# A Survey of Inverse Surface Design From Light Transport Behavior Specification

Gustavo Patow and Xavier Pueyo

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

**Abstract**

*Inverse surface design problems from light transport behavior specification usually represent extremely complex and costly processes, but their importance is well known. In particular, they are very interesting for lighting and luminaire design, in which it is usually difficult to test design decisions on a physical model in order to avoid costly mistakes. In this survey we present the main ideas behind these kinds of problems, characterize them, and summarize existing work in the area, revealing problems that remain open and possible areas of further research.*

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

## 1. Introduction

Inverse problems are emerging as an important research topic for the graphics community due to their importance in a wide range of application fields including lighting engineering and lighting design. These problems are usually of extreme complexity. Although progress in rendering to date has mainly focused on improving the accuracy of the physical simulation of light transport and developing algorithms with better performance, some attention has been given to the problems related to inverse analysis, leading to interesting results. Other interesting results come from Applied Mathematics as well as Optical and Thermal Engineering.

Traditional direct problems in rendering involve computing the radiance distribution in an environment that is completely known a priori (geometry and materials). Inverse Rendering Problems refer to all the problems in which, as opposed to what happens in traditional direct rendering problems, several aspects of the scene are unknown. Some aspects that might be unknown are light source positions and/or their orientations, luminaries' emittances, surfaces reflectances and the shape or position and orientation of the surfaces (and other reflective/refractive elements) of the luminaries in the scene. The Rendering Inverse Problems are not well-posed: the solution does not depend continuously on the data, which means that small errors in measurements (input data) may cause large errors in the solution (see [HMH95]).

As proposed by Marschner [Mar98] and described in section 2.1, Inverse Rendering Problems can be classified into three classes: *Inverse Emittance*, *Inverse Reflectometry* and *Inverse Geometry*. In the present paper we study the contributions for Inverse Surface Design and characterize them in the class of Inverse Geometry Problems. It is important to note that, although the problem is often generically called *Inverse Reflector Design*, it encompasses the design of both reflective and refractive elements in an optical system. For a survey of *Inverse Emittance* and *Inverse Reflectometry* the reader can refer to [PP03].

The problem of Inverse Surface Design from light transport behavior specification, as focused in this survey, can be stated as follows: given a light source (bulb) with a known light intensity distribution, a surface should be constructed in such a way that a prescribed illumination intensity is obtained on a prescribed region in space, after reflection/refraction at the surface, see Figure 1. This prescribed distribution can be given either as a near- or a far-field distribution: the first given is directional and spatial distribution (although generally it is defined only as irradiance on a certain plane), and the second given is only in purely directional distribution terms. This later case can be thought of as a limiting case when the plane to be illuminated moves "infinitely" *far* away from the source. In general, radiance distributions are not defined in the continuum, but in some set of directions (far-field) or points in space (near-field), although they are later extended to the continuum by interpolation.

There are works found in the literature that present important problems and results, which have so many similarities that can generically be enclosed within the problem studied here. Among them we find works on the measurement of the human cornea (which acts as a reflector) and luminaire design (street lamps, car headlights, etc.). In addition, it is important to review the works on Inverse Design of Refractors, where the surface to be found is refractive instead of reflective. For this case, the design of progressive lenses represents a very important application.

The paper is organized as follows: In Section 2 the theoretical background is given, related problems are briefly commented, and a classification based on the rendering equation is presented. Next, in Section 3, analytical methods for surface design are explained, followed in Section 4 by the numerical solutions to the problem. In both cases, the methods presented are organized according to the treatment of the light transport problem they use, and then according to the type of surface used. Finally, in Section 5 the conclusions and open lines of research are presented.

## 2. Inverse Geometry Problems: Surface Design

### 2.1. Theoretical Background

The behavior of light transport is characterized by the properties of the particles (photons) when traversing the environment. The most fundamental quantity in global illumination is radiance $L(\mathbf{r}, \omega)$ which is defined as the power radiated at a given point $\mathbf{r}$ in a given direction $\omega$ per unit of projected area perpendicular to that direction per unit solid angle for a given frequency ($Watt\, m^{-2}\, sr^{-1}$).

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

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

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

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

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

which accounts for the scattering of incident radiant energy. The measures $d\omega_i$ and $d\mu(\omega_i)$ are related by $d\mu(\omega_i) = \cos\theta d\omega_i$. Here $h$ is a field radiance function, that corresponds to all incident light. Next, we can define the *field radiance operator* $\hat{G}$, that transforms an exiting light distribution into the incident light distribution that results from surfaces illuminating one another:

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

where $\nu(\mathbf{r}, \omega)$ is the *visible surface function* and is defined [Arv95b] as $\nu(\mathbf{r}, \omega) \equiv inf\{x > 0 : \mathbf{r} + x\omega \in \text{Surfaces in the environment}, \infty\}$, and $\mathbf{p}(\mathbf{r}; \omega) \equiv \mathbf{r} + \nu(\mathbf{r}, \omega)\omega$ the *ray casting function.*

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

$$L = L_e + \hat{K}\hat{G}L \tag{2}$$

Following the outlines in Stephen Marschner's PhD. thesis introduction [Mar98], we can classify the different papers on inverse lighting problems according to which quantities of the above equation are unknown:

- **If** $L_e$ is unknown, and $\hat{K}$, $\hat{G}$ and $L$ or part of it, are known, we have a problem of *inverse lighting.*
- **If** $\hat{K}$ is unknown, and $\hat{G}$, $L_e$ and part of $L$ are known, we must solve for information about $\hat{K}$. This problem can, in general, be called *inverse reflectometry*, and a particular case is the one called *image-based reflectometry* in [Mar98], where images are used as input to the information about $L$.

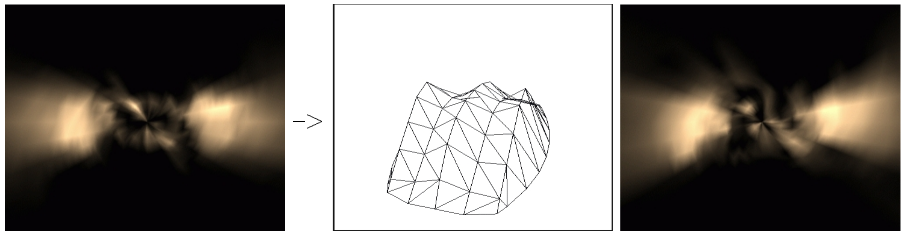

**Figure 1:** *Example of an Inverse Surface Design Problem: The desired light distribution is shown on the left, and the algorithm should produce the surface shown in the middle, which generates (in this case by reflection) the light distribution shown on the right.*

- **If** $\hat{G}$ is unknown, we have an *inverse geometry* problem, which encompasses the long-studied computer vision *shape from shading* problem, the *reflector design* inverse problem and the *recovering objects from photographs* problem. See Figure 1.

On the other hand, direct problems are those which, given known values for $L_e$, $\hat{K}$ and $\hat{G}$, solve for $L$.

### 2.2. Related Problems

In this survey we will review the work done on the family of problems enclosed in the last item of the previous section, when elements of the geometry are unknown, i.e. an *inverse geometry* problem, we will deal with the Inverse Surface Design problem. Other closely related problems are:

- The *shape from shading* problem, although similar in nature to the problem presented here, it is fundamentally different in that at each sampling position on the image (at the observer location) we can safely assume that only a small subset of points on the surface project onto it, and thus only a small subset of the surface contributes to its value (often, it is considered that only a single point contributes). Instead, in the problem treated here, there is usually a very important influence from all the points on the surface and the sampling positions (both in the far- and near-field problems), giving us a strong correlation between surface control points and sampling positions. A review of the *shape from shading* problem is beyond the scope of the present work, but the interested reader is referred to [HB89] or [SL97] for a description of the field.
- Another field closely related to the one reviewed here is *Optical Design*, which is the process of describing the refracting or reflecting elements in an optical system so that they meet a set of performance specifications. Typically, the performance specifications concern the imaging characteristics of the system, such as the resolution, magnification, numerical aperture, and field of view. Other requirements may include tolerances, size, weight, and cost. The result of an optical design project is typically a prescription or database that lists the materials and shapes of the optical elements required. This is different from the problem treated in our survey, baceause in Optical Design problems the generic shape of the surfaces is generally known in advance, and can only be modified by global parameters such as lens thickness or radii, but not its *shape*. For more information on the subject, please refer to [KM00] [Soi02] [Sha97] and [Smi90]. To the best of our knowledge, most of the existing commercial software belong to this field, using either the design-direct simulation-restart principle [Jen01] [Org02a], or based on a local optimization [Org02b] [Cor] [Eng], or a global optimization approach [Vas98] [SO]. A comprehensive list of commercial software can be found in [Opt].
- Although the *Design of Reflector Antennas* is closely related to the problem treated here, most of the works in the field have unique features that make it quite difficult to generalize their results to incoherent radiation as produced by most light sources, like coherent radiation, which implies taking into account aspects such as phase cancellation and other interference effects [Kil00].
- It is also important to note the difference between computations of the shape of a source (see [PP03]), and the problem of surface design. The latter computes the shape of a reflecting or transmitting surface from light transport behavior, while the former computes the shape of the emitting surface itself (the bulb).

