

Relighting Photographs of Tree Canopies

Marcio Cabral, Nicolas Bonneel, Sylvain Lefebvre and George Drettakis

Abstract—We present an image-based approach to relighting photographs of tree canopies. Our goal is to minimize capture overhead; thus the only input required is a set of photographs of the tree taken at a *single* time of day, while allowing relighting at any other time. We first analyze lighting in a tree canopy both theoretically and using simulations. From this analysis, we observe that tree canopy lighting is similar to volumetric illumination. We assume a single-scattering volumetric lighting model for tree canopies, and diffuse leaf reflectance; we validate our assumptions with synthetic renderings.

We create a volumetric representation of the tree from 10-12 images taken at a single time of day and use a single-scattering participating media lighting model. An analytical sun and sky illumination model provides consistent representation of lighting for the captured input and unknown target times. We relight the input image by applying a ratio of the target and input time lighting representations. We compute this representation efficiently by simultaneously coding transmittance from the sky and to the eye in spherical harmonics. We validate our method by relighting images of synthetic trees and comparing to path-traced solutions. We also present results for photographs, validating with time-lapse ground truth sequences.

Index Terms—Image based rendering, relighting

1 INTRODUCTION AND MOTIVATION

TREES are a very compelling and important component of many outdoor images. Matting photos of trees is common practice, e.g., in photomontages for urban planning projects or in other image editing applications. One important problem in such contexts is that the lighting conditions of the tree do not correspond to the lighting of the final image (see Figure 1(b)-(c)). We introduce a novel approach which allows us to take a small set of photos of a tree at a *single* time of day, and then relight the tree canopy with a different, target lighting condition (see Figure 1(d)-(e) and Figures 14, 15, 21).

Relighting photographs is a long-standing goal of computer graphics. In many previous methods (e.g., [1], [2], [3], [4], [5], [6]) geometric and photometric reconstruction and/or capture are required, followed by different kinds of inverse (global) illumination computations. More recently, relighting research has concentrated on complex, often expensive, capture setups (e.g., [7], [8]) typically involving multiple light sources and cameras.

It is unclear how tree canopies can be relit with such methods. Despite recent work on image-based tree reconstruction [9], [10], pixel-accurate geometric reconstruction of tree canopies is very hard. Even if a synthetic tree were to be reconstructed from images, it is unlikely that photometric, geometric and lighting calibration would be sufficiently accurate to convincingly mix the synthetic tree with the photograph. On the other hand, solutions based on laboratory capture setups are

inappropriate for trees. To our knowledge, no previous method has been proposed that allows relighting of tree canopy photos, using a *single* lighting condition as input.

Given the above, our goals are as follows:

- 1) Allow relighting of a given tree canopy from a given input photograph;
- 2) Provide an *image-based* method, i.e., that only takes photographs as input. No geometry or material information is required for the tree canopy;
- 3) Require photographs to be taken *only at a single time of day*; no input is thus required for the target relit time.

These goals have the advantage of allowing us to relight at any target time, given input photographs at a single time of day. However, a consistent analytic model must exist to represent illumination at the input time (which is captured), and at the target time (for which no data is required). As a result, methods based e.g., on light probes, which would require capturing lighting both at input and target times, are incompatible with our goals.

We do not deal with relighting of the rest of the photograph, such as shadow removal and relighting of parts of the scene other than the canopy. In future work, our results could potentially be integrated with one of the previous approaches mentioned above to handle this more general problem.

The proposed tree canopy relighting method is based on a volumetric single-scattering participating media approximation; it is image-based and does not use or generate leaf geometry or normals to achieve relighting; in this sense it should not be confused with a traditional rendering algorithm.

