REFLECTANCE ESTIMATION FROM URBAN TERRESTRIAL IMAGES: VALIDATION OF A SYMBOLIC RAY-TRACING METHOD ON SYNTHETIC DATA

Fabien Coubard [a], Mathieu Brédif [a], Nicolas Paparoditis [a], Xavier Briottet [b]

[a] MATIS, Institut Géographique National, 73 avenue de Paris F-94160 SAINT MANDÉ, France
(fabien.coubard, mathieu.bredif, nicolas.paparoditis)@ign.fr
[b] DOTA, ONERA, 2 avenue Edouard Belin 31055 TOULOUSE, France

Working Groups III/4, III/5

KEY WORDS: radiometry, reflectance, BRDF, ray-tracing, terrestrial images

ABSTRACT:

Terrestrial geolocalized images are nowadays widely used on the Internet, mainly in urban areas, through immersion services such as Google Street View. On the long run, we seek to enhance the visualization of these images; for that purpose, radiometric corrections must be performed to free them from illumination conditions at the time of acquisition. Given the simultaneously acquired 3D geometric model of the scene with LIDAR or vision techniques, we face an inverse problem where the illumination and the geometry of the scene are known and the reflectance of the scene is to be estimated. Our main contribution is the introduction of a symbolic ray-tracing rendering to generate parametric images, for quick evaluation and comparison with the acquired images. The proposed approach is then based on an iterative estimation of the reflectance parameters of the materials, using a single rendering pre-processing. We validate the method on synthetic data with linear BRDF models and discuss the limitations of the proposed approach with more general non-linear BRDF models.

1 INTRODUCTION

1.1 Context

The development of digital cameras has led to new possibilities for radiometric processes: the linear response of the sensor with respect to incoming radiance gives radiometer abilities to the camera. So understanding the physical processes that lead to the formation of digital images is a key to developing many applications, for professionals as well as for the general public. For instance, terrestrial images are nowadays widely available on the Internet: Google Street View in 2007, Microsoft Bing Streetside in 2009 (North America only), or the research project iTowns (Devaux and Paparoditis, 2010). These images have very high spatial resolution (a few centimeters); so the radiative phenomena must be modelized with the same resolution. In this work, we focus on images acquired from terrestrial vehicles, with large spectral bands in the visible domain. The seamless visualization of such images has to deal with the dependence of the pixel values on the illumination conditions at the time of acquisition. For instance, the illuminated and shadowed areas are determined by that time, and can be a disturbance for the viewer. On the long run, the objective is to free images from their illumination conditions, enabling applications such as relighting or augmented reality. That can be done through the estimation of the intrinsic color properties of the scene surface: the reflectance.

1.2 Related work

The estimation of one or several radiometric parameters of a scene, from a set of observations (mainly images), is referred to as inverse rendering in computer graphics. Patow and Pueyo (2003) propose a survey of these methods. As part of it, inverse reflectometry aims at retrieving reflectance properties of objects from one or several images. The hypotheses and input data of these works can be very different: global or only direct illumination, controlled or uncontrolled light sources, invariant or textured reflectance map, single or multi-angular. Machida et al. (2007) propose an inverse method using photon mapping, with a parametric model of BRDF (Bidirectional Reflectance Distribution Function), from images with very directional light source with known position. Yu et al. (1999) use inverse radiosity in a multi-scale hierarchical method, from a large set of images from different camera positions. Boivin and Gagalowicz (2001) estimate BRDF from a single image. Lensch et al. (2003) estimate spatially-varying BRDF without considering reflections between objects. The main goal of these methods is often to produce augmented reality from real images, such as addition of new objects or relighting under user-defined lighting conditions.

In the remote sensing community, some authors follow a more physical approach, looking for accurate estimation of properties, mainly analyzing outdoor aerial or satellite images for classification purposes. A physical modelization of the different radiance terms in urban images has been proposed by Miesch et al. (2000). Using multi-view aerial images, Martinoty (2005) uses the same modelization for classification of roof materials with BRDF parameters as criteria, and Lachérade et al. (2008) estimates the albedo (in the Lambertian case) for accurate synthesis of satellite images under any atmospheric conditions.

