

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
dagush@iia.udg.es, 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.

Keywords: computer graphics, inverse problems, global illumination, lighting design, luminaire design.

ACM CSS: 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 [61], and an excellent survey on inverse placement of cameras, curves and objects for animation can be found in [27] 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. Traditional 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 [23]. The inverse rendering problems lack at least one of the Hadamard [19] 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 [23]).

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 [20], 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 [20] 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 [8, 56]. 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 ω 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), θ the angle between the surface normal at \mathbf{r} and ω , \mathcal{S}_i the hemisphere of incoming directions with respect to \mathbf{r} and ω_i an incoming direction.

This classical governing equation can be concisely expressed as a linear operator equation [1, 2]. 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 [2] 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 [36], 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 [36], 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 [46].

Equation (1) can be regarded in a signal processing framework [51] under the restrictions of distant illumination,


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

Inverse Lighting (L_e)	[52], [37], [36], [7], [55], [35] [23],[14], [12], [32], [15], [38] [40], [51], [59],[58],[57] [42],[47],[18]
Inverse Reflectometry (\hat{K})	[9], [3], [14], [12], [62], [32], [31], [26], [28], [44], [41], [13], [34], [45, 43], [36, 39], [54], [60], [10], [63], [36], [51]
Combined Problems (L_e and \hat{K})	[29], [51], [58], [59]

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 [51] 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 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.

**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**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.

	Emittances	Positioning
General	[52],[37], [36]	[7], [55], [35]
Radiosity	[23], [14], [12], [32], [15], [38]	
Monte Carlo	[40]	
Local	[51], [59], [58], [57]	[42],[47], [18]

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 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 Φ_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 \Re with some measured intensity values α_j in the scene or in a screen, we arrive at an expression of the form

$$\alpha_j = \Re(\Phi) = \sum_i^n u_i \Re(\Phi_i) \quad (3)$$

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, ρ_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 trivial, 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 [37, 36] studied the ill-conditioning of the inverse lighting problem in the case of diffuse surfaces, and later Ramamoorthi and Hanrahan [51] 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.* [52], where the user defines the light features by “spraying” color onto surfaces. Later, Stephen Marschner and Donald Greenberg [37, 36] 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 α_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 decomposition (SVD) to allow the regularization parameter to be adjusted interactively. Here, the α_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, ρ_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.* [23], 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 [25] (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.* [15] and Morales *et al.* [38] 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 [14] also base their formulation on the radiosity approximation, directly fitting the element emissions to the observed values (α_i). Drettakis, Robert and Bougnoux [12] improved the work of Fournier *et al.* by setting up a hierarchical radiosity system. Loscos *et al.* [32] 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 [40]. 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 [51] 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.* [59] 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 [58], 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 [51] 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.* [57] 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 [11] 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 [49, 48].
- **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 [50]. 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 [8],

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

where θ_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 [5] formula.

3.2.1. General LSPP Formulations

Costa *et al.* [7] 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 [17]. 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 [49, 48]. 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, Shackled and Lichinski [55] 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 [53]. The main difference with [7] 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.* [35]. 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, [42]. 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 [5] shading as expressed by Blinn [4]. 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 determining 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 [16] generated by it must be used [45]. 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 [47] 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 [22], 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




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

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 threshold τ . 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 [42] 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 [18], 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 [52], [37], [36] and [51], 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. ([51] 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 [40] and [52], [37], [36], mainly because the first paper presents an algorithm that samples on the surfaces pointwise and the others are BRDF independent.
- Only França *et al.* [15] and Morales *et al.* [38] 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 PDE.
- Among the reviewed papers, only [52], [37], [36] and [29] 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.* [52] 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 [37, 36], 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.* [59, 58] 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 [51] 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 [7], [55] and [35] 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.* [7] 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 [55] and Costa *et al.* [7] 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.* [35], 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 [47] with the work by Pouin and Fournier

Table 3: Classification of inverse reflectometry problems organized with respect the treatment they give to Equation (1): general, radiosity, Monte Carlo and local.

General		[9], [3]
Radiosity		[14], [12], [62], [32], [31]
Local	Texture-based	[26], [28], [44], [41], [13], [34]
	Angular-based	[45, 43], [36, 39], [54], [60], [10]
	General	[63], [36], [51]

[42] 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 [18] 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{G} , L_e and part of L are 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:

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

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

$$\rho_i = (B_i - L_{e_i}) / \sum_j F_{ij} B_j.$$




Figure 3: Diagram for inverse reflectometry problems.