We build our approach based on an intuitive argument of qualitative similarity between volumetric single scattering and the effect of lighting distribution in a tree canopy. We use synthetic renderings of trees to illustrate

• M. Cabral, N. Bonneel, S. Lefebvre and G. Drettakis are with REVES / INRIA Sophia Antipolis, France. Nicolas Bonneel is currently at UBC in Vancouver, Canada.
E-mail: {firstname.lastname}@inria.fr

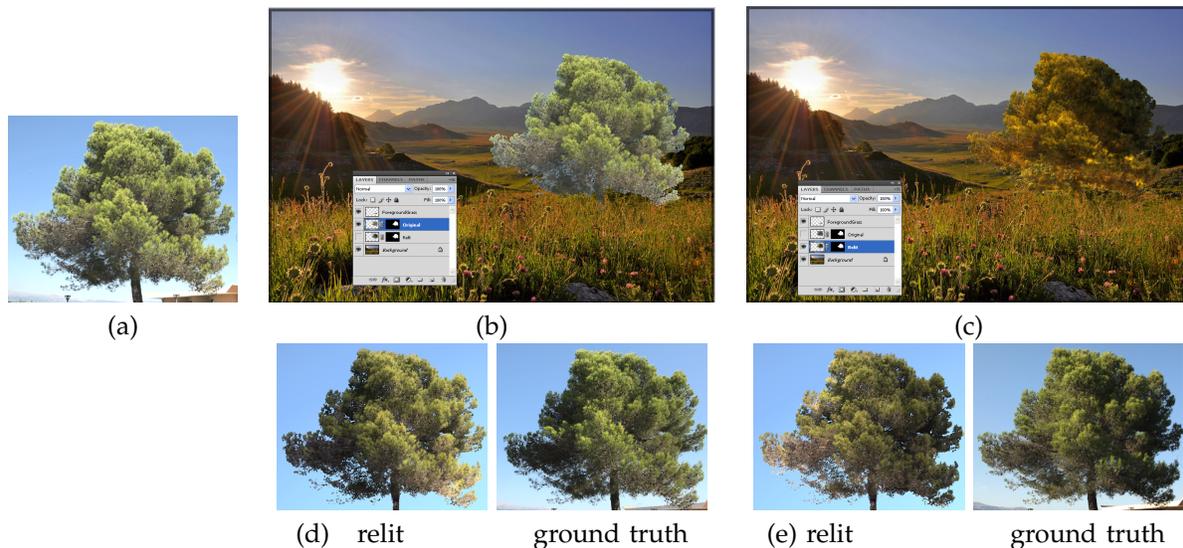


Fig. 1. (a) One of the input images taken at noon. (b) A snapshot of an image editing program interface, with the layer containing a tree, and the target photo. Inserting the tree directly in the image is unsatisfactory, clearly revealing the lighting inconsistency. (c) Using our approach, the relit tree fits much better with the target lighting condition, in particular for shadowing in the canopy. (d)-(e) Canopies relit with our approach using only photos taken at noon (a), and the ground truth photographs at the target relighting times (taken for purposes of comparison only, and not used by the algorithm). Please note that we only use (a) as input, together with a set of images taken at the same time around the tree. The target lighting conditions (i.e., (c), (d) and (e)) are generated automatically by our method.

and validate this principle. This intuition is then used to estimate a per pixel radiance value for the input photograph and then estimate an equivalent value at the target time, allowing us to achieve relighting. Our goal is to provide a convincing relighting algorithm for tree canopies, and as such our approximations do not constitute a physically accurate model. Nonetheless, as we shall see, our assumptions provide satisfactory results for relighting. In particular, we have three main contributions:

- We present an analysis of lighting in a tree canopy, starting from first principles, and using physically-based synthetic renderings to guide our study. Synthetic trees are used only to generate renderings for analysis, and do not constitute part of our contribution. The analysis indicates that in terms of overall lighting behavior we can assume diffuse reflectance for leaves, and approximate lighting behavior in a tree canopy using a single-scattering isotropic participating media model.
- Based on this analysis, we develop a relighting approach based on a volumetric lighting model, and an analytical sky/sun model [11]. This is achieved by computing approximate representations for irradiance both at the time of the input photograph and that of the desired lighting condition, and relighting using their ratio.
- Finally, we develop a fast relighting algorithm using spherical harmonics to combine the effect of transmittance from the sky sampling directions and the path to the eye.