1.3 Proposed approach

In this work, the input data is a set of terrestrial images shot from a mobile-mapping vehicle. We also suppose that we have a 3D geometric model of the scene, segmented in areas that are homogeneous in reflectance. That segmentation could be obtained with the LIDAR backscattering data (independent of illumination) simultaneously acquired by the vehicle, but that topic is not tackled in this work (see (Müller et al., 2007) for an example of facade modeling). We also consider the data are acquired under good atmospheric conditions which are supposed to be known through observable variables for aerosol (type, concentration) and gas (water vapour content). The method aims at enabling radiometric corrections of terrestrial images (shadow removal, relighting, etc), and intends to be as general as possible, though adapted to outdoor urban scenes. Furthermore, we do not look for purely image-based methods, such as tonemapping; but

we instead try to use a physical model of the image radiometry, in order to explore the potential of physically-based methods. To summarize, our method uses the following input data: known but uncontrolled light sources (sun, sky), a 3D geometric model of the scene segmented into regions of homogeneous reflectance properties, and single or multi-view calibrated images. It takes into account the reflections between objects and the bidirectional behaviour of materials (up to the limitations discussed in section 4.3). This paper presents the physical background of that study, in section 2, and the simulation method used in section 3. Then the inversion algorithm is presented in section 4, as well as the motivations for symbolic ray-tracing. Results on a synthetic scene are then shown in section 5, as well as the computational limitations of the inversion method.

2 RADIATIVE TRANSFER IN URBAN SCENES

An image is the result of the propagation of light rays from the light sources to the sensor. This propagation of a light ray is affected by the medium it goes through and the surfaces it hits. In the case of an outdoor urban scene: the source is the sun, the medium is the atmosphere, and the surfaces are the objects of the scene: buildings, soil... So the modelization of the radiometric pixel values requires the modelization of the radiative transfer in the atmosphere and the reflection on the objects of the scene.

2.1 Participating medium

The solar light interacts with the atmospheric particles (gas molemolecules or aerosol) located at point Q. Part of its energy is absorbed and scattered, in a proportion given by the attenuation coefficient σ_e. Along a unit vector $\vec{\omega}$, the radiance going out of the particle at wavelength λ is given by the radiative transfer equation (Meyzonnette and Lépine, 1999):

$$L(Q+\delta\vec{\omega}, \vec{\omega}, \lambda) = L(Q, \vec{v}, \lambda) - \sigma_e\left(L(Q, \vec{\omega}, \lambda) - J(Q, \vec{\omega}, \lambda)\right)\delta$$
(1)

where δ is a small length, and J is the source function that takes into account the scattering of the light hitting the particle from other directions than $\vec{\omega}$ and volumetric light emission.

2.2 Interaction with surfaces

When hitting a surface at point P, the light is reflected according to its Bidirectional Reflectance Distribution Function (BRDF), that links the outgoing radiance L_{out} in a given direction of reflection $\vec{\omega}_{out}$ to the radiance distribution L_{in} coming from the incident directions $\vec{\omega}_{in}$ (Nicodemus et al., 1977). Considering a non-emitting surface, and denoting by Ω_P the hemisphere directed by the surface normal $\vec{n}_p$, this relation is (Kajiya, 1986):

$$L_{out}(P, \vec{\omega}_{out}, \lambda) =$$
$$\int_{\vec{\omega}_{in}\in\Omega_P} brdf_P(\vec{\omega}_{in}, \vec{\omega}_{out}, \lambda) \cdot L_{in}(P, -\vec{\omega}_{in}, \lambda)\langle\vec{\omega}_{in}\cdot\vec{n}_p\rangle d\vec{\omega}_{in}$$
(2)

The high dimension of the parameter space of BRDF distributions prevents its accurate modelization with a limited number of parameters. Many models exist, each fitted for certain materials; for instance the models of Torrance and Sparrow (1967), Cook and Torrance (1981) and Ward (1992) are widely used in computer graphics (Yu et al., 1999; Machida et al., 2007). For reasons of space-complexity of the inversion method (see section 4.3), and at the expense of physical reliability, we propose to limit our search space to the particular class of kernel BRDFs.

We denote a **kernel BRDF** as a linear combination of parameterless BRDF basis functions $f_i^{seg(P)}(P, \vec{\omega}_{in}, \vec{\omega}_{out})$ weighted by geometry-agnostic functions $\rho^{seg(P)}(\lambda)$ which are the unknown constants for each scene surface segment of homogeneous reflectance ($seg(P)$ denotes the index of the scene segment containing P).

$$brdf_P(\vec{\omega}_{in}, \vec{\omega}_{out}, \lambda) = \sum_i \rho_i^{seg(P)}(\lambda) \cdot f_i^{seg(P)}(P, \vec{\omega}_{in}, \vec{\omega}_{out})$$
(3)

