INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Image-Based Rendering from a Single Image

Samuel Boivin — André Gagalowicz

N° 4098

Janvier 2001

THÈME 3

ISSN 0249-6399 ISRN INRIA/RR--4098--FR+ENG

Image-Based Rendering from a Single Image

Samuel Boivin , André Gagalowicz

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet MIRAGES

Rapport de recherche n° 4098 — Janvier 2001 — 31 pages

Abstract: In this report, we present a new method to recover an approximation of the bidirectional reflectance distribution function (BRDF) of the surfaces present in a real scene. This is done from a single photograph and a 3D geometric model of the scene. The result is a full model of the reflectance properties of all surfaces, which can be rendered under novel illumination conditions, with, for example, viewpoint modification and the addition of new synthetic objects. Our technique produces a reflectance model using a small number of parameters. Using a new hierarchical algorithm assumptions are done for each surface about their BRDF, which becomes more and more complex during the rerendering iterations. We treat each type of BRDF separately, applying a specific algorithm to take into account its own characteristics and its parameters. The number of parameters varies with the complexity of the model.
When an approximation of the BRDF of all surfaces has been found, high quality rerenderings are produced using a global illumination software. We present several synthetic images that are compared to the original ones, as well as some possible applications in augmented reality.

Key-words: Computer Graphics, Computer Vision, Illumination, Image-Based Rendering, Radiosity, Ray Tracing, Reflectance Recovery, Rendering, Inverse Rendering, Rerendering, BRDF, Augmented Reality

Unité de recherche INRIA Rocquencourt
Domaine de Voluceau, Rocquencourt, BP 105, 78153 Le Chesnay Cedex (France)
Téléphone : +33 1 39 63 55 11 — Télécopie : +33 1 39 63 53 30

Rendu Inverse à partir d'une Seule Image Réelle

Résumé : Dans ce rapport de recherche, nous présentons une nouvelle méthode pour la reconstruction d'une approximation de la fonction de distribution de réflexion bidirectionnelle (BRDF). Les données nécessaires à une telle reconstruction sont décrites par une seule photographie de la scène et un modèle géométrique 3D de cette scène. Le résultat produit est un modèle complet des propriétés de réflexion de toutes les surfaces, incluant la possibilité de recalculer une nouvelle image synthétique, soit sous les conditions initiales, soit sous de nouvelles conditions (nouvel éclairement, nouveau point de vue, ajout/ablation d'objets, etc.). Notre technique emploie un modèle de BRDF simple, basé sur cinq paramètres au maximum, pour calculer les réflectances des objets, qui deviennent de plus en plus complexes au fur et à mesure des itérations de rendu inverse. Chaque type de surface extraite de l'image réelle est donc traitée séparément, en appliquant un algorithme spécifique pour leur calculer des propriétés de réflectance. Le nombre de paramètres de réflectance varie avec la complexité de la BRDF choisie pour une surface.
Lorsqu'une approximation de la BRDF de chaque surface de la scène a été trouvée, il est possible de calculer une nouvelle image synthétique en utilisant un logiciel d'illumination globale. Nous présentons ici plusieurs images synthétiques qui sont directement comparées aux images réelles, ainsi que des possibilités d'applications dans le domaine de la réalite augmentée.

Mots-clés : Vision par Ordinateur, Synthèse d'Images, Illumination, Rendu à base d'Images Réelles, Radiosité, Lancer de Rayons, Recouvrement de Réflectances, Rendu Réaliste, Rendu Inverse, BRDF, Réalité Augmentée

1 Introduction

1.1 Overview of the problem

Since its origin, Computer Graphics has aimed to simulate an illusion of reality. Rendering algorithms have been developed specifically to generate near-perfect images under realistic illumination conditions. It is often difficult to say if such images are realistic or not because there is no real reference such as a photograph. Moreover, the application may need to create novel viewpoints and/or novel illumination conditions from a sparse set of photographs. This is difficult to achieve without using image-based modeling and rendering algorithms. For example, suppose we want to insert a new synthetic object on top of a real anisotropic mirror inside a real scene. This operation clearly requires to take into account the interaction between the new object and its environment (especially this mirror). This is impossible to do, if we do not have an approximation of the reflectance properties of the real surfaces in the image. Therefore specific algorithms are necessary to recover these reflectance properties from the real images.
Many authors have contributed to the resolution of this problem [20, 24, 31, 30, 32, 25, 26, 33, 7, 39, 22, 23, 29, 28, 14, 11, 27]. The algorithms that they have produced vary greatly and not all can be re-used for our applications. Considerable work has been done for the reflectance estimation of an isolated object in particular illumination conditions [20, 24, 31, 30, 32, 25, 26] . Although these techniques often bring very detailed reflectance information (i.e. a full BRDF sometimes), their goal is more to replace the use of an expensive gonioreflectometer rather than to be able to change the viewpoint and/or the illumination. Recently, several methods have been developed to take into account the interaction between objects inside a real scene, from a sparse set of photographs [7, 39, 22, 23]. Fournier [14] proposed a different approach but with the use of a single image. However, his technique was limited to perfect diffuse environments and was not be able to take into account specular surfaces. Our method has the similar ambition to recover an approximation of the BRDF of the surfaces from a single image, including the processing of specular, isotropic or anisotropic surfaces. This is extremely difficult to achieve because it is not possible to compute a full BRDF correctly without having several images.

We propose a hierarchical and iterative technique that computes the best possible approximation of a real image, using the error computed between the rerendered image and the real one. Each of the new images is generated by making more and more complex assumptions about the reflectance properties of the real surfaces. It is rendered by a global illumination software that takes into account these reflectance changes (see figure 1). The main advantages of our approach are: it does not need any special device to capture the real image (a classical camera is enough), and it estimates the reflectances of all types of surfaces (including anisotropic mirrors) from a single image without any particular constraint for the viewpoint position, the light sources[1] or the objects orientation. The goal of our method is

[1] In fact, the emittances of the light sources are supposed to be known. However, if it is not the case [14] propose a method to recover them automatically.

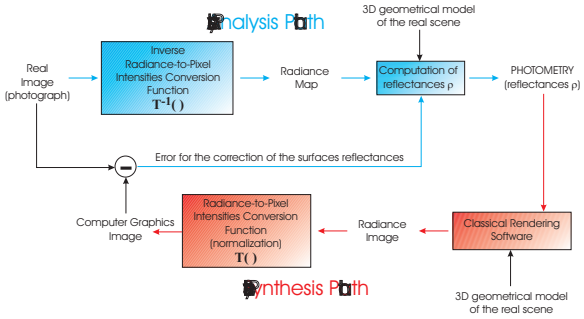

Figure 1: **General Principle of our Method** this figure shows the global scheme of the inverse rendering process. Initial data are: one real image and a 3D geometrical model of the scene.

to recover an approximation of the BRDF of the surfaces, and to compute the best synthetic image preserving the real properties of the scene.

1.2 Organization of the report

The report is organized as follows. In the next section, we discuss previous work related to image-based rendering. Section 3 describes the bases and the tools necessary to our algorithm. In particular we give the notion of group which solves the problem of objects that are not directly seen in the image, and the BRDF model that we use. We also give a short description of the input data and the rendering software that we have developed. Section 4 describes the algorithm in full details using the previously discussed tools. In particular, we explain the methods to process each case of reflectance property separately. In section 5 , we describe one of the advantages inherent in our methodology: the possibility to analyze some surfaces that are not directly seen in the real image, but indirectly through a mirror. Section 6 completes the technical discussion by explaining the optimizations that we have implemented to accelarate the rerendering process. Section 7 shows several results of rerendering, including images containing all kinds of photometric properties. Some applications are given in the domain of augmented reality, including rendering of new images under novel viewpoints, novel illumination conditions and the insertion/removal of objects. The last section gives some conclusions and future research directions.

2 Background and Related Work

All the techniques and ideas in this report have been made possible by works about photorealistic rendering including global illumination and ray tracing, image-based modeling and BRDF modeling. However, the most concerned domains deal with *inverse rendering*,

image-based rendering and *reflectance recovery*. We can split the *reflectance recovery* algorithms into three parts: direct measure of reflectances on the object using a specific device [36, 19, 2, 6], the extraction of reflectances from a set of images [20, 24, 31, 30, 32, 25, 26, 33, 7, 39, 40, 22, 23], and the extraction of reflectances from a single image [29, 28, 14, 11, 27]. The last two parts may be subdivided into two categories, depending on whether the method takes into account energetic interreflections (using a global illumination algorithm for example) or not.