We present validation results on synthetic trees (Figures 11, 12), for which exact ground truth exists, and for (sparse) time-lapse sequences of three different real trees (see Figure 1(c)-(e) and Figures 14, 15 and 21)

2 PREVIOUS WORK

The goal of changing the lighting conditions of a photograph has been the focus of much past research in Computer Graphics. Initial approaches were based on constructing an approximate model of the scene, estimating reflectance parameters and then performing a relighting calculation sometimes based on global illumination computations (e.g., [1], [2], [12], [5], [4]). The environments used in these approaches did not involve complex geometries such as trees, since they are hard to reconstruct from images. An interesting recent solution to image-based relighting taps into the sheer volume of images and lighting conditions available on the web [13]. While very promising, it has similar limitations concerning required geometric reconstruction in the case of trees.

Lighting in tree canopies has been previously studied in computer graphics and other fields. Soler et al. [14] describe a hierarchical instantiation scheme for radiosity rendering of vegetation that takes advantage of self-similarities that exists in plants; their solution is on entirely synthetic models, with exact geometry and materials available. In the field of agronomy, de Castro and Fetcher [15] model lighting in a tree canopy by analyzing illumination arriving at individual grid cells overlaid on a vegetation. They validate this model with real measurements on a coarse grid and in an “artificial”

plantation. While this model has some similarities to our approach, their focus is on producing average illumination values at much larger scales, suitable for agronomy. As a result this approach cannot be directly applied for image relighting. In our case, we will be using a participating media lighting model. For a comprehensive survey, please see [16]; our notation is inspired from this survey. In addition, we assume that the medium is isotropic and that we have single scattering.

The idea of modeling the tree canopy as a volume has been explored by Boulanger et al. [17]. The authors approximate visibility between leaves within the tree canopy using a volumetric approach in order to render the final result. However they rely on actual geometry and normals to compute visibility, as opposed to our method which is image-based and volumetric.

There has also been a significant amount of recent work on tree modeling from images. Most approaches (e.g., [18], [9], [10]) use the images to guide the generation of a completely synthetic model; an alternative is to create an image-based representation using a voxel grid with billboards [19]. While it would be possible to use the resulting purely synthetic tree models for relighting in the spirit of [20], [4], [6], it would be necessary to extract leaf reflectance parameters from the images. While some methods exist for large scale reflectance capture in radiometry [21] it is unclear whether these could be used here. Even if reflectance were available, the final result would not completely match the image since the resulting synthetic models are not usually pixel accurate. Thus our goal of relighting the tree in the actual photograph would not be achieved.

Our relighting formulation is based on ratios of images that depict different environment conditions to compute lighting at a target time. These methods have been extensively used in the literature for solving inverse lighting problems [22], [23], [24], [25].

A recent trend in relighting research is based on complex capture setups, such as the Light Stage [7] and followup work, e.g., [8]. These results have produced impressive results of stunning realism, and have been extensively used in film production. Our approach is orthogonal to such capture setups, since our goal is to have a simple capture workflow using a simple digital camera. In addition, most such laboratory setups [7] are inappropriate for capturing trees, due to size and transportability issues.

To our knowledge, and as can be seen from the previous discussion, no previous method can relight tree canopies using only single lighting condition photographs as input. Evidently, we build on previous techniques. We use a volumetric representation of the canopy, and a ratio-based relighting technique such as those discussed above. Our volumetric lighting estimation is based on techniques developed in volumetric and tree rendering methods discussed previously. However, since no geometry nor material properties are explicitly used our approach is not a “pure” rendering technique.

