constant on patches + lambertian surfaces and isotropic media.	Discretization. Truncation of singular values.	TSVD
MMOH97	Direct + Finite Volume method for media	Finite Elements. P1 Spherical Harmonics Approximation. Radiosity constant on patches + lambertian surfaces and isotropic media.	MTSVD
MC97 Mar98	Indirect (precomputed renderings for each light source)	Distant light sources only. Linear relationship between light sources and the mapping to image values.	Linear 1st order regularized Least Squares method with Generalized SVD for acceleration
DRB97, LFD <sup>+</sup> 99, FGR93	Indirect	Diffuse Surfaces with constant radiosity. Known Emitting Patches with known geometry.	Least Squares. DRB97 and LFD <sup>+</sup> 99 build a hierarchical radiosity system to speed up computations.
PF92	Indirect	Blinn's illumination model.	Analytical Optimization
PRJ97	Indirect	none	NLCO
CSF99	Indirect	none	SA (with possible final local search step)
SSI99b SSI99c	Direct	Distant sources, No interreflections, uniform reflectance, camera & known geom	CLS
SSI99a	Direct	Common features in stereo pair, no interreflections, uniform reflectance, camera & known geom	CLS
SL01	Indirect	None (although implemented as local illum)	Steepest descent
MAB <sup>+</sup> 97	Indirect	None	Selections dispersion and manual user selection
Gut00	Indirect	Known geometry and diffuse reflectances	Levenberg-Marquardt
RH01b	1st Indirect, 2nd Direct	Local illum. Simplified BRDF (diff + specular)	1) Linear LS with cutoff regularization. 2) Subtract diff comp of BRDF and recover angular-based illum resulting from specular BRDF.

Table 5: Classification of inverse rendering methods: Inverse Lighting. (NLCO = Non-Linear Constrained Optimization. MTSVD = Modified Truncated Singular Value Decomposition. LS = Least-Squares. CLS = Constrained LS).

Work	Type of approach	Approximations	Method used
PF95	Indirect	none	NLCO or NLCLS
Mar98	Indirect	BRDF w/ few spatially varying params. Discrete Sampl.	LS BRDF fitting (from photographs).
Mar98 MWL <sup>+</sup> 99	Direct	$f_r(r, \omega, \omega_i) \equiv f_r(\omega, \omega_i)$ . Point light sources.	Measurement from photographs.
YDMH99	Mixed	Low param BRDF & Scene decomp. in regs. w/ const. directional refl. props.	Iterative BRDF fitting: NLLS steps.
DeB98	Indirect	Scene subdivided in far & local regions.	Iterative BRDF fitting.
POF98	Direct	Purely Lambertian surfaces.	Color: sampling the reproj. texels in the corresp. imgs.
YM98	Indirect	Pseudo-BRDFs. 3 sources: Sun (parallel), Sky (fitted) and Landscape (Env. Map). Spec lobe from sky & landscape const on large regs	Diff comp: direct measure. Sky & landscape spec lobe: LS. Sun spec lobe: Nelder-Mead w/ SA.
FGR93 DRB97	Direct	Diffuse Surfaces. Does not distinguish between colour and shadows.	Reflectance proportional to average radiosity
LFJ <sup>+</sup> 99	Direct	Assumes known exact emittance from surfaces	As DRB97, but refl. is avrg over imgs. confidence weights.
LDR00	Indirect	Diffuse surf	As DRB97 + blocked light from rad. sol. + heuristic corr. based on illum. on ref. patch.
OSRW97	Direct	No highlights in most images	Multiresolution textures by recursive reprojection
DHT <sup>+</sup> 00 SW197	Indirect	Assumes spec color = light color. No changes in ind. illum.	Separation of diff & spec parts. Direct recov of BRDF params
RH01b	Indirect	First Terms in Sph Harm series	Simplex method for nonlinear params & linear fitting for rest
BC01	Indirect	Hyerarchically tests simple to complex BRDF models	Simple Models: Iterations with image-based error. Complex: Simplex Method.
IS91	Indirect	Constant BRDF over regions (restricted Torrance-Sparrow BRDF)	Iterative LS fitting
KC94	Indirect	Requires Several Photographs	BRDF sample inversion (nohighlight: LS; highlight: separable LS with Reg)
S196	Mixed	Separable diffuse & specular parts	Separation of diff & spec and fitting of Torrance-Sparrow by Levenberg-Marquardt
SW197	Mixed	Local illum. Convex Objects. Interpolated BRDF params	1st measured, 2nd fitted Torrance-Sparrow
DvGNK99	Direct	Totally known lighting	Measurements by Spectrophotometer
LKG <sup>+</sup> 01	Indirect	Full knowledge of Lights, camera & object geom	1st BRDF sample classification, 2nd Levenberg-Marquardt method for each samples cluster

Table 6: Classification of inverse rendering methods: Inverse Reflectometry.

NLCO = Non-Linear Constrained Optimization. LS = Least Squares. SA = Simulates Annealing. NLCLS = Non-Linear Constrained Least-Squares. NLLS = Non-linear Least squares.

	Constraints imposed	Initial guess / Initial approx.	Objective function (?) / direct inversion method	Tests Performed
SDS <sup>+</sup> 93 (CLS)	All weights $\geq 0$ .	Any value	Minimize: $\ \sum_i \omega_i \phi_i - g^{exp}\ $ where $\phi_i$ is the effect of the $i$ -th light source.	Complex env. with 19000 polys, 27000 vert. and 12 (14) light sources.
OH95	None	Initial emittance distribution proposed for unknowns.	Inverse MC method (requires high number of samples)	Rectangular 2D env. with one surf with flux & heat unknown.
HMH95	# surfaces w/ flux and emissivity known = # surfaces w/ both unknown.	Use initial SVD solution to guide algorithm.	Invert algebraic system using MTSVD.	Rectangular 2D env. with one surf with flux & heat unknown.
FOH98	# of vol. elem. with both emissivities and source strength unknown = # of surf. elem. with known flux and emissivity	Use initial SVD solution to guide algorithm.	MTSVD for resulting algebraic system (both media and surfaces discretized).	Rectangular 2D env. filled with isotropic hom. media, with one surf with flux & heat unknown.
MMOH97	None	PDE of media with "cold wall" boundary conditions. Use SVD solution to guide MTSVD algorithm.	Solve Volume PDE by Finite Volume method. Invert resulting algebraic system for surfaces only using MTSVD.	Rectangular 2D env. filled with isotropic hom. media, with one surf with flux & heat unknown.
MG97, Mar98	Positive coefficients.	Any value	Linear combination of basis images to match a given photograph.	Relighting: Photos + 3D model to modify a given photo.
FGR93, DRB9 LDR00	$\Sigma$ emittances fits estim total intensity	Radiosities from images	Fitting of the sources power	Acquired real video scenes.
PF92	None	Any value	$(\mathbf{N} \cdot \mathbf{H})^n$ , for dir $\mathbf{L}$ and glossiness $n$ . Use shadow volumes for final position.	Simple primitives. e.g.: cone under triangular light.
PR97 (NLCO)	- in cone defined by each point - same side of normal at point - right side of half-cone WRT cone axis - Umbra: $\forall$ Light Vertex (LV), Sketched Point (SP), $LV \in$ cone SP. - Penumbra: $\forall LV, \exists SP, LV \in$ cone SP. - within cone defined by each sketched point. - same side of normal at sketched point. - right side of half-cone WRT cone axis.	Above ellipsoid centered on occluder and radius $r$ along direction from center of mass of the sketched points to ellipsoid center. Just above an ellipsoid $E$ centered on occluder along dir. from mass center of the sketched points to $E$ center Any value	Maximize: $\sum (p_i - p_i)^2 =$ distance from every sketched points $p_i$ to light position $p_l$ . Maximize: $\sum (p_i - p_i)^2 =$ distance from every sketched points	Several primitives: 0.003 sec/Sketched point - interactive speed.
CSF99	used-defined (added as penalties in the objective function)	User-provided	Minimize: Value of Point Light within the implicit cone equation.	Several primitives: 0.01 sec/Sketched point when source is large WRT cones. 0.015 sec/sketched point.
SSI99b, SSI99c	Positive coefs.	Lambertian coef. estim. for shadow surface	User-provided by a script.	1) room, luminarie & table with homogeneous desired illum. 2) office (10550 obj), goal: reduce glare, shadow boundaries & redundancy.
SSI99a	None	Any Value	Linear System of Eqs: sources vs. pixels in photographs	simple polyhedra on tablecloth.
SL01	none	Light dirs = centroid of normals of viewed surfaces. Intensities = average desired to viewed intensities ratio Sha01	Find bright features in each onmi-image and correlate. Map remaining as image. Perception-based image quality metric (extent in communication scene info)	Indoor & outdoor real scenes isolated complex geometrical objects
MAB <sup>+</sup> 97	None	User-defined initial lighting	Manual user optimization	Complex scenes
RH01b	None	Any value	1)LS for Spherical Harmonics coeffs. 2.1) subtract diff comp from refl field 2.2) recover high-res angular-space illum	Complex Objects
Gm00	Sources inside regions	Intersections of $m$ directional lights calc. from diffuse surfaces	Least-squares error between the surface radiances and the measured rad.	Simple Cornell Box-like scenes

Table 7: Summary of the mathematical aspects of the inverse lighting approaches presented: Inverse Lighting, SVD = Singular Value Decomposition, MTSVD = Modified Truncated SVD, MC = Monte Carlo, PDE = Partial Differential Equation.

Work	Constraints imposed	Initial guess / Initial approx.	Objective function (?) / direct inversion method	Tests Performed
PT95 with $n \leq points$	Values $\geq 0$ .	a) w/ small glossiness ( $n$ ). b) conv sol as init guess for next it.	User-defined terms in the BRDF.	Phong-Blinn BRDF: 1st Ambient, 2nd Diffuse, 3rd specular
PT95 with $n \geq 3$ points	Values $\geq 0$ . Weight points with most important contribution.	Any value	Distance between BRDF and user-provided samples.	Phong-Blinn BRDF.
Mar98	None	Any value	Weighted Squared error between BRDF and measured samples.	Synthetic sphere w/ textures for calib. Squash. Rock.
Mar98, MWL <sup>†99</sup>	None	Any value	Direct measurement of $\omega, \omega_i$ , obtaining $f_r(\omega, \omega_i)$	Gray cyl for verif. Squash. Human skin. (300000 sample/channel)
YDMH99	None	rough estimate of BRDF	Different BRDF isotropic and anisotropic cases. Plastics different from metals.	Synthetic test: room w/ mutual illum. Real test: room w/ furniture, metallic balls and coloured posters.
De98	None	Initial guess (user)	Diff. between rendered local scene and photos.	Real complex scenes w/ simple synthetic objects.
POF98	None	Any value	All polys textures fitted to one another, corresp. colors combined WRT set of crit. to form a unique text.	320 x 400 textures for book from real photographs.
YM98	None	Any value	Diff: measure, Landscape spec: LS diff pseudo-BRDF & recorded rad. Sun spec: robust param recovery <sup>PTV92</sup> .	Relighting of existing, real tower.
FG93 DRB97	None	Any value	Avg reflect proportional to radiosity	Acquired video image scene
LDR00	None	Any Value	1st: eval. as <sup>PTB97</sup> . 2nd: Add blocked light from rad. sol. 3rd: Heur. correc: find unoccluded ref. patch for color ref.	Acquired real video Image scene
OSRW97	None	Any Value	Multiresolution quadtree textures re-covered by recursive reprojection.	Pictures & simple textured objects
DHT <sup>†00</sup>	None	Spec color = source color. Diff color = median RGB measurements.	Direct for spec. Lambertian fit.	Change viewpoint, illum & matching real-world illum.
SWI97	None	Spec color = source color. Diff color = separated & fitted to Torrance - Sparrow.	Direct for spec. Lambertian fit.	Coffee cup.
RH01b	None	Any value	First Terms in Sph Harm series: BRDF = Diffuse + slow-varying terms + fast-varying specular lobe.	Real-world models.
BG01	None	Diffuse Surfaces used first. Then more complex BRDFs.	Ratio of average of object image radiance to object rendered image.	Office Scene
TS91	Constant BRDF over regions	Intensity pixel values with normal incident light	RMS: 1) observed brightness and orientation from range finder, 2) Log(spec comp) and log(BRDF spec part)	Simple Objects
KC94	Diffuse coef smooth on range-deficient(rd) regions, specular over non-spec & rd, roughness everywhere	Any	Linear LS error in BRDF values	Textured spheres
SI96	None	Any	not specified (fit of specular comp)	Complex objects
DVGK99	None	Any	Measurement of BRDF samples with diff. comb. of viewing & lighting	Database of materials
LKG <sup>†01</sup>	None	Params from avg BRDF	RMSE b/ sampled & predicted values	texture recovery from complex objects

Table 8: Summary of the mathematical aspects of the inverse lighting approaches presented: Inverse Reflectometry. SVD = Singular Value Decomposition. MTSVD = Modified Truncated SVD. MC = Monte Carlo. PDE = Partial Differential Equation.

## References

- Arv95a James Arvo. *Analytic Methods for Simulated Light Transport*. Ph.D. thesis, December 1995.
- Arv95b James Arvo. The Role of Functional Analysis in Global Illumination. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 115–126, New York, NY, 1995. Springer-Verlag.
- BG01. Samuel Boivin and Andre Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001)*, pages 107 – 116, August 2001.
- Bli77. James F. Blinn. Models of light reflection for computer synthesized pictures. In *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '77 Proceedings)*, pages 192–198, 1977.
- BT75. P. Bui-T. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, June 1975.
- CSF98. Antonio Cardoso Costa, A. Sousa, and F. Ferreira. Design de iluminação. In *8th Portuguese workshop on Computer Graphics*, Coimbra, Portugal, 1998. in Portuguese.
- CSF99. A. Costa, A. Sousa, and F. Ferreira. Lighting design: A goal based approach using optimization. In D. Lichinski and G. Ward Larson, editors, *Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering)*, pages 317–328, New York, NY, June 1999. Springer-Verlag.
- CW93. Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Boston, MA, 1993.
- Deb98. Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '98 Proceedings)*, pages 189–198, 1998.
- DHT<sup>+</sup>00. Debevec, T. Hawkins, C. Tchou, H-P Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2000)*, pages 145–156, 2000.
- DM97. Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '97 Proceedings)*, pages 369–378, 1997.
- DRB97. George Drettakis, Luc Robert, and Sylvain Bougnoux. Interactive common illumination for computer augmented reality. In Julie Dorsey and Phillip Slusallek, editors, *Rendering Techniques '97 (Proceedings of the Eighth Eurographics Workshop on Rendering)*, pages 45–56, New York, NY, 1997. Springer Wien.
- DvGNK99. Justin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999.
- FGR93. Alain Fournier, Atjeng S. Gunawan, and Chris Romanzin. Common Illumination Between Real and Computer Generated Scenes. In *Proceedings of Graphics Interface '93*, pages 254–262, San Francisco, CA, May 1993. Morgan Kaufmann.
- FOH98. França, M. Oguma, and J. R. Howell. Inverse radiative heat transfer with nongray, nonisothermal participating media. In R.A. Nelson, T. Chopin, and S.T. Thynell, editors, *ASME HTD*, volume 361-5, pages 145–151. IMECE, November 1998.
- FvDFH00. James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics, Principles and Practice, Second Edition*. Addison-Wesley, Reading, Massachusetts, 1990.
- Gla95. Andrew S. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Francisco, CA, 1995.
- Gui00. Erwan Guillou. *Simulation d'environnements complexes non lambertiens à partir d'images: Application à la réalité augmentée*. PhD thesis, Mathématiques. Informatique, Signal et électronique et Télécommunications. IFSIC/IRISA, 2000. in french.
- Had02. J. Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. In *Bull. Univ. Of Princeton*, pages 49–52, 1902. In French.
- HB89. Berthold K. P. Horn and Michael J. Brooks. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
- Hec90. Paul Heckbert. Adaptive Radiosity Textures for Bidirectional Ray Tracing. In *Computer Graphics (ACM SIGGRAPH '90 Proceedings)*, volume 24, pages 145–154, August 1990.
- HH90. P. Hanrahan and P. Haeberli. Direct wisiwyw painting and texturing on 3d shapes. *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '90 Proceedings)*, 24:215–223, 1990.
- HMH95. Harutunian, J. C. Morales, and J. R. Howell. Radiation exchange within an enclosure of diffuse-gray surfaces: The inverse problem. In *Inverse Problems in Heat Transfer, ASME/AIChE National Heat Transfer Conference*, Portland, 1995.
- HSA91. Pat Hanrahan, David Salzman, and Larry Aupperle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics (ACM SIGGRAPH '91 Proceedings)*, volume 25, pages 197–206, July 1991.
- HSS92. P.C. Hansen, T. Sekii, and H. Shibahashi. The modified truncated svd method for regularization in general form. *SIAM - J. Sci. Stat Comput.*, 13(5):1142–1150, 1992.
- IS91. Katsushi Ikeuchi and K. Sato. Determining reflectance properties of an object using range and brightness images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1139 – 1153, November 1991.
- Kas92. Michael Kass. Inverse problems in computer graphics. In Nadia Magnenat Thalmann and Daniel Thalmann, editors, *Creating and Animating the Virtual World (Computer Animation '92, Fourth Workshop on Computer Animation)*. Springer-Verlag Tokyo, May 1992.
- KC94. G. Kay and T. Caelli. Inverting an illumination model for range and intensity maps. *Image Understanding*, 59:183–201, 1994.
- KPC93. John K. Kawai, James S. Painter, and Michael F. Cohen. Radiooptimization - Goal Based Rendering. In *Computer Graphics Proceedings, Annual Conference Series, 1993 (ACM SIGGRAPH '93 Proceedings)*, pages 147–154, 1993.
- LD00. C. Loscos and G. Drettakis. Low-cost photometric calibration for interactive relighting. In *First French-British International Workshop on Virtual Reality*, Brest, France, 2000.
- LDR00. Celine Loscos, George Drettakis, and Luc Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, 6(4):289–305, October-December 2000.
- LD<sup>+</sup>99. Loscos, M. C. Frasson, G. Drettakis, B. Walter, X. Grainer, and P. Poulin. Interactive virtual relighting and remodeling of real scenes. In *Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering)*, pages 329–340, New York, NY, 1999. Springer Wien.
- LD<sup>+</sup>99. Eric P. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. In *Computer Graphics (ACM SIGGRAPH '97 Proceedings)*, volume 31, pages 117–126, 1997.
- LKG<sup>+</sup>01. Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatially varying materials. Research Report MPI-I-2001-4-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, March 2001.

- MAB<sup>+</sup>97Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '97 Proceedings)*, pages 389–400, 1997.
- Mar98.SyStephen R. Marschner. *Inverse Rendering in Computer Graphics*. PhD thesis, Program of Computer Graphics, Cornell University, Ithaca, NY, 1998.
- MG97.Stephen R. Marschner and Donald P. Greenberg. Inverse lighting for photography. In *Proceedings of the IS&T/SID Fifth Color Imaging Conference*, pages 262–265, Scottsdale, AZ, November 1997. Society for Imaging Science and Technology.
- MMOH7C. Morales, M. Matsumura, M. Oguma, and J. R. Howell. Computation of inverse radiative heat transfer within enclosures. In *Proc. 1997 ASME National Heat Transfer Conference*, Baltimore, August 1997.
- MWL<sup>+</sup>9Stephen R. Marschner, Stephen H. Westin, Eric P. F. Lafortune, Kenneth E. Torrance, and Donald P. Greenberg. Image-based brdf measurement including human skin. In D. Lichinski and G. Ward Larson, editors, *Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering)*, pages 131–144, New York, NY, June 1999. Springer-Verlag.
- OH95. M. Oguma and J. R. Howell. Solution of the two-dimensional blackbody inverse radiation problem by inverse monte carlo method. In *ASME/JSME Thermal Engineering Conference*, volume 3, pages 243–250. ASME, 1995.
- OSRW97. Ofek, E. Shilat, A. Rappoport, and M. Werman. Multiresolution textures from image sequences. *IEEE Computer Graphics and Applications*, 17(2):18–29, 1997.
- PF92. Pierre Poulin and Alain Fournier. Lights from highlights and shadows. *Computer Graphics*, 25(2):31–38, March 1992.
- PF95. P. Poulin and A. Fournier. Painting surface characteristics. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 119–129, New York, NY, 1995. Springer-Verlag.
- POF98Pierre Poulin, Mathieu Ouimet, and Marie-Claude Frasson. Interactively modeling with photogrammetry. In *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)*, pages 93–104, New York, NY, June 1998. Springer Wien.
- Pou93.Pierre Poulin. *Shading and Inverse Shading from Direct Illumination*. PhD thesis, Department of Computer Science, University of British Columbia, Vancouver, British Columbia, 1993. Available from <http://www.cs.ubc.ca/nest/imager/th.html>.
- PP00. G. Patow and X. Pueyo. A survey on reflector design problems. Technical Report TR-IIIA 00-16-RR, Universitat de Girona, 2000.
- PRJ97Pierre Poulin, Karim Ratib, and Marco Jacques. Sketching shadows and highlights to position lights. In *Proceedings of Computer Graphics International 97*, pages 56–63. IEEE Computer Society, June 1997.
- PT91. W. H. Press and S.A. Teukolsky. Simulated annealing optimization over continuous spaces. *Computers in Physics*, pages 426–429, Jul/Aug 1991.
- PTVF9W. H. Press, S.A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.
- RH01a.Ravi Ramamoorthi and Pat Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *Journal of the Optical Society of America (JOSA A)*, pages 2448–2459, October 2001.
- RH01bRavi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001)*, pages 117–128, August 2001.
- SDS<sup>+</sup>9Chris Schoeneman, Julie Dorsey, Brian Smits, James Arvo, and Donald Greenberg. Painting With Light. In *Computer Graphics Proceedings, Annual Conference Series, 1993 (ACM SIGGRAPH '93 Proceedings)*, pages 143–146, 1993.
- Sha01.R. Shacked. *Automatic lighting design using a perceptual quality metric*. PhD thesis, The Hebrew University of Jerusalem, Israel, 2001. <http://www.cs.huji.ac.il/~danix/ldesign>.
- SI96. Y. Sato and K. Ikeuchi. Reflectance analysis for 3d computer graphics model generation. *CVGIP Graphical Models and Image Processing*, 58(5):437–451, 1996.
- SL01. R. Shacked and D. Lischinski. Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum (Eurographics 2001)*, 20(3):215 – 227, 2001.
- SP94. Francois Sillion and Claude Puech. *Radiosity and Global Illumination*. Morgan Kaufmann, San Francisco, CA, 1994.
- SSI99aI. Sato, Y. Sato, and K. Ikeuchi. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics*, 5(1):1–12, 1999.
- SSI99bI. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *Proceedings of IEEE ICCV'99*, pages 875–882, 1999.
- SSI99cI. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from shadows. In *Proceedings of IEEE Conf. on Computer Vision and Pattern Recognition'99*, pages 306–312, 1999.
- SWI97Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '97 Proceedings)*, pages 379–387, 1997.
- WW92Alan Watt and Mark Watt. *Advanced Rendering and Animation Techniques: Theory and Practice*. Addison-Wesley, 1992.
- YDMH9Zhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *Computer Graphics Proceedings, Annual Conference Series (Proc. ACM SIGGRAPH '99 Proceeding)*, pages 215–224, August 1999.
- YM98. Yizhou Yu and Jitendra Malik. Recovering photometric properties of architectural scenes from photographs. In Michael Cohen, editor, *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '98 Proceedings)*, pages 207–218. Addison Wesley, July 1998.