### 2.3. Inverse Surface Design Problems

We are interested in constructing of reflective/refractive surface shapes from prescribed optical properties of the luminaries and/or bulbs (far-field or near-field radiance distributions) and geometrical constraints. See Figure 3. From a mathematical

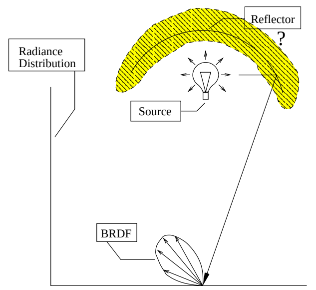


**Figure 2:** *Inverse Surface Design Problem.*

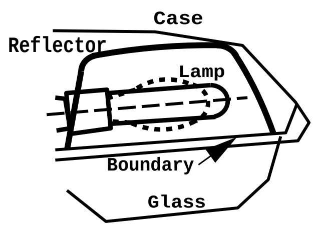


**Figure 3:** *An optical set for reflector design.*

point of view, the problem to solve belongs to the category of non-linear inverse problems. In the case of manufacturing design problems, restrictions on the shape imposed by the needs of the manufacturing process (Figure 3) should also be taken into account, which generally results in adding constraints to the problem.

The basic problem is building a surface shape from the description of the optical properties of the emitting light bulb and the desired light distribution the optical set (surface plus bulb) must provide. A user-provided tolerance is also given. Comparison of outgoing light distributions (the ideal desired one and the computed one) requires establishing a distance measure, which in turn will provide a logical target function for any optimization procedure: minimize the distance between the light distribution given by the current calculation surface (in general, a reflective surface) and the ideal, user-provided desired light distribution. The theoretical formulation of this problem leads to a non-linear partial differential equation of the Monge-Ampère type, as described in the literature [WN75] [Wes83] [EN91].

In general, we can say that all the works are based on the same scheme, which is illustrated in Figure 4: Finding the minimum of the function that describes the error in the outgoing light distribution for the optical set in the space of possible surfaces, with respect to the prescribed, user-given light distribution. As is known, optimization procedures require an *iterative* procedure that repeatedly evaluates the objective function and determines the minimum from those measurements. In our case, the function to be evaluated always consists of two parts: the simulation of the light propagation from the light bulb to the registration area (as mentioned above, far-field or near-field), and finally an evaluation of the distance between the calculated outgoing light distribution and the user-provided desired distribution. The initial surface for the optimization must be, in general, manually provided.

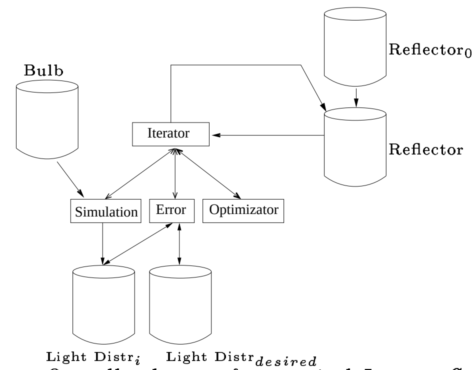


**Figure 4:** *Overall scheme of numerical Inverse Surface Design methods.*

In general, we can say that the works reviewed here share the same elements which may be used to build an initial characterization (See Table 1). In particular, the different works studied can be characterized according to the treatment they give to the light transport from the light bulb to the evaluation region, the shape representation chosen, the optimization method, the type of BRDF used for the surface, the representation of the light emitter, the distance measure used and the nature of the problem focused on (far-field vs. near-field problems)

- **The function mapping light from the bulb to the evaluation space, usually represented by the light propagation algorithm**. Basically, the purpose of this step is to compute the light transport from the light bulb, reflecting or refracting on the surface, and arriving to the final, desired registration region. Although most of the studied algorithms use a local illumination-based scheme (light from the source bounces/refracts once on the surface and "goes" directly to the receiving region), inter-reflections are an important part of the simulation and should be considered. In general, we can observe that there are three main approaches to this particular aspect:
  - *"One-to-one incoming-to-outgoing rays" local illumination*: This is a severe approximation that

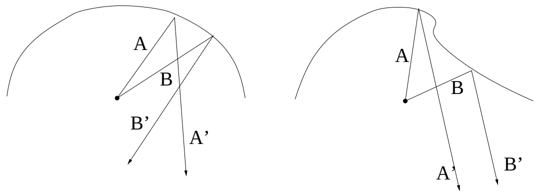


**Figure 5:** *The one-to-one assumption: The left reflector satisfies the condition, since each ray that arrives on the reflector bounces in a* different*, primed direction, but the reflector on the right does not since rays labeled A and B bounce in the* same *direction (A' = B').*

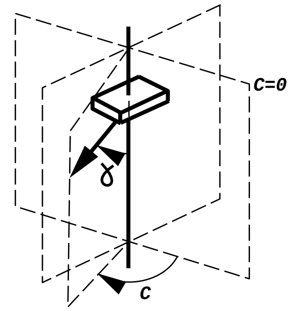


**Figure 6:** *The* $C - \gamma$ *coordinate system.*

assumes there is a one-to-one relation between incident rays and outgoing rays: no two rays can be reflected in the same outgoing direction (see Figure 5). This is an approximation mostly used for analytical works as it greatly simplifies the calculations needed for the convergence demonstrations.

  – *Local illumination*: An illumination scheme in which rays from the light bulb bounce once, accordingly to the surface BRDF, towards the measurement region. This implies that no inter-reflections are computed, and in general, visibility issues (generated by the bulb and supporting elements or by the reflector itself) are disregarded.
  – *Global illumination*: In this computation scheme, irradiance is followed as it bounces in its way out of the optical system. In general, this requires both computing multi-reflections as well as taking into account the visibility problem, as rays might be blocked by the surface itself, the bulb and/or supporting elements. Now, depending on the surface BRDF, many options can be considered, from radiosity solutions to a full treatment of Equation 1.

- **The shape definition for the surface, and if applicable, the restrictions imposed on the space of possible achievable shapes**. Different articles use different kinds of models for the surface representation, ranging from simple geometric primitives (for example, pieces of quadrics) or combinations of them, up to polygonal-based or NURBS-based definitions. Different choices have a direct impact on the number and type of the optimizable degrees of freedom, as well as modifying the light propagation algorithm to some extent.
- **The surface material used**. In most reviewed works it is a perfect specular BRDF. Other models could be used, but with a significant impact on accuracy as the BRDF goes diffuse. This clearly influences the light propagation method.
- **Representation of the light emitter (bulb)**. The light emitter can be represented either as a point source or as an emitting surface (closer to reality). It can be an isotropic source that emits equally in every direction (in general an approximation that is too simplistic) or anisotropic, which is the common choice. Once again, this mainly affects the light propagation algorithm through the system.
- **The definition of the distance between the desired outgoing light distribution and the corresponding one for a given surface**. Light distributions are usually represented in some sort of discrete set of directions and positions, so all works use the $\ell^2$ norm to define the distance between radiance distributions. See [CW93] Section 10.1.4, General Luminaries. An example of such a distribution can be found in Figure 7.
- **The optimization method used**. Here is where the papers discussed differ most, because the non-linear inverse problem faced can be solved in different ways: by a global approach or by starting the algorithms close enough to the desired solution in such a way that a local optimization algorithm would lead to the correct solution. Other approaches impose restrictions on the achievable shape in order to guarantee local convergence.
- **The problem faced can be in its far-field or near-field form, although some works deal with both kinds of approaches**. As mentioned above, this affects the way the light propagation is evaluated and, of course, the error measure being used. Although most of the papers found in the literature make a clear distinction between both problems, the resulting algorithms/methods can, in general, be easily adapted to solve any one of the two types, thus blurring the importance of the distinction from a practical point of view (See Tables 3 and 5).