2.1 Reflectance Recovery using a Specific Device

Ward [36] proposed to directly measure the reflectances of an object, using a low-cost device. Ward introduced a device to estimate the five parameters of his anisotropic BRDF model, that he developed for these purposes. Karner et al. [19] presented a simplification of this device.
Baribeau, Rioux and Godin [2] described a method for measuring three reflectance parameters of several objects inside a scene. The diffuse reflectance , the Fresnel Term and the roughness of the objects are estimated using a polychromatic laser range sensor. However, this method is limited to uniform reflectance properties over each object.
[6] suggest using a device containing a robotic manipulator and CCD camera to allow simultaneous measurement of the BTF (Bidirectional Texture Function) and the BRDF of large samples (about 10cm $\times$ 10cm).

2.2 Reflectance Recovery from Several Images

2.2.1 Methods without Global Illumination

Kay and Caelli [20] described a method to compute the surface reflectances using the Torrance-Sparrow light reflection model [34]. They used a depth map and four or eight images obtained with different point light sources. By increasing the number of intensity images, they estimated the parameters of the Torrance-Sparrow's model, reduced to three terms: the diffuse reflection coefficient k_d, the specular reflection coefficient k_s and the rugosity factor c.
Lu and Little [24] did not use any reflection model, but directly estimated the reflectances from the pixel intensities. Nineteen black and white images were captured using a custom device that turns around the object. For each incident angle of light, they built a reflection function, depending on the maximum pixel intensity in the image.
[31] and [30] proposed to register a range map (to get a 3D geometric model) and a set of images of an object, using a 360 degrees rotation device. Next, they extracted the pixel intensities from the images and from the 3D model of the object reprojected onto the images by a Z-buffer algorithm. Theses parameters were used to compute the diffuse component and the specular component separately, i.e. the k_d term of the Lambert's model and the k_s and c terms of a simplified Torrance-Sparrow's reflection model.
[32] needed 120 color images and 12 range maps to compute the Torrance-Sparrow's pa-

rameters, separating the diffuse and the specular component. They recovered the BRDF of highly textured objects (this was impossible to do with previous techniques presented in [20, 24, 31]), and proposed the creation of new images under novel viewpoints and novel illumination conditions.
[25] directly estimated the Lafortune's et al. BRDF [21] of an object from a set of images (30). To obtain the BRDF, the radiance received by the pixels from the object is divided by the irradiance received by this object from the light source. He applied this computation to the rerendering of objects under novel illumination conditions.
Finally, Wong et al. [33] described a method that recovers the reflectance of each pixel of an image, considered as a set of small facets, each one having its own BRDF. The BRDFs are estimated from a set of images taken under different viewpoint and illumination conditions, as the ratio of the pixel intensity divided by the light source intensity. Wong et al. applied their method to the re-illumination of the scene with new light sources.

2.2.2 Methods with Global Illumination

[7] used global illumination for augmented reality applications. To insert new objects inside a real image, he needed to take into account interreflections and compute the reflectances of the surfaces in the part of the scene influenced by this insertion. He created a geometrical 3D model of this part of the scene (called the *local scene*), and calculated manually the reflectance parameters of all the modeled objects. Each of the BRDF parameters are changed by the user iteratively until the rerendered image becomes close enough to the original one.
Yu et al. [39] proposed a complete solution for the recovery of the surfaces BRDF from a sparse set of images captured with a camera (twelve of the 150 images were taken specifically to get specular highlights on surfaces). They built 40 radiance maps for the estimation of the reflectance parameters and the computation of the radiance-to-pixel intensities conversion function (camera transfer function) [8]. Using an image-based modeling software such as *Facade* [9], a 3D geometrical model of the scene was built from the set of images. All these data were then utilized to recover the BRDF of the modeled surfaces. Their method minimized the error on the parameters of the Ward's anisotropic BRDF model [36] to estimate the best possible BRDF for each object. This work was applied to the insertion of new objects in the scene, to the modification of the illumination conditions and to the rendering of a new scene under novel viewpoints. However, this method only works if at least one specular highlight is visible on an object. Otherwise this object is simulated as perfect diffuse.
Loscos et al. [22] proposed a method based on an original idea from Alain Fournier [14]. Their algorithm recovers the diffuse reflectances of the surfaces inside a set of photographs of a scene, taking into account the textures of the objects (each surface has to be unshadowed in at least one image of the set). They applied their technique, to insert/remove objects and to modify the lighting conditions of the original scene (insertion of a new light source for example). More recently, Loscos et al. [23] extended this technique by removing the constraint of the unshadowed surfaces. To improve the results, they transformed their reflectance recovery algorithm into an iterative process. However, the method remained limited to perfect diffuse surfaces (the mirrors are considered to be diffuse textured objects for example).

2.3 Reflectance Recovery from a Single Image

2.3.1 Methods without Global Illumination

K. Sato and K. Ikeuchi [29] described an algorithm for the reflectance recovery of an isolated object from a single image and a 3D geometrical model of this object. They applied some constraints on the light source position and the camera parameters. In addition, they simplified the Torrance-Sparrow reflection model. This way, they estimated separately the diffuse component and the specular component to recover the uniform reflectance of the surface.
More recently, I. Sato et al. [28] proposed to recover the BRDF of an object, using the shadows generated by the surfaces of the scene. They used a single omnidirectional image of the environment and a 3D geometrical description of the surfaces. They developed a 6-step iterative algorithm to minimize the error between the real and the synthetic image with respect to the BRDF parameters of the surfaces.

2.3.2 Methods with Global Illumination

Our approach belongs to this category.
A pioneering work in this domain was completed by Fournier et al. [14] in 1993. He proposed to rerender an original image using a 3D representation of the scene (including the positions of the light source and the camera parameters) and a single image of this scene. All the surfaces are considered as perfect diffuse, and they used their reprojection on the real image to estimate their reflectances. A radiosity-based algorithm then computes an image applying these reflectances to a progressive radiosity technique [4] to obtain a new synthetic image.
An extension of the previous method was developed by Drettakis et al. [11]. They proposed an interactive version of the initial paper and added a vision algorithm for the camera calibration and the 3D geometrical model automatic positioning. They described a slightly different technique for the estimation of the reflectances of the surfaces and they used a hierarchical radiosity algorithm [18] to compute a new synthetic image close to the real one.
An approach similar to Fournier et al.'s was chosen in [27]. It included a feedback that compares the real image to the synthetic one. Gagalowicz [27] described a technique to generate a new synthetic image from a single one (except the 3D geometrical model, which was built from two stereo images) using an iterative method that minimizes the error between the real image and the synthetic one. However, this technique is limited to a pure lambertian approximation of the surface reflectances.

3 Elements of Reflectance Recovery

3.1 The Notion of Group

The inputs of our reflectance recovery algorithm are separated into two categories, the 3D geometrical model of the scene and a single image of this scene captured with a standard

camera. This method is based on the extraction of the object reflectances from the pixels covered by the projection of these objects in the image (as described later in section 4).
Using a single image to recover all the surface reflectances of the scene raised several problems related to the geometrical model and the size of the projection of the objects in the image. First of all, there are generally a lot of surfaces that are not directly visible in the real image. It is then extremely difficult (sometimes impossible) to compute their reflectances because no information is available about them. This is not important if the position of the observer is never changed. However, it is usual to modify this position especially in augmented reality applications. Therefore, we introduce the notion of *group* of objects and surfaces. These *groups* specify the objects and the surfaces which have the same reflectance properties. This is a very fast manual operation left to the user after or during the geometrical modeling process. For example, in figure 2, the 'red cube' was modeled as an object containing six planar surfaces which have the same reflectance properties. Our reflectance algorithm will then use this description to propagate the estimated reflectance from the three visible faces of the cube to the three other ones.
This *group* notion often solves the second modeling problem which could happen during the reflectance estimation. Indeed, the area covered by the projection of some objects in the real image could be too small to give a good approximation to the reflectance of these objects. Therefore, if the user joins these objects with others which have the same reflectance and a bigger projection area in the real image, it becomes possible to obtain a better approximation of their reflectance. However, if there are no other bigger objects, a very rough approximation of the reflectance will be computed for these small objects, and the resulting image may be biased. This problem is inherent in all image-based rendering methods [7, 39, 40, 22, 23, 14, 11, 27] which use the area covered by the projection of an object in the real image to determine its reflectance. Nervertheless, as our method uses a feedback through the comparison between the real and synthetic image, bias is considerably reduced.

3.2 Reflectance Model and Data Description

For the past several years, the construction of a 3D geometrical model from a single image or a set of images has been widely investigated and is known as *image-based modeling* (see [9] for an overview of these methods). In image-based modeling, the 3D geometrical model is built interactively using *Alias|Wavefront*'s *Maya* modeler. The positioning of the full 3D geometrical model of figure 2 took around six hours to be complete, including the recovery of the camera parameters and the light sources positions. Typically, for the camera parameters we use the Dementhon and Davis [10] technique combined with a Simplex minimization method [17]. The light sources have been modeled approximately (because of their complex geometry) and they have been placed manually with a precision of $\pm$ 5cm[2]. Our photometric recovery method is based on the use of Ward's reflectance model [36]. We chose the same BRDF model as [39, 40] because of its small number of parameters and its ability to simulate

[2]Our technique can be used regardless of how the geometry is acquired.

anisotropic surfaces. This model only requires the knowledge of five parameters for a complex BRDF: ρ_d the diffuse reflectance, ρ_s the specular reflectance, $\vec{x}$ the anisotropy direction (called the *brushed direction*) and the anisotropic roughness parameters α_x and α_y (see [36] for a detailed description of this BRDF model). Furthermore, this model avoids the costly computation of the Fresnel term which has been replaced by a normalization factor.