# Classification of illumination methods for mixed reality

Katrien Jacobs<sup>†</sup> and Céline Loscos<sup>‡</sup>

Department of Computer Science, University College London, UK

---

## Abstract

A mixed reality (MR) represents an environment composed both by real and virtual objects. MR applications are used more and more, for instance in surgery, architecture, cultural heritage, entertainment, etc. For some of these applications it is important to merge the real and virtual elements using consistent illumination. This report proposes a classification of illumination methods for MR applications that aim at generating a merged environment in which illumination and shadows are consistent. Three different illumination methods can be identified: common illumination, relighting and methods based on inverse illumination. In this report a classification of the illumination methods for MR is given based on their input requirements: the amount of geometry and radiance known of the real environment. This led us to define four categories of methods that vary depending on the type of geometric model used for representing the real scene, and the different radiance information available for each point of the real scene. Various methods are described within their category.

The classification points out that in general the quality of the illumination interactions increases with the amount of input information available. On the other hand, the accessibility of the method decreases since its pre-processing time increases to gather the extra information. Recent developed techniques managed to compensate unknown data with clever techniques using an iterative algorithm, hardware illumination or recent progress in stereovision. Finally a review of illumination techniques for MR is given with a discussion on important properties such as the possibility of interactivity or the amount of complexity in the simulated illumination.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: augmented reality, mixed reality, common illumination, relighting, inverse illumination

---

## 1. Introduction

To understand the concept *mixed reality* it is necessary to classify the different types of environments that can be generated with a computer. Milgram et al. [MK94, OT99] present such classification based on the amount and type of virtual and real elements that constitute the resulting world. In their classification, all possible environments form one continuum called *reality-virtuality continuum (RV)*, see Figure 1. In this continuum, four worlds can be identified that have an outspoken character. These four worlds lie next to each other in the RV continuum and might even overlap. The first and most straightforward of these is the real world without any addition of virtual elements; it will be referred to as *reality* and it lies on the left end of the RV contin-

uum. In the second This world is referred to with the term *augmented reality (AR)* [Azu95][ABB\*01][BKM99]. In an opposite scenario, the world consists of a virtual environment, augmented with real elements. This world is consequently called an *augmented virtuality (AV)*. The last and fourth world doesn't contain any real elements and is therefore labelled as a *virtual environment (VE)*; it lies on the right end of the RV continuum. The term *Mixed Reality (MR)* refers to those worlds that are a mix of virtual and real elements, or, MR spans the RV continuum. In general methods that are developed for AR focus on real-time applications. Therefore they usually differ from methods that are specifically designed for MR applications whose focus can be on non real-time applications. This report will discuss the various existing illumination methods for MR applications in general.


Two different classes of AR exist; they differ in the realisation of the AR [MK94]. The first class groups the methods

---

<sup>†</sup> K.Jacobs@cs.ucl.ac.uk

<sup>‡</sup> C.Loscos@cs.ucl.ac.uk





**Figure 1:** Simplified representation of a Reality-Continuum [MK94][OT99]. Courtesy of Milgram et al.

for semi-transparent or see-through displays, examples are [SCT\*94][BGWK03]. This first class contains two different see-through display methods. The first method (optical AR method [Azu95]) projects the virtual objects on a transparent background, most likely the glasses of goggles. The second method (video AR method [Azu95]) uses a head-mounted display: a head-mounted camera records the environment and this image is projected inside the display together with the virtual objects. The second class of AR replaces the expensive see-through devices with a non-immersive display; it is usually called a computer-augmented reality (CAR). The quality of the immersion is higher for the first class than for the second. Nevertheless see-through devices are not always required by the application: urban planning, architecture and some applications in the entertainment industry are satisfied with the second type of CAR display methods. In earlier approaches of AR, virtual objects were positioned on top of a real environment. Calibration and registration are difficult processes and for long the focus lied upon taking into account the possible occlusion and collision effects, while no further adaptations on real and virtual objects were carried out. In other words, after the inclusion, no resulting shadows were generated and no lighting changes were put through. An AR system of such kind does not yield a high level of realism, consistency between objects is restricted to geometric aspects. Nowadays three illumination techniques can be identified that attempt to raise the quality of AR and in general MR: *common illumination*, *relighting* and *inverse illumination for relighting or common illumination*. These techniques vary in the quality of the illumination and in the consistency obtained between the illumination of the real and virtual objects.

The most straightforward method of these three results in the addition of shadows in the MR environments. Generating shadows is just as important as taking into account occlusions, since they help situating the objects in the scene and give information about the distance between different objects [SUC95]. A higher level of realism can also be obtained when the local light interaction between real and virtual objects is incorporated in the MR scene. Simulating such effects results in *common illumination*. An example of an application that uses common illumination to improve the MR can be found in the movie industry. Special effects in movies make an effort to mix lighting effects and reflections as realistic as possible, resulting in brilliant graphical

effects in recent movies such as *Jurassic Park*, *Harry Potter* and *The Lord of the Rings* trilogy. In these movies computer-generated effects are blended entirely with the real footage; usually this is carried out by hand.

Some methods allow to change the original illumination, hereby influencing the appearance of virtual and real objects. An example of an application domain for this method is architecture. Being able to virtually change the lighting conditions in the real scene makes it possible to see the impact of a new building in a street under different lighting conditions without the need of recording the real environment under all these different conditions. Another application area is crime investigation [HGM00]: a recording of a scene at a certain time can be changed to the illumination at a different daytime, making it possible to visualise the perception of the criminal at the time of the crime. Techniques that focus on virtually changing the illumination of an existing scene are simply known as *relighting techniques*.

The techniques brought together in a third category are based on more complex estimations of the reflectances in the environment in order to provide more accurate results. The process of estimating the reflectances (bidirectional reflectance distribution function or BRDFs) from an existing lighting system is called *inverse illumination*. It was originally developed to give more realism in computer graphics. Reflectance properties of objects were estimated in order to reproduce a more realistic simulation of virtual scenes. In the context of MR, inverse illumination techniques aim at making a correct estimate of the photometric properties of the objects in the scene. While other techniques search for acceptable solutions for the new illumination problem, inverse illumination makes it possible to produce relit images that aim to be an *exact* replica of the real conditions. A full discussion of the current state of the art of inverse illumination techniques can be found in [PP03], while Ramamoorthi and Marschner [RM02] present a tutorial on some of the leading research work in this area.

At the moment, fairly good techniques exist that can relight an augmented scene with a different illumination. It is getting more difficult to differentiate between virtual objects and real objects. The main limitation of most techniques is the tedious pre-processing time and the slow update rate, which excludes real-time applications. When a geometric model of the scene is required, the user will have to create one, usually in a semi-manual and error-prone manner. The scene update rate is often too slow to allow real-time user interaction, even with the current progress in computer hardware and software. The research focus is moving towards using hardware for the calculation instead of software to accelerate computation. Early results are promising, but more research needs to be carried out in this area. This report does not review all existing work. Instead it concentrates on illumination techniques in MR that are meant for large environments. When optimized and extended, these techniques

will be widely applicable in real-time applications, for instance see-through display in AR. Several techniques exist for relighting human faces [Mar98][DHT\*00], or that focus on local objects or simple scenes [ALCS03][SWI97]. These techniques are classified mainly in the domain of inverse illumination as the emphasis was placed on this aspect in the referenced papers. Although these techniques are designed for small objects they can be used to build extremely useful and strong methods for illumination in MR but they will not be further discussed in this paper.

This report discusses the state of the art of those techniques that strive to solve the problem of illumination in MR environments and gives an objective evaluation of their quality. Section 2 describes in more detail the context of this review and the assessment of the criteria on which the classification is based. Section 3 gives a structured overview of all the illumination methods that were developed for MR. A further qualitative evaluation and discussion of the different techniques is given in section 4. Finally section 5 gives conclusions and presents the necessary future work for illumination methods for MR.

## 2. Problem Assessment

### 2.1. Objective and difficulties

The classes described above are not necessarily designed to lure the users into believing that what they see is real. For instance VR often aims at trying to create the perception of a real world, without necessarily using convincing real imagery. Some AR systems merely add data displays to real scenes, making no attempt to mix the two seamlessly. This report considers MR scenes that *do* try to convince the users of *believing* that a real world is surrounding them and will use this as a measure to assess the quality of the method.


An MR is convincingly real when it is impossible to separate the virtual elements from the real elements in the resulting environment. Four critical success factors were identified that need to be present in the MR in order to be convincingly real:

- **After including the virtual object(s), the resulting scene needs to have a consistent shadow configuration**[SUC95]. The main difficulty to obey this requirement is to find the correct appearance of the new shadows: their position in the scene, shape and colour. Sometimes these are estimated, but they can be calculated exactly if the geometry of the scene, the illumination characteristics and the material properties of all objects in the scene are known.
- **The virtual object(s) must look natural.** A cartoon-like virtual object is easily detectable and therefore efforts have been made to model objects that look realistic. One successful technique is *image-based modelling*, in which objects are rendered with textures based on real images.
- **The illumination of the virtual object(s) needs to resemble the illumination of the real objects.** There are two possible methodologies to achieve this requirement. Either the illumination pattern of the real scene is known, which in turn is used to illuminate the virtual object or all material properties of all objects in the scene are known or estimated, which allows the entire scene to be relighted with a consistent known illumination pattern.
- **If the user can interact with the MR environment, it is clearly important that all update computations occur in real-time.** Any delay in the interaction will remind the user of the fact that what is seen is unreal [MW93]. The requirement of a real-time system is one of the most difficult to achieve, especially when no pre-processing time is allowed.

### 2.2. Assessment of existing techniques

The ultimate objective of the aforementioned techniques is defined by the amount of realism perceived by the user. This inherent subjectivity complicates an objective assessment of the various techniques. In this section a few quality criteria are listed that will be used in section 4 to assess the presented methods:

- **Amount of realism.** In some cases it is impossible to evaluate the amount of realism without using a statistical measure. For instance, a test audience can evaluate the technique, if the test group is large enough, a statistical value can be derived from the group evaluation. Alternatively, if the inserted virtual object is an exact replica of an existing real object, it is possible to give an exact value of the amount of realism in the produced scene. It suffices to compare the generated scene with an image of the real object in the same scene. The difference between the two gives a measure of the level of realism.
- **Input requirements.** It is expected that the more input data is available, the higher the quality of the end result will be. On the other hand, the usability of the system reduces with the complexity of the input data. Possible input data are: the geometry, the light position, the illumination pattern and the material properties. This report gives a classification of the various techniques based on their input requirement.
- **Processing time.** the time needed to create the end result is another important characteristic of the method. To offer the user a highly realistic interactive environment, the computations need to be done in real-time. Unfortunately, this is very hard to achieve. If geometric and material properties of a scene need to be known, it is unavoidable that some pre-processing time needs to be incorporated. In general the usability of the proposed techniques depends on the amount of pre-processing time needed and the computation speed of the illumination technique.
- **Level of automation:** if the technique under consideration requires a considerable amount of manual interaction



**Figure 2:** Results for Sato et al. [SSI99]. The top row shows results for an indoor scene, the bottom row for an outdoor scene. The images on the left are the input images, the images on the right illustrate the resulting MR. Soft shadows are produced using local common illumination. Courtesy of Sato et al.

while processing the input data, the technique is less interesting than one that is automated.

- **Level of interaction:** a technique can be judged based on its dynamic character: the possibility to change the viewpoint of the camera, or the possibility to let the user interact with the environment for instance by moving around objects. A higher degree of interaction gives a greater usability of the method.

### 2.3. Methodology

The various existing techniques can be grouped into three different classes, based on the methodology used to solve the problem. They were already listed in the introduction and are further discussed in this section:

1. **Common illumination:** to this category belong all methods that provide a certain level of illumination blending, like the addition of shadows projected from real objects on virtual objects and shadows cast by virtual objects on real objects. These techniques do not allow any modification of the current illumination of the scene. Two different types of common illumination can be considered: *local* and *global* common illumination, referring to the type of illumination simulated. For local common illumination, there is usually no requirement of any BRDF information. For global illumination, it is often important to have an estimate of the material properties of the real

objects. The accuracy of this type of techniques depends on the accuracy of the known geometric model of the real scene. In Figure 2 an example is given of a rendering using global common illumination [SSI99].

2. **Relighting after light removal:** relighting techniques make it possible to change the illumination of the scene in two steps. Firstly, the current illumination effects of the real scene are analysed and possibly removed. Secondly, new illumination effects (shadows, intensity changes, addition of a new light, indirect lighting effects, etc.) are generated based on a new illumination pattern. These methods do not necessarily require an exact knowledge of the BRDF values of the real scene objects as for some methods the focus lies on generating a scene that *looks* realistic. These techniques require in general a detailed geometric model of the real scene. An example of a relighted scene using global illumination techniques [LDR00] is given in Figure 3.
3. **Physically based illumination:** this last category encloses those methods that make an attempt to retrieve the photometric properties of all objects in the scene often referred to by the term *inverse illumination*. These methods estimate BRDF values as correctly as possible as well as the emittance and positions of the light sources in the real scene. The BRDF values can be estimated using a goniometer [War92] or can be calculated based on the photometric equilibrium equations [Mar98] [SWI97].



**Figure 3:** Results for Loscos et al. [LDR00]. The image on the left hand side represents the real scene. The image on the right hand side shows the relighted synthetic scene for which real light sources have been virtually turned off and a virtual light source is inserted. Global common illumination updates are performed at interactive rates using an adapted version of [DS97]. Courtesy of Loscos et al.



**Figure 4:** Results for Yu et al. [YDMH99]. Left: the original input scene. Right: the result of illuminating the original scene using a different illumination pattern. The specular and diffuse parameters of the real objects are calculated. Courtesy of Yu et al.

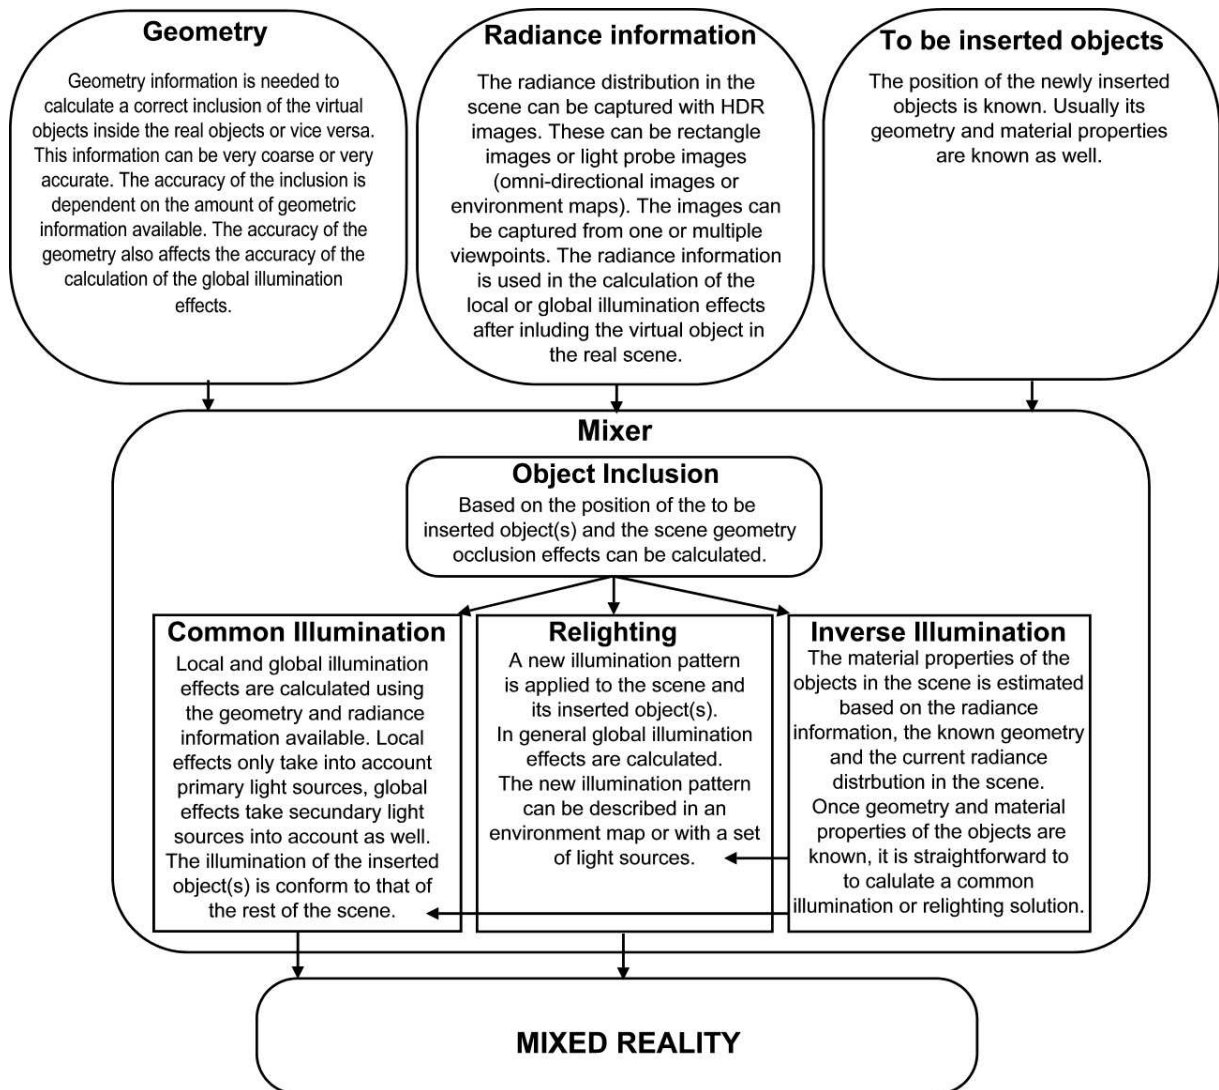
The BRDF information can be used for both common illumination or relighting methods. The accurate BRDF estimation often permits to perform a complete and realistic relighting, which takes both reflections and global illumination techniques into account. Patow et al. [PP03] give an in-depth overview of inverse illumination techniques. An example of inverse global illumination [YDMH99] is illustrated in Figure 4.

### 3. Classification of Illumination Methods for Mixed Reality

MR brings together those applications that create a new environment, around or in front of a user, containing both real and virtual elements. Sections 2.1 and 2.2 formulated the objectives, the difficulties encountered and the assessment criteria of the MR systems. One of these criteria, the type of input requirements, regulates the accessibility and accuracy of the technique. This criteria will be used to classify the different methods. An overview of the input information needed to calculate a common illumination, relighting or inverse illumination solution is given in Figure 5.