  In general, when dealing with the *far-field* problem,

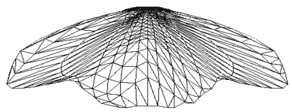

**Figure 7:** *An example of a real outgoing radiance distribution used in industry.*

| | Theoretical | Numerical |
|---|---|---|
| Local Illum | [WN75] [Wes83] [BW78] [Oli89] [Wan96] [Oli02] [KO97] [KOvT98] [Oli03] | [CKO99] [KO03] [EN91] [Neu94] [Neu97] [KO98] [EN91] [KN96] [Neu97] [HBKM95a] [Hal96] [HBKM95b] [HBKM96] [Hal96] [LSS98] |
| Global Illum | | [DCC99b] [DCC01] [DCC99a] [PPV04] |

**Table 1:** *Initial classification of the studied papers.*

it is a very common choice to use a spherical coordinates system as a discrete representation for the outgoing radiance distributions, like the well known $C-\gamma$ system[CM97], that represents a standard in the lighting engineering industry (see Figure 6).

In this survey we will build our classification around three main aspects from among those mentioned above: the analytical vs. numerical nature of the papers, the light propagating algorithm/method used to compute the outgoing radiance distributions and the shape definition used for the different computations.

## 3. Analytical Methods for Inverse Surface Design

This section deals with the theoretical analysis works done on the problem of Inverse Surface Design. Basically, they formulate the problem in precise mathematic terms using differential geometry, although adding important constraints on the formulation in order to make the problem theoretically tractable. For example, all of the reviewed works assume that the surface is perfectly specular, and many assume that

| | Rotational Symm Approx | Non-rotational Symm Approx |
|---|---|---|
| 1-1 in-to-out rays | [WN75] [Wes83] | |
| Local Illumination | [BW78] [Oli89] | [Wan96] [KO97] [KOvT98] [Oli02] [Oli03] |

**Table 2:** *Sub-classification of analytical papers on Inverse Surface Design.*

| | near-field | far-field |
|---|---|---|
| 1-1 in-to-out Local Illumination | | [WN75] [Wes83] |
| Local Illumination | [BW78] [Oli89] [KO97] [KOvT98] [Oli02] [Oli03] | [Wan96] |

**Table 3:** *Alternative classification of analytical papers on Inverse Surface Design.*

there is a one-to-one relation between incident rays and outgoing directions, as mentioned above.

Below, we present the classification of theoretical papers according to the type of treatment given to the light propagation step: either they use a 1-to-1 correspondence between incoming and outgoing rays or a local illumination approach, vs. the type of surface they consider in their approaches. A summary can be found in Table 2.

However, it is also possible to present an alternative sub-classification of those works with respect to a far-field treatment vs. a near-field approach, as shown in Table 3, a distinction often made by the authors of the reviewed works.

### 3.1. "One-to-one incoming-to-outgoing rays" local illumination

In this section we will review the contributions that treat the 1-to-1 approximation for local illumination as described above, in Section 2.1. This means that no two outgoing rays have the same orientation, a constraint usually referred to as the "one-to-one correspondence between incoming and outgoing rays" approximation (See Figure 5). To our understanding, this is too restrictive in practice. In general, these papers are based on the requirement of a rotational symmetric approximation for the studied surfaces, thus greatly reducing the space of the studied surfaces.

One of the first papers that dealt with this problem in its far-field approach is by Wescott et al. [WN75]. The authors presented an analytical formulation of the problem under the assumptions that the reflector surface is perfectly specular and rotationally symmetric. The authors studied the solutions for the case of fields with even azimuthal symmetry by using a spherical coordinate representation for them, and they also gave proof of existence and uniqueness.

Numerical solutions for these rotationally symmetric-restricted surfaces were also studied by Wescott [Wes83]. Both works, [WN75] and [Wes83], are focused on the design of reflector antennas in the particular case where incoherent light is reflected.

### 3.2. Local Illumination

The papers reviewed here treat the local illumination problem where rays are fired from the source, bounced once on the surface and reach their destination without any visibility or inter-reflection calculations. In addition, all of them use a perfect specular surface for their computations.

#### 3.2.1. Rotational Symmetric Surfaces

As mentioned above, many works impose the restriction that the surfaces used must be rotationally symmetric, thus reducing demonstrations to a simpler 2D problem.

In [BW78], Brickell et al. study the problem for the near-field distribution. In this case the distribution was considered on a flat object, arriving to a differential equation in complex form of Monge-Ampère type.

Oliker, in 1989 [Oli89] [Oli03], reformulated the same problem for rotationally symmetric reflectors in simpler differential geometry terms and without resorting to using complex structures. This resulted in a more general expression valid for the case of curved objects being illuminated in a prescribed way (near-field problem). He also proved the existence and uniqueness of the rotationally symmetric solution for the radially symmetric case.

#### 3.2.2. Non-Rotational Symmetric Surfaces

All the analytical works studied here are able to generalize the sort of surfaces they use for their studies to be non-rotationally symmetric, thus showing a significantimprovement in the generality of the solutions found.

Following almost the same assumptions as Wescott [WN75] (i.e. perfect specular surfaces) Xu-Jia Wang [Wan96] studied the existence, uniqueness and smoothness of the solution for the general problem in the far-field approximation. His presentation was based on a differential geometry formulation of the problem, which resulted in a clearer expression to work with. The author also showed that the regularity of the solutions fails in even the simplest cases.

In [KO97], Kochengin and Oliker considered the general problem with near-field scattering data without *a priori* assumptions regarding any rotational symmetry. Thus, they formulated the problems in differential geometry terms and established the existence of a weak solution. See also [Oli02] and [Oli03].

In another work, Kochengin, Oliker and von Tempski [KOvT98] studied the problem of finding a convex surface R which refracts a given anisotropic bundle of rays from a source in such a way that the refracted rays are incident on a specified set of points in space and produce a specified intensity distribution there. The refracting surface was found by taking the boundary of families of intersecting hyperboloids, each characterized by its polar radius. The light source is placed at the common focus of all hyperboloids, it is subdivided and rays are counted if they arrive at an objective point or direction [Koc].

## 4. Numerical Methods for Inverse Surface Design

The works reviewed in this section deal with numerical solutions of the proposed problem. We can sub-classify them according to how they treat the light propagation problem, showing results with both local-illumination and global-illumination algorithms. Then, we will characterize them based on the type of representation chosen for the surface (See table 4).

An alternative classification can be found in Table 5, according to the type of problem faced (far-field or near-field) vs. the type of light propagation method used, a classification often made in the literature.

### 4.1. Local Illumination

As stated previously, the works presented in this section share the common approximation of dealing with a simplified form of local illumination light propagation algorithms, where no inter-reflections or visibility issues are considered. They can be further classified according to the sort of representation chosen for the surface: the first set uses the intersection of simple primitives to define a convex surface, while the second one defines the surface with a spline-based representation.

| | Local Illumination | Global Illumination |
|---|---|---|
| Convex Surface | [KO98] [Oli03] [CKO99] [KO03] | |
| Splines | [EN91] [KN96] [Neu94] [Neu97] [HBKM95a] [Hal96] [HBKM95b] [HBKM96] [LSS98] | [DCC99b] [DCC99a] [DCC01] |
| Polygonal Surface | | [PPV04] |

**Table 4:** *Further classification of numerical papers on Inverse Surface Design.*

| | near-field | far-field |
|---|---|---|
| Local Illumination | [KO98] [Oli03] [EN91] [KN96] [Neu97] [HBKM95a] [Hal96] [HBKM95b] [HBKM96] [LSS98] | [CKO99] [KO03] [EN91] [Neu94] [Neu97] |
| Global Illumination | [DCC99b] [DCC99a] [DCC01] | [PPV04] |

**Table 5:** *Alternative classification of numerical papers on Inverse Surface Design.*

#### 4.1.1. Convex Surfaces

In general, we can say that this section includes the works where a convex surface resulting from the intersection of a set of primitives is found. As we will see, these primitives are usually paraboloids, ellipsoids or hyperboloids depending on the type of problem faced: reflective far-field or near-field problems, or refractive ones.