For instance, the following kernel BRDF blends a Lambertian model (isotropy of the outgoing radiance; corresponding to a perfectly diffuse behaviour), weighted by ρ_d, and a perfectly specular model (mirror behaviour), weighted by ρ_s :

$$brdf_P(\vec{\omega}_{in}, \vec{\omega}_{out}, \lambda) = \frac{\rho_d(\lambda)}{\pi} + \rho_s(\lambda).\delta_{\{bisect(\vec{\omega}_{in}, \vec{\omega}_{out})=\vec{n}_P\}}$$
(4)

2.3 The specific case of urban terrestrial images

High spatial resolution urban terrestrial imagery imply specific conditions that prevent us from using the hypotheses commonly assumed in aerial and satellite remote sensing. First, the objects of the scene cannot be considered as equally distant from the camera, so the scene cannot be modelized as a plane (as in the radiometric equalization proposed by Chandelier and Martinoty (2009)). This also means that the reflections between objects have to be taken into account. However, the reflected light may be neglected after a few reflections, because its radiance decreases with the number of reflections (due to the geometric attenuation, the transmittance, the reflectance). The importance of reflections decreases also with the square of the distance of the reflecting surface. Lachérade et al. (2008) show that it may be neglected after two reflections.

The Lambertian hypothesis can be valid with low or medium resolution images, because each pixel integrates the contribution of a relatively large surface: the roughness, that is relative to the resolution, usually is high enough to observe an almost perfectly diffuse reflectance. This implies that the BRDF does not depend on the angles of incidence and reflection, so it makes it possible to compute the irradiance on each surface without saving the information of incident angle for each ray. This might be very convenient for further inversion (Lachérade et al., 2008). Unfortunately, the Lambertian hypothesis is no longer valid when the image has a resolution of a few centimeters (Martinoty, 2005).

The transmittance along each rectilinear portion of a ray can be computed simply in the case of a homogeneous medium without scattering, by multiplying the radiance by $e^{-\sigma_e d}$ where σ_e is the attenuation coefficient and d the length of the rectilinear portion. But in general, there is scattering and inhomogeneous medium, and we have to stochastically divert rays while they propagate. That complete computation, though possible, might be very costly. However, a simplification can be adopted: the atmosphere within the convex hull of the 3D model is almost perfectly transparent, due to the low distances between sensor and objects. This hypothesis may be numerically validated; the extinction coefficient σ_e is linked to the horizontal visibility V (that is the greatest distance at which an object can be seen): $e^{-\sigma_e V} = 0.02$. For an atmosphere with a visibility of 23 km, the transmittance at a distance of 30 m is 0.995. So the absorption and the scattering inside the scene may be neglected. Particularly, it implies that fog or fumes at the street level are not taken into account. Furthermore, the size of the scene can be considered as negligible compared to distance at which occurs the atmospheric

scattering; the downward radiance of the sky coming from a given direction can be considered as constant on the whole scene.

3 DIRECT PROBLEM: IMAGE SYNTHESIS

3.1 Ray-tracing

Computing the total radiance analytically is untractable in the general case; we have to use numerical methods to estimate the radiance at the sensor level (Miesch et al., 2000). These rendering methods have been developed for realistic image synthesis; see (Dutre et al., 2006) for a description of the different rendering methods. In this work, we chose ray-tracing method because of its genericity: it can simulate complex BRDF models, and takes into account any scattering or reflection or absorption phenomenon (within the limits of geometrical optics). Furthermore, it allows any kind of geometry, since it only uses the local normals to the surfaces and ray/surface intersection tests; no discretization of surfaces is needed.

The principle of ray-tracing is to use the local equations of radiative transfer (1) and (2) to compute the exitant radiance for every point hit by a light ray. To compute the radiance at the sensor level for a pixel of the image, we throw many rays from the camera in the FOV of the pixel. When hitting a surface, a ray is reflected in a stochastically sampled direction, until it hits a light source. This produces a light path C, made of 3D line segments, which carries a radiance $L(C)$ and a sampling probability $pdf(C)$. The Monte-Carlo estimate (5) gathers the radiances $L(C_k)$ carried by all the paths C_k generated by rays thrown from a pixel FOV into the radiance value of the pixel $L_{sensor}^{synthetic}$, and may even be slightly modified to take into account the point spread function (PSF) of the sensor (Pharr and Humphreys, 2004):

$$\frac{1}{N} \sum_{k=1}^{N} \frac{L(C_k)}{pdf(C_k)} \underset{N \to \infty}{\to} L_{sensor}^{synthetic} \qquad (5)$$

3.2 Implementation details

The criteria of choice of a ray-tracer are: an unbiased estimation, a reasonable computation time (real-time is however out of scope), open source software to allow modifications, a good control through physically pertinent parameters. For these reasons, we chose to use the GPL physically-based renderer LuxRender, that is based on PBRT (Physically-Based Ray-Tracing) (Pharr and Humphreys, 2004).