The classification put forward firstly takes into account the required geometric model of the real scene, starting with the techniques that require no geometric model and ending with techniques that require a precise geometric model. In this report a geometric model is defined as a reconstruction of a part of the (or the entire) real scene with significant detail. The pre-processing workload for techniques that extract a basic geometric model, e.g. the depth at a low resolution, is significantly lower than those methods that do require a high-level geometric model. Therefore techniques using basic geometric information are classified under that group of methods that do not require a geometric model, as this will give a better indication of the amount of pre-processing time required for each different class.

Two different approaches exist to reconstruct a geometric model of the real scene. Either the scene is scanned with a scanning device [Nyl] [MNP\*99] [3Ds] or it is reconstructed using stereovision [HGC92][Har97][Fau92][Fau93]. The first option of using a scanning device gives a precise geometric model but is expensive and tedious. Often the model will capture too much detail, which is not always necessary and is difficult to manage for real-time applications. Ob-



**Figure 5:** An overview of the dataflow in illumination calculation for MR. Three input sources are considered: the scene geometry, the scene radiance and information about the to be inserted object(s). These are used to calculate the illumination solution. The accuracy of the integration depends on the amount of input information available.

jects such as trees and objects with a highly specular surface are for some scanning techniques difficult to model accurately. Instead of using a scanning device, modelling techniques based on stereovision can be used to reconstruct the geometry of a scene. Most methods described in this survey requiring a 3D model of the scene make use of this low cost solution. In general, the 3D reconstruction requires at least two images from two different viewpoints. However, the entire geometry of a scene cannot be captured with one image pair only, this would create gaps in the known geometry. Usually more than one image pair is required for a complete geometric reconstruction. The ease

at which this reconstruction can take place depends on the amount of information that is available for the camera(s) used. If the internal and external camera parameters are known, the reconstruction is easier. Internal parameters can be estimated in a relatively simple way. Recording the external parameters is more difficult and involves a precise and tedious capture process. Fortunately, the fundamental matrix of the stereovision system can be estimated based on the internal parameters only, if at least eight corresponding points are known [HGC92][Har97]. This was reduced to six known points [Fau92][Fau93] for a calibration relative to a scale factor, which is often sufficient. Having the fundamen-

tal matrix can ease the reconstruction but does not make it trivial. Research led to different types of systems: non constraint systems [FRL\*98][SWI97] and constraint systems [POF98][DTM96][DBY98][MYTG94]. Good commercial reconstruction software [Rea][Met][Eos][Int] exists, but most of them lack the option of reconstructing complex shapes and large environments. In general, we can conclude that systems requiring geometric information demand a long pre-processing time, and are not guaranteed to get an accurate geometric model. It is really important to recognize the geometric acquisition as a difficult problem, that still requires much more research efforts.

Our classification of methods is, parallel with the classification based on the geometric information, based on this amount of image data needed to reconstruct a MR environment. Hereby excluding the image data needed for retrieving geometric information. More precisely, the classification will use the amount of different input images used for the rendering or calculation process. These processes might use more than one image (or texture containing radiance information) for the same point in the real scene, for instance for the BRDF estimation.

Some projects adopted the concepts of High Dynamic Range Images (HDR images) [Lar91] which can be computed using techniques such as [DM97][MN99]. Each HDR image is generated based on a set of images taken from the same viewpoint of the same scene, but with a different exposure. The end-result is one image containing radiance values instead of ordinary RGB values. In other words, radiance values are not clamped in the RGB space. It may be argued that methods using HDR images should be classified under that class with methods that use more than one image for each point of the scene. However, this report considers that a HDR image provides one radiance value per point, and methods that use only one HDR image for a certain point of the scene, are therefore classified as methods requiring only one input image. Similarly, techniques that require a few or many HDR images are classified as methods using respectively a few or many images.

The following classification is used throughout the rest of this section:

1. **Model of the real scene unknown, one image known** (section 3.1): this category lists those techniques that do not require any model of the real scene, except for some low-level geometry like depth information. Any necessary radiance information of a certain point in the real scene is extracted from one single image.
2. **Model of the real scene known, one image known** (section 3.2): a geometric model of the real scene is available. Any necessary radiance information is extracted from one image only.
3. **Model of the real scene known, few images known** (section 3.3): again a geometric model of the scene is re-

quired. For a certain point in the scene, radiance information is available from a few different images.

4. **Model of the real scene known, many images known** (section 3.4): this class groups those techniques that require both a detailed geometric model of the real scene and radiance information from a large set of different images.

The rest of this section lists the most significant methods based on the above mentioned classification and briefly discusses their methodology. A discussion of the techniques based on the assessment criteria mentioned in section 2.2 is given in section 4.

### 3.1. Model of real scene unknown, one image known

To this challenging category, in terms of output quality, belong those methods that require very little relevant information about the real scene. Since no geometric model of the scene is available it might be necessary to calculate depth information of the scene, to allow a correct inclusion of the virtual objects, or some lighting information. For this group, all radiance information is extracted from one single image.

Nakamae et al. [NHIN86] were the first to propose a method for composing photographs with virtual elements. Input photographs are calibrated and a very simple geometric model of the real scene is extracted. The viewpoints of the photograph and the virtual scene are aligned to ensure an appropriate registration of the virtual objects within the photographed elements. The sun is positioned within the system according to the time and date when the picture was taken. The sun intensity and an ambient term are estimated from two polygons in the image. The illumination on the virtual elements is estimated and adjusted to satisfy the illumination in the original photograph. The composition is done pixel by pixel and at that stage it is possible to add fog. All parameters are very inaccurate and therefore the results are limited in accuracy. However, they were one of the first to mention the importance of using a radiometric model to improve the image composition. Figure 6 displays gives an example of the obtained results.

Techniques exist in computer graphics that use *environment maps* to render objects in a scene. They were introduced to approximate reflections for interactive rendering [BN76][Gre86][VF94]. These techniques can also be used to assist the rendering of glossy reflections [CON99][HS99][KM00] by pre-filtering a map with a fixed reflection model or a BRDF. At this moment, graphics cards extensions support the real-time use of environment maps, this encourages its use even more. Graphics cards now support cube maps [NVi]; ATI [ATI02] presented at SIGGRAPH 2003 a demonstration of a real-time application for high resolution. Environment maps can be used to represent the real scene in a MR environment as a panorama and the information from these images can be used to simulate re-





**Figure 6:** Results for Nakamae et al. [NHIN86]. The first image depicts the background scene, the second image shows the augmented scene where the illumination of the augmented objects are matched to their surroundings and the shadows are cast accordingly. Courtesy of Nakamae et al.

reflections on a vertical object positioned at the center of the environment map [Che95].

Agusanto et al. [ALCS03] exploited the idea of environment maps to provide reflections in AR. They use HDR images of the environment captured by a light probe to create the environment map. These maps are filtered off-line to decompose the diffuse from the glossy components. The rendering is then performed with a multi-pass rendering algorithm that exploits hardware capabilities. After some pre-processing, like the inclusion of shadows, they present results for MR environments rendered on a desktop. An impressive aspect of their work is that the method also works for real-time AR. The implementation of their method works with ARToolKit [KBBM99] and the results show interactive reflections from the real scene on virtual objects at interactive frame rate. An example of such a projection is given in Figure 7. Although it should be feasible, they have not yet provided a shadow algorithm for the AR application.

Sato et al. [SSI99] adopt a technique that extends the use of environment maps to perform common illumination. In their method, it is assumed that no geometry is known a-priori. However, at least a few images are known from different but very restricted and known viewpoints, which can be used to estimate a very simple geometry of the scene and the position of the light sources. The obtained geometry does not offer a reliable occlusion detection and the positions of the virtual object are therefore restricted to lie in front of all real objects in the real scene. After this low-level geometric reconstruction, a set of omni-directional images are captured with varying shutter speed. From these images, a radiance distribution is calculated, which in turn is mapped onto the geometry. To calculate the shadows and the local illumination a ray casting technique is adopted. The radiance values of the virtual objects are calculated using the information known about the light sources, the radiance values of the real scene, the geometry and the BRDF values of the virtual



**Figure 7:** Results for Agusanto et al. [ALCS03]. The virtual objects are rendered with skin textures. The left object is blended with a diffuse map and no soft shadows. The right object is blended with a glossy map and with soft shadows. Courtesy of Agusanto et al.

objects. To simulate the shadows cast by virtual objects on real objects, the radiance values of those points in the scene that lie in shadow are scaled. The simulated soft shadows look realistic, see Figure 2. Their geometric estimate is poor and therefore usability of the method is limited and the positions of the virtual objects are restricted. Nevertheless the method produces convincing local common illumination.

### 3.2. Model of real scene known, one image known

Most of the existing illumination methods assume that a geometric model of the scene is available. The more detailed the geometric model is, the more reliable the occlusion detection will be. Although not all techniques explain where this model should come from, it is doubtful that a perfect geometric model can ever be acquired and this should be taken

into account when evaluating a specific method. In this section a discussion is given of those methods that take a certain 3D geometric model of the real scene as input and extract radiance information from one single image. All methods that belong to this category are further divided into three groups based on the type of illumination they produce:

- local illumination for AR applications: 3.2.2
- common illumination: 3.2.1
- relighting: 3.2.3

### 3.2.1. Local illumination for AR

As mentioned before, AR has long been an area wherein people focused on registration and calibration as these are still difficult problems to solve in that area. However, a few papers tried to introduce shadows in their systems, to show how well the registration was done and to improve the rendering quality. Recent improvements in graphics hardware for rendering shadows made it possible to perform real-time rendering of shadows on well-registered systems where the geometry is known. Early work was presented by State et al. [SHC\*94] in which virtual objects are inserted in the see-through real scene. A real light source is moved around and tracked, and shadows of the virtual object due to this real light source are virtually cast onto real objects by using the shadow mapping technique [Bli88]. In this case, the light source is assumed to be a point light source. It was very promising that some researchers in AR were interested in using local common illumination in their systems, but it was followed by a long period in which no innovative material emerged. Only recently, additional work of Haller et al. [HDH03] and Jacobs et al. [JAL\*05] was carried out to add shadows from a virtual object onto real objects. Both methods use shadow volumes, and in order to get good quality results knowledge about the scene geometry is essential. However, Jacobs et al. start from a rough approximation of the geometry and use image processing techniques to compensate for approximation errors. Other methods exist [BGWK03] that we will not discuss here, since they will not be applicable in general MR systems because these systems would require the capture of a large environment.

### 3.2.2. Common illumination


Jancene et al. [JNP\*95] use a different approach to illuminate the virtual objects, they base their method, called RES (Reality Enriched by Synthesis), on the principle of composition. The objective is to add virtual animated objects in a calibrated video sequence. The final video is a composition of the original video sequence with a virtual video sequence that contains both virtual objects and a representation of the real objects. The geometry of the real object is reconstructed a-priori so that for each frame in the video the geometry is known. The rendering in the virtual sequence is performed using ray tracing. It is possible to modify the reflectance properties of real objects. Shadows are simulated

in the virtual sequence, the impact of the shadows in the final video is acquired by modifying the original video with an attenuation factor. An occlusion mask is created to reflect occlusion between virtual and real objects. This method came quite early in the history of common illumination and video composition. Even though it is not applicable for real-time applications, it allows local common illumination and virtual modification of the reflectance properties of real objects. The images on the left in Figure 8 illustrate the original scene, the images on the right illustrate the composition.

Gibson and Murta [GM00] present a common illumination method, using images taken from one viewpoint that succeeds in producing MR images at interactive rates, by using hardware accelerated rendering techniques. Apart from constructing the geometry of the scene, the pre-processing involves creating a set of radiance maps based on an omnidirectional HDR image of the entire scene. New virtual objects are rendered via a spherical mapping algorithm, that maps the combination of these radiance maps onto the virtual object under consideration. Later shadows are added using a two step algorithm. To simulate the shadows, a set of  $M$  light sources are identified, which imitate the true, unknown illumination in the scene. Each light source is assigned a position and two parameters  $\alpha_i$  and  $I_i$ , which define the colour of the shadow. For each light source, a shadow map is calculated using efficient hardware calculations (z-buffer). Shadow mapping is an intensively technique supported by the graphics hardware that helps to create shadows in a fast and efficient way. The shadows created with shadow maps are in nature hard shadows and therefore unsuitable for realistic shadow generation. Gibson and Murta combine the  $M$  shadow maps in a specific way, using the above-mentioned parameters and now the system succeeds in simulating soft shadows, looking almost identical to the solutions obtained with a more computational and traditional ray-casting algorithm, see Figure 9. The system of  $M$  light sources needs to be defined so that it represents a close replica to the current illumination system, an increase in number of light sources affects the rendering time. To demonstrate their method, Gibson and Murta used eight light sources to simulate an indoor environment. The position and the parameters of the light sources are defined via an optimisation algorithm, which needs to be executed only once for each different scene.

Debevec [Deb98] presents a more advanced common illumination technique that estimates the BRDF values for a small part of the scene. It is argued that if a virtual object is inserted into the scene, only a small fraction of the scene experiences an influence from that inclusion. Relighting techniques using inverse illumination therefore only require the BRDF values of those points that lie in this fraction. Since for most applications it is possible to know the position of the virtual objects, Debevec uses this position to divide the entire scene into two parts: the *local scene* and the *distant scene*. The local scene is that fraction of the scene whose





**Figure 8:** Results for Jancene et al. [JNP\*95]. The images on the left hand side show the original scene and the registration of the cardboard box within this scene. The images on the right hand side show two screen shots from the video sequence in which a virtual dynamic green ball and static pink cube have been added to the original scene. The reflection of the green ball is visible on the board behind it. Courtesy of Jancene et al.



**Figure 9:** Results for Gibson et al. [GM00]. Comparison of a ray-traced (left) and a hardware generated image (right). The ray-traced image was generated using RADIANCE [War94], the hardware generated image made use of the rendering method described in [GM00]. The generation of the ray-traced image took approximately 2 hours, the generation of the hardware rendered image took place at nearly 10 frames-per-second. Courtesy of Gibson et al.

appearance might alter after inclusion and the BRDF of the materials in that part need to be estimated. On the other hand, the distant scene is that part of the scene that undergoes no physical alteration after inclusion. A schematic overview of the division in local and distant scene and their corresponding influences is presented in Figure 10. The local scene is restricted to be diffuse only; the distant scene has no restrictions. An omni-directional HDR image is captured using a mirrored ball. The resulting *light probe image* is used to present the illumination in the real scene. Based on the geometric model, the light probe image and the division into local and distant scene, the BRDF values in the local scene are estimated. The calculations are straightforward, since only diffuse BRDF values are considered. A *differential rendering* technique was developed to reduce the possible inconsistencies in the geometric model and the (specular) error on the BRDF estimates to an acceptable level. The rendering is a two pass mechanism. First, the augmented scene is rendered using a global illumination technique, the result is

denoted by  $LS_{obj}$ . Next the scene is rendered using the same global illumination technique, without including the virtual objects, denoted by  $LS_{noobj}$ . If the input scene is represented by  $LS_b$ , then the difference between  $LS_b$  and  $LS_{noobj}$  is exactly the error that results from an incorrect BRDF estimation. The differential rendering therefore calculates the final output rendering  $LS_{final}$  as:

$$LS_{final} = LS_b + (LS_{obj} - LS_{noobj})$$

This differential rendering technique removes most of the inaccuracies and in a certain way it is similar to the one of Jancene et al. [JNP\*95] presented above. The results of this technique are promising, see Figure 10, but it still suffers from a few deficiencies. Firstly, only diffuse parameters of the local scene are estimated, this introduces an error that should be compensated by the differential rendering. Secondly, the viewpoint can be altered but the technique is too slow to work at interactive rates. If the rendering could be




**Figure 10:** Debevec et al. [Deb98]. Left: a diagram illustrating the relation between the different components presented in [Deb98]. The real scene is divided into a local scene and a distant scene. The illumination from the distant scene influences the local scene and the virtual objects. The virtual objects influence the local scene. The local scene and the virtual objects do not have an influence on the distant scene. Middle: an image of the real scene. Right: an example of the differential rendering technique for an indoor scene after inserting virtual objects. Diffuse effects are simulated. Courtesy of Debevec et al.

accelerated using low cost graphics hardware, it could be possible to achieve interactive update rates for the MR.

Gibson et al. [GCHH03] developed a method to create soft shadows using a set of shadow maps. They created a rapid shadow generation algorithm to calculate and visualize the shadows in a scene after the material properties of the scene are calculated. A proper estimate of both the geometry and the radiance information in the real scene needs to be available. It is assumed that the BRDF for all materials is diffuse. This diffuse BRDF is estimated using geometry and radiance information (one radiance image per 3D point). In their method, the scene is divided into two parts: one part contains all patches in a scene that are visible from the camera, called the *receiver patches* and another part contains those patches in the scene that have a significant radiance, called the *source patches*. Then they organize these patches to build a shaft hierarchy between the receiver patches and the source patches. The shaft hierarchy contains information on which patches block receiver patches from other source patches. Next they render the scene from a certain viewpoint. This rendering is a two-pass mechanism. In a first pass, they go through the shaft hierarchy to see which source patches partially or completely illuminate a receiver patch. Once these source patches are identified, they set the radiance of each receiver patch to the sum of all irradiance coming from these source patches, without taking occlusions into account. The second rendering pass, takes the shadows in consideration. To calculate the portion of blocked light, they use the shadow mapping technique. In fact, they create a shadow map for each source patch. At each receiver patch, these maps are then combined and subtracted from the radiance value that was rendered in the first pass. This technique is capable of producing soft shadows in a fast and efficient way. In Figure 11 examples are given of synthetic scenes rendered using the above described method. Renderings of the same synthetic scenes using a ray tracing method are given as well. The images in the last column are photographic reference images.

Another set of methods were built to exploit the structure of a radiosity method. Fournier et al. made pioneering work in this direction [FGR93]. When this method was developed, facilities for modelling a geometric model from a real scene were not available. To overcome this issue, Fournier et al. decided to replace the geometry of the objects in the real scene by their bounding box, and an image of the object was applied on each of the faces of the box. An example of such a model is shown in Figure 12. To set up the scene for global common illumination computation, faces of the boxes representing the real objects are divided into patches. Using the information contained in the radiance textures, a diffuse local reflectance is computed by averaging pixels covered by each patch. Light source exitances are estimated and the radiosity of the patches are set as an average of the per pixel radiance covered by each patch. After insertion of the virtual objects and the virtual light sources in the model of the real scene, new radiosity values are computed for the elements in the scene using *progressive radiosity* [CCWG88]. The rendering is carried out by modifying the intensity of each patch with the ratio obtained by dividing the new radiosity by the original one. In Figure 12 an illustration of the result of this method is given. The results of this technique look promising but it suffers from the lack of a detailed geometry. This leads to misaligned shadows and other types of mismatching between real and virtual objects. The technique is slow and will not allow real-time interaction. Nevertheless, this pioneering method has influenced subsequent research work, e.g. Drettakis et al. [DRB97] and Loscos et al. [LDR00] as presented in the remainder of this section.

Drettakis et al. [DRB97] present a method that builds on Fournier et al. [FGR93], but uses a finer model of the real scene. The same equations are used to estimate the light sources emittance, the reflectance of the patches and the original radiosity. Drettakis et al. make use of the more recent hierarchical radiosity method hierarchical [HSA91] accelerated by using clustering [RPV93][Sil95][SAG94].



**Figure 11:** Results for Gibson et al. [GCHH03]. A comparison of the rendering quality for three different scenes. The images in the left column are produced using the system presented in [GCHH03]. The images in the middle column are rendered using ray tracing. The images in the right column are photographic reference images. Courtesy of Gibson et al.

Based on [DS97] a hierarchy of shafts is built from the real scene model, which allows a local understanding when virtual objects are added. This permits an easy identification of all patches that need to undergo a radiosity alteration due to the influence of the newly added object. The advantage of this shaft hierarchy is that it permits interactive updates of the illumination in the augmented scene when virtual objects move. The final display is made similarly to the method of Fournier et al. [FGR93]: the intensity of the patches is modified with the ratio defined by the modified radiosity divided by the original radiosity. This type of rendering is fast, compared to a ray tracing method, as it uses the hardware capability to render textured polygons. This method provides global common illumination with possible interaction. Unfortunately, the technique does not allow changing either the current illumination or the current viewpoint. In Figure 13 a

screen shot is given of the 3D reconstruction and an example of the resulting MR using the above explained method.

### 3.2.3. Relighting

In Loscos et al. [LDR00], relighting is made possible, while keeping the framework set by Fournier et al. [FGR93] and Drettakis et al. [DRB97]. The scene parameters are extracted in the same way but all calculations are extended to the use of HDR images [Los99] as well. Since this technique focuses on relighting, a specific subdivision of the real scene is made to detect as much direct shadows as possible. The radiosity of each element is modified to simulate non-blocked radiosity, in other words, to erase the shadows from the textures. A factor is computed using the radiosity method without taking the visibility in consideration. Then the new radiosity value is used to update the texture. Approximations



**Figure 12:** Results for Fournier et al. [FGR93]. Left: wire-frame image, all objects in the scene are represented by a box, that narrowly fits the object. Middle: Image information is mapped on the boxes (note that for the ball, a more complex shape was used). Right: an example of a MR, the book on top of another book, lying on a table is virtual. Also a virtual light source is added. The global common illumination effects are generated with an adaptive progressive radiosity algorithm. Courtesy of Fournier et al.