In [KO98] Kochengin and Oliker continued their previous theoretical work [KO97] by presenting a numerical solution for a discrete version of the problem based on the demonstrations in [KO97]. By using the geometric optics property that any ray starting at one focal point of an ellipsoid arrives at the other after one bounce, they build a reflector as the convex body formed by the boundary of the intersection of ellipsoids sharing one foci and with their other foci on a point of the surface where a prescribed intensity is

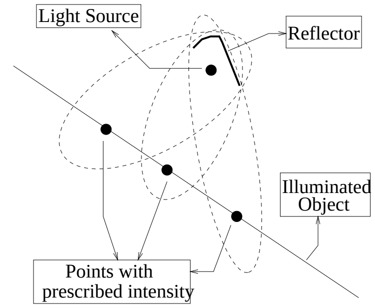


**Figure 8:** *Reflector construction as the boundary of the intersection of confocal ellipsoids [KO98]. Notice that the other foci of each ellipsoid lies on a point with prescribed intensity on the illuminated object. Compare with Figure 9.*

specified (Figure 8). The light source is placed at the common focal point, which results in a light propagation scheme that maps families of rays that impinge on an ellipsoidal piece onto a given destination foci. There are as many ellipsoids as points on the prescribed intensity distribution. As each ellipsoid can be described by its polar radius, the algorithm presents an iterative scheme that minimizes the difference between light arriving at each ellipsoid foci and the light prescribed there. In each iteration a new reflector is built from the previous one until there is no change from one reflector to the next one. Each reflector is built by decreasing each polar radius in turn from its previous value until the light difference for that prescribed point is below a certain threshold. If a polar radius already satisfies that condition, the old value is used instead. The authors show that the algorithm presents a time complexity of $O(K^4 lnK)$, with $K$ as the number of points with prescribed illumination on the target object, and of course, the number of ellipsoids to work with. A detailed study of the mathematical properties of such a resulting surface can be found in [Oli02], [Oli03].

For the discrete far-field problem, Caffarelli et al. [CKO99] present a very similar approach to the one presented by Kochengin and Oliker [KO98] for the near-field problem, but this time based on using paraboloids [Oli02] [Oli03]. They exploit the optical property of perfect paraboloidal reflectors that the light which starts at the focal point leaves the reflector in a direction parallel to its axis. Therefore, by having one paraboloid for each prescribed direction in the far-field region, they can build the final reflector as the boundary surface of the intersection of all confocal paraboloids (Figure 9). As in [KO98], the light source is placed at the focal point, thus mapping families of light rays that arrive at a paraboloidal

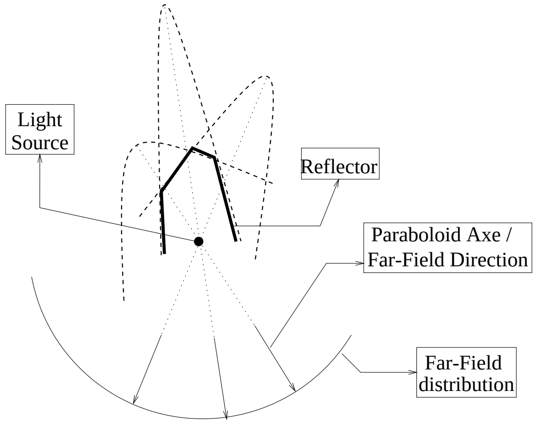


**Figure 9:** *Reflector construction as the boundary of the intersection of confocal paraboloids [CKO99]. Notice that the axis of each paraboloid lies on a direction with prescribed intensity in the far-field region. Compare with Figure 8.*

region onto a given direction in the far-field distribution. The algorithm is the same as described for that paper, but switching at the end to a final Newton-type method for faster convergence. Since the paper deals with the far-field problem, the global size of the reflector is unimportant, so it fixes the focal radius of one paraboloid and optimizes the remaining ones. The authors show that the algorithm converges and it does so at least linearly and, with a proper starting point, the Newton step gives super-linear behavior. Nevertheless, the same authors later [KO03] showed that their algorithm has three disadvantages with respect to its convergence: The last iterations show a significant decrease in their convergence, the convergence becomes worse as the number of paraboloids increases and the error for the paraboloid with a fixed focal radius is much bigger than the error for the remaining directions.

In [KO03] the above mentioned “brute force” method [CKO99] is compared with the Nelder-Mead simplex method [PTVF92]. The main disadvantage of the Nelder-Mead algorithm is that it does not guarantee convergence to a solution, but its convergence depends on how “nicely” the initial reflector is built. They conclude that using their original algorithm to find a first approximation (with only one iteration) and then using it as input to the Nelder-Mead method, reduces all the convergence shortcomings of the former method, providing faster and better convergence behavior. The fact that the error distribution is now much more uniform than in [KO98] is particularly noteworthy.

#### 4.1.2. Surfaces defined as splines

Defining surfaces as splines provides several advantages, like a good tradeoff between global control and good flexibility in the range of achievable shapes, but of course, there are also some drawbacks if the designer is interested in non-smooth surfaces, as points/lines of $C^0$ continuity are hard to generate in an automated way.

In [EN91], Engl and Neubauer face both sorts of problems (far-field and near-field) without the need to map areas on the reflector to points on the near-field plane as the previous article did, although they also use perfect reflectors without any visibility or inter-reflection computations. They also assumed that the starting reflector would satisfy the constraint that no points/directions in the target region would be met by more than one light ray, a condition that is very difficult to satisfy for a real industry problem. In this context and for the near-field case, they derive the associated Monge-Ampère equation. The paper reports a first unsuccessful numerical attempt by using Newton's method with line search for the non-linear partial differential equation, which leads to a sequence of linear elliptic boundary value problems. Unfortunately, this approach does not work well because, as soon as the non-linear problem is unsolvable in the finite element space used, Newton's method becomes useless. Thus, they resort to an optimization approach: they represent the reflector as cubic tensor product splines with a given boundary curve, which is appropriate for the manufacturing process (compare this with the ellipsoid-based representation in [KO98]). They performed a conjugate gradient minimization (Powell's method) of the $L^2$-error in the non-linear equation, taking constraints into account. Side conditions include that the boundary of the reflector can be fixed or variable within prescribed bounds, box constraints on the B-spline coefficients were imposed and a non-unattainable points condition on the reflector was imposed. Finally, the problem of finding an initial reflector to start their algorithm was left to further research, since it is very difficult to find a good starting reflector.

Later, Neubauer presented [Neu94] [Neu97] a solution for the far-field problem without the assumption that parallel outgoing rays will not occur for the starting reflector. Again, a bicubic B-spline representation is used and the spline coefficients are determined to be the best in the least squares sense. The sphere of outgoing directions is divided into 3200 subdomains and the solution is formulated as the minimization of the MS error between the obtained irradiances on the far-field sphere and the user-prescribed irradiances. The light propagation step is computed by Gaussian quadrature based on $4 \times 4$ nodes in the reflector parameterization domain. For each node the incoming and outgoing directions are computed, and each outgoing ray is assigned the corresponding value divided by 16. Since the distribution of light is required to be differentiable (to be able to analytically compute the derivatives), each outgoing ray has its energy "distributed" on the next four neighboring subdomains. Weights are added to the least squares formulation to influence accuracy in different subregions. The side conditions used are that the boundary curve of the reflector can either be fixed or free, box constraints on the B-spline coefficients and the inner product between the incoming ray direction and the surface normal are not allowed to change sign on the whole reflector surface. The minimization is solved iteratively by a projected conjugate gradient method developed by Powell, and the line search is performed via quadratic interpolation. Stability questions are briefly commented. As an example, a reflector generating uniform distribution on an infinitely distant plane is used with 169 spline coefficients (507 variables) and the deviation from the desired light distribution is reduced by a factor of 3 (take into account that this problem has many local minima).

In [KN96] [Neu97] the near-field problem is faced, which has a similar derivation of the integro-differential equation as the previous work [EN91]. Again, as in their earlier work[Neu94] for the far-field problem, a bicubic B-spline representation is used for the reflector surface, and the light propagation step is computed by tracing rays from the source to vertices on the surface, which were obtained from a regular grid in the spline parametric space. To achieve differentiability of the algorithm with respect to the spline coefficients, the light values are distributed onto the neighboring subregions defined in the near-field object in a differentiable way. The minimization procedure seeks to optimize a weighted least squares error with an extra added term that accounts for the light falling *outside* the near-field region. The optimization algorithm chosen is the same as in [Neu94].

In [HBKM95a] [Hal96] Halstead presents a different problem: The reconstruction of a three-dimensional surface model from an image of this surface illuminated with a structured pattern of light. The surface is ideally specular. The authors apply their algorithm to the measurements of the human cornea, which are made by a device called a *videokeratograph* (Figure 10a). The input to the algorithm is the set of feature positions (important features located on the image plane) extracted from a single image (Figure 10b). These positions are in the coordinate system of the image plane, and each feature must already be linked to an identifiable feature in the source pattern. The algorithm is expected to output the reconstructed cornea in the form of a continuous function describing the surface's position. For this function a biquintic B-spline is used, and its coefficients have to be found by the algorithm. The algorithm uses constrained optimization (Levengerg-Marquardt method) to solve the non-linear least squares problem and the objective is a function of the surface variables which measure the error between the surface and the original cornea. Its formulation greatly affects the efficiency of the solution process. The authors discovered that using the comparison of the real image to the image generated synthetically as the error measure is not a good option since useful error information is gained only around the boundaries between the structured regions. However, since each feature is associated with a corresponding point or set of points in the source pattern, a good error definition could be the square of the shortest distance between the intersection point of a ray from the image (reflected at the cornea surface) with the source pattern and the set of corresponding source points (Figure 11), the objective function is the sum of errors of all the image features.