We next start with a careful analysis of tree canopy relighting. We use synthetic global illumination renderings to study lighting in tree canopies, which leads us to adopt a volumetric rendering model (Section 3). Based on this analysis, we develop a volumetric tree canopy relighting method (Section 4) and an efficient relighting algorithm based on spherical harmonics. Our method is explained in more detail in Section 5.

3 AN ANALYSIS OF TREE CANOPY LIGHTING

As discussed above, no previous method is able to achieve the goal of relighting single lighting condition photographs of tree canopies. This is due to the inherent complexity of tree canopy geometry and illumination. The analysis presented in this section

shows that a number of simplifying assumptions can be made concerning lighting of tree canopies, providing us with the qualitative intuition required for our fast and efficient relighting algorithm

In particular, using theoretical arguments and synthetic simulation, we will show that a diffuse reflectance assumption, and a single scattering volumetric lighting approximation (Section 4) are qualitatively acceptable, enabling the use of a ratio-based relighting technique, in the spirit of [22], [23], [24], [25].

In the following, we will be using synthetic images to provide intuition about lighting, and to validate our approximations and assumptions.

Note however that the geometry and normals of these trees are never used in any way to generate the results of our method; using synthetic trees is just a convenient way to generate images while controlling various parameters.

These synthetic images are rendered using PBRT [26], a physically-based rendering system. For these images we used photon-mapping (5M photons) with final gather (128 samples/pixel and 128 directional samples), as a low-noise global illumination solution to best match reality. Note that each image requires several hours to compute. We use diffuse reflectance for the leaves, which are also slightly translucent for Figure 11. The sun is modeled as a directional light source and the sky is modeled using the Preetham model [11].

3.1 Canopy Lighting: diffuse assumption

Consider an image of a tree canopy (e.g., Figure 2(a)); in what follows we will only be referring to pixels which are on the tree canopy itself. Without loss of generality, we assume that for a given pixel x we have a single corresponding 3D point x_l on a leaf in the real tree. The reflectance of x_l is $\rho(\omega_i, \omega_o)$, where ω_i is the incoming direction and ω_o is the outgoing direction. We assume two light sources, the sun, with radiance L_{sun} and the sky, which is a hemispherical source. We can sample the sky in a given direction ω , giving radiance $L_{sky}(\omega)$. We also denote $L_{ind}(\omega)$ the radiance due to indirect illumination, and θ the angle of ω with the surface

normal; Ω is the positive hemisphere of directions. The radiance arriving at pixel x from point x_l , is thus:

$$L(x) = \int_{\Omega} (\rho(\omega, \omega_o) L_{sky}(\omega, x) \cos(\theta) + \rho(\omega, \omega_o) L_{ind}(\omega, x) \cos(\theta)) d\omega + \rho(\omega_{sun}, \omega_o) L_{sun}(x) \cos(\theta_{sun}) \quad (1)$$

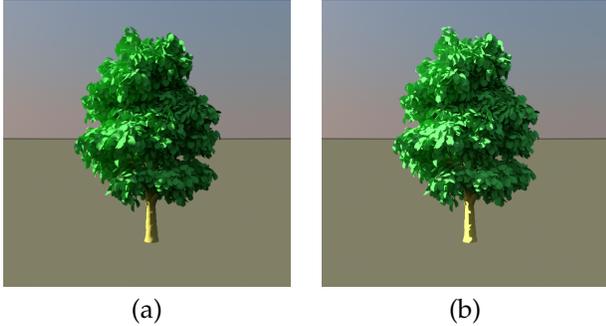


Fig. 2. (a) A globally illuminated image of a tree. (b) the same tree, in which the dependency on the cosine to the normal is ignored in all lighting computations.