In order to fasten the convergence of the estimation (5), LuxRender uses importance sampling to drive the stochastic choice of the reflection direction; the BRDF values are used as a probability density function (pdf). When evaluating the radiance coming directly from the light sources, the pdf takes into account the BRDF and the angular distribution of light sources (multiple importance sampling (Veach, 1997)).

To provide an unbiased estimate, LuxRender emulates the sampling of paths with an unbounded number of reflections (using a Russian roulette technique). We however chose to limit the maximum number of allowed reflections N_{max} to 3 to prevent the combinatorial explosion of our technique, driven by the observation in section 2.3 that the reflected light may be neglected after a few interactions.

As discussed in section 2.3, the downward radiance from the sky can be represented as a function of the direction. An environment map is thus sufficient to describe the radiance coming from the
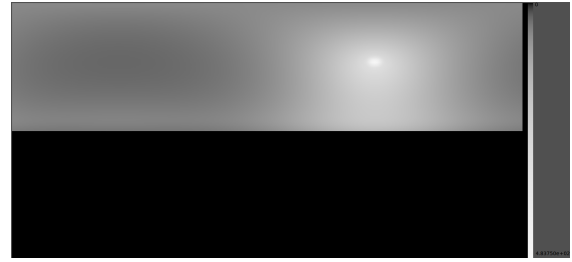

Figure 1: An example of environment map computed with 6S, at wavelength $550nm$. The solar halo is visible, but this environment map does not represent the direct solar radiance.
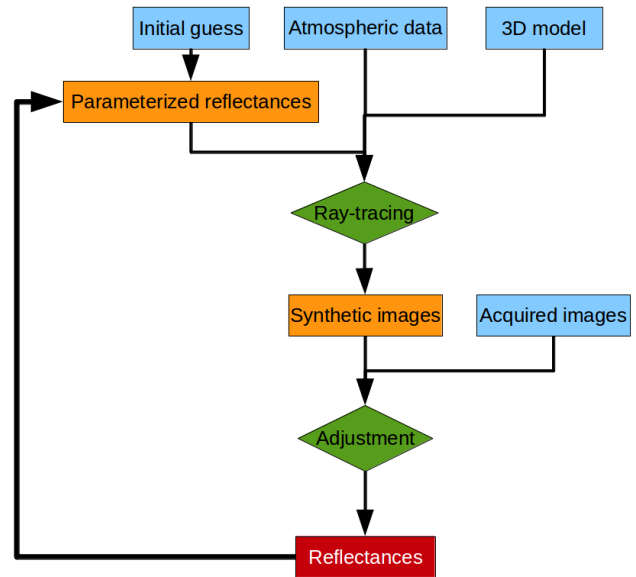

Figure 2: Principle of reflectance inversion.

sky, i.e. a Look-Up Table (LUT) giving for each direction the value of the downward radiance from the sky (figure 1). That downward radiance can be computed by a radiative transfer code, such as 6S (Vermote et al., 2002). That environment map and the irradiance at top of atmosphere and the optical thickness of the atmosphere are sufficient to model the light sources and the participating medium in our case.