Later, in [HBKM95b] [HBKM96] [Hal96], the error was computed differently: To evaluate the error of a single feature they again trace a ray from the feature (in the real image) through the nodal point (the camera acts as pinhole camera) to the surface. Then, they compute at which point they would like the reflected ray to intersect the source pattern. Finally, they keep the closest point in the set of corresponding features, and compute the normal that would reflect the incoming ray to the desired source point. The error is measured in terms of the difference between this normal and the current surface normal. The sum of the error terms of all features gives the objective function. This last definition reduced the error evaluation cost from hours to minutes. Each time the variation of the mean angle between normals changes a little in successive iterations, the surface is refined doubling the number of patches to optimize. Note that the above backward ray-tracing computations are done without any visibility or inter-reflection computations, but the extra evaluations are not needed due to the particular geometric setting used. Thus, the main difference from the previously mentioned approaches depends on the backwards ray tracing algorithm used, mainly due to the geometric nature of the stated problem. Another difference concerns error definition, based on a feature-match difference measurement.

In [LSS98], Loos, Slusallek and Seidel deal with a similar problem: designing Progressive Lenses for defects of the human visual system (a refraction problem). They use the results of wavefront tracing for re-focusing the eye while rendering and for deriving

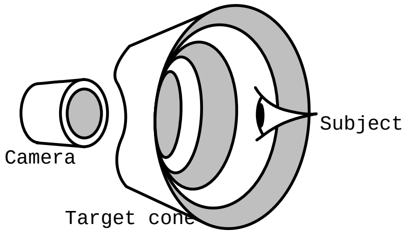


a)

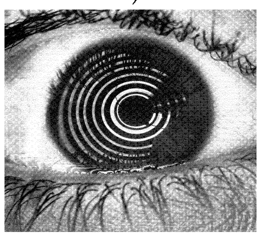

b)

**Figure 10:** *a) Videokeratograph system and b) Captured image [HBKM96, Hal96]. Reprinted with the permission of Brian A. Barsky, UC Berkeley.*

an accurate error functional that describes the desired properties and optical error across the lens. The algorithm performs a minimization of this error, yielding an optimal free-form refractive lens surface. The rendering algorithm finds, at first, the point in the environment the eye should focus on by tracing a single ray from the center of the pupil to the center of the pixel. Then, wavefront tracing is performed and the accommodation required to get a sharp image is computed. Finally, distribution ray tracing is used to compute the final pixel color. The error functional for describing the optimal progressive lens is based on wavefront tracing and the effective astigmatism, a generalization of traditional error measures used in optics. The lens is represented only as the front surface, with a fixed toroidal or spherical back surface. This back surface is discretized in about $100 \times 100$ sample points. The optimization is performed by a variant of the well-known Newton-iteration method, showing that no more than two iterations were needed. The main difference with the other works reviewed in this section concerns the new error definition, which is an extremely specific, optics-based error measure.

### 4.2. Global Illumination

The papers studied in this section share the important feature of facing the Surface Design Problem in the context of global illumination treatment of light propagation. They can be classified according to the representation chosen for the surface: splines or simple polygons.

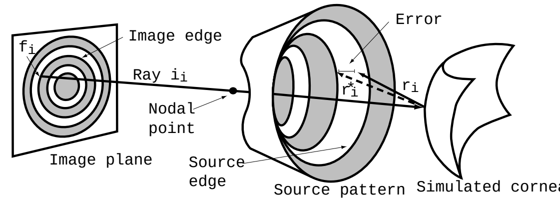


**Figure 11:** *By using backward ray-tracing, in [HBKM96, Hal96] the authors determine the error in surface shape for a given image sample point. Reprinted with the permission of Brian A. Barsky, UC Berkeley.*

#### 4.2.1. Surfaces defined as splines

As mentioned above, splines show several good characteristics for global optimization, and the works studied in this section share the common approach of simulating global light propagation in a two-dimensional setting, thus only requiring 2D splines to represent the surfaces to optimize.

Doyle, Corcoran and Connell present [DCC99b] [DCC99a] [DCC01] an evolution strategy (a variation on the Genetic Algorithm) for 2D luminary design with point[DCC99b] and extended[DCC99a] light sources, later completed with a careful analysis of the objective function[DCC01] (called Merit Function in the papers). As such, they defined a three term objective function consisting of a term that represents the absolute element-wise difference between the target and resultant distributions in an objective region, and a term representing a penalty for the rays that lie outside the desired region, plus a third term that represents a measure of the difference between the reflected and the desired light power. The 2D surface definition is a cubic Bézier curve whose four control points are the genes of the chromosomes for the optimization evolution algorithm. Light propagation is performed by a recursive ray tracer with the number of bounces used as a stopping criterion (rays with more bounces than allowed are "lost"). Extended light sources were defined as a set of Lambertian-like (cosine distribution) point sources located at regular intervals on the perimeter of an arbitrarily sized circle. Each point source has its main axis in the circle radial direction. In this case, a penalty for the total light power re-intercepted by the extended source was added to the objective function. It is important to note that the main difference from all the other works mentioned is the usage of both a global 2D illumination algorithm and a global optimization technique like genetic programming for finding the surface-defining coefficients.

#### 4.2.2. Polygonal Surfaces

In the work by Patow et al.[PPV04], the problem of the far-field case taking into account multiple bounces of light inside a reflector with occlusion and a generic BRDF was studied. They use a Monte-Carlo based light propagation algorithm that traces rays from the light source (punctual in their studies, but the presented algorithm easily accommodates extended light sources as well) and follows them until they get out of the optical system. The optimization algorithm used is a global "brute force method", that conducts tests of the performance of each member of a family of polygonal reflectors obtained by iteratively combining the addition of an increment to each of the vertices. For this to be useful, the size of this generated family of reflectors has to be kept manageable. Obviously, the desired accuracy given by these increments added to the vertices (trying to ensure that the family contains at least one reflector close to the desired target), and the size of such a family are closely related. The vertices are sorted according to the contribution they make to the overall illumination error, and more effort (i.e. more samples) is put into the vertices with worst error. The results show good convergence of the algorithm, but the reported times are slow.

## 5. Discussion

Here we will summarize the conclusions of the surveyed work, to underline common problems, characterize solution approaches and present the open issues.

As seen, Inverse Surface Design papers can be grouped into two sets: analytical and numerical works. In Table 7 we see that the analytical works reviewed can in turn be grouped into methods that deal with the rotationally symmetric approximation ([WN75], [BW78] and [Oli89]) and those that don't ([Wan96], [KO97] and [KOvT98]). The former works present the equation formulation for the problem of reflector design in the rotationally symmetric simplification, and show existence and uniqueness of the solution. The latter present a more general equation formulation, showing existence, uniqueness and smoothness of the solution. In particular, [KO97] and [KOvT98] present a constructive demonstration in terms of weak convergence to the solution. It is also possible to group of the studied works according to the type of assumption they make, far-field or near-field, which is often found in the literature.