**Figure 13:** Results for Drettakis et al. [DRB97]. In the left image a screen shot is given of the 3D reconstruction of the real scene. The right image gives an example of the MR, the floating box is the virtual object. The virtual objects can be moved at interactive rate while keeping the global illumination effects. This is carried out by using an adaptation of hierarchical shafts for hierarchical radiosity [DS97]. Courtesy of Drettakis et al.

of the estimation and of the input data lead to an inexact modification of the texture. In a second step, another factor is applied to automatically correct the imprecisions. This is done by using a reference patch that reflects the desired result. Once this is done, the new textures are used instead of the original ones, and reflectance and original radiosity values are updated accordingly. Shadows can be simulated using the factor of the newly computed radiosity solution divided by the original radiosity (without shadows). This technique also extends the method presented in [DS97] for the insertion of virtual lights. In the system of Loscos et al. [LDR00], it is possible to virtually modify the intensity of

real light sources, to insert virtual objects that can be dynamically moved and to insert virtual light sources. The problem that comes with inserting new lights or increasing light source intensity is that the value of the factor computed between the new radiosity value, divided by the original radiosity, may be greater than one. In that case, multi-pass rendering is used to enable the visualisation of brighter illumination. This method allows interactivity and is fairly rapid in the pre-processing computation. However, the obtained results are inaccurate as the illumination of the real scene is not fully estimated. Firstly, because lit areas are not altered at all, and secondly, because it concentrates on the diffuse





**Figure 14:** Results for Boivin et al. [BG01]. The top left image illustrates the original scene. The top right image is a relighted synthetic image. Diffuse and specular effects are simulated using an optimisation algorithm. The bottom left image illustrates the possibility of changing the viewpoint by grouping objects with similar properties. The bottom right image illustrates the relighting of the original scene with a different illumination pattern. Courtesy of Boivin et al.

component only. An example of the results is shown in Figure 3 using HDR images as an input.

Although it doesn't seem feasible to estimate specular components of the BRDF from one single image, Boivin et al. [BG01] present a technique that re-renders diffuse and specular effects based on radiance information from one single image and a full geometric model of the scene, including the light source positioning and the camera properties. With a hierarchical and iterative technique they estimate the reflectance parameters in the scene. In this method, the reflectance model of Ward [War92] is used, which presents the entire BRDF with either 3 (isotropic materials) or 5 (anisotropic materials) different parameters. The BRDF estimation process starts by assuming that the BRDF values are all diffuse. A synthetic scene is rendered using the geometry, the current BRDF estimate and global illumination techniques. If the difference between the real scene and the syn-

thetic scene is too large, the BRDF values are re-estimated using a more complex BRDF model. First specular effects are added and a roughness factor is estimated using an time-consuming optimisation process. Later anisotropic effects are introduced and the optimisation continues until a reasonable synthetic scene is acquired. This is very similar to the way parameters are estimated in [YDMH99]. However, in this case, only one input image is used, and anisotropic parameters are estimated as well. The method of Boivin et al. relies on one single image to capture all photometric information. The advantage of such an approach is that the image capturing is relatively easy; the disadvantage is that only partial geometric information is available: there is no information for those surfaces that are not visible in the image. Nevertheless, the proposed technique allows changing the viewpoint. If a sufficiently large portion of a certain object is visible in the image, the reflectance properties of the miss-

ing parts of the object are calculated based on this portion. Grouping objects with similar reflectance properties makes this process more robust. On the other hand, this requires that not only the geometry needs to be known, but also a partitioning of the scene into objects with similar reflectance properties, which greatly compromises the operability of this technique. Although optimised, the rendering algorithm is computationally expensive and therefore only a non real-time solution can be obtained. In Figure 14 an illustration is given of the output results of the described method.

### 3.3. Model of real scene known, few images known

If more information about the radiance of the points in the scene is available, a better BRDF estimate can be acquired. The radiance perceived at a certain point depends on the viewing angle, on the angle of incident light and the BRDF. Hence, it is possible to gain more information about the BRDF of a certain point in the scene if radiance information is available from images captured from a different viewing angle. Alternatively, if the viewpoint is kept the same but the position of the light sources is changed, extra BRDF information is captured as well. In this section, the methods are discussed that make use of this extra information.

Loscos et al. [LFD\*99] developed a system that allows relighting, as well as virtual light source insertion, dynamic virtual objects inclusion and real object removal. They identified that it is difficult to estimate reflectance values in shadow regions due to saturation and because this estimate depends on the quality of the indirect light estimation. This is compensated for by adding extra photographs captured under different lighting. The geometry of the real scene is modelled from photographs. This geometric model is textured using one of the images, taken from the different viewpoints. A set of pictures is then taken from this chosen viewpoint while a light source is moved around the scene to modify the illumination. These pictures can be HDR images as used in [Los99]. Loscos et al. decided to mix a ray-casting approach to compute the local illumination and a radiosity approach to compute the indirect lighting. Two sets of reflectances are thus computed. First diffuse reflectance values are computed for each pixel of the viewing window. This is done with a weighted average of the reflectance evaluated with each input image differently lit. The applied weight is based on whether the 3D point associated to the pixel is in shadow relative to the light source position, and also whether the radiance value captured is saturated. The reflectance values are then used to initialise a radiosity system similar to those in [DRB97][LDR00]. This reflectance can be refined by an iterative algorithm [LFD\*99]. With this reflectance, Loscos et al. are able to relight the scene using global illumination. Pixel values are updated by adding the local illumination value, computed by ray casting, to the indirect illumination value, computed by hierarchical radiosity using a rough subdivision of the scene. Local modifications are

made after the insertion or moving of virtual objects by selecting the area of the window where local illumination will be affected. Indirect illumination is modified by adapting the technique of [DS97]. Similarly, virtual light sources can be added, and intensity of real light sources can be modified. A very interesting application of this method is the removal of real objects. The unknown information previously masked by the object is filled using automatic texture synthesis of a sample of the image of the reflectance values of the previously hidden object. The results show that the relighting and the interaction with virtual objects can be achieved in an interactive time. Image examples of the results are shown in Figure 15. The produced results are good but could be improved by considering specular effects. Due to the nature of the image capture process, it would be very difficult to apply this technique on real outdoor scenes.

A different approach taken by Gibson et al. [GHH01] results in another relighting method, in which the reflectance of the material is roughly estimated based on a restricted amount of geometry and radiance information of the scene. In theory, only geometry and radiance information is needed for those parts of the scene that will be visible in the final relighted MR. In their approach a photometric reconstruction algorithm is put forward, that is capable of estimating reflectance and illumination for a scene if only incomplete information is available. To achieve this the direct illumination is modeled as coming from unknown light sources using virtual light sources, see Figure 16. The aim is not to produce an accurate illumination model, but rather a model that produces a similar illumination as in the original scene. The model used is a spherical illumination surface: a set of small area light sources that surrounds the known geometry. The parameters of this surface, the position and emission of the light sources, are estimated using an iterative minimization algorithm. Based on this model, the reflectance of the materials in the scene are estimated. The MR scene is rendered using a ray tracing algorithm. User interaction is impossible at real-time update rate but nevertheless the method illustrates the possibility of getting fairly realistic mixed realities, without limiting input requirements. This method is original, interesting and very practical to adapt to many situations where information on a real scene is partially known. Imprecisions and ambiguities are compensated for, resulting in a more accurate simulation of the existing illumination. An example of a rendered scene and its comparable real scene are given in Figure 17.

### 3.4. Model of real scene known, many images known

This category collects those techniques that require the most input information. Not only the geometry is known but also radiance information under many different geometric setups. Two significant methods could be identified that belong to this category of MR methods. They were selected from a broad set of techniques on inverse illumination because



**Figure 15:** Results for Loscos et al. [LFD\*99]. The left image is one of the input images of the real scene. The middle image is a relighted image of the real scene, using the calculated BRDF values. The left image illustrates the removal of an object (the door), the insertion of a new virtual object (the chair) and the insertion of a virtual light source. All manipulations are carried out at interactive update rates. The illumination is updated locally with ray casting. The consistency of the indirect illumination is kept using an adaptation of [DS97]. Courtesy of Loscos et al.



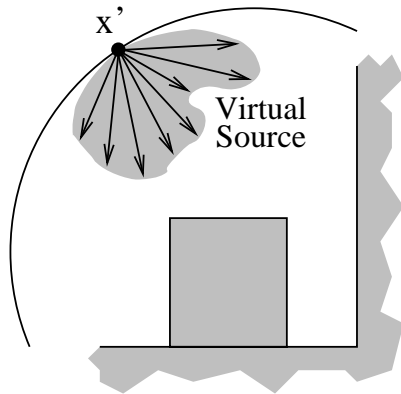
**Figure 17:** Results for Gibson et al. [GHH01]. The left images illustrates the reconstructed scene from a novel viewpoint. The image in the middle is a synthetic image illuminated with virtual light sources. The right image illustrates the addition of virtual objects. Both specular and diffuse effects are simulated. Courtesy of Gibson et al.

they provide a solution for a large group of objects, which is essential for MR. The first inverse illumination method [YDMH99] focuses on the BRDF estimation, using many HDR images from different viewpoints. The second [YM98] allows to relight outdoor scenes. The remainder of this section briefly discusses these two techniques.

Yu et al. [YDMH99] use a low parametric reflectance model, which allows the diffuse reflectance to vary arbitrarily across the surface while non-diffuse characteristics remain constant across a certain region. The input to their system is the geometry of the scene, a set of HDR images and the position of the direct light sources. An inverse radiosity method is applied to recover the diffuse albedo. The other two parameters in the reflectance model of Ward [War92], the roughness and the specular component, are estimated by a non-linear optimisation. For the estimation of the specular BRDF, it is assumed that many HDR images are available from a different set of viewpoints. The estimation makes use of the position of the light sources and the possible highlights they may produce on a surface due to specular effects. It is therefore helpful to capture images of the scene with a various number of light sources, since this might increase the number of specular highlights. This precise estimate of the

BRDF values in the scene allows to remove all illumination in the scene and a new illumination pattern can be applied. To render the scene they make use of Ward's RADIANCE system [War94]. No further steps were taken to speed up the rendering process. Figure 4 illustrates the results obtained for augmented images compared to photographs of the real scene. This technique is interesting for MR because it provides an algorithm to estimate an accurate complex BRDF of a complex real scene, resulting in an accurate representation of the illumination.

Yu and Malik [YM98] present a technique that allows relighting for outdoor scenes based on inverse illumination. As it is impossible to retrieve the geometry of the entire scene, they separate the scene into four parts: the local model, the sun, the sky and the surrounding environment. The illumination sources are the sun, the sky, and the surrounding environment. Luminance due to the sun and the sky are estimated based on a set of input images. At least two photographs per surface of the local model are captured, which should show two different lighting conditions (directly and not directly lit by the sun). The local model is subdivided into small surfaces. Based on these two photographs, two pseudo-BRDF values are estimated per surface. One relates to the illumina-



**Figure 16:** Technique Gibson et al. [GHH01]. The real illumination is approximated by a illumination surface. This illumination surface is covered by a set of virtual light sources. The parameters of these virtual light sources are estimated such that its effect resembles the real illumination. Courtesy of Gibson et al.

tion from the sun, the other relates to the illumination from the integrated environment (sky plus surrounding environment). A least square solution is then used to approximate the *specular term* for each surface and for each lighting conditions (from the integrated environment and from the sun). This approach uses an approximation of the inverse illumination equation. It illustrates the difficulty of setting up a parameterised MR system for outdoor scenes. At rendering time, different positions of the sun are simulated. After extracting the sun and the local model from the background, sky regions are identified and they are mapped on a mesh supported by a hemisphere. Three parameters control the sky intensity. A first scale factor is applied when simulating sunrise and sunset; it is constant otherwise. The second parameter adjusts the intensity of the sky depending on the position of the mesh on the dome. A last parameter controls the sky intensity depending on the sun's position. Next, the radiance values and the pseudo-BRDFs are used to reproduce the global illumination on the local scene. This method is the first to present the possibility of relighting outdoor scenes. Results of these relighted scenes and a comparison image are shown in Figure 18. Although it is difficult to evaluate the quality of the relighting from the images provided by the authors, the images resemble the real conditions, and this can satisfy most of the MR applications for outdoor environments.

#### 4. Discussion

In section 2.2 we pointed out that the assessment of the various illumination techniques for MR comes with a certain degree of subjectivity. Fortunately there are some aspects that can be evaluated in a rather objective way. Some of these

measures will be used in this section to assess the methods from section 3. Section 4.1 discusses the amount of *pre-processing* required. In section 4.2 an evaluation of the *degree of interactivity* is given and in section 4.3, the methods will be evaluated based on the *quality* of the results. Section 4.4 explains which methods are suitable for outdoor scenes. Finally an overview of the discussed methods is given in section 4.5.

#### 4.1. Pre-processing time

The term *pre-processes* refers to those steps, carried out once, that are required by the method before the merging of real and virtual objects takes place. The geometry reconstruction, image capturing and BRDF estimation, are considered as pre-processing steps.

A few methods do not require a full geometric model of the real scene: Sato et al. [SS199], Nakamae et al. [NHIN86] and Haller et al. [HDH03]. All other methods require a geometric model. Some of these methods do not explain how this model can be constructed, others assume that it is constructed using semi-manual 3D reconstruction software, examples of such software were given in section 3. Using reconstruction software usually results in a low resolution model and is in general error prone, this is due to the fact that no automatic, accurate 3D reconstruction software is yet commercially available. Scanning devices give a better resolution, but these devices are expensive and while the scanning of a small object might be straightforward, the scanning of a larger scene is tedious. As a summary we can say that a perfect geometric model is difficult to acquire and that reconstruction is always a tedious work.

Some methods require radiance information captured from several viewpoints [YDMH99] [GHH01] or under different types of illumination [LFD\*99] [YM98]. Taking several HDR images from different viewpoints and under different illumination delays the image capture time.


Many methods calculate a BRDF estimate, some use a diffuse model, some allow a more complex model. Often the calculation of the BRDF needs to be carried out off-line, due to timing issues and is therefore considered as pre-processing work. Methods that calculate a diffuse-only BRDF are: [Deb98][FGR93][DRB97][LDR00][LFD\*99], methods that allow specular components are: [GHH01][YDMH99][YM98][BG01].

#### 4.2. Level of interactivity

Interactivity means:

- the possibility of navigating objects or viewpoints in the scene,
- the effort made to get an interactive rendering,
- the possibility to modify reflectance properties of real objects in the scene,





**Figure 18:** Results for Yu et al. [YM98]. The top row images illustrates the original tower from different viewpoints. The bottom row are synthetic images of the tower from approximately the same viewpoint. The real images were not used to generate the synthetic images, nevertheless the synthetic and real images look very similar. Courtesy of Yu and Malik.

- the possibility to modify the illumination sources in the real scene.

A few methods allow to navigate the virtual objects or the viewpoints. These techniques have either enough BRDF information [BG01][YDMH99][FGR93], enough geometry and illumination information [SSI99][YM98] or use a different approach [ALCS03][SHC\*94][JNP\*95].

Only a few of the methods operate in true real-time (RT) [ALCS03][SHC\*94][GCHH03], others are near real-time (near RT) [LDR00][LFD\*99][DRB97] but most of them are non real-time (NRT). However, it should be noted that some methods were developed years ago, when computer hardware and software were much slower than nowadays. Also, it should be pointed out that some methods did not made a special attempt in producing interactive systems. With a few modifications, it should be possible to speed up most of the described systems.

Some methods that specifically tried to speed up the computations are worth mentioning. Agusanto et al. [ALCS03] exploited the idea of environment mapping while State et al. [SHC\*94] used shadow mapping and Haller et al. [HDH03] shadow volumes. Gibson et al. [GM00] developed a new technique to simulate soft shadows at interactive rates and Drettakis et al. [DRB97], Loscos et al. [LDR00] and Loscos et al. [LFD\*99] made use of a hierarchical radiosity algo-

rithm, that decreased the computation time to interactive rates as well. Gibson et al. [GCHH03] used shadow volumes.

Most methods that calculate the BRDF values are in principle capable of changing the BRDF values into something new. This can be used to modify the appearances of real objects in the scene. Relighting methods can use this BRDF information to relight a scene using a different illumination pattern. Table 1 gives an overview of the various different types of illumination the discussed methods allow.

### 4.3. Evaluation of the quality

Some of the described methods evaluated the quality of their method using one or more of the following evaluation methods:

- a comparison is made between a photograph reference of the real scene and a synthetic version of the same scene,
- the BRDF is measured using a device and these results are compared with the calculated BRDF values.

Gibson et al. [GCHH03] compare their shadow rendering technique with a ray traced rendering and an image of the real scene, see Figure 11. They are capable of producing realistic and similar shadows as in the real image and at a faster time than the ray traced rendering. In [GM00] the presented extended shadow mapping is compared with a ray traced version using the same input parameters, see Figure 9. There are

some differences between the two synthetic scenes, but the generated shadows look realistic.

Boivin et al. [BG01] extract a full BRDF model and compare their rendering with an original image of the real scene, see Figure 14. In [YM98] the diffuse and specular components are calculated; the resulting rendering is compared with an original image of the real scene. Likewise, [LDR00] [LFD\*99] estimate the diffuse BRDF and compare a synthetic rendering with a original image of the real scene (see Figures 3 and 15). In both methods, the rendering occurs at interactive update rates.

Similarly, Gibson et al. [GHH01], see Figure 17, compare an original and synthetic image and find that the error between the two images decreases drastically in the first three iterations. Both diffuse and specular reflectances are modelled.

Yu et al. [YDMH99], see Figure 4, estimate diffuse and specular BRDF values and compare these with measured BRDF values of objects in the scene. The estimates and the true values are similar.

We can also compare methods that use both specular and diffuse BRDF values for the rendering with those that have a more restrictive understanding of the BRDF. It is understood that systems based on a more complete BRDF model result in an MR of a higher quality than those based on diffuse BRDF values only or those that do not estimate BRDF values at all. For some methods, only a subjective user perceptible assessment can be made.

#### 4.4. Usability on indoor and outdoor scenes

The reader may have noticed that most techniques were tested on indoor scenes. Outdoor scenes are more complex than indoor scenes. Not only is the geometry more difficult to model, the illumination is difficult to extract as well. Outdoor illumination is time and weather dependent and difficult to model and simulate. Only three methods from section 3 explicitly used an outdoor scene to test their method [Deb98][GCHH03][SS199][YM98] but this does not imply that the other methods are not suitable for outdoor scenes. For instance one might argue that all methods that use environment maps are capable of capturing the outdoor illumination. But some caution is in place when interpreting this statement [JL04]. If HDR images are used to capture the environment map, which is in general the case, one needs to bare two things in mind. Firstly, the intensity of the sun is in general too bright to be captured in a HDR image without saturation, even at very fast shutter speeds. Secondly, if the sky is clouded and the clouds drift in the sky, there will inevitable be some movement in the low dynamic images used to compile the HDR image, making them worthless. It should be clear, that the extension from indoor to outdoor scenery is not straightforward. The current state of the art of MR shows no good solutions for the outdoor scenes.

#### 4.5. Overview

Table 1 gives an overview of all methods discussed in section 3. For each method, the overview discusses the following aspects:

- **Geometric model of the scene:** whether or not the method requires a geometric model of the scene.
- **Number of different images:** the number of different images needed per point in the scene, to calculate the MR.
- **Methodology:** the methodology used to create the MR. In section 2.3 three different approaches were discussed:

1. common illumination,
2. relighting,
3. inverse Illumination.

Further to this division, a distinction is made between local and global illumination techniques.

- **Rendering:** the rendering method used to compose the MR. Possible answers are: ray-casting, ray-tracing, radiosity, etc.
- **Computation time:** the *update* time of the method is real-time (RT), non real-time (NRT) or near real-time (near RT).

**Table 1:** Overview of illumination methods for mixed reality.

	Geometric model	# Images	Methodology	Rendering	Computation time	Section
[ALCS03]	no	one	global common illumination	environment maps, multipass rendering	RT	3.1
[NHIN86]	no	one	local common illumination	ray casting	NRT	3.1
[SSI99]	no	one	global common illumination	ray casting	NRT	3.1
[SHC*94]	yes	one	local common illumination	shadow mapping	RT	3.2
[HDH03]	yes	one	local common illumination	shadow volumes	RT	3.2
[JNP*95]	yes	one	local common illumination	ray tracing	NRT	3.2
[GCHH03]	yes	one	global common illumination	shadow mapping	RT	3.2
[Deb98]	yes	one	global common illumination	differential rendering + ray tracing	NRT	3.2
[GM00]	yes	one	global common illumination	extended shadow mapping	near RT	3.2
[FGR93]	yes	one	global common illumination	radiosity + ray casting	NRT	3.2
[DRB97]	yes	one	relighting using global illumination	hierarchical radiosity algorithm	near RT	3.2
[LDR00]	yes	one	relighting using global illumination	hierarchical radiosity algorithm	near RT	3.2
[BG01]	yes	one	inverse global illumination	ray tracing	NRT	3.2
[LFD*99]	yes	few	relighting using global illumination	hierarchical radiosity algorithm + ray casting	near RT	3.3
[GHH01]	yes	few	relighting using global illumination	ray tracing	NRT	3.3
[YDMH99]	yes	many	inverse global illumination	ray tracing	NRT	3.4
[YM98]	yes	many	inverse global illumination	ray tracing	NRT	3.4

## 5. Conclusions and Future work

In the past few years research has been motivated to consider a new type of simulation: the simulation of a new reality called *mixed reality* which refers to the concept of mixing a real scene with virtual elements. Mixed reality has now become very important for various applications ranging from entertainment with movie post-production and games, architecture, cultural heritage, education and training, etc. Several problems arise when composing reality with virtual elements. A first problem is how to automatically calibrate the relative position and the occlusion of virtual objects with virtual ones. As mentioned in this report, this has been addressed successfully in two different ways. One can use a scanning device or one can use real-time stereovision to extract depth and shape information. A second problem is how to illuminate the virtual objects consistently with the original light conditions in the real scene. Research papers appeared already in the late eighties to answer this last need of the industry, but it is only recently, within the last ten years, that the international community made a more significant effort to provide more automated solutions for computing the illumination in mixed reality.