4 INVERSE PROBLEM: REFLECTANCE ESTIMATION

4.1 General principle and plain algorithm

We face an inverse problem which direct problem has been discussed in the previous section. We consider that the unknown properties are the reflectances of the materials that compose the scene. That means that lighting conditions are supposed to be known, as well as the geometric 3D model (including the geometric calibration of the camera). The estimation of the reflectance from the images is a complex inverse problem in the general case (non-Lambertian materials, any geometry), because the total irradiance at a point P depends on the reflectance of the other surface elements of the scene, through the reflected irradiance. What is more, the directional effects of reflectance have to be modelized by a parametric function, the BRDF, that is not linear. So we use a minimization algorithm to estimate the parameters of the BRDF function of every material, in an iterative way, as shown on figure 2. The function to be minimized is the difference between
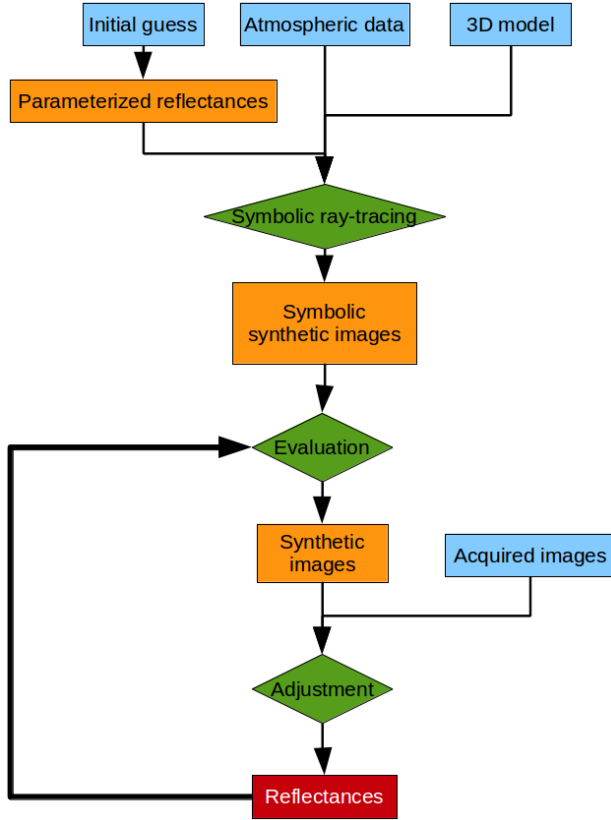
Figure 3: Inversion algorithm with symbolic ray-tracing.

acquired and synthetic sensor images:

$$\epsilon = \left(\sum_{image\ j} \sum_{pixel\ i} [L_{sensor}^{synthetic}(i,j,\boldsymbol{\rho}) - L_{sensor}^{acquired}(i,j)]^p \right)^{\frac{1}{p}} \tag{6}$$

This difference is the cost function of the set of parameters $\boldsymbol{\rho} = (\rho_i)_i$ of the BRDF models (*e.g.* albedos of the Lambertian materials). The input data is composed of the atmospheric properties (defining lighting conditions), the segmented geometric 3D model and initial reflectance guesses. They may be based on prior knowledge of the scene materials or on a preliminary inversion under a Lambertian hypothesis.

4.2 Symbolic ray-tracing for inversion

A major difficulty appears while minimizing (6): it requires the evaluation of the radiance image with different sets of parameters, naively leading to numerous computationally intensive ray-tracing renderings. For instance, for a gradient descent optimization relying on finite difference estimates, images not only have to be rendered for the current parameter values, but also for each dimension of the parameter vector $\boldsymbol{\rho}$. We can however observe that for a given camera position, scene geometry and illumination, the radiance can be expressed as a function of the BRDF parameters. Thus, by keeping the rendered radiance image as a function of the BRDFs parameters instead of a numerical value, the costly ray-tracing rendering may only be performed once. That is what we refer to as **symbolic ray-tracing**. For instance, if we assume that the materials of the scene are Lambertian, their BRDFs are each described by a single parameter: the albedos ρ_j, and $\dim(\boldsymbol{\rho})$ counts these Lambertian materials. Thanks to equations (1) and

(2), the radiance carried by any light path C is of the form:

$$L(\boldsymbol{\rho}, C) = \alpha \prod_{j=1}^{\dim(\boldsymbol{\rho})} \rho_j^{e_j} \quad , \quad \text{with} \quad \begin{cases} \alpha \in \mathbb{R}, \quad e_j \in \mathbb{N} \\ \sum_j e_j \leq N_{max} \end{cases} \tag{7}$$

where α is a coefficient depending on the geometry, the atmospheric transmission and the power of the light source where the path originates. The exponent e_j is equal to the number of reflections of the path on material j (and thus 0 if that material is not hit by the path C). The radiance of a pixel being the Monte-Carlo sum (5) of the weighted contributions of many light paths, it may be expressed as a multivariate polynomial function:

$$L_{sensor}^{synthetic}(\boldsymbol{\rho}) = \sum_k \beta_k \cdot L(\boldsymbol{\rho}, C_k) = \sum_k \gamma_k \prod_{j=1}^{\dim(\boldsymbol{\rho})} \rho_j^{e_{j,k}} \tag{8}$$

It is easy to show that if the scene reflectances are all kernel BRDFs and not only Lambertian, the pixel radiance $L_{sensor}^{synthetic}$ is still a multivariate polynomial (8), where $\dim(\boldsymbol{\rho})$ is now the total number of BRDF kernel terms f_i (3).