Papers on numerical methods, summarized in Table 6, can be characterized by:

- *Shape definition*: [KO98], [CKO99] and [KO03] use the boundary of the intersection of a series of simple primitives (ellipsoids or paraboloids) to define the surface. [EN91], [Neu94, Neu97], [KN96, Neu97], [HBKM95a, Hal96], [HBKM95b, HBKM96, Hal96] and [LSS98] use tensor product of splines (bicubic or biquintic). A simple polygon-based shape definition was used in [PPV04].
- *Light Propagation method*: We see that almost all the reviewed papers, with the exceptions of [HBKM95a, Hal96] and [HBKM95b, HBKM96, Hal96], basically use light ray tracing [Arv86] to compute the light propagation at each iteration of their optimization method. The mentioned exceptions use traditional ray tracing [Whi80], basically because they are computing the shape of the human cornea and working with a convex surface shape, and all the reflected rays reach the *videokeratograph*. To the best of our knowledge, almost all works which study the far-field problem only use local illumination schema to solve the light propagation step, with the only exception being [PPV04].
- *Type of problem faced*: It can be either a near-field formulation of the problem or a far-field one. In the former, the desired light distribution is defined on a registration area with distances comparable to the ones involved in the optical set, while in the later it is defined at a region infinitely far away from it.
  - near-field distribution: [EN91], [KO98], [KN96], [Neu97], [HBKM95a], [Hal96], [HBKM95b], [HBKM96], [Hal96], [LSS98], [DCC99b], [DCC99a] and [DCC01].
  - far-field distribution:[EN91], [CKO99], [KO03], [Neu94], [Neu97] and [PPV04].
- *Surface material assumptions*: We see that almost all the reviewed works only use the perfect specular BRDF; in addition [KO98], [CKO99] and [KO03] and [EN91] assume that there are no two outgoing rays with the same orientation. This basically implies that every two different points on the reflector surface will reflect light in two different directions. The other papers surveyed do not use this highly restrictive assumption. To the best of our knowledge, the only paper that uses a generic BRDF is [PPV04], sampled by a Monte Carlo scheme.
- *Optimization method*: [KO98], [CKO99] and [KO03] use and compare their own custom method (named "Brute Force Method") with traditional methods like Newton or Nelder-Mead. [EN91], [Neu94], [Neu97] and [KN96] [Neu97] use the Projected Conjugate Gradient method; [HBKM95a], [Hal96] use the Levengerg-Marquardt method, and [HBKM95b], [HBKM96], [Hal96] use Least Squares as the chosen optimization method. [LSS98] uses a modified Newton-iteration method. On the other hand, Doyle et al. [DCC99b] [DCC99a] [DCC01] use a Genetic Algorithm for the optimization, and

[PPV04] uses a custom global "Brute Force" algorithm. To the authors knowledge, choosing one optimization method or another is quite arbitrary, and no method has proven to be the best for this sort of optimization problem.

We have also seen that almost all works use a local illumination approach, with the further restrictions of using pure specular surfaces and the highly restrictive condition that no inter-reflections occur, with the only exception in the works by Doyle et al.[DCC99b] [DCC99a] [DCC01], which use global illumination with a perfect specular BRDF model with a threshold based stopping criteria, and [PPV04], in which a more generic global illumination algorithm is used.

- *Inter-reflections*: The fact that almost no article inter-reflections are computed is especially important when the reflector is highly specular and its surface concave, as in most papers reviewed, since a very important contribution to the final light distribution may be underestimated. This is not applicable to [HBKM95a] [Hal96] and [HBKM95b] [HBKM96] [Hal96] due to the convexity of the surface used to approximate the human cornea, and to [LSS98] which computed the refractive surface of a lens.
- *BRDF*: As mentioned, most of the commented works deal only with pure specular surfaces. To understand what happens if other BRDFs are used, we can consider that Equation 1 can be regarded in a signal processing framework [RH01] under the restrictions of no inter-reflections (as most of the papers assume), isotropic BRDFs, known geometry (besides the reflector) and camera parameters. Then, the reflected light field integral is regarded as the convolution of two signals: the bidirectional reflectance function and the incident lighting; i.e. by filtering the illumination using the BRDF. Then, inverse rendering can be simply viewed as a deconvolution of the two signals. This framework [RH01] led the authors to conclude that BRDF recovery is well-conditioned (in a mathematical sense) when lighting contains high frequencies (e.g. directional sources) and is ill-conditioned for soft lighting. Counterwise, inverse lighting is well-conditioned for BRDFs with high-frequency components (specular peaks) and ill-conditioned for diffuse surfaces. The same analysis can be carried on for the inverse problem of Surface Design, leading to the conclusion that convergence of the algorithms would become worse as the BRDF gets more diffuse: this is one of the main limitations for solving a generic-BRDF problem. Experimental results [PPV04] confirm this analysis.
- *Local vs. Global Illumination*: Another factor to take into account is that almost none of the different papers treat the full global illumination equation, Equation 2, limiting themselves to using of a simpler local-illumination version based on a simplification of the equation being used by considering only point light sources and without considering inter-reflections. As most of the reviewed papers omit the treatment of inter-reflections, the resulting algorithms provide solutions that are not readily applicable in real-life situations. Multiple reflections would not only introduce more complexity in the light propagation method evaluation (requiring more computing time to get accurate results), but would also severely increase the associated variance for non-specular surfaces, thus lowering convergence, as can be seen in [PPV04]. Other studies that take into account global illumination effects are presented in [DCC99b] [DCC99a] [DCC01], but, as they use a perfect specular surface, the results do not show any serious increase in the associated variance, and thus do not affect the way their algorithms converge.
- *Visibility*: It is also important to mention the treatment of the visibility in the different approaches reviewed: when computing the radiance with the above equations, the visibility problem consists of detecting if there are any blockers between the source and the surface being illuminated, and not adding their contribution in that case. The same is true for the paths from the surface to the eye or the region where the final radiance computations are needed. In general, this would be solved by adding visibility calculations to most approaches, simply by not accounting for occluded light. This would affect the error measure and steer the optimization, adding a strong non-linearity which would cause serious convergence problems for local optimization based algorithms. In those cases, a stimulated annealing approach ([PTVF92]), a genetic algorithm ([DCC01] [DCC99b] [DCC99a]) or even a brute force approach would be needed [PPV04]. This can be seen in the work by Doyle, Corcoran and Connell [DCC99a], where the occlusion introduced by the extended light source is introduced. This resulted in a restriction of the search space of possible reflectors, which forced the authors to introduce several changes to the reflector shape and size in order to compensate for this.

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

- Letting an algorithm automatically choose between some possibilities the more suited material for a surface to get a given illumination.
- Another open problem can be found in determining the shape of surfaces by using the full rendering equation, since all treatments up to now deal with

highly simplified versions of the problem. Additionally, more general BRDFs should be considered.

- One promising line of research is the one introduced by Costa et al. [CSF98], in which they show that the potential equation can be used to simplify the problem from considering the whole environment to only a small subset of it, thus reducing the number of possible degrees of freedom to be taken into account. In spite of this, the general problem of obtaining inverse surface shapes from general shading in an environment remains untouched.
- Determining scene shapes other than sources and surfaces from direct illumination information could be an interesting problem to solve, in spite of its inherent high complexity. For example, it would be interesting to find the shape needed in a room to get a desired illumination effect, both by its direct contribution as a light-blocking element, and by the distortions introduced by its presence in the light propagation through the environment.
- Calculating shapes in the presence of participating media has never been considered, despite their importance for design in any given adverse conditions.

## 6. Acknowledgments

We would like to thank Pierre Poulin for his insights, Brian A. Barsky for giving us permission to reprint his figures, Sergey A. Kochengin for the extra explanations about his work and Jack Howell for his reprints and for proofreading an earlier version of this paper. Partially done under grant TIC2001-2226-C0202 of CICYT (Spain) and 2001SGR0296 of DURSI of the Generalitat de Catalunya.