Figure 2: Example of a real image with the superposition of its 3D reconstructed geometrical model (in white)

When the 3D geometrical model (objects, camera and light sources positions) and the photometric model (reflectances and light sources intensity) are determined, it is possible to render a synthetic image using a classical rendering software such as *Radiance* [37]. In this report, we developed our own rendering software called *Phoenix* to obtain a high-performance computing power and to take advantage of the specific architecture of the Silicon Graphics workstations used[3]. *Phoenix* is a global illumination software. It computes the form factors of a progressive radiosity system [4] using a 64 bits A-Buffer [3, 13] mapped on each face of the hemicube [5]. This increases the resolution of each face of the hemicube by a factor of 64 with a negligible increase in computation time, with respect to a classical Z-Buffer software. Moreover, *Phoenix* uses advanced *OpenGL* programming techniques called *offscreen rendering* to compute the index buffers (or *item buffers* [11]) necessary for the extraction of the pixel intensities from the original image and the synthetic one. Each number in the index buffer indicates either a group number, or an object number, depending on whether we need to compute the reflectance of a group or of an object.

[3]This work was carried out on a SGI Octane SI 2× R12000 300Mhz.

4 Inverse Rendering from a Single Image

4.1 Overview of the Algorithm

The core of our technique is incremental and hierarchical (see figure 3). It is incremental because the surface reflectances are going to evolve to their optimum value. It is hierarchical because the general algorithm forces the surfaces BRDF to be more and more complex if the error between the real and the synthetic image does not decrease for these surfaces. This algorithm is iterative and will proceed to successive corrections of the surface reflectances by minimizing the error between the real and the synthetic image. Indeed, each computed error for a group of objects having the same photometric properties drives the correction of their reflectance. Our technique successively applies the selected assumption on the group reflectances until the error became smaller than a user-defined threshold. The notion of threshold and how to fix its value to give them will be discussed in the section 6.

We start the algorithm with the *perfect diffuse* case without considering texture (the diffuse reflectance of a group is computed averaging the radiances covered by its projection in the real image). All the surfaces are then considered as perfectly lambertian, and the global illumination rendering software (*Phoenix* in this case) computes an new approximation of the image. If the difference between the real and the synthetic image for a group is greater than a fixed threshold on all the group projection, then the reflectance of this group is considered as *perfect specular* for the next rerendering iteration. If, after *Phoenix* has recomputed a new image using the new assumption, the error for this group remains big, then its reflectance is simulated as *non-perfect specular*. We apply the same principle to change again the group reflectance to a *both diffuse and specular* one. Until then, all the surfaces were considered with no roughness term (only a ρ_d and a ρ_s were estimated). In the next assumption, if the difference between the two images still produces big errors, they are considered as isotropic and a rugosity factor (α) has to be evaluated. This assumption is extended to anisotropic properties if the user-defined threshold for the error has not been reached. If all assumptions have failed, the group is supposed to be highly textured. Since only a single image is available, it is extremely difficult and sometimes impossible to create a combination between this texture and other reflectance properties (a glossy textured surface for example). This situation is discussed in paragraph 4.7.

4.2 Case of perfect diffuse surfaces

One of the most simple cases of reflectances is the *perfect diffuse* one. During the first inverse rendering iteration, all the objects of the scene are simulated as perfectly diffuse. A diffuse reflectance (ρ_d) is then computed for each group, as the average of radiances covered by the projection of the groups in the original image. This technique is different from [11, 14] because we do not pay attention to the texture of the surfaces. It is interesting to note that some textured surface may be simulated using a pure diffuse reflectance (as shown in figure 14), to create a good visual approximation. This method is very different from [11, 14] because it is not limited to the computation of the average reflectance to produce the new

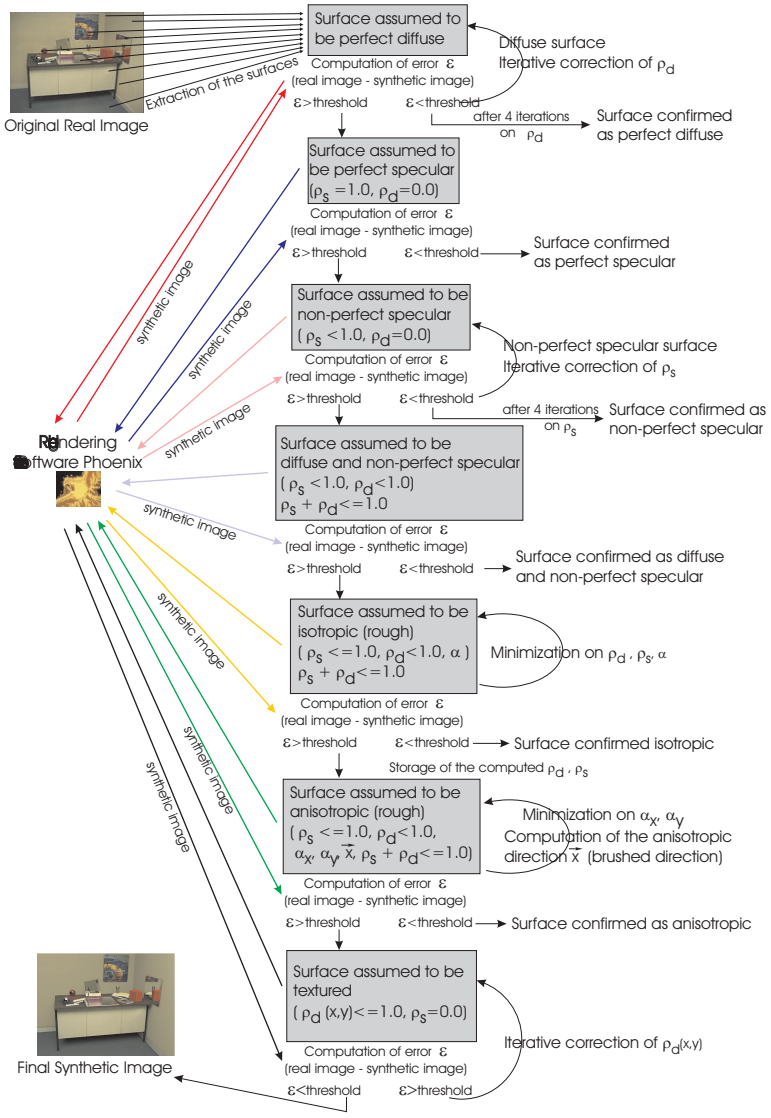

Figure 3: General iterative and hierarchical algorithm for reflectance recovery. Each surface of the scene is analyzed separately, depending on the assumption about its reflectance (perfect diffuse, perfect specular, etc.). If the assumption is false (the error between the real and the synthetic image is big), then the surface reflectance is assumed to be more complex (hierarchical principle). If the assumption is correct then the surface reflectance is modified accordingly in order to minimize the error between the two images (iterative principle).

final synthetic image. We correct this reflectance iteratively until the error between the original and the rerendered image becomes small. For an object, this error is computed as the ratio between the average of the radiances covered by the projection of the groups in the original image, and the average of the radiances covered by the projection of the groups in the synthetic image (see equation 1). These radiances have been obtained using the inverse of the camera transfer function that was simulated as a γ correction function with a γ value of 2.2 according to [35]. However a more sophisticated algorithm could be applied if we had more than one photograph of our scene [8] .