The most common way of knowing B_i and L_{e_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 ρ^* and estimating the surface reflectance as

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

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 [51] 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 [9], 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 [3] 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.* [14], assumes 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 [12], 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.* [62] 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 [11] 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 [32] Loscos *et al.* a different approach for reflectance recovery is presented: a set of images from a fixed viewpoint but with controlled, varying illumination (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 [31] and an algorithm for interactive re-lighting was introduced, based on a preprocessing step that reconstructs geometry and creates un-occluded illumination textures (thus taking into account the effect of shadows). Unfortunately, only diffuse surfaces can be considered. The creation of these un-occluded 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 un-occluded 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 [30] 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 does not 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.* [62], 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 ω and ω_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 ω and ω_i only, disregarding spatial variations.

4.3.1. Texture-based BRDF Recovery

In 1991, Ikeuchi and Sato [26], 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 [28] 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.* [44] 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.* [41] 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.* [13] 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.* [34] 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 [33] 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 [45, 43] 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 [36, 39] 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 [36] 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 [39] 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.* [10] 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.* [60] 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 compute the BRDF parameters. This work is an extension of a method [54] 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.* [36, 39], 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 [63] 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 [11]. 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 [33]. 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 [37] described in section 3, Marschner's thesis presents two inverse reflectometry problems, the first one being *Photographic Texture Measurement* [36]. 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 [51] 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 [63], nor does it require highly controlled lighting conditions (generally by careful active positioning of a single source) like Marschner *et al.* [37]. 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 Galalowicz [3] presents the clear advantage over other image-based methods (like Debevec *et al.* [9], Yu *et al.* [62], Loscos *et al.* [32], Fournier *et al.* [14] and Drettakis *et al.* [12]) 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.* [62] do. On the other hand, observe that the method by Yu *et al.* [62] and the one by Boivin and Galalowicz [3] work on

full scenes, while Debevec [9] only estimated material properties on a part of the scene.

- With respect to radiosity-based approaches, it is important to note that in [12] 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.* [32] 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.* [62] 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.* [44] and the one by Poulin and Fournier [42] are the only interactive methods that allow some sort of texture recovery. Instead, Ofek *et al.* [41] 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.* [13] 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.* [34] 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.* [10] and Sato *et al.* [60] for BRDF recovery is that, in the second step in the Sato *et al.* [60] work, the diffuse part is fitted to the diffuse term in a modified Torrance-Sparrow model. Unlike Marschner *et al.* [36, 39], 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 [51] 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

Table 4: Inverse reflectometry papers that use images from different points of view (PoV) as input.

	Several PoV	Single PoV
Local	[54], [60], [41], [28], [63], [13], [10], [34]	[26], [43], [45], [36], [39], [44], [51]
General	[9]	[3]
Radiosity	[62], [32], [31]	[12], [14]

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 [9] [3] [62] [14] [12] [3] [32] [30] [31] 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 [44, 10, 60] 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 [36, 39], which take the information directly from the user-provided photographs, without needing an intermediate optimization process to get the results.
- Instead, [62] 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:
 - 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 [3] 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 [51], 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 [29] 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 [24] 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.* [59] 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 [58] 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 [51] 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 [58], 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 [29], 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 [51] being able to be considered as an extension of Sato *et al.* [59, 58] 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 3).

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 [52], [29], [37, 36], [43], [36], [62].
- *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 [23, 15, 38]. Another approach [40] 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 [54], [60] and [63].

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. ([37], [36] and [51], 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, [23] and [38].

With respect to the inverse reflectometry problems, we see in Table 3 that [9] and [3] 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 [44] presents a direct approach, and together with [36], [39], [14], [12], [32], [41] and [13] are the purely direct approaches among those surveyed. On the other hand, [62], [54] and [60] 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 [62] 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.