Although it is tempting to compare techniques relatively to the quality of the results achieved, this report classifies them depending on the context and the goal of the method. Firstly, it is of course easier to compute illumination for mixed reality if a 3D model is available. Secondly, it may be that only a few images of the real scene are available from different viewpoints, and some available with different lighting conditions. The more images are available for the illumination extraction, the easier the computation procedure becomes. On the contrary, the fewer images are available the more difficult it is to perform an accurate estimation and therefore simulation. Consequently it was decided that it would be fairer and more interesting to compare techniques using similar types of data. Four different categories were identified based on the amount of geometric detail and radiance information available. Different manners to compare the illumination techniques used for mixed reality were presented as well. For example, the methods were compared based on the type of the illumination achieved: local or global, diffuse or complex illumination effects. It was also pointed out if relighting was possible and if the user could interact with the scene.

An ideal conclusion of this report would state which technique is the most perfect one. However it is impossible to assess the methods without knowing the application at hand. It is therefore very difficult to describe the ideal method. It should be a real-time and automatic method with no pre-processing requirements. It would allow any type of virtual interaction: modification of lighting, removal of real objects, modification of material properties of real objects and addition of virtual elements. And the rendering quality would perfectly match with the real one. Research is head-

ing towards this, and it is likely that this technology will become more accessible in future years. Progress in stereovision techniques, in automatic calibration, registration and in computer graphics will help in the progression in illumination for mixed reality. More automatic reconstruction methods of the geometry are needed, that will also model more complex details. Progress in augmented reality is heading towards systems being able to recognise shape and depth without markers. Computer graphics research needs to provide more precise description of reflection models and rendering software needs to be adapted to these more complex materials. Little work has been done in modelling the behaviour of light sources, which are often assumed diffuse. It will be important for future work to consider more complex lighting in order to find a better estimate for the illumination in mixed reality. Finally, most of the methods have been designed for indoor environments that are easier to control. Outdoor environments present a real challenge, both in the capture and in the simulation. It is expected that more work for outdoor environments will appear in the near future.

## 6. Acknowledgements

We would like to thank Simon Gibson, Greg Ward, Anthony Steed, Mel Slater, Erik De Witte, Hila Ritter Widerfeld and Bernhard Spanlang for reviewing this paper.

This report was partly funded by CREATE (IST-2001-34231), a 3-year RTD project funded by the 5th Framework Information Society Technologies (IST) Programme of the European Union and a UCL graduate school research scholarship.

## References

- [3Ds] 3DSCANNERS.: online. [www.3dscanners.com](http://www.3dscanners.com).
- [ABB\*01] AZUMA R., BAILLOT Y., BEHRINGER R., FEINER S., JULIER S., MACINTYRE B.: Recent advances in augmented reality. *IEEE Computer Graphics Applications* 21, 6 (2001), 34–47.
- [ALCS03] AGUSANTO K., LI L., CHUANGUI Z., SING N. W.: Photorealistic rendering for augmented reality using environment illumination. In *proceedings of IEEE/ACM International Symposium on Augmented and Mixed Reality (ISMAR '03)* (October 2003), vol. 21, pp. 208–216.
- [ATI02] ATI: Car paint. online, 2002. [www.ati.com/developer/demos/r9700.html](http://www.ati.com/developer/demos/r9700.html).
- [Azu95] AZUMA R.: A survey of augmented reality. In *ACM Siggraph '95, Course Notes #9: Developing Advanced Virtual Reality Applications* (August 1995), pp. 1–38.
- [BG01] BOIVIN S., GAGALOWICZ A.: Image-based

- rendering of diffuse, specular and glossy surfaces from a single image. In *proceedings ACM Siggraph '01 (Computer Graphics)* (2001), ACM Press, pp. 107–116.
- [BGWK03] BIMBER O., GRUNDHEIMER A., WETZSTEIN G., KNODEL S.: Consistent illumination within optical see-through augmented environments. In *proceedings of IEEE/ACM International Symposium on Augmented and Mixed Reality (ISMAR'03)* (2003), ACM Press, pp. 198–207.
- [BKM99] BEHRINGER R., KLINKER G., MIZELL D. W. (Eds.): *Augmented reality: placing artificial objects in real scenes* (1999).
- [Bli88] BLINN J.: Me and my (fake) shadow. *IEEE Computer Graphics Applications* 8, 1 (1988), 82–86.
- [BN76] BLINN J. F., NEWELL M. E.: Texture and reflection in computer generated images. *Communications of the ACM* 19, 10 (1976), 542–547.
- [CCWG88] COHEN M. F., CHEN S. E., WALLACE J. R., GREENBERG D. P.: A progressive refinement approach to fast radiosity image generation. In *proceedings of ACM Siggraph '88 (Computer Graphics)* (1988), Annual Conference Series, ACM Press, pp. 75–84.
- [Che95] CHEN S. E.: QuickTime VR — an image-based approach to virtual environment navigation. In *proceedings of ACM Siggraph '95 (Computer Graphics)* (1995), vol. 29 of *Annual Conference Series*, pp. 29–38.
- [CON99] CABRAL B., OLANO M., NEMEC P.: Reflection space image based rendering. In *proceedings of ACM Siggraph '99 (Computer Graphics)* (1999), Annual Conference Series, ACM Press/Addison-Wesley Publishing Co., pp. 165–170.
- [DBY98] DEBEVEC P., BORSHUKOV G., YU Y.: Efficient view-dependent image-based rendering with projective texture-mapping. In *proceedings of the 9th Eurographics Workshop on Rendering (Rendering Techniques '98)* (June 1998), Springer-Verlag, (Ed.), pp. 105–116.
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *proceedings of ACM Siggraph '98 (Computer Graphics)* (1998), vol. 32 of *Annual Conference Series*, pp. 189–198.
- [DHT\*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *proceedings of ACM Siggraph '00 (Computer Graphics)* (2000), ACM Press/Addison-Wesley Publishing Co., pp. 145–156.
- [DM97] DEBEVEC P. E., MALIK J.: Recovering high dynamic range radiance maps from photographs. In *proceedings of ACM Siggraph '97 (Computer Graphics)* (1997), vol. 31 of *Annual Conference Series*, pp. 369–378.
- [DRB97] DRETTAKIS G., ROBERT L., BUGNOUX S.: Interactive common illumination for computer augmented reality. In *proceedings of the 8th Eurographics Workshop on Rendering (Rendering Techniques '97)* (Saint Etienne, France, June 1997).
- [DS97] DRETTAKIS G., SILLION F. X.: Interactive update of global illumination using a line-space hierarchy. In *proceedings of ACM Siggraph '97 (Computer Graphics)* (August 1997), vol. 31 of *Annual Conference Series*, pp. 57–64.
- [DTM96] DEBEVEC P. E., TAYLOR C. J., MALIK J.: Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *proceedings of ACM Siggraph '96 (Computer Graphics)* (1996), vol. 30 of *Annual Conference Series*, pp. 11–20.
- [Eos] EOS SYSTEMS INC: Photomodeller. online. [www.photomodeler.com](http://www.photomodeler.com).
- [Fau92] FAUGERAS O.: What can be seen in three dimensions with an uncalibrated stereo rig. In *Proceedings of the 2nd European Conference on Computer Vision* (May 1992), Sandini G., (Ed.), vol. 588, pp. 563–578.
- [Fau93] FAUGERAS O.: *Three-dimensional computer vision - A geometric viewpoint*. MIT press, 1993.
- [FGR93] FOURNIER A., GUNAWAN A. S., ROMANZIN C.: Common illumination between real and computer generated scenes. In *proceedings of Graphics Interface '93* (Toronto, Canada, May 1993), pp. 254–262.
- [FRL\*98] FAUGERAS O., ROBERT L., LAVEAU S., CSURKA G., ZELLER C., GAUCLIN C., ZOGHLAMI I.: 3-d reconstruction of urban scenes from image sequences. *CVGIP : Image Understanding* 69, 3 (1998), 292–309.
- [GCHH03] GIBSON S., COOK J., HOWARD T., HUBBOLD R.: Rapid shadow generation in real-world lighting environments. In *proceedings*

- of the 13th Eurographics workshop on Rendering (Rendering Techniques '03) (2003), Eurographics Association, pp. 219–229.
- [GHH01] GIBSON S., HOWARD T., HUBBOLD R.: Flexible image-based photometric reconstruction using virtual light sources. In *proceedings of Eurographics 2001* (September 2001). Manchester, UK.
- [GM00] GIBSON S., MURTA A.: Interactive rendering with real-world illumination. In *proceedings of 11th Eurographics Workshop on Rendering (Rendering Techniques '00)* (June 2000). Springer Wien.
- [Gre86] GREENE N.: Environment mapping and other applications of world projections. *IEEE Computer Graphics Applications* 6, 11 (1986), 21–29.
- [Har97] HARTLEY R.: In defence of the eight point algorithm. *IEEE International Conference on Computer Vision* 19, 6 (June 1997), 580–593.
- [HDH03] HALLER M., DRAB S., HARTMANN W.: A real-time shadow approach for an augmented reality application using shadow volumes. In *proceedings of the ACM symposium on Virtual reality software and technology (VRST '03)* (2003).
- [HGC92] HARTLEY R., GUPTA R., CHANG T.: Stereo from uncalibrated cameras. In *proceedings of Computer Vision and Pattern Recognition* (1992), IEEE Computer Society Press, pp. 761–764.
- [HGM00] HOWARD T., GIBSON S., MURTA A.: Virtual environments for scene of crime reconstruction and analysis. In *proceedings of SPIE Electronic Imaging 2000* (Jan 2000), vol. 3960.
- [HS99] HEIDRICH W., SEIDEL H.-P.: Realistic, hardware-accelerated shading and lighting. In *proceedings of ACM Siggraph '99 (Computer Graphics)* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 171–178.
- [HSA91] HANRAHAN P., SALZMAN D., AUPPERLE L.: A rapid hierarchical radiosity algorithm. In *proceedings of ACM Siggraph '91 (Computer Graphics)* (1991), vol. 25 of *Annual Conference Series*, pp. 197–206.
- [Int] INTEGRA: Renoir. online. [www.integra.co.jp/eng/products/renoir](http://www.integra.co.jp/eng/products/renoir).
- [JAL\*05] JACOBS K., ANGUS C., LOSCOS C., NAHMIA S. J.-D., RECHE A., STEED A.: Automatic consistent shadow generation for augmented reality. In *proceedings of Graphics Interface '05* (2005).
- [JL04] JACOBS K., LOSCOS C.: Relighting outdoor scenes, July 2004. BMVA one day symposium: Vision, Video and Graphics.
- [JNP\*95] JANCENE P., NEYRET F., PROVOT X., TAREL J.-P., VEZIEN J.-M., MEILHAC C., VERROUST A.: Computing the interactions between real and virtual objects in video sequences. In *proceedings of 2nd IEEE Workshop on Networked Realities* (October 1995), pp. 27–40.
- [KBBM99] KATO H., BILLINGHURST M., BLANDING R., MAY R.: *ARToolkit*. Tech. rep., Hiroshima City University, 1999.
- [KM00] KAUTZ J., MCCOOL M. D.: Approximation of glossy reflection with prefiltered environment maps. In *proceedings of Graphics Interface '00* (2000), pp. 119–126.
- [Lar91] LARSON G. W.: Real pixels. In *Graphics Gems II* (1991), Arvo J., (Ed.), pp. 80–83.
- [LDR00] LOSCOS C., DRETTAKIS G., ROBERT L.: Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics* 6, 3 (July-September 2000), 289–305.
- [LFD\*99] LOSCOS C., FRASSON M.-C., DRETTAKIS G., WALTER B., GRANIER X., POULIN P.: Interactive virtual relighting and remodeling of real scenes. In *proceedings of 10th Eurographics Workshop on Rendering (Rendering Techniques '99)* (New York, NY, Jun 1999), Lischinski D., Larson G., (Eds.), vol. 10, Springer-Verlag/Wien, pp. 235–246.
- [Los99] LOSCOS C.: *Interactive relighting and remodelling of real scenes for augmented reality*. PhD thesis, iMAGIS-GRAVIR/IMAG-INRIA, 1999. PhD thesis.
- [Mar98] MARSCHNER S. R.: *Inverse rendering in computer graphics*. PhD thesis, Department of Computer Graphics, Cornell University, Ithaca, NY, 1998. Program of Computer Graphics, Cornell University, Ithaca, NY.
- [Met] METACREATIONS: Canoma. online. [www.metacreations.com/products/canoma](http://www.metacreations.com/products/canoma).
- [MK94] MILGRAM P., KISHINO F.: A taxonomy of mixed reality visual displays. *IEICE Transactions on Information Systems* E77-D, 12 (December 1994), 1321–1329.
- [MN99] MITSUNAGA T., NAYAR S. K.: Radiometric self calibration. In *proceedings of IEEE*

- Conference on Computer Vision and Pattern Recognition* (June 1999), Fort Collins, pp. 374–380.
- [MNP\*99] MCALLISTER D., NYLAND L., POPESCU V., LASTRA A., MCCUE C.: Real-time rendering of real-world environments. In *proceedings Eurographics Workshop on Rendering (Rendering Techniques '99)* (1999), pp. 145–160.
- [MW93] MACKENZIE I. S., WARE C.: Lag as a determinant of human performance in interactive systems. In *proceedings of the SIGCHI conference on Human factors in computing systems* (1993), ACM Press, pp. 488–493.
- [MYTG94] MEAS-YEDID V., TAREL J.-P., GAGALOWICZ A.: Calibration métrique faible et construction interactive de modèles 3d de scènes. In *Congrès Reconnaissance des Formes et Intelligence Artificielle* (Paris, France, 1994), AFCET.
- [NHIN86] NAKAMAE E., HARADA K., ISHIZAKI T., NISHITA T.: A montage method: the overlaying of the computer generated images onto a background photograph. In *proceedings of ACM Siggraph '86 (Computer Graphics)* (1986), ACM Press, pp. 207–214.
- [NVi] NVIDIA: Cube environment mapping. online. [www.nvidia.com/object/feature\\_cube.html](http://www.nvidia.com/object/feature_cube.html).
- [Nyl] NYLAND L. S.: Capturing dense environmental range information with a panning, scanning laser rangefinder. [www.cs.unc.edu/ibr/projects/rangefinder](http://www.cs.unc.edu/ibr/projects/rangefinder).
- [OT99] OHTA Y., TAMURA H.: *Mixed Reality - merging real and virtual worlds*. Ohmsha and Springer-Verlag, 1999. Chapter 1, by Paul Milgram and Herman Colquhoun Jr.
- [POF98] POULIN P., OUIMET M., FRASSON M. C.: Interactively modeling with photogrammetry. In *proceedings of Eurographics Workshop on Rendering (Rendering Techniques '98)* (June 1998), pp. 93–104.
- [PP03] PATOW G., PUEYO X.: A survey on inverse rendering problems. *Computer Graphics Forum* 22, 4 (2003), 663–687.
- [Rea] REALVIZ: Image modeller. online. [www.realviz.com](http://www.realviz.com).
- [RM02] RAMAMOORTHY R., MARSCHNER S.: Acquiring material models using inverse rendering. In *ACM Siggraph '02, Course Notes #39* (2002). Organizer: Ravi Ramamoorthi, Steve Marschner, Lecturers: Samuel Boivin, George Drettakis, Hendrik P. A. Lensch, Yizhou Yu.
- [RPV93] RUSHMEIER H. E., PATTERSON C., VEERASAMY A.: Geometric simplification for indirect illumination calculations. In *proceedings of Graphics Interface '93* (May 1993), pp. 227–236.
- [SAG94] SMITS B., ARVO J., GREENBERG D.: A clustering algorithm for radiosity in complex environments. In *proceedings of ACM Siggraph '94 (Computer Graphics)* (1994), vol. 28 of *Annual Conference Series*, pp. 435–442.
- [SCT\*94] STATE A., CHEN D. T., TECTOR C., BRANDT A., CHEN H., OHBUCHI R., BAJURA M., FUCHS H.: *Case Study: Observing a volume rendered fetus within a pregnant patient*. Tech. Rep. TR94-034, University of North Carolina, Chapel Hill, 18, 1994.
- [SHC\*94] STATE A., HIROTA G., CHEN D. T., GARRETT W. F., LIVINGSTON M. A.: Superior augmented reality registration by integrating landmark tracking and magnetic tracking. In *proceedings of ACM Siggraph '94 (Computer Graphics)* (1994), Rushmeier H., (Ed.), vol. 30 of *Annual Conference Series*, pp. 429–438.
- [Sil95] SILLION F. X.: A unified hierarchical algorithm for global illumination with scattering volumes and object clusters. *IEEE Transactions on Visualization and Computer Graphics* 1, 3 (1995), 240–254.
- [SSI99] SATO I., SATO Y., IKEUCHI K.: Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics* 5, 1 (1999), 1–12.
- [SUC95] SLATER M., USOH M., CHRYSANTHOU Y.: The influence of dynamic shadows on presence in immersive virtual environments. In *Selected papers of the Eurographics workshops on Virtual environments '95* (1995), Springer-Verlag, pp. 8–21.
- [SWI97] SATO Y., WHEELER M. D., IKEUCHI K.: Object shape and reflectance modeling from observation. In *proceedings of ACM Siggraph '97 (Computer Graphics)* (1997), vol. 31 of *Annual Conference Series*, pp. 379–388.
- [VF94] VOORHIES D., FORAN J.: Reflection vector shading hardware. In *proceedings of ACM Siggraph '94* (1994), ACM Press, pp. 163–166.
- [War92] WARD G. J.: Measuring and modeling anisotropic reflection. In *proceedings of ACM Siggraph '92 (Computer Graphics)* (1992), ACM Press, pp. 265–272.

- [War94] WARD G. J.: The radiance lighting simulation and rendering system. In *proceedings of ACM Siggraph '94 (Computer Graphics)* (1994), ACM Press, pp. 459–472.
- [YDMH99] YU Y., DEBEVEC P., MALIK J., HAWKINS T.: Inverse global illumination: recovering reflectance models of real scenes from photographs. In *proceedings of ACM Siggraph '99 (Computer Graphics)* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 215–224.
- [YM98] YU Y., MALIK J.: Recovering photometric properties of architectural scenes from photographs. In *proceedings of ACM Siggraph '98 (Computer Graphics)* (1998), ACM Press, pp. 207–217.



# A Survey of Inverse Surface Design From Light Transport Behavior Specification

Gustavo Patow and Xavier Pueyo

Grup de Gràfics de Girona  
Institut d'Informàtica i Aplicacions,  
Universitat de Girona,  
Campus de Montilivi,  
E-17003 Girona, Spain.  
E-mail: {dagush,xavier}@iia.udg.es

---

## Abstract

*Inverse surface design problems from light transport behavior specification usually represent extremely complex and costly processes, but their importance is well known. In particular, they are very interesting for lighting and luminaire design, in which it is usually difficult to test design decisions on a physical model in order to avoid costly mistakes. In this survey we present the main ideas behind these kinds of problems, characterize them, and summarize existing work in the area, revealing problems that remain open and possible areas of further research.*

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture I.4.7 [Image Processing and Computer Vision]: Feature Measurement

---

## 1. Introduction

Inverse problems are emerging as an important research topic for the graphics community due to their importance in a wide range of application fields including lighting engineering and lighting design. These problems are usually of extreme complexity. Although progress in rendering to date has mainly focused on improving the accuracy of the physical simulation of light transport and developing algorithms with better performance, some attention has been given to the problems related to inverse analysis, leading to interesting results. Other interesting results come from Applied Mathematics as well as Optical and Thermal Engineering.

Traditional direct problems in rendering involve computing the radiance distribution in an environment that is completely known a priori (geometry and materials). Inverse Rendering Problems refer to all the problems in which, as opposed to what happens in tra-

ditional direct rendering problems, several aspects of the scene are unknown. Some aspects that might be unknown are light source positions and/or their orientations, luminaries' emittances, surfaces reflectances and the shape or position and orientation of the surfaces (and other reflective/refractive elements) of the luminaries in the scene. The Rendering Inverse Problems are not well-posed: the solution does not depend continuously on the data, which means that small errors in measurements (input data) may cause large errors in the solution (see [HMH95]).