| | near/ far | Shape Definition | BRDF | Light Emitter | Global/ local illum | Light Propagation | Optimization Method | Approx. |
|---|---|---|---|---|---|---|---|---|
| [WN75] | far | Rotationally Symm. | Perfect Specular | Isotropic point light | Local | 1-1 | Theoretical | Shpe with even azimuthal symm. |
| [Oli89] | near | Rotationally Symm. | Perfect Specular | Point light | Local | 1-1 | Theoretical | Shpe with even azimuthal symm. |
| [Wan96] | far | Generic | Perfect Specular | Point light | Local | Theoretical ray tracing | Theoretical | |
| [KO97] | near | Families of confocal ellipsoids | Perfect Specular | Point light | Local | Maps sets of rays impinging on a ellipsoid onto a direction | Theoretical | Convex Shape |
| [KOvT98] | near | Families of confocal hyperboloids | Perfect REFRACTIVE | Point light | Local | Theoretical tracing of refractive rays | Theoretical | Convex Shape |
| [KO98] | near | $\partial(\bigcap$ ellipsoids) | Perfect Specular | Point Source | Local | Light ray tracing. Source is subdivided and rays are counted if arrive to objective point/ dir. | BFM | ˜I, ˜V, convex shape |
| [CKO99] | far | $\partial(\bigcap$ paraboloids) | | | | | BFM + Newton | ˜I, ˜V, convex shape |
| [KO03] | far | $\partial(\bigcap$ paraboloids) | | | | | Nelder-Mead | ˜I, ˜V, convex shape |
| [EN91] | both | Cubic spline tensor product | Perfect Specular | Point Source | Local | Not stated explicitly. Probably as in [Neu94] [Neu97] | Projected Conjugate Gradient | ˜I, ˜V, No parallel outgoing rays |
| [Org02b] | both | Segmented reflector | Perfect Specular | Point Source | Local | ray tracing | Optimization by user | ˜I, ˜V |
| [Neu94] [Neu97] | far | Bicubic B-splines | Perfect Specular | Point Source | Local | Gaussian quad. on refl. param. domain. For each point, in & out rays are found and out-ray is "distributed" on Far-Field subdomains. | Projected Conjugate Gradient with line search via quadratic interpolation. | ˜I, ˜V |
| [KN96] [Neu97] | near | Bicubic B-splines | Perfect Specular | Point Source | Local | Gaussian quad. on refl. param. domain. For each point, in & out rays are found and out-ray is "distributed" on Near-Field subdomains. | Projected Conjugate Gradient with line search via quadratic interpolation. | ˜I, ˜V |
| [HBKM95a] [Hal96] | near | Biquintic B-splines | Perfect Specular | Point Source | Local | Backwards ray-tracing from the feature (in the real image) through the nodal point to the specular surface, then reflected and traced to the pattern source. | Levengerg-Marquardt method. | ˜I, ˜V |
| [HBKM95b] [HBKM96] [Hal96] | | | | | | | Least Squares | |
| [LSS98] | near | bicubic NURBS | Perfect REFRACTIVE | Point source | Local | ray tracing to locate objects, wavefront tracing to re-focus the eye + distribution ray-tracing | Modif. Newton-it | |
| [DCC99b] [DCC01] | near | 2D Bèzier curve | Perfect Specular | Point Source | 2D Global | Ray tracing | Evolution Strategy | 2D |
| [DCC99a] | near | 2D Bèzier curve | Perfect Specular | Extended lambertian Light Source | 2D Global | Ray tracing from uniform samples on source | Evolution Strategy | 2D |
| [PPV04] | far | Polygonal Surf | General | Point (Extended too) | Global | Monte Carlo Light Tracing | "Brute Force" | |

Table 6: Papers with numerical solutions to the Reflector Design problem. (continued).

**Appendix A:** Appendix

In this appendix we analyze the mathematical aspects of each surveyed paper's solution method. Table 8 explains the constraints imposed on the solution method (for optimization approaches) for Inverse Geometry problems. In this case the restrictions are incorporated into the optimization algorithm. The third column shows the starting point used by the different methods. Finally, the fourth column presents the objective function to be optimized in the case of optimization approaches.

| | near-field / far-field | Demonstrates / Shows | Restrictions / approximations |
|---|---|---|---|
| [WN75] | far | Eq. formul. ∃ & uniq. for RS. | PS, 1-1, ˜I, ˜V, RS |
| [BW78] | near | Eq. formul. ∃ & uniq. for RS. Dem. that paraboloid is sol. for const. phase func. | PS, 1-1, ˜I, ˜V, RS |
| [Oli89] | near | ∃ & uniq. for RS. | PS, 1-1, ˜I, ˜V, RS |
| [Wan96] | far | ∃, uniq. & smothness | PS, ˜I, ˜V |
| [KO97] | near | Probl. def. ∃ of weak sol. for the class of convex sol. | PS, ˜I, ˜V |
| [KOvT98] | near (refractions) | Probl. def. ∃ of weak sol. for the class of convex sol. | PS, ˜I, ˜V |

**Table 7:** *Theoretical papers on Inverse Surface Design. In this table "RS" stands for Rotationally Symmetric approximation, "PS" means that the surfaces are perfectly specular or perfectly refractive, "1-1" that the incoming-ray/outgoing-ray uniqueness relationship approximation is used, "˜I" means that interreflections are disregarded and "˜V" means that no visibility computation is done. The second column refers to the type of problem faced: near-field or far-field problem.*

| Paper | Constraints imposed | Initial guess / Initial approx. | Objective function | Tests Performed | comments |
|---|---|---|---|---|---|
| [KO98] | None | $M = \max_{x \in Surface} \lvert x \rvert$ then first polar radius $= 16M$ and all other p.r. $= 64M$, i.e. reflector is first ellipsoid | Minimize each error ($\ell^\infty$) separately | Plane uniformly illuminated with 4, 10 and 25 ellipsoids. | Shows convergence & $O(K^4 ln K)$ |
| [CKO99] | First polar radius is fixed | First Polar Radius = 1, rest chosen such that radiance at a point $\leq$ prescribed energy + $\epsilon$ | Minimize each error ($\ell^\infty$) separately | Uniform illum with 7, 19 paraboloids. | $O(linear)$ |
| [KO03] | First polar radius is fixed for all reflectors | Simplex vertices in neighborhood of initial reflector | Minimize simplex energy | Uniform illum with 17, 19 and 61 paraboloids. | $O < linear$ |
| [EN91] | Boundary fixed or variable. Box constr. on B-spline coefs. Non-unattainable points on refl. | User provided with no parellel outgoing rays for first reflector: Too difficult to provide! | Least squares error between the current light distribution from the reflector and the desired distribution. | Synthetic examples: 5 × 5 nodes, good results. Real Ex. (Poor behavior due to bad starting point): Rot. sym. refl. w/ unif. illum. on pl. at ∞: $10 \times 10$ nodes, 48 iterations (40 minutes). | Only coarse solutions: too high evaluation cost. |
| [Neu94] [Neu97] | Boundary fixed or variable. Box constr. on B-spline coefs. Surface normal is not allowed to change sign on the refl. surface. | User's (manufacturer) provided. | | Rot. sym. refl. w/ unif. illum. on subreg. of pl. at ∞: $10 \times 10$ nodes, 360its (6 days w/ analytic grads.), error reduced 3 times. | |
| [KN96] [Neu97] | Boundary fixed or variable. Box constr. on B-spline coefs. Surface normal is not allowed to change sign on the refl. surf. No reflected ray must go away from near-field reg. | User's (manufacturer) provided. | | Synthetic Ex: 10 × 10 nodes, near planar subreg., error reduced from 189% to 3%. Real Ex: near plannar subreg., 149 its (5.5 hours), error reduced 3 times. | |
| [HBKM95a] [Hal96] | None | Little effect: paraboloid | Error $= \sum$ \|feature pos - backwards trac. ray pos\|$^2$ | Synthetic ellipsoid: 1 to 8 × 8 patches. $RMSE \leq 8.5 \times 10^{-6} mm$. 5760 samples at image. | Error $= \sum$ \|feature pos - backwards trac. ray pos\|$^2$ |
| [HBKM95b] [HBKM96] [Hal96] | None | Guess at the shape of the cornea | Error $= \sum$ dist(surface normal, desired normal) | Synthetic data sets: RMS of 0.0092 microns! Real Data sets: RMS of 0.9-1.5 microns. All: 8 × 8 final patches, over 5000 image-samples. | Error $= \sum dist$(surface normal, desired normal) |
| [LSS98] | None | front-surf with const dist form back-surf | Error functional that describes the aberrations in a progressive lens | bi-cubic tensor product of NURBS, took 66.5 secs in a SGI Onix w/ R10000. | |
| [DCC99b] [DCC99a] [DCC01] | None | Random Initial Population | Diff between target and resultant distrib at obj region plus lost rays plus difference between reflected and desired power[DCC01] | 2D: parabolic & elliptic reflectors | |
| [PPV04] | None | Generic Initial Refl | Least squares error between the current light distribution from the reflector and the desired distribution. | Different Configurations. 10 days on a P-IV PC. | |

Table 8: Summary of the mathematical aspects of the inverse lighting approaches presented: Inverse Geometry (continued).

## References


[Arv86] James R. Arvo. Backward Ray Tracing. In *ACM SIGGRAPH '86 Course Notes - Developments in Ray Tracing*, volume 12, August 1986.

[Arv95a] James Arvo. *Analytic Methods for Simulated Light Transport*. Ph.D. thesis, December 1995.

[Arv95b] James Arvo. The Role of Functional Analysis in Global Illumination. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 115–126, New York, NY, 1995. Springer-Verlag.

[BW78] F. Brickell and Brian S. Westcott. Phase and power distribution on plane apertures of reflector antennas. *J. Phys. A: Math. Gen*, 11:777–789, 1978.

[CKO99] Luis A. Caffarelli, Sergey A. Kochengin, and Vladimir I. Oliker. On the numerical solution of the problem of reflector design with given far-flield scattering data. *Contemporary Methematics*, 226, 1999.

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

[Cor] ZEMAX Development Corporation. Zemax software for optical design. Technical report, ZEMAX Development Corporation, http://www.zeemax.com.