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
[52]	Indirect	Lambertian surfaces. Final impl: local illum.	CLS
[40]	Direct	Lambertian surfaces. Uniform spatial and angular subdiv.	Inverse Monte Carlo
[23]	Direct	Lambertian surfaces + Radiosity constant on patches. Truncation of singular values.	MTSVD
[15]	Direct + Finite Volume method for media + Radiosity constant on patches + lambertian surfaces and isotropic media.	Discretization. Truncation of singular values.	TSVD
[38]	Direct + Finite Volume method for media	Finite Elements, P1 Spherical Harmonics Approximation. Radiosity constant on patches + lambertian surfaces and isotropic media.	MTSVD
[37], [36]	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
[12], [32], [14]	Indirect	Diffuse Surfaces with constant radiosity. Known Emitting Patches with known geometry.	Least Squares. [12] and [32] build a hierarchical radiosity system to speed up computations.
[42]	Indirect	Blinn's illumination model.	Analytical Optimization
[47]	Indirect	none	NLCO
[7]	Indirect	none	SA (with possible final local search step)
[58], [59]	Direct	Distant sources, No interreflections, uniform reflectance, camera & known geom	CLS
[57]	Direct	Common features in stereo pair, no interreflections, uniform reflectance, camera & known geom	CLS
[55]	Indirect	None (although implemented as local illum)	Steepest descent
[35]	Indirect	None	Selections dispersion and manual user selection
[18]	Indirect	Known geometry and diffuse reflectances	Levenberg-Marquardt
[51]	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 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.

Work	Type of approach	Approximations	Method used
[43]	Indirect	none	NLCO or NLCLS
[36]	Indirect	BRDF w/ few spatially varying params. Discrete Sampl.	LS BRDF fitting (from photographs).
[36], [39]	Direct	$f_r(\omega, \omega_i) \equiv f_r(\omega, \omega_i)$. Point light sources.	Measurement form photographs.
[62]	Mixed	Low param BRDF & Scene decomp. in regs. w/ const. directional refl. props.	Iterative BRDF fitting; NLLS steps.
[9]	Indirect	Scene subdivided in far & local regions.	Iterative BRDF fitting.
[44]	Direct	Purely Lambertian surfaces.	Color: sampling the reproj. texels in the corresp. imgs.
[63]	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.
[14], [12]	Direct	Diffuse Surfaces. Does not distinguish between colour and shadows.	Reflectance proportional to average radiosity
[32]	Direct	Assumes known exact emittance from surfaces	As [12], but refl. is avrg over imgs. confidence weights.
[31]	Indirect	Diffuse surf	As [12] + blocked light from rad. sol. + heuristic corr. based on illum. on ref. patch.
[41]	Direct	No highlights in most images	Multiresolution textures by recursive reprojection
[10], [60]	Indirect	Assumes spec color = light color. No changes in ind. illum.	Separation of diff & spec parts. Direct recov of BRDF params
[51]	Indirect	First Terms in Sph Harm series	Simplex method for nonlinear params & linear fitting for rest
[3]	Indirect	Hyerarchically tests simple to complex BRDF models	Simple Models: Iterations with image-based error. Complex: Simplex Method.
[26]	Indirect	Constant BRDF over regions (restricted Torrance-Sparrow BRDF)	Iterative LS fitting
[28]	Indirect	Requires Several Photographs	BRDF sample inversion (nohighlight: LS; highlight: separable LS with Reg)
[54]	Mixed	Separable diffuse & specular parts	Separation of diff & spec and fitting of Torrance-Sparrow by Levenberg-Marquardt
[60]	Mixed	Local illum. Convex Objects. Interpolated BRDF params	1st measured. 2nd fitted Torrance-Sparrow
[13]	Direct	Totally known lighting	Measurements by Spectrophotometer
[34]	Indirect	Full knowledge of Lights, camera & object geom	1st BRDF sample classification, 2nd Levenberg-Marquardt method for each samples cluster

Table 7: Summary of the mathematical aspects of the inverse lighting approaches presented: *inverse lighting*, $SVD = \text{singular value decomposition}$, $MTSVD = \text{modified truncated SVD}$, $MC = \text{Monte Carlo}$, $PDE = \text{partial differential equation}$.