As proposed by Marschner [Mar98] and described in section 2.1, Inverse Rendering Problems can be classified into three classes: *Inverse Emittance*, *Inverse Reflectometry* and *Inverse Geometry*. In the present paper we study the contributions for Inverse Surface Design and characterize them in the class of Inverse Geometry Problems. It is important to note that, although the problem is often generically called *Inverse Reflector Design*, it encompasses the design of both re-

flective and refractive elements in an optical system. For a survey of *Inverse Emittance* and *Inverse Reflectometry* the reader can refer to [PP03].

The problem of Inverse Surface Design from light transport behavior specification, as focused in this survey, can be stated as follows: given a light source (bulb) with a known light intensity distribution, a surface should be constructed in such a way that a prescribed illumination intensity is obtained on a prescribed region in space, after reflection/refraction at the surface, see Figure 1. This prescribed distribution can be given either as a near- or a far-field distribution: the first given is directional and spatial distribution (although generally it is defined only as irradiance on a certain plane), and the second given is only in purely directional distribution terms. This later case can be thought of as a limiting case when the plane to be illuminated moves “infinitely” far away from the source. In general, radiance distributions are not defined in the continuum, but in some set of directions (far-field) or points in space (near-field), although they are later extended to the continuum by interpolation.

There are works found in the literature that present important problems and results, which have so many similarities that can generically be enclosed within the problem studied here. Among them we find works on the measurement of the human cornea (which acts as a reflector) and luminaire design (street lamps, car headlights, etc.). In addition, it is important to review the works on Inverse Design of Refractors, where the surface to be found is refractive instead of reflective. For this case, the design of progressive lenses represents a very important application.

The paper is organized as follows: In Section 2 the theoretical background is given, related problems are briefly commented, and a classification based on the rendering equation is presented. Next, in Section 3, analytical methods for surface design are explained, followed in Section 4 by the numerical solutions to the problem. In both cases, the methods presented are organized according to the treatment of the light transport problem they use, and then according to the type of surface used. Finally, in Section 5 the conclusions and open lines of research are presented.

## 2. Inverse Geometry Problems: Surface Design

### 2.1. Theoretical Background

The behavior of light transport is characterized by the properties of the particles (photons) when traversing the environment. The most fundamental quantity in global illumination is radiance  $L(\mathbf{r}, \omega)$  which is defined as the power radiated at a given point  $\mathbf{r}$  in a given

direction  $\omega$  per unit of projected area perpendicular to that direction per unit solid angle for a given frequency ( $Watt m^{-2} sr^{-1}$ ).

The boundary conditions of the integral form of the transport equation are expressed as:

$$L(\mathbf{r}, \omega) = L_e(\mathbf{r}, \omega) + \int_{S_i} f_r(\mathbf{r}, \omega_i \rightarrow \omega) L(\mathbf{r}', \omega_i) \cos \theta d\omega_i \quad (1)$$

for points  $\mathbf{r}$  in surfaces,  $f_r$  being the bidirectional reflection (and/or transmission) distribution function (BRDF),  $\theta$  the angle between the normal at  $\mathbf{r}$  and  $\omega_i$ ,  $\mathbf{r}'$  the point that is visible to  $\mathbf{r}$  in the direction  $\omega_i$ ,  $S_i$  the hemisphere of incoming directions with respect to  $\mathbf{r}$  and  $\omega_i$  an incoming direction.

This classical governing equation can be expressed concisely as a linear operator equation [Arv95a, Arv95b]. First, define the *local reflection operator*  $\hat{K}$  by

$$(\hat{K}h)(\mathbf{r}, \omega) \equiv \int_{S_i} k(\mathbf{r}; \omega_i \rightarrow \omega) h(\mathbf{r}, \omega_i) d\mu(\omega_i)$$

which accounts for the scattering of incident radiant energy. The measures  $d\omega_i$  and  $d\mu(\omega_i)$  are related by  $d\mu(\omega_i) = \cos \theta d\omega_i$ . Here  $h$  is a field radiance function, that corresponds to all incident light. Next, we can define the *field radiance operator*  $\hat{G}$ , that transforms an exiting light distribution into the incident light distribution that results from surfaces illuminating one another:

$$(\hat{G}h)(\mathbf{r}, \omega) \equiv \begin{cases} h(\mathbf{p}(\mathbf{r}; -\omega), \omega) & \text{when } \nu(\mathbf{r}, \omega) < \infty \\ 0 & \text{otherwise} \end{cases}$$


where  $\nu(\mathbf{r}, \omega)$  is the *visible surface function* and is defined [Arv95b] as  $\nu(\mathbf{r}, \omega) \equiv \inf\{x > 0 : \mathbf{r} + x\omega \in \text{Surfaces in the environment}, \infty\}$ , and  $\mathbf{p}(\mathbf{r}; \omega) \equiv \mathbf{r} + \nu(\mathbf{r}, \omega)\omega$  the *ray casting function*.

Defining these operators we can factor out the implicit function  $\mathbf{r}'(\mathbf{r}, \omega)$  from the integral Equation 1 and we may write:

$$L = L_e + \hat{K}\hat{G}L \quad (2)$$

Following the outlines in Stephen Marschner’s PhD. thesis introduction [Mar98], we can classify the different papers on inverse lighting problems according to which quantities of the above equation are unknown:

- If  $L_e$  is unknown, and  $\hat{K}$ ,  $\hat{G}$  and  $L$  or part of it, are known, we have a problem of *inverse lighting*.
- If  $\hat{K}$  is unknown, and  $\hat{G}$ ,  $L_e$  and part of  $L$  are known, we must solve for information about  $\hat{K}$ . This problem can, in general, be called *inverse reflectometry*, and a particular case is the one called *image-based reflectometry* in [Mar98], where images are used as input to the information about  $L$ .



**Figure 1:** Example of an Inverse Surface Design Problem: The desired light distribution is shown on the left, and the algorithm should produce the surface shown in the middle, which generates (in this case by reflection) the light distribution shown on the right.

- If  $\hat{G}$  is unknown, we have an *inverse geometry* problem, which encompasses the long-studied computer vision *shape from shading* problem, the *reflector design* inverse problem and the *recovering objects from photographs* problem. See Figure 1.

On the other hand, direct problems are those which, given known values for  $L_e$ ,  $\hat{K}$  and  $\hat{G}$ , solve for  $L$ .

## 2.2. Related Problems

In this survey we will review the work done on the family of problems enclosed in the last item of the previous section, when elements of the geometry are unknown, i.e. an *inverse geometry* problem, we will deal with the Inverse Surface Design problem. Other closely related problems are:

- The *shape from shading* problem, although similar in nature to the problem presented here, it is fundamentally different in that at each sampling position on the image (at the observer location) we can safely assume that only a small subset of points on the surface project onto it, and thus only a small subset of the surface contributes to its value (often, it is considered that only a single point contributes). Instead, in the problem treated here, there is usually a very important influence from all the points on the surface and the sampling positions (both in the far- and near-field problems), giving us a strong correlation between surface control points and sampling positions. A review of the *shape from shading* problem is beyond the scope of the present work, but the interested reader is referred to [HB89] or [SL97] for a description of the field.
- Another field closely related to the one reviewed here is *Optical Design*, which is the process of describing the refracting or reflecting elements in an optical system so that they meet a set of performance specifications. Typically, the performance specifications concern the imaging characteristics of the system, such as the resolution, magnification,

numerical aperture, and field of view. Other requirements may include tolerances, size, weight, and cost. The result of an optical design project is typically a prescription or database that lists the materials and shapes of the optical elements required. This is different from the problem treated in our survey, because in Optical Design problems the generic shape of the surfaces is generally known in advance, and can only be modified by global parameters such as lens thickness or radii, but not its *shape*. For more information on the subject, please refer to [KM00] [Soi02] [Sha97] and [Smi90]. To the best of our knowledge, most of the existing commercial software belong to this field, using either the design-direct simulation-restart principle [Jen01] [Org02a], or based on a local optimization [Org02b] [Cor] [Eng], or a global optimization approach [Vas98] [SO]. A comprehensive list of commercial software can be found in [Opt].

- Although the *Design of Reflector Antennas* is closely related to the problem treated here, most of the works in the field have unique features that make it quite difficult to generalize their results to incoherent radiation as produced by most light sources, like coherent radiation, which implies taking into account aspects such as phase cancellation and other interference effects [Kil00].
- It is also important to note the difference between computations of the shape of a source (see [PP03]), and the problem of surface design. The latter computes the shape of a reflecting or transmitting surface from light transport behavior, while the former computes the shape of the emitting surface itself (the bulb).

## 2.3. Inverse Surface Design Problems

We are interested in constructing of reflective/refractive surface shapes from prescribed optical properties of the luminaries and/or bulbs (far-field or near-field radiance distributions) and geometrical constraints. See Figure 3. From a mathematical




Figure 2: Inverse Surface Design Problem.




Figure 3: An optical set for reflector design.

point of view, the problem to solve belongs to the category of non-linear inverse problems. In the case of manufacturing design problems, restrictions on the shape imposed by the needs of the manufacturing process (Figure 3) should also be taken into account, which generally results in adding constraints to the problem.

The basic problem is building a surface shape from the description of the optical properties of the emitting light bulb and the desired light distribution the optical set (surface plus bulb) must provide. A user-provided tolerance is also given. Comparison of outgoing light distributions (the ideal desired one and the computed one) requires establishing a distance measure, which in turn will provide a logical target function for any optimization procedure: minimize the distance between the light distribution given by the current calculation surface (in general, a reflective surface) and the ideal, user-provided desired light distribution. The theoretical formulation of this problem leads to a non-linear partial differential equation of the Monge-Ampère type, as described in the literature [WN75] [Wes83] [EN91].

In general, we can say that all the works are based on the same scheme, which is illustrated in Figure 4: Finding the minimum of the function that describes the error in the outgoing light distribution for the optical set in the space of possible surfaces, with respect to the prescribed, user-given light distribution. As is known, optimization procedures require an *iter-*

*ative* procedure that repeatedly evaluates the objective function and determines the minimum from those measurements. In our case, the function to be evaluated always consists of two parts: the simulation of the light propagation from the light bulb to the registration area (as mentioned above, far-field or near-field), and finally an evaluation of the distance between the calculated outgoing light distribution and the user-provided desired distribution. The initial surface for the optimization must be, in general, manually provided.





Figure 4: Overall scheme of numerical Inverse Surface Design methods.

In general, we can say that the works reviewed here share the same elements which may be used to build an initial characterization (See Table 1). In particular, the different works studied can be characterized according to the treatment they give to the light transport from the light bulb to the evaluation region, the shape representation chosen, the optimization method, the type of BRDF used for the surface, the representation of the light emitter, the distance measure used and the nature of the problem focused on (far-field vs. near-field problems)

- **The function mapping light from the bulb to the evaluation space, usually represented by the light propagation algorithm.** Basically, the purpose of this step is to compute the light transport from the light bulb, reflecting or refracting on the surface, and arriving to the final, desired registration region. Although most of the studied algorithms use a local illumination-based scheme (light from the source bounces/refracts once on the surface and “goes” directly to the receiving region), inter-reflections are an important part of the simulation and should be considered. In general, we can observe that there are three main approaches to this particular aspect:


- “One-to-one incoming-to-outgoing rays” local illumination: This is a severe approximation that



**Figure 5:** *The one-to-one assumption: The left reflector satisfies the condition, since each ray that arrives on the reflector bounces in a different, primed direction, but the reflector on the right does not since rays labeled A and B bounce in the same direction ( $A' = B'$ ).*

assumes there is a one-to-one relation between incident rays and outgoing rays: no two rays can be reflected in the same outgoing direction (see Figure 5). This is an approximation mostly used for analytical works as it greatly simplifies the calculations needed for the convergence demonstrations.

- *Local illumination:* An illumination scheme in which rays from the light bulb bounce once, accordingly to the surface BRDF, towards the measurement region. This implies that no inter-reflections are computed, and in general, visibility issues (generated by the bulb and supporting elements or by the reflector itself) are disregarded.
- *Global illumination:* In this computation scheme, irradiance is followed as it bounces in its way out of the optical system. In general, this requires both computing multi-reflections as well as taking into account the visibility problem, as rays might be blocked by the surface itself, the bulb and/or supporting elements. Now, depending on the surface BRDF, many options can be considered, from radiosity solutions to a full treatment of Equation 1.
- **The shape definition for the surface, and if applicable, the restrictions imposed on the space of possible achievable shapes.** Different articles use different kinds of models for the surface representation, ranging from simple geometric primitives (for example, pieces of quadrics) or combinations of them, up to polygonal-based or NURBS-based definitions. Different choices have a direct impact on the number and type of the optimizable degrees of freedom, as well as modifying the light propagation algorithm to some extent.
- **The surface material used.** In most reviewed works it is a perfect specular BRDF. Other models could be used, but with a significant impact on




**Figure 6:** *The  $C - \gamma$  coordinate system.*

accuracy as the BRDF goes diffuse. This clearly influences the light propagation method.

- **Representation of the light emitter (bulb).** The light emitter can be represented either as a point source or as an emitting surface (closer to reality). It can be an isotropic source that emits equally in every direction (in general an approximation that is too simplistic) or anisotropic, which is the common choice. Once again, this mainly affects the light propagation algorithm through the system.
- **The definition of the distance between the desired outgoing light distribution and the corresponding one for a given surface.** Light distributions are usually represented in some sort of discrete set of directions and positions, so all works use the  $\ell^2$  norm to define the distance between radiance distributions. See [CW93] Section 10.1.4, General Luminaries. An example of such a distribution can be found in Figure 7.
- **The optimization method used.** Here is where the papers discussed differ most, because the non-linear inverse problem faced can be solved in different ways: by a global approach or by starting the algorithms close enough to the desired solution in such a way that a local optimization algorithm would lead to the correct solution. Other approaches impose restrictions on the achievable shape in order to guarantee local convergence.
- **The problem faced can be in its far-field or near-field form, although some works deal with both kinds of approaches.** As mentioned above, this affects the way the light propagation is evaluated and, of course, the error measure being used. Although most of the papers found in the literature make a clear distinction between both problems, the resulting algorithms/methods can, in general, be easily adapted to solve any one of the two types, thus blurring the importance of the distinction from a practical point of view (See Tables 3 and 5).

In general, when dealing with the *far-field* problem,



**Figure 7:** An example of a real outgoing radiance distribution used in industry.

	Theoretical	Numerical
Local Illum	[WN75] [Wes83] [BW78] [Oli89] [Wan96] [Oli02] [KO97] [KOvT98] [Oli03]	[CKO99] [KO03] [EN91] [Neu94] [Neu97] [KO98] [EN91] [KN96] [Neu97] [HBKM95a] [Hal96] [HBKM95b] [HBKM96] [Hal96] [LSS98]
Global Illum		[DCC99b] [DCC01] [DCC99a] [PPV04]

**Table 1:** Initial classification of the studied papers.

it is a very common choice to use a spherical coordinates system as a discrete representation for the outgoing radiance distributions, like the well known  $C - \gamma$  system [CM97], that represents a standard in the lighting engineering industry (see Figure 6).

In this survey we will build our classification around three main aspects from among those mentioned above: the analytical vs. numerical nature of the papers, the light propagating algorithm/method used to compute the outgoing radiance distributions and the shape definition used for the different computations.

### 3. Analytical Methods for Inverse Surface Design

This section deals with the theoretical analysis works done on the problem of Inverse Surface Design. Basically, they formulate the problem in precise mathematical terms using differential geometry, although adding important constraints on the formulation in order to make the problem theoretically tractable. For example, all of the reviewed works assume that the surface is perfectly specular, and many assume that

	Rotational Symm Approx	Non-rotational Symm Approx
1-1 in-to-out rays	[WN75] [Wes83]	
Local Illumination	[BW78] [Oli89]	[Wan96] [KO97] [KOvT98] [Oli02] [Oli03]

**Table 2:** Sub-classification of analytical papers on Inverse Surface Design.

	near-field	far-field
1-1 in-to-out Local Illumination		[WN75] [Wes83]
Local Illumination	[BW78] [Oli89] [KO97] [KOvT98] [Oli02] [Oli03]	[Wan96]

**Table 3:** Alternative classification of analytical papers on Inverse Surface Design.

there is a one-to-one relation between incident rays and outgoing directions, as mentioned above.

Below, we present the classification of theoretical papers according to the type of treatment given to the light propagation step: either they use a 1-to-1 correspondence between incoming and outgoing rays or a local illumination approach, vs. the type of surface they consider in their approaches. A summary can be found in Table 2.

However, it is also possible to present an alternative sub-classification of those works with respect to a far-field treatment vs. a near-field approach, as shown in Table 3, a distinction often made by the authors of the reviewed works.

#### 3.1. “One-to-one incoming-to-outgoing rays” local illumination

In this section we will review the contributions that treat the 1-to-1 approximation for local illumination as described above, in Section 2.1. This means that no two outgoing rays have the same orientation, a constraint usually referred to as the “one-to-one correspondence between incoming and outgoing rays” approximation (See Figure 5). To our understanding, this is too restrictive in practice. In general, these papers are based on the requirement of a rotational symmetric approximation for the studied surfaces, thus greatly reducing the space of the studied surfaces.

One of the first papers that dealt with this problem in its far-field approach is by Wescott et al. [WN75]. The authors presented an analytical formulation of the problem under the assumptions that the reflector surface is perfectly specular and rotationally symmetric. The authors studied the solutions for the case of fields with even azimuthal symmetry by using a spherical coordinate representation for them, and they also gave proof of existence and uniqueness.

Numerical solutions for these rotationally symmetric-restricted surfaces were also studied by Wescott [Wes83]. Both works, [WN75] and [Wes83], are focused on the design of reflector antennas in the particular case where incoherent light is reflected.

### 3.2. Local Illumination

The papers reviewed here treat the local illumination problem where rays are fired from the source, bounced once on the surface and reach their destination without any visibility or inter-reflection calculations. In addition, all of them use a perfect specular surface for their computations.

#### 3.2.1. Rotational Symmetric Surfaces

As mentioned above, many works impose the restriction that the surfaces used must be rotationally symmetric, thus reducing demonstrations to a simpler 2D problem.

In [BW78], Brickell et al. study the problem for the near-field distribution. In this case the distribution was considered on a flat object, arriving to a differential equation in complex form of Monge-Ampère type.

Oliker, in 1989 [Oli89] [Oli03], reformulated the same problem for rotationally symmetric reflectors in simpler differential geometry terms and without resorting to using complex structures. This resulted in a more general expression valid for the case of curved objects being illuminated in a prescribed way (near-field problem). He also proved the existence and uniqueness of the rotationally symmetric solution for the radially symmetric case.

#### 3.2.2. Non-Rotational Symmetric Surfaces

All the analytical works studied here are able to generalize the sort of surfaces they use for their studies to be non-rotationally symmetric, thus showing a significant improvement in the generality of the solutions found.

Following almost the same assumptions as Wescott [WN75] (i.e. perfect specular surfaces) Xu-Jia Wang [Wan96] studied the existence, uniqueness and

smoothness of the solution for the general problem in the far-field approximation. His presentation was based on a differential geometry formulation of the problem, which resulted in a clearer expression to work with. The author also showed that the regularity of the solutions fails in even the simplest cases.

In [KO97], Kochengin and Oliker considered the general problem with near-field scattering data without *a priori* assumptions regarding any rotational symmetry. Thus, they formulated the problems in differential geometry terms and established the existence of a weak solution. See also [Oli02] and [Oli03].

In another work, Kochengin, Oliker and von Temp-ski [KOvT98] studied the problem of finding a convex surface  $R$  which refracts a given anisotropic bundle of rays from a source in such a way that the refracted rays are incident on a specified set of points in space and produce a specified intensity distribution there. The refracting surface was found by taking the boundary of families of intersecting hyperboloids, each characterized by its polar radius. The light source is placed at the common focus of all hyperboloids, it is subdivided and rays are counted if they arrive at an objective point or direction [Koc].

## 4. Numerical Methods for Inverse Surface Design

The works reviewed in this section deal with numerical solutions of the proposed problem. We can sub-classify them according to how they treat the light propagation problem, showing results with both local-illumination and global-illumination algorithms. Then, we will characterize them based on the type of representation chosen for the surface (See table 4).

An alternative classification can be found in Table 5, according to the type of problem faced (far-field or near-field) vs. the type of light propagation method used, a classification often made in the literature.

### 4.1. Local Illumination

As stated previously, the works presented in this section share the common approximation of dealing with a simplified form of local illumination light propagation algorithms, where no inter-reflections or visibility issues are considered. They can be further classified according to the sort of representation chosen for the surface: the first set uses the intersection of simple primitives to define a convex surface, while the second one defines the surface with a spline-based representation.