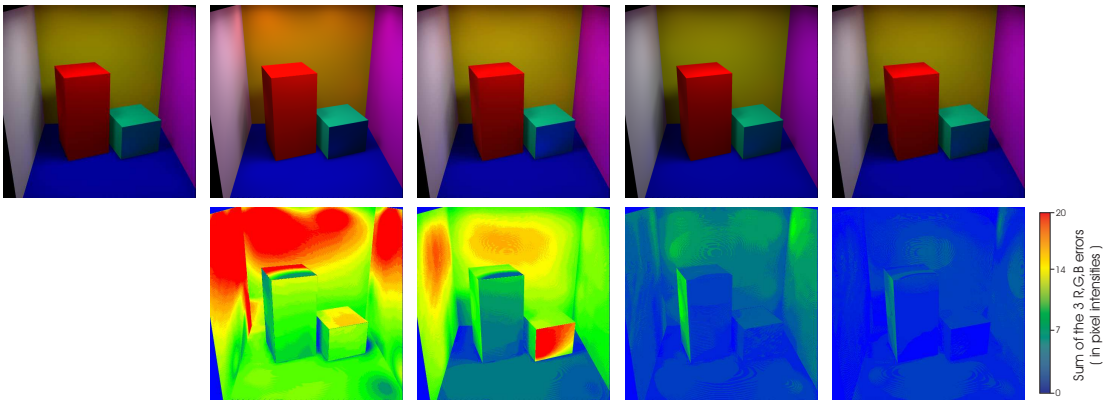

Figure 4: In the top row, from left to right: the original synthetic image (top left) generated using a rendering software was rerendered for 4 iterations (the next four images). The differences between this original image and the regenerated images are shown in the bottom row and displayed using a specific error colormap (at the bottom right). We observe a regular decrease of the error from left to right.

$$\widehat{\varepsilon_j} = \frac{\widehat{B_{org_j}})}{\widehat{B_{new_j}}} = \frac{T^{-1}(\widehat{P_{org_j}})}{T^{-1}(\widehat{P_{new_j}})} \tag{1}$$

where : $\widehat{B_{org_j}}$ and $\widehat{P_{org_j}}$ are respectively the average of the radiances and the pixels covered by the projection of object j in the original image.
$\widehat{B_{new_j}}$ and $\widehat{P_{new_j}}$ are respectively the average of the radiances and the pixels covered by the projection of object j in the synthetic image.
$T()$ is the camera transfer function (a γ correction function here).

Since the diffuse reflectance ρ_{dj} of object j is proportional to the average radiance $\widehat{B_j}$, the iterative correction of the ρ_{dj} can be written as:

$$\rho_{di_{k+1}} = \rho_{di_k} \times \widehat{\varepsilon_i} \tag{2}$$

$$\rho_{di_{k+1}} = \rho_{di_k} \times \frac{\sum_{j=1}^{n_i} f(\widehat{\varepsilon_j}) \cdot (\widehat{\varepsilon_j} \times m_j)}{\underbrace{\sum_{j=1}^{n_i} f(\widehat{\varepsilon_j}) \cdot m_j}_{\neq 0}} \tag{3}$$

$$\text{and } f(\widehat{\varepsilon_j}) = \begin{cases} 0 & \text{if } \widehat{\varepsilon_j} \geq (1+\lambda) \cdot md \\ 1 & \text{else} \end{cases}$$

with :
- k, the current rerendering iteration number.
- $\widehat{\varepsilon_i}$, the total error between the original and the synthetic image for group i.
- $\widehat{\varepsilon_j}$, the total error between the original and the synthetic image for object j.
- n_i, the number of objects for group i.
- md, the median (select the middle value of the sorted samples).
- λ, the authorized dispersion criteria.
- m_j, the number of pixels covered by the projection of object j.

The function $f()$ eliminates problems generated by smaller objects for which the error is very important, because they are more sensitive to the image noise (their projection in the image cover a small amount of pixels). An example of iterative correction of ρ_d is provided by figure 4 on a very simple synthetic scene, nevertheless containing high color bleeding effects (see how the green cube is influenced by the blue floor for example).

As textures are not taken into account in this section, we only consider a diffuse reflectance parameter ρ_d. It could be interesting and maybe faster to directly inverse the radiosity equation as suggested by [39, 40]. If we know the radiances, the emittances and the full geometry (i.e. the form factors), it is possible to directly solve the radiosity equation [16] for the reflectances. However, this is not so simple, because we work with a single image. Consequently, there are some surfaces that are not directly visible in the original

image. Therefore, their radiosities are unknown and it is impossible to guess their values. Thus, we can not inverse the radiosity equation.

4.3 Case of perfect and non-perfect specular surfaces

If the previous diffuse hypothesis about the surface reflectance failed, it is now considered as a*perfect mirror*. It is the easiest case to solve because the diffuse reflectance of a perfect mirror has a null value ($\rho_d = 0$) and its specular reflectance is equal to 1 ($\rho_s = 1$). It is worth noting that there is no need to iterate on the specular reflectance and a new synthetic image can be directly rendered. On the other hand, the reflectance for a non-perfect specular object has to be iteratively modified to obtain an optimum ρ_s. The iterative correction of ρ_s is similiar to equation 3, except ρ_d has to be replaced by ρ_s. An example of the use of the hierarchical algorithm on a scene containing both diffuse, non-perfect specular surfaces is shown in figure 5.

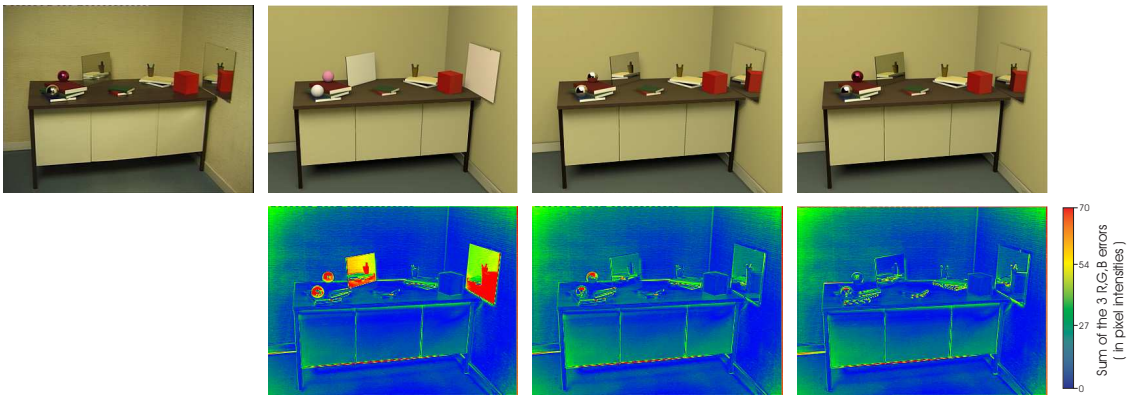

Figure 5: Simulation of hierarchical inverse rendering, where the top row from left to right consists of the real image captured with a camera, the synthetic image with a pure diffuse assumption (first iteration), the synthetic image with perfect diffuse and perfect specular assumptions (fifth iteration) and the synthetic image with pure diffuse and non-perfect specular surfaces (seventh iteration). On the bottom row, we can see the error images corresponding to the difference between the real and the synthetic image.

4.4 Case of both diffuse and specular surfaces with no rugosity factor

At this point of the algorithm, surfaces with big errors are now considered as both diffuse and specular ($\rho_d \neq 0$ and $\rho_s \neq 0$) but still with no roughness.
The differences between the real image and the synthetic are minimized as a function of ρ_d

and ρ_s:

$$(T^{-1}(I_{synth}) - T^{-1}(I_{org}))^2 = \sum_{i=1}^{nbg} (\rho_d \cdot B_d + \rho_s \cdot B_s - T^{-1}(I_{org}))^2 \quad (4)$$

with nbg, the number of pixels covered by the group projection.

This minimization has an analytical solution for each wavelength R, G, B:

$$\begin{pmatrix} \rho_d \\ \rho_s \end{pmatrix} = \begin{pmatrix} \sum_{nbg} B_d T^{-1}(I_{org}) \\ \sum_{nbg} B_s T^{-1}(I_{org}) \end{pmatrix} \begin{pmatrix} \sum_{nbg} B_d^2 & \sum_{nbg} B_d B_s \\ \sum_{nbg} B_d B_s & \sum_{nbg} B_s^2 \end{pmatrix}^{-1} \quad (5)$$

In practice, such surfaces in real cases are very rare but not impossible. For example, the top face of the desk in the figure 14 presents some photometric properties very close to this approximation.

4.5 Case of isotropic surfaces

Until now, all the surfaces were supposed to be without roughness. In the case of an isotropic surface, the diffuse reflectance ρ_d, the specular reflectance ρ_s and a rugosity coefficient α have to be recovered according to Ward's BRDF model. A first idea is to use a classical minimization algorithm to solve for these three paremeters. However, the error function (difference between the real and the synthetic image) for an anisotropic surface is radically different if ρ_d is varying in $]0.0; 1.0[$ (figure 6) or if ρ_d has a null value (figure 7). Directly minimizing the error function for ρ_d, ρ_s and α in the interval $[0.0; 1.0[$ is thus not possible. We propose to miminize the error function using two separate error functions: one for the interval $]0.0; 1.0[$ and the other for the $\rho_d = 0$ particular case. The minimization algorithm (we use the Simplex method [17] for the two minimizations) that provides the smallest error will determine the final value of ρ_d, ρ_s and α.