	Constraints imposed	Initial guess / Initial approach.	Objective function (?) / direct inversion method	Tests Performed
[52] (CLS)	All weights ≥ 0 .	Any value	Minimize: $\ \sum_i \omega_i \phi_i - g^{e/p}\ $ where ϕ_i is the effect of the i -th light source.	Complex env. with 19000 polys, 27000 vert. and 12 (14) light sources.
[40]	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.
[23]	# 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.
[15]	# 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.
[38]	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.
[37, 36]	Positive coefficients.	Any value	Linear combination of basis images to match a given photograph.	Relighting: Photos + 3D model to modify a given photo.
[14, 12], [31]	Σ emittances fits estim total intensity	Radiosities from images	Fitting of the sources power	Acquired real video scenes.
[42]	None	Any value	$(N \cdot H)^n$, for dir L and glossiness n . Use shadow volumes for final position.	Simple primitives. e.g.: cone under triangular light.
[47] (NLCO) Lights from shad sketch	-in cone defined by each point -same side of normal at 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.	Maximize: $\sum (p_i - P_i)^2 = \text{distance from every sketched points } p_i \text{ to light position } P_i$.	Several primitives: 0.003 sec/Sketched point - interactive speed.
-umbra/ penumbra	-Umbra: $\forall \text{ Light Vertex (LV), Sketched Point (SP), LV} \in \text{cone SP}$ -Penumbra: $\forall \text{ LV, } \exists \text{ SP, LV} \in \text{cone SP}$	Just above an ellipsoid E centered on occluder along dir: from mass center of the sketched points to E center	Maximize: $\sum (p_i - P_i)^2 = \text{distance from every sketched points}$	Several primitives: 0.01 sec/Sketched point when source is large WRT cones.
-highlights	-within cone defined by each sketched point. -same side of normal at sketched point. -right side of half-cone WRT cone axis.	Any value	Minimize: Value of Point Light within the implicit cone equation.	0.015 sec/sketched point.
[7]	used-defined (added as penalties in the objective function)	User-provided	User-provided by a script.	1) room, luminarie & table with homogeneous desired illum. 2) office (10550 obj), goal: reduce glare, shadow boundaries & redundancy.
[58, 59]	Positive coeffs.	Lambertian coef. estim. for shadow surface	Linear System of Eqs: sources vs. pixels in photographs	simple polyhedra on tablecloth.
[57]	None	Any Value	Find bright features in each omni-image and correlate. Map remaining as image.	Indoor & outdoor real scenes
[55]	none	Light dirs = centroid of normals of viewed surfaces. Intensities = average desired to viewed intensities ratio [53]	Perception-based image quality metric (extent in communication scene info)	isolated complex geometrical objects
[35]	None	User-defined initial lighting	Manual user optimization	Complex scenes
[51]	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
[18]	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 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.

Work	Constraints imposed	Initial guess / Initial approx.	Objective function (?) / direct inversion method	Tests Performed
[43] · with $n \leq 3$ points	Values ≥ 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
[43] · with $n \geq 3$ points	Values ≥ 0 . Weight points with most important contribution.	Any value	Distance between BRDF and user-provided samples.	Phong-Blinn BRDF.
[36]	None	Any value	Weighted Squared error between BRDF and measured samples.	Synthetic sphere w/ textures for calib. Squash. Rock.
[36, 39]	None	Any value	Direct measurement of ω, ω_i , obtaining $f_r(\omega, \omega_i)$	Gray cyl for verif. Squash. Human skin. (300000 sample/channel)
[62]	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.
[9]	None	Initial guess (user)	Diff. between rendered local scene and photos.	Real complex scenes w/ simple synthetic objects.
[44]	None	Any value	All polys textures fitted to one another. corresp. colors combined WRT set of crit. to form a unique text.	320 × 400 textures for book from real photographs.
[63]	None	Any value	Diff: measure. Landscape spec: LS diff pseudo-BRDF & recorded rad. Sun spec: robust param recovery [49].	Relighting of existing, real tower.
[14], [12]	None	Any value	Avg reflect proportional to radiosity	Aquired video image scene
[31]	None	Any Value	1st: eval. as i [12]. 2nd: Add blocked light from rad. sol. 3rd: Heur correc: find unoccluded ref. patch for color ref.	Aquired real video Image scene
[41]	None	Any Value	Multiresolution quadtree textures recovered by recursive reprojection.	Pictures & simple textured objects
[10]	None	Spec color = source color: Diff color = median RGB measurements.	Direct for spec. Lambertian fit.	Change viewpoint, illum & matching real-world illum.
[60]	None	Spec color = source color: Diff color = separated & fitted to Torrance-Sparrow.	Direct for spec. Lambertian fit.	Cofee cup.
[51]	None	Any value	First Terms in Sph Harm series: BRDF = Diffuse + slow-varying terms + fast-varying specular lobe.	Real-world models.
[3]	None	Diffuse Surfaces used first. Then more complex BRDFs.	Ratio of average of object image radiances to object rendered image.	Office Scene
[26]	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
[28]	Diffuse coef smooth on range-deficient(rd) regions, specular over non-spec & rd, roughness everywhere	Any	Linear LS error in BRDF values	Textured spheres
[54]	None	Any	not specified (fit of specular comp)	Complex objects
[13]	None	Any	Measurement of BRDF samples with diff. comb. of viewing & lighting	Database of materials
[34]	None	Params from avg BRDF	RMSE b/ sampled & predicted values	texture recovery from complex objects