Once the symbolic image has been computed, it is easier to minimize (6) with a standard minimization algorithm: the cost function and its derivative can be evaluated directly and quickly for any set of parameters. The inversion algorithm only needs a single ray-tracing rendering, as shown on figure 3. Furthermore, the minimization process avoids the instability induced by the change in the stochastic noise that would be implied by the several rendering passes of the plain algorithm (figure 2). Besides drastically speeding up the optimization process, symbolic ray-tracing thus improves its numerical stability.

4.3 Symbolic ray-tracing tractability

The number N_e of $(e_j)_j$ series satisfying constraints (7) is the number of N_{max} combinations with repetitions of $\dim(\boldsymbol{\rho}) + 1$ elements, known as the multiset coefficient $\binom{\dim(\boldsymbol{\rho})+N_{max}}{N_{max}}$. The space complexity of storing the multivariate polynomial function $L_{sensor}^{synthetic}(\boldsymbol{\rho})$ (8) is thus bounded by N_e since each of its monomial term corresponds to a unique series $(e_j)_j$. A monomial weight is thus simply a sum of γ_k values corresponding to identical series $(e_{j,k})_j$. For instance, for a scene with 10 different materials and at most 3 reflections, the maximal number of monomial terms in $L_{sensor}^{synthetic}(\boldsymbol{\rho})$ is $N_e = 286$.

In the general case of non-kernel BRDFs, this bounded size refactoring is however not possible, and thus the space-complexity of $L_{sensor}^{synthetic}(\boldsymbol{\rho})$ grows with each new sampled path of the Monte Carlo estimate (5). For instance the models from (Torrance and Sparrow, 1967), (Cook and Torrance, 1981) or (Ward, 1992), whose exponential term, containing both a geometric coefficient and a model parameter, prevents the factorization of the $L(C_k)$ terms in (5). As Monte-Carlo integration relies on many rays to decrease the stochastic noise; using non factorizable BRDF models becomes untractable, which drove our restriction to kernel BRDFs.

4.4 Implementation details

We implemented the symbolic ray-tracing in LuxRender, by keeping the expression of the reflectance every time a ray hits the surface of an object. The symbolic expressions are handled through the GPL library GiNaC. For minimizing (6), we chose a conjugate gradient method, using the GPL scientific library GSL. Furthermore, we chose $p = 2$ for the norm of the difference (6), because it is the lowest degree such that ϵ^p is a multivariate polynomial (hence without absolute differences and differentiable).

(a) Reference image

(b) Final Reconstructed image (Lambertian BRDF)

(c) Image with initial parameters

(d) Reconstructed image after 20 iterations

(e) Final reconstructed image (1000 iterations)

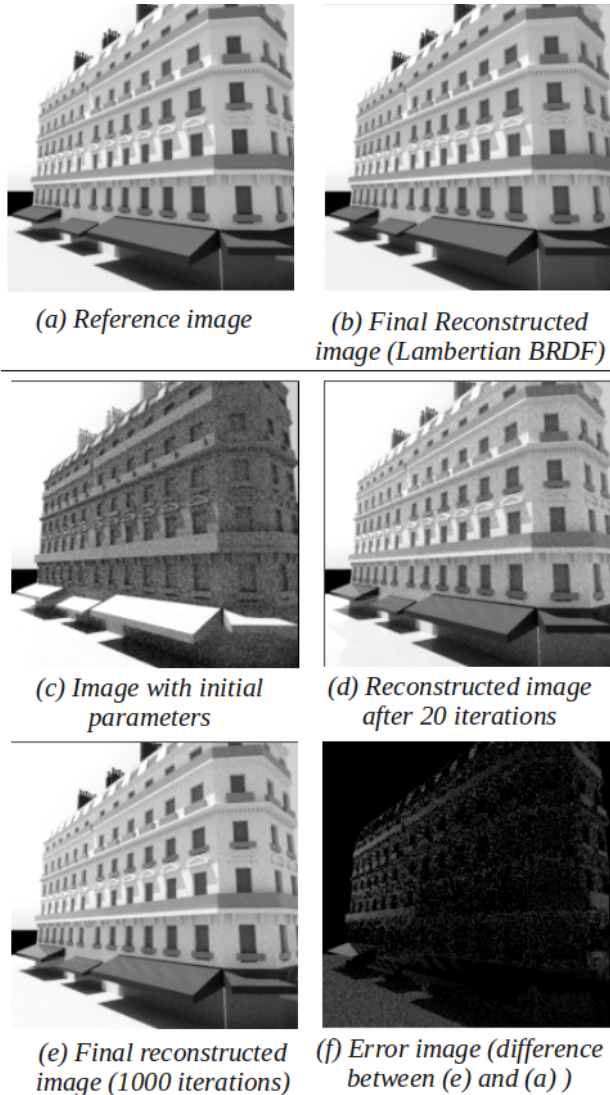(f) Error image (difference between (e) and (a))

Figure 4: Reference image (a), reconstruction with Lambertian BRDF (b), and convergence of the estimation with kernel BRDF.