Our first assumption is that we can assume leaf reflectance ρ to be diffuse. This may seem to be too strong a simplification; however, there are two reasons which justify this choice. First, if we assume that the leaves are randomly oriented in a uniform distribution, in a manner analogous to microfacets for surfaces, at a given scale, where each pixel captures one or more leaves, the canopy can be seen as a diffuse surface. Second, as noted by [17], more than half the total radiance in an image of a tree comes from the sky (which is an hemispherical source) and indirect light. In addition, the direct sunlight, while contributing a high percentage of radiance (about half), only affects a small number of pixels (see also Figure 3). These observations indicate that specular effects, at a certain scale, do not always have a very strong effect on the image of a tree. For leaves with a strong preferred orientation and high specularity, these assumptions no longer hold, and this could adversely affect the quality of our results.

To further examine the effect of these assumptions on the performance of our approach, we varied the reflectance parameters of leaves in synthetic data sets, to have high specularity and then high translucency and studied the effect on the performance of our algorithm. Details of these studies are presented in Section 6.2.

If we assume that ρ is diffuse, we can simplify Eq. 1 by dropping dependencies on angle θ . Since we are integrating over the hemisphere Ω , we are now treating irradiance values, which we note with the symbol E . Radiance received at pixel x thus becomes:

$$L(x) = \rho(E_{sky}(x) + E_{sun}(x) + E_{ind}(x)) = \rho E \quad (2)$$

where E is the total irradiance received at x_l , and $E_{sky}(x)$, $E_{sun}(x)$, $E_{ind}(x)$ are the irradiance due to

sky, sun and indirect illumination respectively. It is important to note that we could still derive Eq. (2) without dropping dependencies on angle: we do so to emphasize the fact that we do not have any information on normals in our volumetric approximation. For clarity, $E_{sky}(x)$, $E_{sun}(x)$, $E_{ind}(x)$ will be called E_{sky} , E_{sun} , E_{ind} from now on.

One way to evaluate the validity of discarding the cosine term in the volumetric model assumption is to examine the importance, with respect to lighting, of normals compared to visibility for tree canopies. To do this, we compute an image of a tree (shown in Figure 2(a)) using standard PBRT rendering. We then compute an image of the same tree with the same lighting parameters, but we ignore the cosine with the normal in the computation of lighting (Eq. 1) (see Figure 2(b)). As we can see the difference in the images is minimal, except for leaves which receive direct light.

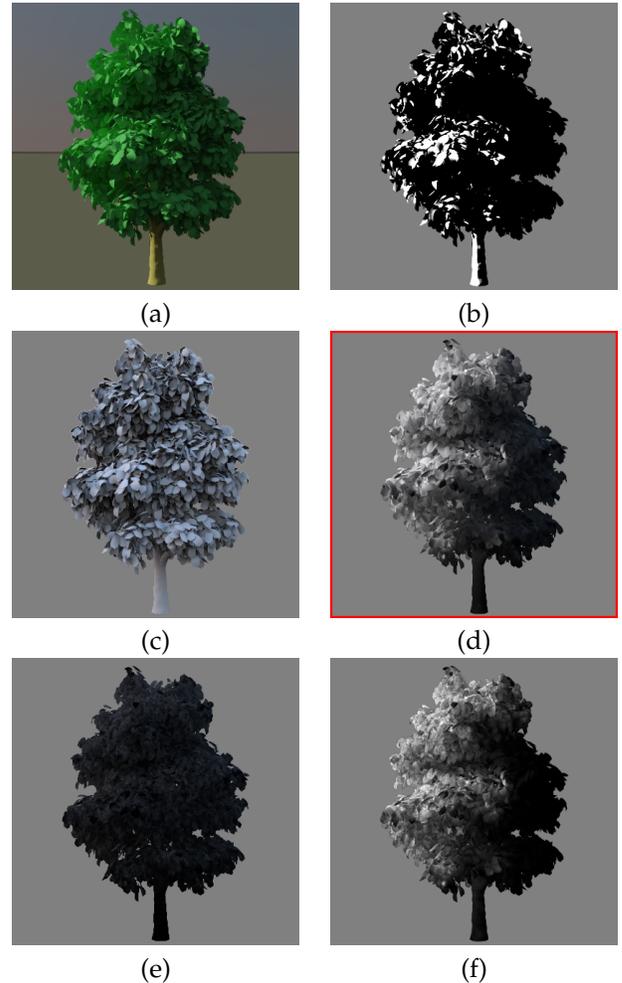


Fig. 3. (a) A globally illuminated tree, using photon mapping, computed with PBRT. Images (b)-(f) are computed with leaf reflectance equal to (1,1,1) to allow comparison with our participating media model. (b) Direct sunlight. (c) Direct skylight. (d) Indirect illumination. (e) Indirect illumination (sky only) (f) Indirect illumination (sun only)