One of the disadvantages of the method is that it could take a lot of time to minimize such functions. Indeed, these isotropic surfaces use ray-tracing [1] techniques for their correct simulation. Even if optimizations techniques greatly accelerate the rendering [15, 12], it still could take around one hour and fifty minutes to recover the ρ_d, ρ_s and α values (using ten bounced rays for each primary ray (nine per pixel) that reached a glossy surface). In fact, the optimum values of ρ_d and ρ_s are found in only two minutes because the resulting value does not need to be obtained with a precision better than $1 \cdot 10^{-2}$ (the visual difference became imperceptible). On the other hand, α requires a determination with a $1 \cdot 10^{-4}$ precision (according to [36], the α parameters may vary between 0.001 for a quite perfect specular surface to 0.2 for a mostly diffuse surface).
Figure 8 shows the result of these minimizations: the aluminium surface (in the center of image) has been simulated as isotropic, and an optimum value of $\rho_d = 0.0$ and $\rho_s = 1.0$ has been found. However the error image shows that maybe a better approximation seems

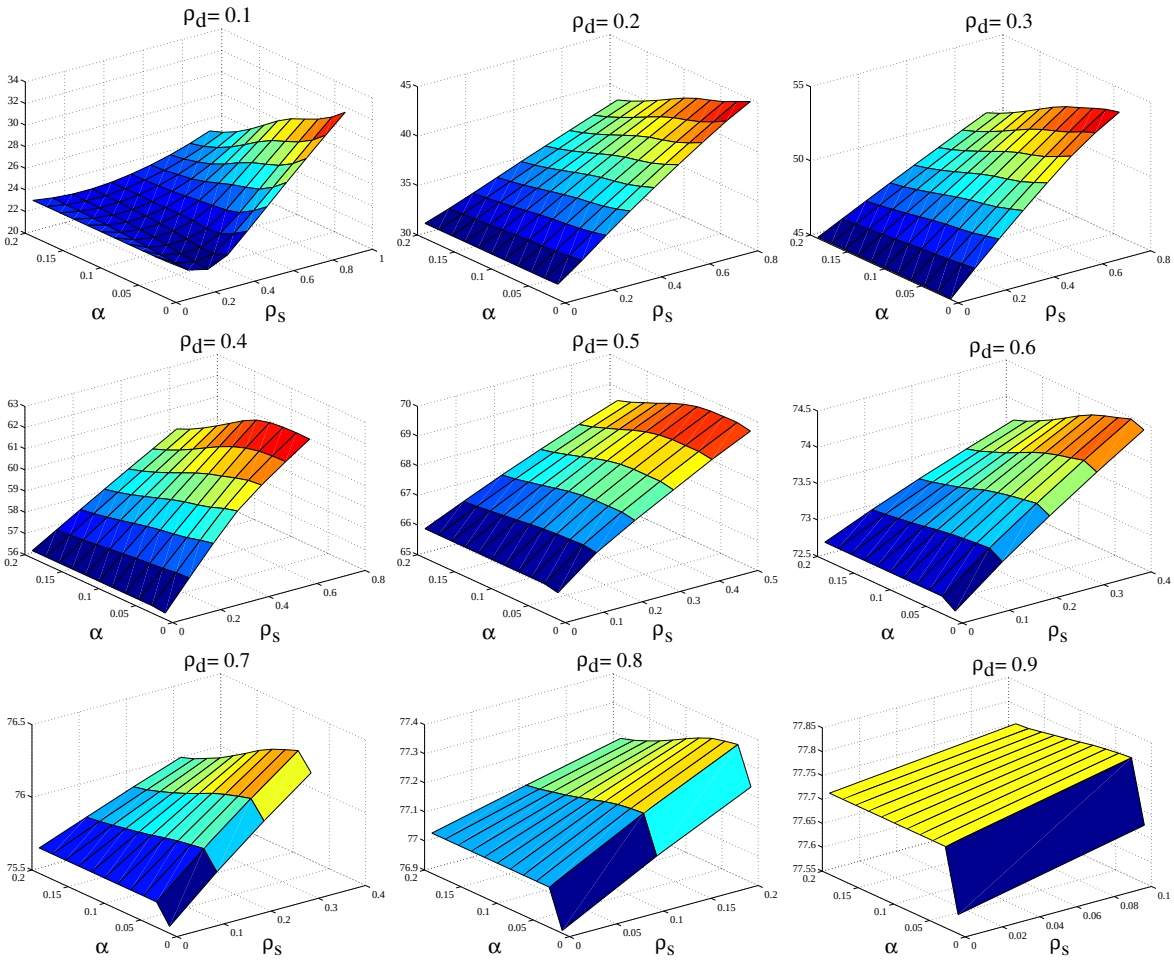

Figure 6: Error function (synthetic image - real image), for a fixed diffuse reflectance, with respect to variations of the isotropic values α, and ρ_s specular reflectance. The evolving steps are 0.018 for α and 0.1 for ρ_s.

to be possible for this particular surface. The error remains important in the extent of the specular reflection area of the two books on this surface. Therefore a more complex BRDF is needed and the algorithms tries now to simulate the surface as an anisotropic one.

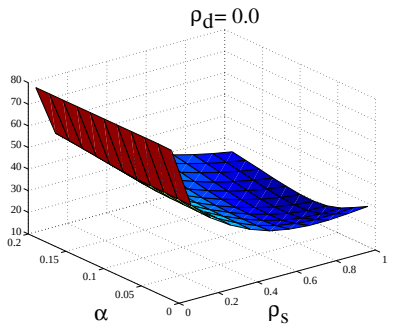

Figure 7: Error function (synthetic image - real image), for $\rho_d = 0$ with respect to variations of α (isotropy value), and ρ_s (specular reflectance).

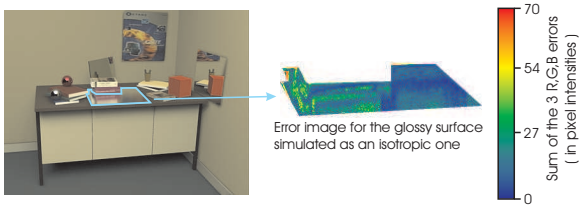

Figure 8: Approximation of the aluminium surface (anisotropic) of the real image (left) by an isotropic surface in the synthetic image (center). The error between these two images for the aluminium surface is visible in the right image. We remark that the error is still important in the area of the specular reflection of the books. The red pixels correspond to a high error but they are not significant because they are coming from an approximative positioning of the 3D geometrical model ont the image, especially on the edges of the objects.

4.6 Case of anisotropic surfaces

This is clearly the most complicated case of our algorithm because the anisotropic model of Ward requires to minimize a function of five parameters: the diffuse reflectance ρ_d, the specular reflectance ρ_s, the anisotropy direction $\vec{x}$ (or *brushed direction* [36]) and the rugosity factors α_x, α_y. However, it is possible to keep the previous ρ_d and ρ_s values computed for the isotropic case: the error functions (see figure 6 and 7) show that the ρ_s parameter is not correlated to the α parameter, because these functions are quite constant with respect

to α. We may then suppose that the ρ_d and ρ_s do no differ from the isotropic case to the anisotropic one.
The error function to minimize has now three parameters left (see figure 9). We remark on this figure that for a given rotating angle of the vector $\vec{x}$ and varying values of α_x and α_y, this error function presents several minima. This confirms that a standard minization algorithm will probably not find a global minimum.
Regarding these considerations, we have computed the four images corresponding to the four smallest minima found by a Simplex minimization algorithm (figure 10). It is interesting to note that the rerendered images remain far from the original one and that the error is bigger than for the isotropic case. This brings us to the conclusion that the data for the minimization procedure are not well conditioned. Therefore, we propose to determine the anisotropy vector $\vec{x}$ directly from the real image.

In a first step, we consider the anisotropic surface as a perfect mirror and compute a synthetic image. Next, we estimate the difference between the real image and the synthetic one to visualize the part of the anisotropic mirror where the specular reflection is "extended". This area corresponds to an attenuation of the specular reflection, and this effect is always very important in the direction perpendicular to the brushed direction (or anisotropy direction). In a second step, we compute an index buffer for this mirror of all the surfaces visible through it. We then look for a reference surface that has the biggest reflection area on the anisotropic surface, while being as close as possible to it. This surface is then selected in a such manner that the ratio $\frac{\text{Area(reflected surface)}}{d(S,P)}$ is maximized (with d(S,P), the euclidean distance between the center of gravity of the selected surface and the center of gravity of the anisotropic mirror). The motivation of this choice resides in the fact that surfaces very far from the anisotropic object exhibit a reflection pattern that is too small or too noisy to be usable for the recovery of the brushed direction. In a third step, the anisotropy direction is sampled creating $\vec{x}$ vectors around the normal to the anisotropic surface. Each of these sampled directions determine a direction to traverse the error image and compute the average of the standard error deviations computed in the error image. Finally, the algorithm selects the direction for which this average value is the smallest one (see figure 11). The figure 12 summarizes the complete procedure.
Once the anisotropy direction $\vec{x}$ has been recovered, a Simplex minimization algorithm is used to estimate the rugosity parameters α_x and α_y. Typically, for the synthetic image in the lower right corner of the figure 14, it tooks 50 iterations and $2h30$ to recover the full BRDF of the anisotropic surface. The algorithm found an optimum anisotropy vector for a rotation angle of 0 degrees and then minimized the error function of the upper left corner of the figure 9. The estimated values of α_x and α_y were 0.01 and 0.062 respectively.