5 RESULTS OF THE INVERSION METHOD WITH SYNTHETIC DATA

5.1 Scene description and methodology

The inversion algorithm is tested on a synthetic urban scene (figure 4), at wavelength $\lambda = 550nm$. The geometry is described by a very detailed 3D model, and the material have a Lambertian BRDF or a BRDF with Lambertian and specular kernels (4) (for walls and windows). The environment map describing atmospheric scattering is computed with 6S. A reference image of the synthetic scene is simulated with LuxRender. A second simulation is made with symbolic ray-tracing, creating a symbolic image from the same point of view. The numerical and symbolic images cannot be simulated with the same ray-tracing pass, because it would lead to correlation between the stochastic noise due to the Monte-Carlo integration; that could enhance the inversion results without qualifying the method itself. Furthermore, the symbolic ray-tracing uses the numerical values of the BRDF parameters to drive importance sampling, so these values must be set arbitrarily.

Object	Reference albedo	Estimated albedo	Error
Roof	0.2	0.285	0.085
Wall	0.8	0.761	-0.039
Balcony	0.3	0.291	-0.009
Asphalt	0.2	0.201	0.001
Awning	0.05	0.0496	-0.0004
Window	0.9	0.230	-0.670

Table 1: Inversion results on a synthetic urban scene, with a single image, assuming a Lambertian BRDF model

Object	Param.	Ref. value	Estim. value	Error
Roof	ρ_d	0.2	0.390	0.190
Wall	ρ_d	0.72	0.762	0.042
	ρ_s	0.08	0.326	0.246
Balcony	ρ_d	0.3	0.284	-0.016
	ρ_s	0	0.373	0.373
Asphalt	ρ_d	0.2	0.218	0.018
	ρ_s	0	0.000	0.000
Awning	ρ_d	0.05	0.046	-0.004
	ρ_s	0	0.300	0.300
Window	ρ_d	0.18	0.185	0.005
	ρ_s	0.72	0.494	-0.226

Table 2: Inversion results on a synthetic urban scene, with a single image, assuming a Lambertian+Specular kernel BRDF model

5.2 Parameter estimation

Lambertian BRDF. We first perform the symbolic simulation assuming that the materials have a Lambertian BRDF. The diffuse part of the reference BRDFs are correctly estimated (table 1), except for materials that are not directly seen in the image (such as roofs), as well as for materials with a high specular component (such as windows).

Lambertian+Specular BRDF. In order to retrieve the specular behaviour of materials, we make a symbolic simulation assuming that all materials have a Lambertian+specular kernel BRDF (4). The results of the parameter estimation are shown in table 2. The total process (symbolic ray-tracing and minimization) takes about 20 min, with images of 200x200 pixels. Though the value itself of the parameters can be estimated with high errors, the image reconstructed with these parameters is visually close to the reference image (figure 4). This is due to the ill-posedness of the inverse problem when trying to estimate the specular component, which is not always directly seen by the camera for each material. For instance there is no point light reflection visible in the windows, but only a large surface light (the sky). There is therefore a strong correlation between the effects of the specular and diffuse components.

6 CONCLUSION AND FUTURE WORK

In this work, we have developed an inversion algorithm to estimate the reflectance of materials of an outdoor scene, assuming known atmospheric conditions, and a known and segmented geometric 3D model. The geometric complexity of the 3D model only impacts the computing time of the symbolic ray tracing during the preprocessing step. The inversion method uses symbolic ray-tracing, therefore it can perform reflectance estimation using a single rendering pass. It proves to yield good results with synthetic data for Lambertian BRDF models, but has difficulties for

estimating the specular components of Lambertian+specular kernel BRDFs. The estimated reflectances of the materials can be used to reconstruct images of the scene from any point of view.

To improve the results, we could use the redundant information of multi-view acquisition (that is the case for terrestrial acquisitions with a vehicle). It cannot be done for complex scenes without improving our current code, because it is highly memory-consuming. However, as the image reconstructed from the estimated parameters with Lambertian models is close to the reference image, the algorithm should be tested on a real set of urban terrestrial images, to qualify the reliability of the Lambertian hypothesis for relighting purpose.