Of the reviewed papers that deal with inverse emittance or combined problems, only [37, 36] and [29] 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 [29] 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 [51] across a scene, and performing correlations of the obtained information (a generalization of the work by Sato *et al.* [59]) 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 [21]. 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.* [6].
- 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 CICYT and 2001SGR-0296 of Generalitat de Catalunya.

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.

References

1. James Arvo. Analytic Methods for Simulated Light Transport, Ph.D. thesis, 1995.
2. James Arvo. The Role of Functional Analysis in Global Illumination. In P. M. Hanrahan and W. Purgathofer (eds), *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, Springer-Verlag, New York, pp. 115–126. 1995.
3. 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)*, pp. 107–116. 2001.
4. James F. Blinn. Models of light reflection for computer synthesized pictures. In *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '77 Proceedings)*, pp. 192–198. 1977.
5. P. Bui-T. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.
6. Antonio Cardoso Costa, A. Sousa and F. Ferreira. Design de iluminação. In *8th Portuguese workshop on Computer Graphics*. Coimbra, Portugal, 1998 in Portuguese
7. A. Costa, A. Sousa and F. Ferreira. Lighting design: A goal based approach using optimization. In D. Lichinski and G. Ward Larson (eds), *Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering)*, Springer-Verlag, New York, NY, pp. 317–328. 1999.

8. Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Boston, MA, 1993.
9. 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)*, pp. 189–198. 1998.
10. P. 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)*, pp. 145–156. 2000.
11. 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)*, pp. 369–378. 1997.
12. George Drettakis, Luc Robert and Sylvain Bugnoux. Interactive common illumination for computer augmented reality. In Julie Dorsey and Philipp Slussalek (eds), *Rendering Techniques '97 (Proc. Eighth Eurographics Workshop on Rendering)*, Springer Wien, New York, NY, pp. 45–56. 1997.
13. Kristin 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.
14. Alain Fournier, Atjeng S. Gunawan and Chris Romanzin. Common Illumination Between Real and Computer Generated Scenes. In *Proceedings of Graphics Interface '93*, Morgan Kaufmann, San Francisco, CA, pp. 254–262. 1993.
15. F. 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 (eds), *ASME HTD*, vol. 361-5, IMECE, pp. 145–151. 1998.
16. 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.
17. Andrew S. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Francisco, CA, 1995.
18. 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. IF-SIC/IRISA, 2000 in french
19. J. Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. In *Bull. Univ. Of Princeton*, pp. 49–52. 1902 In French
20. Berthold K. P. Horn and Michael J. Brooks. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
21. Paul Heckbert. Adaptive Radiosity Textures for Bidirectional Ray Tracing. In *Computer Graphics (ACM SIGGRAPH '90 Proceedings)*, vol. 24, pp. 145–154. 1990.
22. P. Hanrahan and P. Haeberli. Direct wysiwyg painting and texturing on 3d shapes. *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '90 Proceedings)*, 24:215–223, 1990.
23. V. 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.
24. Pat Hanrahan, David Salzman and Larry Auplerle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics (ACM SIGGRAPH '91 Proceedings)*, vol. 25, pp. 197–206. 1991.
25. 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.
26. 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, 1991.
27. Michael Kass. Inverse problems in computer graphics. In Nadia Magnenat Thalmann and Daniel Thalmann (eds), *Creating and Animating the Virtual World (Computer Animation*