4.7 Case of textured surfaces

When the simulation of a surface as anisotropic still produces big errors in the difference image, we proceed to texture extraction.
Extracting the texture from the real image is an easy task that can be realized using the

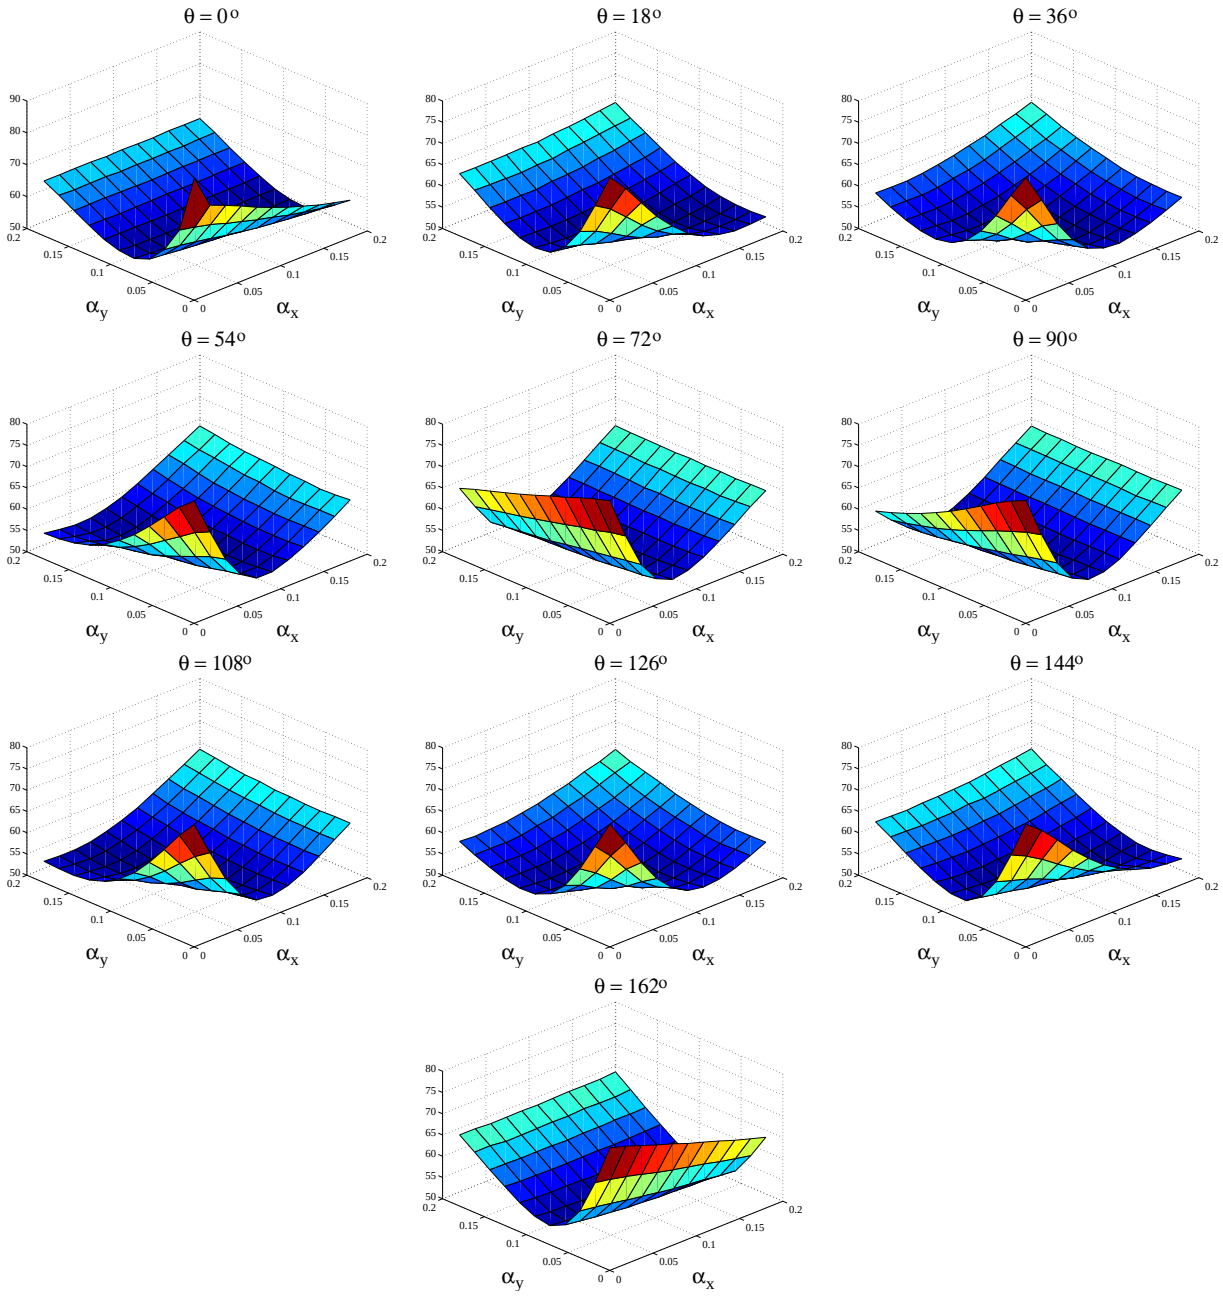

Figure 9: Error function (synthetic image - real image), for different anisotropy directions $\vec{x}$ (the vector is turned around the normal to the surface using a step of 18 degrees) with respect to variations of the roughness parameters α_x, α_y (with a step of 0.018). The diffuse reflectance and the specular reflectance terms have been estimated during the isotropy analysis.

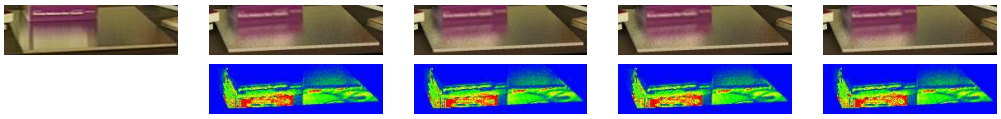

Figure 10: The first image (top left) is the original one (reduced here to the interest area). The next four images have been produced using the four smallest minima found by the minimization algorithm. We can see that all of these images are far from the original one (the vertical black line on the white book (see figure 8) has disappeared from the specular reflection) and that a lot of details have been smoothed. The error colormap remained the same as on figure 8.

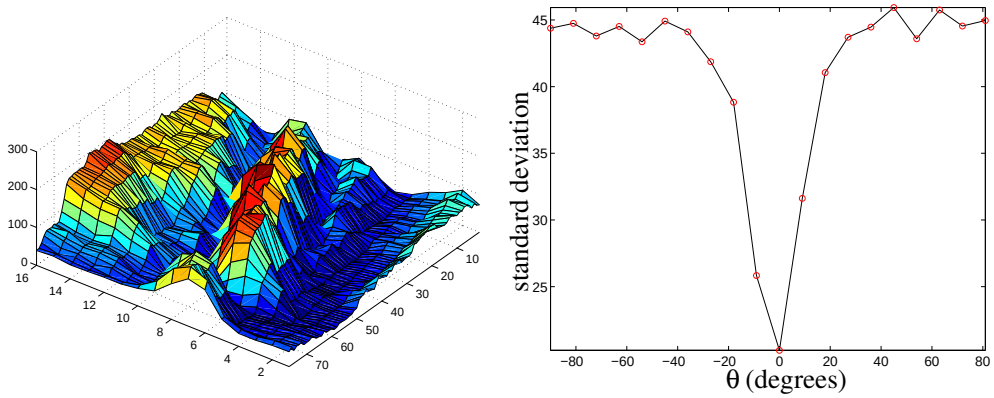

Figure 11: The selected object used here to recover the anisotropy direction is the violet book of the lower left real image of figure 14. The 3D surface (left image) shows the error image for the difference between the perfect specular reflection area of this selected object, and its corresponding area in the real image. The 2D curve (right) shows the average of the standard error deviations computed from the error image along the sampled anisotropy directions (see also figure 12).