[CSF98] Antonio Cardoso Costa, A. Augusto Sousa, and F. Nunes Ferreira. Design de iluminação. In *8th Portuguese workshop on Computer Graphics*, Coimbra, Portugal, 1998. in Portuguese.

[CW93] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Boston, MA, 1993.

[DCC99a] Steven Doyle, David Corcoran, and Jon Connell. Automated mirror design for an extended light source. In *Proceedings of SPIE*, volume 3781, page 94, 1999.

[DCC99b] Steven Doyle, David Corcoran, and Jon Connell. Automated mirror design using an evolution strategy. *Optical Engineering*, 38(2):323–333, 1999.

[DCC01] Steven Doyle, David Corcoran, and Jon Connell. A merit function for automated mirror design. *Journal of the Illuminating Engineering Society*, 30(2):3–11, 2001.

[EN91] Heinz W. Engl and Andreas Neubauer. Reflector design as an inverse problem. In Heiliö M., editor, *Proceedings of the Fifth European Conference on Mathematics in Industry*, pages 13–24, Teubner, Stutgart, 1991.

[Eng] West Coast Engineering. Odp optical design program. Technical report, West Coast Engineerin, http://www.westcoastengineering.com/.

[Hal96] Mark A. Halstead. *Efficient Techniques for Surface Design Using Constrained Optimization*. PhD thesis, University of California, Berkeley, May 1996.

[HB89] Berthold K. P. Horn and Michael J. Brooks. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.

[HBKM95a] Mark A. Halstead, Brian A. Barsky, Stanley A. Klein, and Robert B. Mandell. Geometric modeling of the cornea using videokeratography. In Morton Daelhen, Tom Lyche, and Larry L. Schumaker, editors, *Mathematical Methods for Curves and Surfaces*, pages 213–223. Vanderbilt University Press, Nashville, 1995.

[HBKM95b] Mark A. Halstead, Brian A. Barsky, Stanley A. Klein, and Robert B. Mandell. A spline surface algorithm for reconstruction of corneal topography from a videokeratographic reflection pattern. *Optometry and Vision Science*, 72(11):821–827, November 1995.

[HBKM96] Mark A. Halstead, Brain A. Barsky, Stanley A. Klein, and Robert B. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. *Computer Graphics*, 30(Annual Conference Series):335–342, 1996.

[HMH95] Vigain Harutunian, Juan C. Morales, and John R. Howell. Radiation exchange within an enclosure of diffuse-gray surfaces: The inverse problem. In *Inverse Problems in Heat Transfer, ASME/AIChE National Heat Transfer Conference*, Portland, August 1995.

[Jen01] Daid G. Jenkins. Reflector design using lighttools. Application summary, Optical Research Associates, http://www.opticalres.com, 2001.

[Kil00] Per-Simon Kildal. *Foundations of Antennas - A Unified Approach.* Studentlitteratur, Lund, Sweden, 2000.

[KM00] Bernhard Kress and Patrick Meyrueis. *Digital Difractive Optics.* Ed. John Wiley and Sons, New York/Chicester/Brisbane/Toronto, 2000.

[KN96] Christa König and Andreas Neubauer. Design of a 3d-reflector for the near field problem: a nonlinear inverse problem. *Math. engng. Ind.*, 5(4):269–279, 1996.

[KO97] Sergey A. Kochengin and Vladimir I. Oliker. Determination of reflector surfaces from near-field scattering data. *Inverse Problems*, 13:363–373, 1997.

[KO98] Sergey A. Kochengin and Vladimir I. Oliker. Determination of reflector surfaces from near-field scattering data ii. numerical solution. *Numer. Math.*, 79(4):553–568, 1998.

[KO03] Sergey A. Kochengin and Vladimir I. Oliker. Computational algorithms for constructing reflectors. *Computing and Visualization in Science*, 6:15–21, 2003.

[Koc] Sergey A. Kochengin. Personal Comunication.

[KOvT98] Sergey A. Kochengin, Vladimir I. Oliker, and Oliver von Tempski. On design of reflectors with prespecified distribution of virtual sources and intensities. *Inverse Problems*, 14(3):661–678, 1998.

[LSS98] Joachim Loos, Philipp Slusallek, and Hans-Peter Seidel. Using wavefront tracing for the visualization and optimization of progressive lenses. *Computer Graphics Forum (Eurographics '98)*, 17(3):–, 1998.

[Mar98] Stephen R. Marschner. *Inverse Rendering in Computer Graphics.* PhD thesis, Program of Computer Graphics, Cornell University, Ithaca, NY, 1998.

[Neu94] Andreas Neubauer. The iterative solution of a nonlinear inverse problem from industry: Design of reflectors. In P. J. Laurent, A. Le Méhauté, and L. L. Schumaker, editors, *Curves and Surfaces in Geometric Design*, pages 335–342, Boston, 1994. A. K. Peters.

[Neu97] Andreas Neubauer. Design of 3d-reflectors for near field and far field problems. In L.T. Bieyler, T. F. Coleman, A.R. Conn, and F.N Santosa, editors, *Large Scale Optimization with Applications. Part I: Optimization in Inverse Problems and Design*, volume 92 of *IMA Volumes in Mathematics and its Applications*, pages 101–118. Springer, New York, 1997.

[Oli89] Vladimir I. Oliker. On reconstructiong a reflecting surface from the scattering data in the geometric optics approximation. *Inverse Problems*, 5:51–65, 1989.

[Oli02] Vladimir I. Oliker. On the geometry of convex reflectors. *Banach Center Publications*, 57:155–169, 2002.

[Oli03] Vladimir I. Oliker. Mathematical aspects of design of beam shaping surfaces in geometrical optics. In S. Kromker, R. Rannacher, and F. Tomi, editors, *Trends in Nonlinear Analysis*, pages 193–224, Springer-Verlag, 2003.

[Opt] Optenso. Links on optical design and engineering. Technical report, Optenso: Optical Engineering Software, http://www.optenso.de/links/links.html.

[Org02a] Breault Research Organization. Asap technical summary. Application Summary BRO-0400, Breault Research Organization, http://www.breault.com, March 2002.

[Org02b] Breault Research Organization. Create segmented reflectors quickly with reflectorcad. Feature Note BRO-FN1405, Breault Research Organization, http://www.breault.com, January 2002.

[PP03] Gustavo Patow and Xavier Pueyo. A survey on inverse rendering problems. *Computer Graphics Forum*, 22(4):663–687, 2003.

[PPV04] Gustavo Patow, Xavier Pueyo, and Àlvar Vinacua. Reflector design from radiance distributions. *International Journal of Shape Modelling*, 10(2):211–235, 2004.

[PTVF92] W. H. Press, S.A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C : The Art of Scientific Computing.* Cambridge University Press, 1992.

[RH01] Ravi Ramamoorthi and Pat Hanrahan.

A signal-processing framework for inverse rendering. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001)*, August 2001.

[Sha97] Robert R. Shannon. *The Art and Science of Optical Design.* Cambridge University Pres, 1997.

[SL97] A. James Stewart and Michael S. Langer. Towards accurate recovery of shape from shading under diffuse lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):1020–1025, September 1997.

[Smi90] Warren J. Smith. *Modern Optical Engineering (Second Edition).* McGraw-Hill, 1990.

[SO] Inc. Sinclair Optics. Oslo: Optics software for layout and optimization of optical systems. Technical report, Sinclair Optics, Inc., http://www.sinopt.com/.

[Soi02] Victor A. Soifer. *Methods for Computer esign of Diffractive Optical Elements.* Ed. John Wiley and Sons, New York/Chicester/Brisbane/Toronto, 2002.

[Vas98] Darko Vasiljevic. Optimization of the cooke triplet with the various evolution strategies and the damped least squares. In *Optical Design and Analysis Software, Proceedings of SPIE*, volume 3780, pages 207–215, 1998.

[Wan96] Xu-Jia Wang. On the design of a reflector antenna. *Inverse Problems*, 12:351–375, 1996.

[Wes83] Brian S. Wescott. *Shaped Reflector Antenna Design.* Letchworth: Research Studies Press, 1983.

[Whi80] Turner Whitted. An improved illumination model for shaded display. *CACM*, 23(6):343–9, June 1980. also in Tutorial: Computer Graphics: Image Synthesis, Computer Society Press, Washington, 1988. (Abstract) in Computer Graphics, no. 2 (SIGGRAPH '79 Proceedings), p. 14.

[WN75] Brian S. Westcott and A. P. Norris. Reflector synthesis for generalized far-fields. *J. Phys. A: Math. Gen*, 8(4):521–532, 1975.