From a more theoretical point of view, the space-complexity of the symbolic ray-tracing can be limited by using of a decomposition of BRDF in an adapted function basis, leading to a kernel BRDF model; see (Rusinkiewicz, 1997) for a review on BRDF decomposition. That decomposition can better represent the physical behaviour of materials, that includes non-delta specular and retro-specular lobes. But the high number of parameters introduced by the decomposition leads to an ill-posed problem, and cannot be used in our context.

Another strong hypothesis used in this work is the segmentation of the 3D model into homogeneous regions; textured segments may thus not be processed. For instance, (Yu et al., 1999) consider a spatially-varying diffuse albedo, while the specular part is constant per area. This can lead to a better modelization of the materials of a real scene, which are not homogeneous, but increases the number of parameters to be estimated.

ACKNOWLEDGMENTS

References

Boivin, S. and Gagalowicz, A., 2001. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques, pp. 107–116.

Chandelier, L. and Martinoty, G., 2009. A radiometric aerial triangulation for the equalization of digital aerial images and orthoimages. Photogrammetric Engineering and Remote Sensing.

Cook, R. and Torrance, K., 1981. A reflectance model for computer graphics. In: Proceedings of the 8th annual conference on Computer graphics and interactive techniques, ACM, pp. 307–316.

Devaux, A. and Paparoditis, N., 2010. Increasing Interactivity in Street View Web Navigation Systems. ACM Multimedia.

Dutre, P., Bala, K. and Bekaert, P., 2006. Advanced global illumination. AK Peters Ltd.

Kajiya, J., 1986. The rendering equation. In: Proceedings of the 13th annual conference on Computer graphics and interactive techniques, ACM, pp. 143–150.

Lachérade, S., Miesch, C., Boldo, D., Briottet, X., Valorge, C. and Le Men, H., 2008. ICARE: A physically-based model to correct atmospheric and geometric effects from high spatial and spectral remote sensing images over 3D urban areas. Meteorology and Atmospheric Physics 102(3), pp. 209–222.

Lensch, H., Kautz, J., Goesele, M., Heidrich, W. and Seidel, H., 2003. Image-based reconstruction of spatial appearance and geometric detail. ACM Transactions on Graphics (TOG) 22(2), pp. 234–257.

Machida, T., Takemura, H. and Yokoya, N., 2007. Inverse reflectometry for real objects with diffuse and specular interreflections. Electronics and Communications in Japan (Part II: Electronics) 90(1), pp. 50–60.

Martinoty, G., 2005. Reconnaissance de matériaux sur des images aériennes en multirecouvrement, par identification de fonctions de réflectances bidirectionnelles. PhD thesis, Université Paris 7.

Meyzonnette, J. and Lépine, T., 1999. Bases de radiométrie optique. Cépaduès.

Miesch, C., Briottet, X., Kerr, Y. and Cabot, F., 2000. Radiative transfer solution for rugged and heterogeneous scene observations. Applied Optics 39(36), pp. 6830–6846.

Müller, P., Zeng, G., Wonka, P. and Van Gool, L., 2007. Image-based procedural modeling of facades. ACM Transactions on Graphics 26(3), pp. 85.

Nicodemus, F., Richmond, J., Hsia, J. and Ginsberg, I., 1977. Geometrical considerations and nomenclature for reflectance.

Patow, G. and Pueyo, X., 2003. A survey of inverse rendering problems. In: Computer graphics forum, Vol. 22number 4, pp. 663–687.

Pharr, M. and Humphreys, G., 2004. Physically based rendering: from theory to implementation. Morgan Kaufmann.

Rusinkiewicz, S., 1997. A Survey of BRDF Representation for Computer Graphics. Technical report, Princeton University.

Torrance, K. and Sparrow, E., 1967. Theory for off-specular reflection from roughened surfaces. Journal of the Optical society of America 57(9), pp. 1105–1114.

Veach, E., 1997. Robust Monte Carlo methods for light transport simulation. PhD thesis, Stanford University.

Vermote, E., Tanré, D., Deuzé, J., Herman, M. and Morcette, J., 2002. Second simulation of the satellite signal in the solar spectrum, 6S: an overview. Geoscience and Remote Sensing, IEEE Transactions on 35(3), pp. 675–686.

Ward, G., 1992. Measuring and modeling anisotropic reflection. In: Proceedings of the 19th annual conference on Computer graphics and interactive techniques, ACM, pp. 265–272.

Yu, Y., Debevec, P., Malik, J. and Hawkins, T., 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In: Proceedings of the 26th annual conference on Computer graphics and interactive techniques, pp. 215–224.