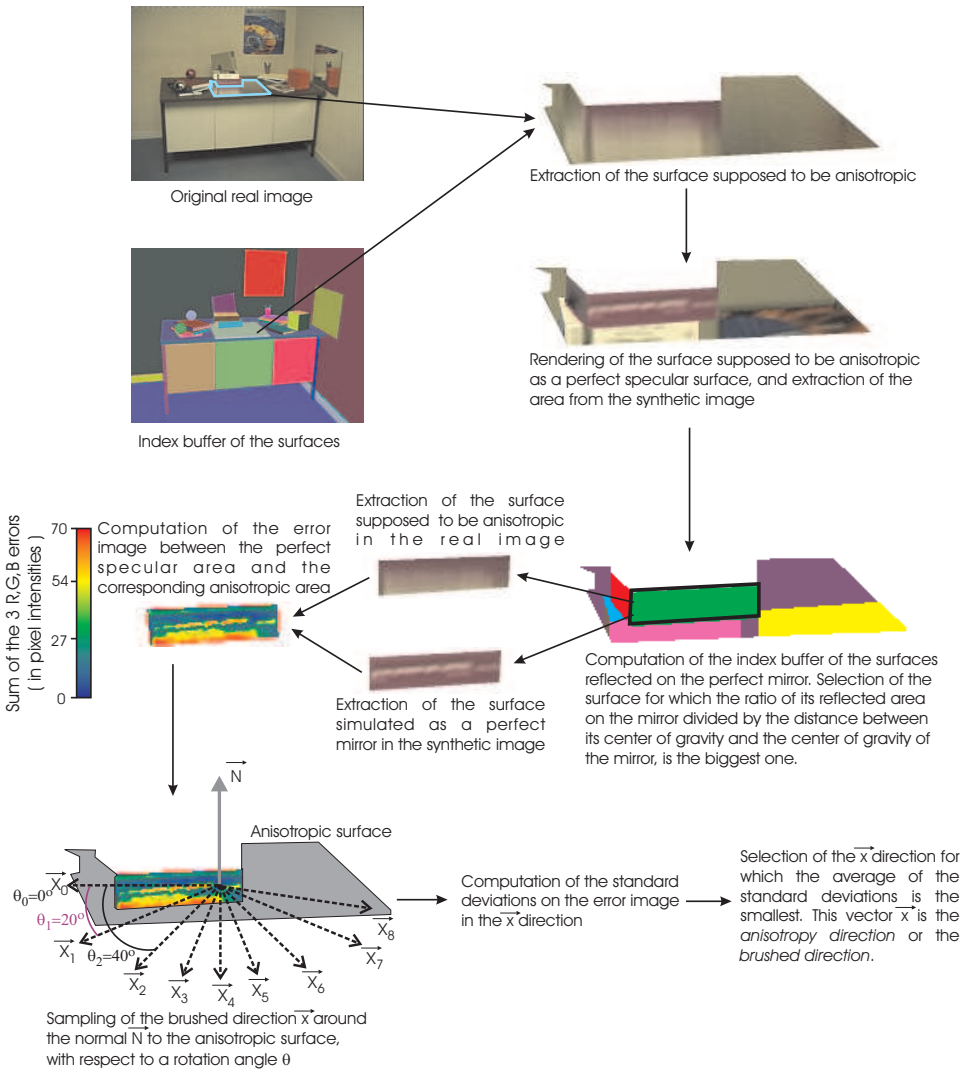

Figure 12: Computation method of the anisotropy direction $\vec{x}$ for a glossy surface.

technique proposed by [38] for example. However, we have to extract this texture while taking into account the fact that it already has received the energy from the light sources, and that the pixels covered by its projection in the real image contain this information. Otherwise, if we send the energy of the light sources to these textures again, they will be over-illuminated. Therefore, we introduce here a notion called *radiosity texture* that balances the extracted texture with an intermediate texture in order to minimize the error between the real and the synthetic image. As for the perfect diffuse reflectance case, this intermediate texture is computed by an iterative method.
At the first iteration, the texture used to rerender the image is the texture directly extracted from the real image. At the second iteration, the texture used to obtain the resulting synthetic image is multiplied by the ratio between the newly extracted texture of this synthetic image and the texture of the real image. This iterative process stops when the user-defined threshold for textured surface has been reached. The textures of the poster and the books in the rerendered images of section 7 have been obtained using this technique.

The problem of this method is that it computes a texture including the shadows, the specular reflections and the highlights. Typically, suppose that we have a marbled floor on which a sphere is reflected. The texture of this floor in the real image then includes the marble characteristics, its reflectance properties and the sphere reflection including its own reflectance properties. How to extract the marble characteristics only and independently of the rest of the scene ? This is an extremely hard problem, and according to [32] no algorithm has been proposed yet to solve it using a single image.

5 Advanced Analysis of Reflectances

Our inverse rendering procedure gives the opportunity to analyze the reflectances of some surfaces that are not directly seen in the original image. Indeed, if a surface is detected and confirmed as a perfect or non-perfect specular one, we can extend our reflectance recovery algorithm to the surfaces that are seen through this mirror in the real image.
First of all, the index buffer of the groups visible through the mirror are computed using a ray tracing algorithm. If there exists a surface in this index buffer that was not visible before in the real image, then its reflectance is computed taking into account the current assumption made for its group reflectance (the surface has the same photometric properties as its group). In the next iteration, this reflectance is balanced by the mirror reflectance (if it is a non-perfect one), and it is then considered for the correction of the group reflectance (see figure 13).
To our knowledge, this is the first time that an image-based rendering technique deliberatley exploits mirror surfaces to enhance the BRDF recovery process in a scene.

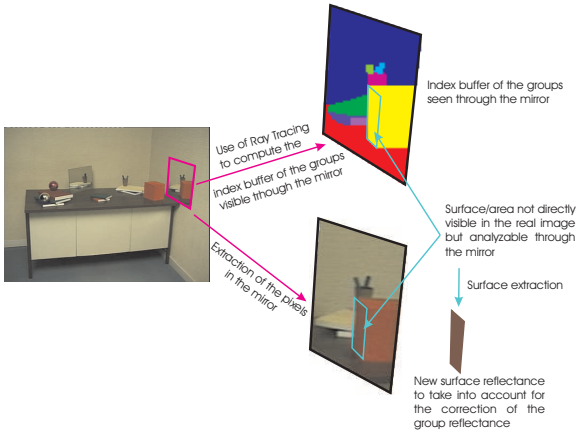

Figure 13: Example of surface not directly seen in the original image. The reflectance of this surface becomes computable through the mirror put on the right wall. This surface belongs to the 'cube' group and now contributes to the estimation of its reflectance. If this face had a specific reflectance (new group), it would be computable only at this point.

6 Optimizations and Determination of Thresholds

Since the reflectance recovery algorithm takes around two hours to simulate an isotropic surface, and two more hours in an anisotropic case, this means that all textured surfaces (which is the final hypothesis after the isotropy and the anisotropy tests) will need four hours to be correctly estimated. This is not acceptable when a lot of surfaces are textured in a real image, but the computing time could be greatly reduced if we can find that the surface is textured before treating the isotropic case. Therefore we introduced a heuristic to solve this problem. It is related to the choice of the thresholds that determine if a surface is correctly simulated. Indeed, after each synthetic image has been produced, it is compared to the real one using a user-defined error threshold to know if the assumptions made about the surface reflectances are correct. For the presented images, the following thresholds were used to produce the results of the section 7. Such thresholds are not critical to the behaviour fo the technique but only affect speed.
From the case of the perfect diffuse assumption up to the isotropic one, the sum of the three R,G,B errors coming from the difference between the real and the synthetic image must have a value smaller than 5%. However, during the non-perfect specular assumption, if the error is greater than 50%, we can directly avoid the isotropic and the anisotropic cases and so greatly increase the performance of the algorithm. We do not have a formal characterization of this optimization, but in practice it seems to work well (see section 7). The isotropic-to-anisotropic threshold has been chosen equal to 1%, to ensure that the algorithm tries the anisotropic case. On the other hand, the threshold used to come to a texture assumption

is equal to 5%. Finally, the last threshold is a global threshold that forces all the groups in the synthetic image to have an error smaller than 5%.

7 Results

All the following synthetic images have been generated using *Phoenix* as rendering and inverse rendering software. The first synthetic image at the top right of figure 14 has been generated in 37 minutes using the hierarchical algorithm, from the left real photograph. Two specular surfaces have been recovered and simulated as non-perfect mirrors. Neither isotropic nor anisotropic hypothesis have been tried thanks to the optimization technique described in section 6, and 14 rerendering iterations were necessary to generate the final image.
The inverse algorithm tooks 4 hours and 40 minutes to produce the image at the bottom right of figure 14. Roughly 4 hours of this time were necessary to recover the anisotropic BRDF of the aluminium surface. The final rendering stage took 32 minutes to render the final image (100 bounced rays have been used for the anisotropic surface).
The images of figure 15 show examples of applications in augmented reality. Some synthetic objects have been added such as a small robot and a luxo-like desk lamp. It is also possible to modify the reflectances easily too. New viewpoints can be generated and new illumination conditions can be created as well.