	Local Illumination	Global Illumination
Convex Surface	[KO98] [Oli03] [CKO99] [KO03]	
Splines	[EN91] [KN96] [Neu94] [Neu97] [HBKM95a] [Hal96] [HBKM95b] [HBKM96] [LSS98]	[DCC99b] [DCC99a] [DCC01]
Polygonal Surface		[PPV04]

**Table 4:** Further classification of numerical papers on Inverse Surface Design.


	near-field	far-field
Local Illumination	[KO98] [Oli03] [EN91] [KN96] [Neu97] [HBKM95a] [Hal96] [HBKM95b] [HBKM96] [LSS98]	[CKO99] [KO03] [EN91] [Neu94] [Neu97]
Global Illumination	[DCC99b] [DCC99a] [DCC01]	[PPV04]

**Table 5:** Alternative classification of numerical papers on Inverse Surface Design.

#### 4.1.1. Convex Surfaces

In general, we can say that this section includes the works where a convex surface resulting from the intersection of a set of primitives is found. As we will see, these primitives are usually paraboloids, ellipsoids or hyperboloids depending on the type of problem faced: reflective far-field or near-field problems, or refractive ones.

In [KO98] Kochengin and Olikier continued their previous theoretical work [KO97] by presenting a numerical solution for a discrete version of the problem based on the demonstrations in [KO97]. By using the geometric optics property that any ray starting at one focal point of an ellipsoid arrives at the other after one bounce, they build a reflector as the convex body formed by the boundary of the intersection of ellipsoids sharing one foci and with their other foci on a point of the surface where a prescribed intensity is

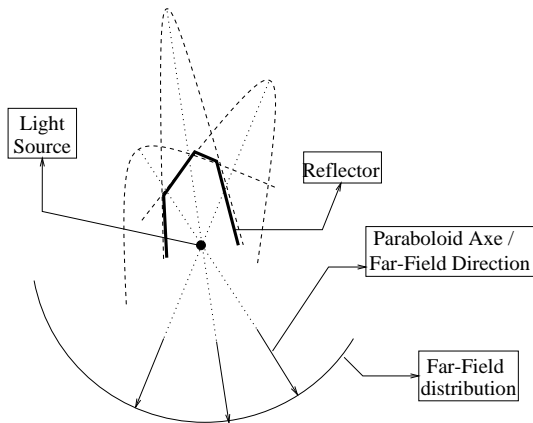


**Figure 8:** Reflector construction as the boundary of the intersection of confocal ellipsoids [KO98]. Notice that the other foci of each ellipsoid lies on a point with prescribed intensity on the illuminated object. Compare with Figure 9.

specified (Figure 8). The light source is placed at the common focal point, which results in a light propagation scheme that maps families of rays that impinge on an ellipsoidal piece onto a given destination foci. There are as many ellipsoids as points on the prescribed intensity distribution. As each ellipsoid can be described by its polar radius, the algorithm presents an iterative scheme that minimizes the difference between light arriving at each ellipsoid foci and the light prescribed there. In each iteration a new reflector is built from the previous one until there is no change from one reflector to the next one. Each reflector is built by decreasing each polar radius in turn from its previous value until the light difference for that prescribed point is below a certain threshold. If a polar radius already satisfies that condition, the old value is used instead. The authors show that the algorithm presents a time complexity of  $O(K^4 \ln K)$ , with  $K$  as the number of points with prescribed illumination on the target object, and of course, the number of ellipsoids to work with. A detailed study of the mathematical properties of such a resulting surface can be found in [Oli02], [Oli03].

For the discrete far-field problem, Caffarelli et al. [CKO99] present a very similar approach to the one presented by Kochengin and Olikier [KO98] for the near-field problem, but this time based on using paraboloids [Oli02] [Oli03]. They exploit the optical property of perfect paraboloidal reflectors that the light which starts at the focal point leaves the reflector in a direction parallel to its axis. Therefore, by having one paraboloid for each prescribed direction in the far-field region, they can build the final reflector as the boundary surface of the intersection of all confocal paraboloids (Figure 9). As in [KO98], the light source is placed at the focal point, thus mapping families of light rays that arrive at a paraboloidal





**Figure 9:** Reflector construction as the boundary of the intersection of confocal parabolooids [CKO99]. Notice that the axis of each parabolooid lies on a direction with prescribed intensity in the far-field region. Compare with Figure 8.

region onto a given direction in the far-field distribution. The algorithm is the same as described for that paper, but switching at the end to a final Newton-type method for faster convergence. Since the paper deals with the far-field problem, the global size of the reflector is unimportant, so it fixes the focal radius of one parabolooid and optimizes the remaining ones. The authors show that the algorithm converges and it does so at least linearly and, with a proper starting point, the Newton step gives super-linear behavior. Nevertheless, the same authors later [KO03] showed that their algorithm has three disadvantages with respect to its convergence: The last iterations show a significant decrease in their convergence, the convergence becomes worse as the number of parabolooids increases and the error for the parabolooid with a fixed focal radius is much bigger than the error for the remaining directions.

In [KO03] the above mentioned “brute force” method [CKO99] is compared with the Nelder-Mead simplex method [PTVF92]. The main disadvantage of the Nelder-Mead algorithm is that it does not guarantee convergence to a solution, but its convergence depends on how “nicely” the initial reflector is built. They conclude that using their original algorithm to find a first approximation (with only one iteration) and then using it as input to the Nelder-Mead method, reduces all the convergence shortcomings of the former method, providing faster and better convergence behavior. The fact that the error distribution is now much more uniform than in [KO98] is particularly noteworthy.

#### 4.1.2. Surfaces defined as splines

Defining surfaces as splines provides several advantages, like a good tradeoff between global control and good flexibility in the range of achievable shapes, but of course, there are also some drawbacks if the designer is interested in non-smooth surfaces, as points/lines of  $C^0$  continuity are hard to generate in an automated way.

In [EN91], Engl and Neubauer face both sorts of problems (far-field and near-field) without the need to map areas on the reflector to points on the near-field plane as the previous article did, although they also use perfect reflectors without any visibility or inter-reflection computations. They also assumed that the starting reflector would satisfy the constraint that no points/directions in the target region would be met by more than one light ray, a condition that is very difficult to satisfy for a real industry problem. In this context and for the near-field case, they derive the associated Monge-Ampère equation. The paper reports a first unsuccessful numerical attempt by using Newton’s method with line search for the non-linear partial differential equation, which leads to a sequence of linear elliptic boundary value problems. Unfortunately, this approach does not work well because, as soon as the non-linear problem is unsolvable in the finite element space used, Newton’s method becomes useless. Thus, they resort to an optimization approach: they represent the reflector as cubic tensor product splines with a given boundary curve, which is appropriate for the manufacturing process (compare this with the ellipsoid-based representation in [KO98]). They performed a conjugate gradient minimization (Powell’s method) of the  $L^2$ -error in the non-linear equation, taking constraints into account. Side conditions include that the boundary of the reflector can be fixed or variable within prescribed bounds, box constraints on the B-spline coefficients were imposed and a non-unattainable points condition on the reflector was imposed. Finally, the problem of finding an initial reflector to start their algorithm was left to further research, since it is very difficult to find a good starting reflector.

Later, Neubauer presented [Neu94] [Neu97] a solution for the far-field problem without the assumption that parallel outgoing rays will not occur for the starting reflector. Again, a bicubic B-spline representation is used and the spline coefficients are determined to be the best in the least squares sense. The sphere of outgoing directions is divided into 3200 subdomains and the solution is formulated as the minimization of the MS error between the obtained irradiances on the far-field sphere and the user-prescribed irradiances. The light propagation step is computed by Gaussian

quadrature based on  $4 \times 4$  nodes in the reflector parameterization domain. For each node the incoming and outgoing directions are computed, and each outgoing ray is assigned the corresponding value divided by 16. Since the distribution of light is required to be differentiable (to be able to analytically compute the derivatives), each outgoing ray has its energy “distributed” on the next four neighboring subdomains. Weights are added to the least squares formulation to influence accuracy in different subregions. The side conditions used are that the boundary curve of the reflector can either be fixed or free, box constraints on the B-spline coefficients and the inner product between the incoming ray direction and the surface normal are not allowed to change sign on the whole reflector surface. The minimization is solved iteratively by a projected conjugate gradient method developed by Powell, and the line search is performed via quadratic interpolation. Stability questions are briefly commented. As an example, a reflector generating uniform distribution on an infinitely distant plane is used with 169 spline coefficients (507 variables) and the deviation from the desired light distribution is reduced by a factor of 3 (take into account that this problem has many local minima).


In [KN96] [Neu97] the near-field problem is faced, which has a similar derivation of the integro-differential equation as the previous work [EN91]. Again, as in their earlier work [Neu94] for the far-field problem, a bicubic B-spline representation is used for the reflector surface, and the light propagation step is computed by tracing rays from the source to vertices on the surface, which were obtained from a regular grid in the spline parametric space. To achieve differentiability of the algorithm with respect to the spline coefficients, the light values are distributed onto the neighboring subregions defined in the near-field object in a differentiable way. The minimization procedure seeks to optimize a weighted least squares error with an extra added term that accounts for the light falling *outside* the near-field region. The optimization algorithm chosen is the same as in [Neu94].

In [HBKM95a] [Hal96] Halstead presents a different problem: The reconstruction of a three-dimensional surface model from an image of this surface illuminated with a structured pattern of light. The surface is ideally specular. The authors apply their algorithm to the measurements of the human cornea, which are made by a device called a *videokeratograph* (Figure 10a). The input to the algorithm is the set of feature positions (important features located on the image plane) extracted from a single image (Figure 10b). These positions are in the coordinate system of the image plane, and each feature must already be linked to an identifiable feature in the source pattern.

The algorithm is expected to output the reconstructed cornea in the form of a continuous function describing the surface’s position. For this function a biquintic B-spline is used, and its coefficients have to be found by the algorithm. The algorithm uses constrained optimization (Levenberg-Marquardt method) to solve the non-linear least squares problem and the objective is a function of the surface variables which measure the error between the surface and the original cornea. Its formulation greatly affects the efficiency of the solution process. The authors discovered that using the comparison of the real image to the image generated synthetically as the error measure is not a good option since useful error information is gained only around the boundaries between the structured regions. However, since each feature is associated with a corresponding point or set of points in the source pattern, a good error definition could be the square of the shortest distance between the intersection point of a ray from the image (reflected at the cornea surface) with the source pattern and the set of corresponding source points (Figure 11), the objective function is the sum of errors of all the image features.

Later, in [HBKM95b] [HBKM96] [Hal96], the error was computed differently: To evaluate the error of a single feature they again trace a ray from the feature (in the real image) through the nodal point (the camera acts as pinhole camera) to the surface. Then, they compute at which point they would like the reflected ray to intersect the source pattern. Finally, they keep the closest point in the set of corresponding features, and compute the normal that would reflect the incoming ray to the desired source point. The error is measured in terms of the difference between this normal and the current surface normal. The sum of the error terms of all features gives the objective function. This last definition reduced the error evaluation cost from hours to minutes. Each time the variation of the mean angle between normals changes a little in successive iterations, the surface is refined doubling the number of patches to optimize. Note that the above backward ray-tracing computations are done without any visibility or inter-reflection computations, but the extra evaluations are not needed due to the particular geometric setting used. Thus, the main difference from the previously mentioned approaches depends on the backwards ray tracing algorithm used, mainly due to the geometric nature of the stated problem. Another difference concerns error definition, based on a feature-match difference measurement.

In [LSS98], Loos, Slusallek and Seidel deal with a similar problem: designing Progressive Lenses for defects of the human visual system (a refraction problem). They use the results of wavefront tracing for re-focusing the eye while rendering and for deriving




**Figure 10:** a) Videokeratograph system and b) Captured image [HBKM96, Hal96]. Reprinted with the permission of Brian A. Barsky, UC Berkeley.

an accurate error functional that describes the desired properties and optical error across the lens. The algorithm performs a minimization of this error, yielding an optimal free-form refractive lens surface. The rendering algorithm finds, at first, the point in the environment the eye should focus on by tracing a single ray from the center of the pupil to the center of the pixel. Then, wavefront tracing is performed and the accommodation required to get a sharp image is computed. Finally, distribution ray tracing is used to compute the final pixel color. The error functional for describing the optimal progressive lens is based on wavefront tracing and the effective astigmatism, a generalization of traditional error measures used in optics. The lens is represented only as the front surface, with a fixed toroidal or spherical back surface. This back surface is discretized in about  $100 \times 100$  sample points. The optimization is performed by a variant of the well-known Newton-iteration method, showing that no more than two iterations were needed. The main difference with the other works reviewed in this section concerns the new error definition, which is an extremely specific, optics-based error measure.

## 4.2. Global Illumination

The papers studied in this section share the important feature of facing the Surface Design Problem in the context of global illumination treatment of light propagation. They can be classified according to the representation chosen for the surface: splines or simple polygons.



**Figure 11:** By using backward ray-tracing, in [HBKM96, Hal96] the authors determine the error in surface shape for a given image sample point. Reprinted with the permission of Brian A. Barsky, UC Berkeley.

### 4.2.1. Surfaces defined as splines

As mentioned above, splines show several good characteristics for global optimization, and the works studied in this section share the common approach of simulating global light propagation in a two-dimensional setting, thus only requiring 2D splines to represent the surfaces to optimize.

Doyle, Corcoran and Connell present [DCC99b] [DCC99a] [DCC01] an evolution strategy (a variation on the Genetic Algorithm) for 2D luminary design with point [DCC99b] and extended [DCC99a] light sources, later completed with a careful analysis of the objective function [DCC01] (called Merit Function in the papers). As such, they defined a three term objective function consisting of a term that represents the absolute element-wise difference between the target and resultant distributions in an objective region, and a term representing a penalty for the rays that lie outside the desired region, plus a third term that represents a measure of the difference between the reflected and the desired light power. The 2D surface definition is a cubic Bézier curve whose four control points are the genes of the chromosomes for the optimization evolution algorithm. Light propagation is performed by a recursive ray tracer with the number of bounces used as a stopping criterion (rays with more bounces than allowed are “lost”). Extended light sources were defined as a set of Lambertian-like (cosine distribution) point sources located at regular intervals on the perimeter of an arbitrarily sized circle. Each point source has its main axis in the circle radial direction. In this case, a penalty for the total light power re-intercepted by the extended source was added to the objective function. It is important to note that the main difference from all the other works mentioned is the usage of both a global 2D illumination algorithm and a global optimization technique like genetic programming for finding the surface-defining coefficients.

### 4.2.2. Polygonal Surfaces

In the work by Patow et al. [PPV04], the problem of the far-field case taking into account multiple bounces of light inside a reflector with occlusion and a generic BRDF was studied. They use a Monte-Carlo based light propagation algorithm that traces rays from the light source (punctual in their studies, but the presented algorithm easily accommodates extended light sources as well) and follows them until they get out of the optical system. The optimization algorithm used is a global “brute force method”, that conducts tests of the performance of each member of a family of polygonal reflectors obtained by iteratively combining the addition of an increment to each of the vertices. For this to be useful, the size of this generated family of reflectors has to be kept manageable. Obviously, the desired accuracy given by these increments added to the vertices (trying to ensure that the family contains at least one reflector close to the desired target), and the size of such a family are closely related. The vertices are sorted according to the contribution they make to the overall illumination error, and more effort (i.e. more samples) is put into the vertices with worst error. The results show good convergence of the algorithm, but the reported times are slow.

## 5. Discussion

Here we will summarize the conclusions of the surveyed work, to underline common problems, characterize solution approaches and present the open issues.

As seen, Inverse Surface Design papers can be grouped into two sets: analytical and numerical works. In Table 7 we see that the analytical works reviewed can in turn be grouped into methods that deal with the rotationally symmetric approximation ([WN75], [BW78] and [Oli89]) and those that don't ([Wan96], [KO97] and [KOvT98]). The former works present the equation formulation for the problem of reflector design in the rotationally symmetric simplification, and show existence and uniqueness of the solution. The latter present a more general equation formulation, showing existence, uniqueness and smoothness of the solution. In particular, [KO97] and [KOvT98] present a constructive demonstration in terms of weak convergence to the solution. It is also possible to group of the studied works according to the type of assumption they make, far-field or near-field, which is often found in the literature.

Papers on numerical methods, summarized in Table 6, can be characterized by:

- *Shape definition*: [KO98], [CKO99] and [KO03] use the boundary of the intersection of a series of simple primitives (ellipsoids or paraboloids) to define the

surface. [EN91], [Neu94, Neu97], [KN96, Neu97], [HBKM95a, Hal96], [HBKM95b, HBKM96, Hal96] and [LSS98] use tensor product of splines (bicubic or biquintic). A simple polygon-based shape definition was used in [PPV04].

- *Light Propagation method*: We see that almost all the reviewed papers, with the exceptions of [HBKM95a, Hal96] and [HBKM95b, HBKM96, Hal96], basically use light ray tracing [Arv86] to compute the light propagation at each iteration of their optimization method. The mentioned exceptions use traditional ray tracing [Whi80], basically because they are computing the shape of the human cornea and working with a convex surface shape, and all the reflected rays reach the *videokeratograph*. To the best of our knowledge, almost all works which study the far-field problem only use local illumination schema to solve the light propagation step, with the only exception being [PPV04].
- *Type of problem faced*: It can be either a near-field formulation of the problem or a far-field one. In the former, the desired light distribution is defined on a registration area with distances comparable to the ones involved in the optical set, while in the later it is defined at a region infinitely far away from it.
  - near-field distribution: [EN91], [KO98], [KN96], [Neu97], [HBKM95a], [Hal96], [HBKM95b], [HBKM96], [Hal96], [LSS98], [DCC99b], [DCC99a] and [DCC01].
  - far-field distribution: [EN91], [CKO99], [KO03], [Neu94], [Neu97] and [PPV04].
- *Surface material assumptions*: We see that almost all the reviewed works only use the perfect specular BRDF; in addition [KO98], [CKO99] and [KO03] and [EN91] assume that there are no two outgoing rays with the same orientation. This basically implies that every two different points on the reflector surface will reflect light in two different directions. The other papers surveyed do not use this highly restrictive assumption. To the best of our knowledge, the only paper that uses a generic BRDF is [PPV04], sampled by a Monte Carlo scheme.
- *Optimization method*: [KO98], [CKO99] and [KO03] use and compare their own custom method (named “Brute Force Method”) with traditional methods like Newton or Nelder-Mead. [EN91], [Neu94], [Neu97] and [KN96] [Neu97] use the Projected Conjugate Gradient method; [HBKM95a], [Hal96] use the Levenberg-Marquardt method, and [HBKM95b], [HBKM96], [Hal96] use Least Squares as the chosen optimization method. [LSS98] uses a modified Newton-iteration method. On the other hand, Doyle et al. [DCC99b] [DCC99a] [DCC01] use a Genetic Algorithm for the optimization, and

[PPV04] uses a custom global “Brute Force” algorithm. To the authors knowledge, choosing one optimization method or another is quite arbitrary, and no method has proven to be the best for this sort of optimization problem.

We have also seen that almost all works use a local illumination approach, with the further restrictions of using pure specular surfaces and the highly restrictive condition that no inter-reflections occur, with the only exception in the works by Doyle et al. [DCC99b] [DCC99a] [DCC01], which use global illumination with a perfect specular BRDF model with a threshold based stopping criteria, and [PPV04], in which a more generic global illumination algorithm is used.

- *Inter-reflections*: The fact that almost no article inter-reflections are computed is especially important when the reflector is highly specular and its surface concave, as in most papers reviewed, since a very important contribution to the final light distribution may be underestimated. This is not applicable to [HBKM95a] [Hal96] and [HBKM95b] [HBKM96] [Hal96] due to the convexity of the surface used to approximate the human cornea, and to [LSS98] which computed the refractive surface of a lens.
- *BRDF*: As mentioned, most of the commented works deal only with pure specular surfaces. To understand what happens if other BRDFs are used, we can consider that Equation 1 can be regarded in a signal processing framework [RH01] under the restrictions of no inter-reflections (as most of the papers assume), isotropic BRDFs, known geometry (besides the reflector) and camera parameters. Then, the reflected light field integral is regarded as the convolution of two signals: the bidirectional reflectance function and the incident lighting; i.e. by filtering the illumination using the BRDF. Then, inverse rendering can be simply viewed as a deconvolution of the two signals. This framework [RH01] led the authors to conclude that BRDF recovery is well-conditioned (in a mathematical sense) when lighting contains high frequencies (e.g. directional sources) and is ill-conditioned for soft lighting. Counterwise, inverse lighting is well-conditioned for BRDFs with high-frequency components (specular peaks) and ill-conditioned for diffuse surfaces. The same analysis can be carried on for the inverse problem of Surface Design, leading to the conclusion that convergence of the algorithms would become worse as the BRDF gets more diffuse: this is one of the main limitations for solving a generic-BRDF problem. Experimental results [PPV04] confirm this analysis.
- *Local vs. Global Illumination*: Another factor to take into account is that almost none of the different

papers treat the full global illumination equation, Equation 2, limiting themselves to using of a simpler local-illumination version based on a simplification of the equation being used by considering only point light sources and without considering inter-reflections. As most of the reviewed papers omit the treatment of inter-reflections, the resulting algorithms provide solutions that are not readily applicable in real-life situations. Multiple reflections would not only introduce more complexity in the light propagation method evaluation (requiring more computing time to get accurate results), but would also severely increase the associated variance for non-specular surfaces, thus lowering convergence, as can be seen in [PPV04]. Other studies that take into account global illumination effects are presented in [DCC99b] [DCC99a] [DCC01], but, as they use a perfect specular surface, the results do not show any serious increase in the associated variance, and thus do not affect the way their algorithms converge.