- '92, *Fourth Workshop on Computer Animation*). Springer-Verlag, Tokyo, 1992.
28. G. Kay and T. Caelli. Inverting an illumination model for range and intensity maps. *Image Understanding*, 59:183–201, 1994.
 29. 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)*, pp. 147–154. 1993.
 30. C. Loscos and G. Drettakis. Low-cost photometric calibration for interactive relighting. In *First French-British International Workshop on Virtual Reality*. Brest, France, 2000.
 31. Celine Loscos, George Drettakis and Luc Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, 6(4):289–305, 2000.
 32. C. 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)*, Springer Wien, New York, NY, pp. 329–340. 1999.
 33. 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)*, vol. 31, pp. 117–126. 1997.
 34. 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, 2001.
 35. J. Marks, 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)*, pp. 389–400. 1997.
 36. Stephen R. Marschner. Inverse Rendering in Computer Graphics, PhD thesis, Program of Computer Graphics, Cornell University, Ithaca, NY, 1998.
 37. Stephen R. Marschner and Donald P. Greenberg. Inverse lighting for photography. In *Proceedings of the IS & T/SID Fifth Color Imaging Conference*, Society for Imaging Science and Technology, Scottsdale, AZ, pp. 262–265. 1997.
 38. J. C. 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, 1997.
 39. Stephen 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 (eds), *Rendering Techniques '99 (Proceedings of the 10th Eurographics Workshop on Rendering)*, Springer-Verlag, New York, NY, pp. 131–144. 1999.
 40. 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*, vol. 3, ASME, pp. 243–250. 1995.
 41. E. Ofek, E. Shilat, A. Rappoport and M. Werman. Multiresolution textures from image sequences. *IEEE Computer Graphics and Applications*, 17(2):18–29, 1997.
 42. Pierre Poulin and Alain Fournier. Lights from highlights and shadows. *Computer Graphics*, 25(2):31–38, 1992.
 43. P. Poulin and A. Fournier. Painting surface characteristics. In P. M. Hanrahan and W. Purghofer (eds), *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, Springer-Verlag, New York, NY, pp. 119–129. 1995.
 44. Pierre Poulin, Mathieu Ouimet and Marie-Claude Frasson. Interactively modeling with photogrammetry. In *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)*, Springer Wien, New York, NY, pp. 93–104. 1998.
 45. Pierre Poulin. Shading and Inverse Shading from Direct Illumination, PhD thesis, Department of Computer Science, University of British Columbia, 1993 Available from <http://www.cs.ubc.ca/nest/imager/th.html>

46. G. Patow and X. Pueyo. A Survey on Reflector Design Problems. Technical Report TR-IIIA 00-16-RR, Universitat de Girona, 2000.
47. Pierre Poulin, Karim Ratib and Marco Jacques. Sketching shadows and highlights to position lights. In *Proceedings of Computer Graphics International 97*, IEEE Computer Society, pp. 56–63. 1997.
48. W. H. Press and S.A. Teukolsky. Simulated annealing optimization over continuous spaces. *Computers in Physics*, 426–429, 1991.
49. 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.
50. 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)*, 2448–2459, 2001.
51. Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001)*, pp. 117–128. 2001.
52. Chris 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)*, pp. 143–146. 1993.
53. 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/l设计>
54. Y. Sato and K. Ikeuchi. Reflectance analysis for 3d computer graphics model generation. *CVGIP Graphical Models and Image Processing*, 58(5):437–451, 1996.
55. R. Shacked and D. Lischinski. Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum (Eurographics 2001)*, 20(3):215–227, 2001.
56. Francois Sillion and Claude Puech. *Radiosity and Global Illumination*. Morgan Kaufmann, San Francisco, CA, 1994.
57. I. 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.
58. I. 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*, pp. 875–882. 1999.
59. I. Sato, Y. Sato and K. Ikeuchi. Illumination distribution from shadows. In *Proceedings of IEEE Conf. on Computer Vision and Pattern Recognition'99*, pp. 306–312. 1999.
60. Y. 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)*, pp. 379–387. 1997.
61. Alan Watt and Mark Watt. *Advanced Rendering and Animation Techniques: Theory and Practice*. Addison-Wesley, 1992.
62. Yizhou 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)*, pp. 215–224. 1999.
63. Yizhou Yu and Jitendra Malik. Recovering photometric properties of architectural scenes from photographs. In Michael Cohen (ed), *Computer Graphics Proceedings, Annual Conference Series (ACM SIGGRAPH '98 Proceedings)*, Addison Wesley, pp. 207–218. 1998.