8 Conclusion and Future Work

In this report, we have presented a new technique allowing to determine an approximation of the reflectance properties of the surfaces of a 3D scene. An incremental and hierarchical algorithm iteratively estimates the diffuse, specular, isotropic and anisotropic reflectance parameters. In a final step, the textured surfaces are considered as a special case of reflectances to be simulated. The method takes as input a single photograph of the scene taken under known illumination conditions as well as a 3D geometric model of the scene. The result is a complete description of the photometric properties of the scene which may be used to produce a photorealistic synthetic image very similar to the real one. We showed that the method is robust and gives the opportunity to visualize the original scene from any new angle, with any illumination conditions and with the addition, removal and modification of objects.
Our work has currently some limitations, especially regarding textured surfaces. Till now, we are not able to discriminate the shadows or highlights from an assumed textured surface. In this regard, it will be interesting to extend our method to these cases, although we think that this is a very difficult problem, if one sticks to the single image assumption.
While many challenges remain, we believe that algorithms for recovering an approximation of the reflectances inside a real scene are an important direction of research for both Computer Vision and Computer Graphics communities. In Computer Vision, it could be possible

Figure 14: Two different examples of synthetic images (right) rerendered from a single real image (left). We remark that the perfect diffuse assumption is realistic enough for many surfaces (including the walls, the floor, the desk, etc.).

Figure 15: Examples of several augmented reality applications. All these new images were rendered using our global illumination software *Phoenix*, which first recovered the surface reflectances from the bottom left image of figure 14. The top left image shows the original scene removing some objects (the feet of the desk and the red cube). Note that the right mirror has taken into account the modification. The right top image shows the original scene rendered under a novel viewpoint. The bottom left image shows the scene with modified photometric properties, and the addition of an object (a small robot). The bottom right image presents the scene under novel illumination conditions with the addition and deletion of objects.

for example to use our method to enhance the positioning of mirrors using a minimization algorithm between the real and the synthetic image. Regarding Computer Graphics, we may extend the reflectance recovery algorithm to objects that have more complex photometric properties such as light beam, small fires, caustics, etc. The hierarchical property of our technique offers a lot of possible extensions.

References

[1] Arthur Appel. Some techniques for shading machine renderings of solids. *AFIPS 1968 Spring Joint Computer Conf.*, 32:37—45, 1968.

[2] R. Baribeau, M. Rioux, and G. Godin. Color reflectance modeling using a polychromatic laser range sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):263–269, February 1992.

[3] Loren Carpenter. The a-buffer, an antialiased hidden surface method. *Computer Graphics (Proceedings of SIGGRAPH 84)*, 18(3):103–108, July 1984. Held in Minneapolis, Minnesota.

[4] M. F. Cohen, D. P. Greenberg, D. S. Immel, and P. J. Brock. An efficient radiosity approach for realistic image synthesis. *IEEE Computer Graphics and Applications*, 6(3):26–35, March 1986.

[5] Michael F. Cohen and Donald P. Greenberg. The Hemi-Cube: A radiosity solution for complex environments. In B. A. Barsky, editor, *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 31–40, August 1985.

[6] 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, January 1999. ISSN 0730-0301.

[7] Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In Michael Cohen, editor, *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, pages 189–198. Addison Wesley, July 1998.

[8] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In Turner Whitted, editor, *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, pages 369–378. Addison Wesley, August 1997.

[9] Paul Ernest Debevec. *Modeling and Rendering Architecture from Photographs.* PhD thesis, University of California, Berkeley, 1996.

[10] D.F DeMenthon and L. Davis. Model-based object pose in 25 lines of code. In *Second European Conference on Computer Vision(ECCV)*, pages 335–343. Springer-Verlag, May 1992.

[11] George Drettakis, Luc Robert, and Sylvain Bougnoux. Interactive common illumination for computer augmented reality. In Julie Dorsey and Philipp Slusallek, editors, *Eurographics Rendering Workshop 1997*, pages 45–56. Springer Wien, June 1997.

[12] Frédo Durand. *Visibilité tridimensionnelle: étude analytique et applications.* Ph.D. thesis, Université Joseph Fourier, Grenoble, France, July 1999.

[13] E. Fiume, A. Fournier, and L. Rudolph. A parallel scan conversion algorithm with anti-aliasing for a general purpose ultracomputer. *Computer Graphics (Proceedings of SIGGRAPH 83)*, 17(3):141–150, July 1983. Held in Detroit, Michigan.

[14] Alain Fournier, Atjeng S. Gunawan, and Chris Romanzin. Common illumination between real and computer generated scenes. In *Graphics Interface '93*, pages 254–262. Canadian Information Processing Society, May 1993. Held in Toronto, Ontario, Canada.

[15] Andrew Glassner. *An Introduction to Ray Tracing.* Academic Press, New York, NY, 1989.

[16] Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, pages 213–222, July 1984.

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

[18] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. *Computer Graphics (Proceedings of SIGGRAPH 91)*, 25(4):197–206, July 1991.

[19] Konrad F. Karner, Heinz Mayer, and Michael Gervautz. An image based measurement system for anisotropic reflection. *Computer Graphics Forum*, 15(3):119–128, August 1996.

[20] G. Kay and T. Caelli. Inverting an illumination model from range and intensity maps. *CGVIP: Image Understanding*, 59:183–201, 1994.

[21] Eric P. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. In *Computer Graphics (ACM SIGGRAPH '97 Proceedings)*, volume 31, pages 117–126, 1997.

[22] C. Loscos, M. C. Frasson, G. Drettakis, B. Walter, X. Grainer, and P. Poulin. Interactive virtual relighting and remodeling of real scenes. Available from www.imagis.imag.fr/Publications RT-0230, Institut National de Recherche en Informatique en Automatique (INRIA), Grenoble, France, April 1999.

[23] Celine Loscos, George Drettakis, and Luc Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, 2000. To appear.

[24] J. Lu and J. Little. Reflectance function estimation and shape recovery from image sequence of rotating object. In *International Conference on Computer Vision*, pages 80–86, June 1995.

[25] Stephen R. Marschner and Donald P. Greenberg. Inverse lighting for photography. In *Proceedings of the Fifth Color Imaging Conference*. Society for Imaging Science and Technology, November 1997.

[26] 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 Dani Lischinski and Greg Ward Larson, editors, *Eurographics Rendering Workshop 1999*. Eurographics, June 1999.

[27] A. Rosenblum. *Data Visualization*, chapter Modeling Complex indoor scenes using an analysis/synthesis framework (André Gagalowicz). Academic Press, 1994.

[28] Imari Sato, Yoichi Sato, and Katsushi Ikeuchi. Illumination distribution from brightness in shadows: Adaptive extimation of illumination distribution with unknown reflectance properties in shadow regions. In *Proceedings of IEEE ICCC'99*, pages 875–882, September 1999.

[29] Kosuke Sato and Katsushi Ikeuchi. Determining reflectance properties of an object using range and brightness images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1139–1153, 1991.

[30] Yoichi Sato and Katsushi Ikeuchi. Temporal-color space analysis of reflection. *Journal of Optical Society of America*, 11(11):2990–3002, November 1994.

[31] Yoichi Sato and Katsushi Ikeuchi. Reflectance analysis for 3d computer graphics model generation. *Graphical Models and Image Processing*, 58(5):437–451, 1996.

[32] Yoichi Sato, Mark D. Wheeler, and Katsushi Ikeuchi. Object shape and reflectance modeling from observation. In Turner Whitted, editor, *Computer Graphics, Proceedings of SIGGRAPH 97*, pages 379–388. Addison Wesley, August 1997.

[33] Siu-Hang Or Tien-Tsin Wong, Pheng-Ann Heng and Wai-Yin Ng. Image-based rendering with controllable illumination. In Julie Dorsey and Phillip Slusallek, editors, *Rendering Techniques '97 (Proceedings of the Eighth Eurographics Workshop on Rendering)*, pages 13–22, New York, NY, 1997. Springer Wien. ISBN 3-211-83001-4.

[34] K.E. Torrance, E.M. Sparrow, and R.C. Birkebak. Polarization, directional distribution, and off-specular peak phenomena in light reflected from roughened surfaces. *Journal of Optical Society of America*, 56(7):916–925, 1966.

[35] Jack Tumblin and Holly Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

[36] Gregory J. Ward. Measuring and modeling anisotropic reflection. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 265–272. ACM Press, July 1992.

[37] Gregory J. Ward. The radiance lighting simulation and rendering system. In Andrew Glassner, editor, *Proceedings of SIGGRAPH 94*, Computer Graphics Proceedings, Annual Conference Series, pages 459–472. ACM Press, July 1994. ISBN 0-89791-667-0. Held in Orlando, Florida.

[38] George Wolberg. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, 1990.

[39] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination : Recovering reflectance models of real scenes from photographs. In A. Rockwood, editor, *Computer Graphics (SIGGRAPH '99 Proceedings)*, volume 19, pages 215–224. Addison Wesley Longman, August 1999.

[40] Yizhou Yu. *Modeling and Editing Real Scenes with Image-Based Techniques*. PhD thesis, University of California, Berkeley, 2000.

Contents

Unité de recherche INRIA Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)
Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399