- *Visibility*: It is also important to mention the treatment of the visibility in the different approaches reviewed: when computing the radiance with the above equations, the visibility problem consists of detecting if there are any blockers between the source and the surface being illuminated, and not adding their contribution in that case. The same is true for the paths from the surface to the eye or the region where the final radiance computations are needed. In general, this would be solved by adding visibility calculations to most approaches, simply by not accounting for occluded light. This would affect the error measure and steer the optimization, adding a strong non-linearity which would cause serious convergence problems for local optimization based algorithms. In those cases, a stimulated annealing approach ([PTVF92]), a genetic algorithm ([DCC01] [DCC99b] [DCC99a]) or even a brute force approach would be needed [PPV04]. This can be seen in the work by Doyle, Corcoran and Connell [DCC99a], where the occlusion introduced by the extended light source is introduced. This resulted in a restriction of the search space of possible reflectors, which forced the authors to introduce several changes to the reflector shape and size in order to compensate for this.

From our survey we can derive that several inverse problems remain open:

- Letting an algorithm automatically choose between some possibilities the more suited material for a surface to get a given illumination.
- Another open problem can be found in determining the shape of surfaces by using the full rendering equation, since all treatments up to now deal with

highly simplified versions of the problem. Additionally, more general BRDFs should be considered.

- One promising line of research is the one introduced by Costa et al. [CSF98], in which they show that the potential equation can be used to simplify the problem from considering the whole environment to only a small subset of it, thus reducing the number of possible degrees of freedom to be taken into account. In spite of this, the general problem of obtaining inverse surface shapes from general shading in an environment remains untouched.
- Determining scene shapes other than sources and surfaces from direct illumination information could be an interesting problem to solve, in spite of its inherent high complexity. For example, it would be interesting to find the shape needed in a room to get a desired illumination effect, both by its direct contribution as a light-blocking element, and by the distortions introduced by its presence in the light propagation through the environment.
- Calculating shapes in the presence of participating media has never been considered, despite their importance for design in any given adverse conditions.

## 6. Acknowledgments

We would like to thank Pierre Poulin for his insights, Brian A. Barsky for giving us permission to reprint his figures, Sergey A. Kochengin for the extra explanations about his work and Jack Howell for his reprints and for proofreading an earlier version of this paper. Partially done under grant TIC2001-2226-C0202 of CICYT (Spain) and 2001SGR0296 of DURSI of the Generalitat de Catalunya.

	near/ far	Shape Definition	BRDF	Light Emitter	Global/ local illum	Light Propagation	Optimization Method	Approx.
[WN75]	far	Rotationally Symm.	Perfect Specular	Isotropic point light	Local	1-1	Theoretical	Shape with even azimuthal symm.
[Oli89]	near	Rotationally Symm.	Perfect Specular	Point light	Local	1-1	Theoretical	Shape with even azimuthal symm.
[Wan96]	far	Generic	Perfect Specular	Point light	Local	Theoretical ray tracing	Theoretical	Convex Shape
[KO97]	near	Families of confocal ellipsoids	Perfect Specular	Point light	Local	Maps sets of rays impinging on a ellipsoid onto a direction	Theoretical	Convex Shape
[KovT98]	near	Families of confocal hyperboloids	Perfect REFRACTIVE	Point light	Local	Theoretical tracing of refractive rays	Theoretical	Convex Shape
[KO98]	near	$\partial(\cap \text{ellipsoids})$	Perfect Specular	Point Source	Local	Light ray tracing. Source is subdivided and rays are counted if arrive to objective point/ dir.	BFM	$\bar{I}, \bar{V}$ , convex shape
[CKO99]	far	$\partial(\cap \text{paraboloids})$					BFM + Newton	$\bar{I}, \bar{V}$ , convex shape
[KO03]	far	$\partial(\cap \text{paraboloids})$					Nelder-Mead	$\bar{I}, \bar{V}$ , convex shape
[EN91]	both	Cubic spline tensor product	Perfect Specular	Point Source	Local	Not stated explicitly. Probably as in [Neu94] [Neu97]	Projected Gradient	$\bar{I}, \bar{V}$ , No parallel outgoing rays
[Org02b]	both	Segmented reflector	Perfect Specular	Point Source	Local	ray tracing	Optimization by user	$\bar{I}, \bar{V}$
[Neu94]	far	Bicubic B-splines	Perfect Specular	Point Source	Local	Gaussian quad. on refl. param. domain. For each point, in & out rays are found and out-ray is "distributed" on Far-Field subdomains.	Projected Gradient with line search via quadratic interpolation.	$\bar{I}, \bar{V}$
[KN96]	near	Bicubic B-splines	Perfect Specular	Point Source	Local	Gaussian quad. on refl. param. domain. For each point, in & out rays are found and out-ray is "distributed" on Near-Field subdomains.	Projected Gradient with line search via quadratic interpolation.	$\bar{I}, \bar{V}$
[Neu97]								
[HBKM95a]	near	Bi-quintic B-splines	Perfect Specular	Point Source	Local	Backwards ray-tracing from the feature (in the real image) through the nodal point to the specular surface, then reflected and traced to the pattern source.	Levenberg-Marquardt method.	$\bar{I}, \bar{V}$
[Ha196]							Least Squares	
[HBKM95b]								
[HBKM96]								
[Ha196]								
[LSS98]	near	bicubic NURBS	Perfect REFRACTIVE	Point source	Local	ray tracing to locate objects, wave-front tracing to re-focus the eye + distribution ray-tracing	Modif. Newton-it	
[DCC99b]	near	2D Bézier curve	Perfect Specular	Point Source	2D Global	Ray tracing	Evolution Strategy	2D
[DCC01]								
[DCC99a]	near	2D Bézier curve	Perfect Specular	Extended lambertian Light Source	2D Global	Ray tracing from uniform samples on source	Evolution Strategy	2D
[PPV04]	far	Polygonal Surf	General	Point (Extended too)	Global	Monte Carlo Light Tracing	"Brute Force"	

Table 6: Papers with numerical solutions to the Reflector Design problem. (continued).

## Appendix A: Appendix

In this appendix we analyze the mathematical aspects of each surveyed paper's solution method. Table 8 explains the constraints imposed on the solution method (for optimization approaches) for Inverse Geometry problems. In this case the restrictions are incorporated into the optimization algorithm. The third column shows the starting point used by the different methods. Finally, the fourth column presents the objective function to be optimized in the case of optimization approaches.



	near-field / far-field	Demonstrates / Shows	Restrictions / approximations
[WN75]	far	Eq. formul. $\exists$ & uniq. for RS.	PS, 1-1, $\neg I$ , $\neg V$ , RS
[BW78]	near	Eq. formul. $\exists$ & uniq. for RS. Dem. that paraboloid is sol. for const. phase func.	PS, 1-1, $\neg I$ , $\neg V$ , RS
[Oli89]	near	$\exists$ & uniq. for RS.	PS, 1-1, $\neg I$ , $\neg V$ , RS
[Wan96]	far	$\exists$ , uniq. & smoothness	PS, $\neg I$ , $\neg V$
[KO97]	near	Probl. def. $\exists$ of weak sol. for the class of convex sol.	PS, $\neg I$ , $\neg V$
[KOV98]	near (refractions)	Probl. def. $\exists$ of weak sol. for the class of convex sol.	PS, $\neg I$ , $\neg V$

**Table 7:** *Theoretical papers on Inverse Surface Design. In this table “RS” stands for Rotationally Symmetric approximation, “PS” means that the surfaces are perfectly specular or perfectly refractive, “1-1” that the incoming-ray/outgoing-ray uniqueness relationship approximation is used, “ $\neg I$ ” means that interreflections are disregarded and “ $\neg V$ ” means that no visibility computation is done. The second column refers to the type of problem faced: near-field or far-field problem.*

Paper	Constraints imposed	Initial guess / Initial approx.	Objective function	Tests Performed	comments
[KO98]	None	$M = \frac{max_x  x }{\sum_{x \in S_{surf}} f_{acc}  x }$ then first polar radius = $16M$ and all other p.r. = $64M$ , i.e. reflector is first ellipsoid	Minimize each error ( $\ell^\infty$ ) separately	Plane uniformly illuminated with 4, 10 and 25 ellipsoids.	Shows convergence & $O(K^4 \ln K)$
[CKO99]	First polar radius is fixed	First Polar Radius = 1, rest chosen such that radiance at a point $\leq$ prescribed energy + $\epsilon$	Minimize each error ( $\ell^\infty$ ) separately	Uniform illum with 7, 19 paraboloids.	$O(\ln near)$
[KO03]	First polar radius is fixed for all reflectors	Simplex vertices in neighborhood of initial reflector	Minimize simplex energy	Uniform illum with 17, 19 and 61 paraboloids.	$O < \ln near$
[EN91]	Boundary fixed or variable. Box constr. on B-spline coeffs. Non-unattainable points on refl.	User provided with no parallel outgoing rays for first reflector. Too difficult to provide!	Least squares error between the current light distribution from the reflector and the desired distribution.	Synthetic examples: $5 \times 5$ nodes, good results. Real Ex. (Poor behavior due to bad starting point): Rot. sym. refl. w/ unif. illum. on pl. at $\infty$ : $10 \times 10$ nodes, 48 iterations (40 minutes).	Only coarse solutions: too high evaluation cost.
[Neu94] [Neu97]	Boundary fixed or variable. Box constr. on B-spline coeffs. Surface normal is not allowed to change sign on the refl. surface.	User's (manufacturer) provided.		Rot. sym. refl. w/ unif. illum. on subreg. of pl. at $\infty$ : $10 \times 10$ nodes, 360its (6 days w/ analytic grads.), error reduced 3 times.	
[KN96] [Neu97]	Boundary fixed or variable. Box constr. on B-spline coeffs. Surface normal is not allowed to change sign on the refl. surf. No reflected ray must go away from near-field reg.	User's (manufacturer) provided.		Synthetic Ex: $10 \times 10$ nodes, near planar subreg., error reduced from 189% to 3%. Real Ex: near planar subreg., 149 its (5.5 hours), error reduced 3 times.	
[HBKM95a] [Hal96]	None	Little effect: paraboloid	$Error = \sum  feature pos - backwards trac. ray pos ^2$	Synthetic ellipsoid: $1$ to $8 \times 8$ patches. $RMS E \leq 8.5 \times 10^{-6} mm$ . 5760 samples at image.	$Error = \sum  feature pos - backwards trac. ray pos ^2$
[HBKM95b] [HBKM96] [Hal96]	None	Guess at the shape of the cornea	$Error = \sum dist(surface normal, desired normal)$	Synthetic data sets: RMS of 0.0092 microns! Real Data sets: RMS of 0.9-1.5 microns. All: $8 \times 8$ final patches, over 5000 image-samples.	$Error = \sum dist(surface normal, desired normal)$
[LSS98]	None	front-surf with const dist form back-surf	Error functional that describes the aberrations in a progressive lens	bi-cubic tensor product of NURBS, took 66.5 secs in a SGI Onix w/ R10000.	
[DCC99b] [DCC99a] [DCC01]	None	Random Initial Population	Diff between target and resultant distrib at obj region plus lost rays plus difference between reflected and desired power [DCC01]	2D, parabolic & elliptic reflectors	
[PPV04]	None	Generic Initial Refl	Least squares error between the current light distribution from the reflector and the desired distribution.	Different Configurations. 10 days on a P-IV PC.	

Table 8: Summary of the mathematical aspects of the inverse lighting approaches presented: Inverse Geometry (continued).

## References

- [Arv86] James R. Arvo. Backward Ray Tracing. In *ACM SIGGRAPH '86 Course Notes - Developments in Ray Tracing*, volume 12, August 1986.
- [Arv95a] James Arvo. *Analytic Methods for Simulated Light Transport*. Ph.D. thesis, December 1995.
- [Arv95b] James Arvo. The Role of Functional Analysis in Global Illumination. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 115–126, New York, NY, 1995. Springer-Verlag.
- [BW78] F. Brickell and Brian S. Westcott. Phase and power distribution on plane apertures of reflector antennas. *J. Phys. A: Math. Gen.*, 11:777–789, 1978.
- [CKO99] Luis A. Caffarelli, Sergey A. Kochengin, and Vladimir I. Olikier. On the numerical solution of the problem of reflector design with given far-field scattering data. *Contemporary Mathematics*, 226, 1999.
- [CM97] J. R. Coaton and A.M. Marsden. *Lamps and Lighting*. Ed. Arnold, London, 1997.
- [Cor] ZEMAX Development Corporation. Zemax software for optical design. Technical report, ZEMAX Development Corporation, <http://www.zemax.com>.
- [CSF98] Antonio Cardoso Costa, A. Augusto Sousa, and F. Nunes Ferreira. Design de iluminação. In *8th Portuguese workshop on Computer Graphics*, Coimbra, Portugal, 1998. in Portuguese.
- [CW93] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Boston, MA, 1993.
- [DCC99a] Steven Doyle, David Corcoran, and Jon Connell. Automated mirror design for an extended light source. In *Proceedings of SPIE*, volume 3781, page 94, 1999.
- [DCC99b] Steven Doyle, David Corcoran, and Jon Connell. Automated mirror design using an evolution strategy. *Optical Engineering*, 38(2):323–333, 1999.
- [DCC01] Steven Doyle, David Corcoran, and Jon Connell. A merit function for automated mirror design. *Journal of the Illuminating Engineering Society*, 30(2):3–11, 2001.
- [EN91] Heinz W. Engl and Andreas Neubauer. Reflector design as an inverse problem. In Heiliö M., editor, *Proceedings of the Fifth European Conference on Mathematics in Industry*, pages 13–24, Teubner, Stuttgart, 1991.
- [Eng] West Coast Engineering. Odp optical design program. Technical report, West Coast Engineerin, <http://www.westcoastengineering.com/>.
- [Hal96] Mark A. Halstead. *Efficient Techniques for Surface Design Using Constrained Optimization*. PhD thesis, University of California, Berkeley, May 1996.
- [HB89] Berthold K. P. Horn and Michael J. Brooks. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
- [HBKM95a] Mark A. Halstead, Brian A. Barsky, Stanley A. Klein, and Robert B. Mandell. Geometric modeling of the cornea using videokeratography. In Morton Daelhen, Tom Lyche, and Larry L. Schumaker, editors, *Mathematical Methods for Curves and Surfaces*, pages 213–223. Vanderbilt University Press, Nashville, 1995.
- [HBKM95b] Mark A. Halstead, Brian A. Barsky, Stanley A. Klein, and Robert B. Mandell. A spline surface algorithm for reconstruction of corneal topography from a videokeratographic reflection pattern. *Optometry and Vision Science*, 72(11):821–827, November 1995.
- [HBKM96] Mark A. Halstead, Brian A. Barsky, Stanley A. Klein, and Robert B. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. *Computer Graphics*, 30(Annual Conference Series):335–342, 1996.
- [HMH95] Vigain Harutunian, Juan C. Morales, and John R. Howell. Radiation exchange within an enclosure of diffuse-gray surfaces: The inverse problem. In *Inverse Problems in Heat Transfer, ASME/AIChE National Heat Transfer Conference*, Portland, August 1995.
- [Jen01] Daid G. Jenkins. Reflector design using lighttools. Application

- summary, Optical Research Associates, <http://www.opticalres.com>, 2001.
- [Kil00] Per-Simon Kildal. *Foundations of Antennas - A Unified Approach*. Studentlitteratur, Lund, Sweden, 2000.
- [KM00] Bernhard Kress and Patrick Meyrueis. *Digital Diffractive Optics*. Ed. John Wiley and Sons, New York/Chicester/Brisbane/Toronto, 2000.
- [KN96] Christa König and Andreas Neubauer. Design of a 3d-reflector for the near field problem: a nonlinear inverse problem. *Math. engng. Ind.*, 5(4):269–279, 1996.
- [KO97] Sergey A. Kochengin and Vladimir I. Oliker. Determination of reflector surfaces from near-field scattering data. *Inverse Problems*, 13:363–373, 1997.
- [KO98] Sergey A. Kochengin and Vladimir I. Oliker. Determination of reflector surfaces from near-field scattering data ii. numerical solution. *Numer. Math.*, 79(4):553–568, 1998.
- [KO03] Sergey A. Kochengin and Vladimir I. Oliker. Computational algorithms for constructing reflectors. *Computing and Visualization in Science*, 6:15–21, 2003.
- [Koc] Sergey A. Kochengin. Personal Communication.
- [KovT98] Sergey A. Kochengin, Vladimir I. Oliker, and Oliver von Tempiski. On design of reflectors with prespecified distribution of virtual sources and intensities. *Inverse Problems*, 14(3):661–678, 1998.
- [LSS98] Joachim Loos, Philipp Slusallek, and Hans-Peter Seidel. Using wavefront tracing for the visualization and optimization of progressive lenses. *Computer Graphics Forum (Eurographics '98)*, 17(3):–, 1998.
- [Mar98] Stephen R. Marschner. *Inverse Rendering in Computer Graphics*. PhD thesis, Program of Computer Graphics, Cornell University, Ithaca, NY, 1998.
- [Neu94] Andreas Neubauer. The iterative solution of a nonlinear inverse problem from industry: Design of reflectors. In P. J. Laurent, A. Le Méhauté, and L. L. Schumaker, editors, *Curves and Surfaces in Geometric Design*, pages 335–342, Boston, 1994. A. K. Peters.
- [Neu97] Andreas Neubauer. Design of 3d-reflectors for near field and far field problems. In L.T. Bieyler, T. F. Coleman, A.R. Conn, and F.N Santosa, editors, *Large Scale Optimization with Applications. Part I: Optimization in Inverse Problems and Design*, volume 92 of *IMA Volumes in Mathematics and its Applications*, pages 101–118. Springer, New York, 1997.
- [Oli89] Vladimir I. Oliker. On reconstructiong a reflecting surface from the scattering data in the geometric optics approximation. *Inverse Problems*, 5:51–65, 1989.
- [Oli02] Vladimir I. Oliker. On the geometry of convex reflectors. *Banach Center Publications*, 57:155–169, 2002.
- [Oli03] Vladimir I. Oliker. Mathematical aspects of design of beam shaping surfaces in geometrical optics. In S. Kromker, R. Rannacher, and F. Tomi, editors, *Trends in Nonlinear Analysis*, pages 193–224, Springer-Verlag, 2003.
- [Opt] Optenso. Links on optical design and engineering. Technical report, Optenso: Optical Engineering Software, <http://www.optenso.de/links/links.html>.
- [Org02a] Breault Research Organization. Asap technical summary. Application Summary BRO-0400, Breault Research Organization, <http://www.breault.com>, March 2002.
- [Org02b] Breault Research Organization. Create segmented reflectors quickly with reflectorcad. Feature Note BRO-FN1405, Breault Research Organization, <http://www.breault.com>, January 2002.
- [PP03] Gustavo Patow and Xavier Pueyo. A survey on inverse rendering problems. *Computer Graphics Forum*, 22(4):663–687, 2003.
- [PPV04] Gustavo Patow, Xavier Pueyo, and Àlvar Vinacua. Reflector design from radiance distributions. *International Journal of Shape Modelling*, 10(2):211–235, 2004.
- [PTVF92] W. H. Press, S.A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C : The Art of Scientific Computing*. Cambridge University Press, 1992.
- [RH01] Ravi Ramamoorthi and Pat Hanrahan.

- A signal-processing framework for inverse rendering. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001)*, August 2001.
- [Sha97] Robert R. Shannon. *The Art and Science of Optical Design*. Cambridge University Press, 1997.
- [SL97] A. James Stewart and Michael S. Langer. Towards accurate recovery of shape from shading under diffuse lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):1020–1025, September 1997.
- [Smi90] Warren J. Smith. *Modern Optical Engineering (Second Edition)*. McGraw-Hill, 1990.
- [SO] Inc. Sinclair Optics. Oslo: Optics software for layout and optimization of optical systems. Technical report, Sinclair Optics, Inc., <http://www.sinopt.com/>.
- [Soi02] Victor A. Soifer. *Methods for Computer Design of Diffractive Optical Elements*. Ed. John Wiley and Sons, New York/Chicester/Brisbane/Toronto, 2002.
- [Vas98] Darko Vasiljevic. Optimization of the cooke triplet with the various evolution strategies and the damped least squares. In *Optical Design and Analysis Software, Proceedings of SPIE*, volume 3780, pages 207–215, 1998.
- [Wan96] Xu-Jia Wang. On the design of a reflector antenna. *Inverse Problems*, 12:351–375, 1996.
- [Wes83] Brian S. Wescott. *Shaped Reflector Antenna Design*. Letchworth: Research Studies Press, 1983.
- [Whi80] Turner Whitted. An improved illumination model for shaded display. *CACM*, 23(6):343–9, June 1980. also in Tutorial: Computer Graphics: Image Synthesis, Computer Society Press, Washington, 1988. (Abstract) in *Computer Graphics*, no. 2 (SIGGRAPH '79 Proceedings), p. 14.
- [WN75] Brian S. Westcott and A. P. Norris. Reflector synthesis for generalized far-fields. *J. Phys. A: Math. Gen.*, 8(4):521–532, 1975.