3.2 Volumetric approximation

A tree canopy is made up of a very large number of leaves. As such, there is a qualitative similarity to a participating medium, in which the leaves have the roles of particles. Our method is based on this idea, which will allow us to approximate irradiance at given lighting conditions, thus enabling tree relighting. As we shall see in Section 5, this will allow us to approximate E_{sky} , E_{sun} and E_{ind} from Equation 2 using an approximation of $L(x)$, shown in Equation 12.

To investigate this qualitative analogy, we will analyze the behavior of light by separating out the various components of Eq. 1. In Figure 3(a) we show an image of a tree with global illumination. We then show the various components of illumination, notably direct sunlight Figure 3(b), direct skylight Figure 3(c) and indirect illumination Figure 3(d) (from both sun and skylight), Figure 3(e) indirect due only to skylight and Figure 3(f) indirect due only to sun light. Since we will be comparing to a participating media model, whose phase function will be monochromatic, we set leaf reflectance equal to (1,1,1). As we shall see later (Section 4), reflectance cancels out in our computations.

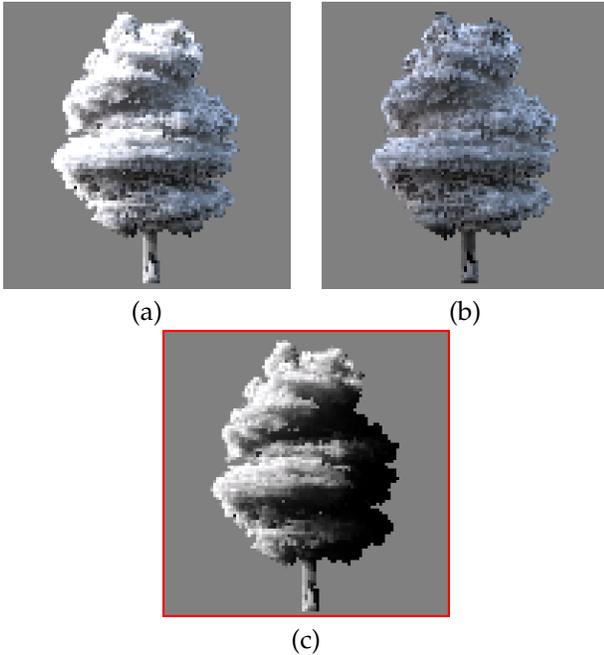


Fig. 4. (a) The image \tilde{E} computed with our single scattering model. (b) Skylight only component and (c) Sunlight only component.

As we can see in Figure 3(b), direct lighting from the sun only affects a small number of pixels since it is only present on unoccluded front-facing leaves (with respect to the sun). Skylight (Figure 3(c)), while present everywhere, contributes less to the overall illumination of the tree canopy. Indirect illumination (Figure 3(d)) represents a significant portion of overall illumination. As expected, the overall aspect of indirect light, especially due to the sun, is similar to that of scattering within a volume. As

we can see in Figure 3(e) and (f) this is mainly due to indirect reflections of sunlight.

We use this qualitative observation as the intuition to derive our volumetric model, which allows us to develop an effective tree canopy relighting approach.

We will be using a standard participating media lighting model, described in Section 4.1. In particular, we shall see that we can compute an image of irradiance for a volumetric representation of the tree canopy. We use [19] (see Section 4.1 for details) to construct a volume, using exact “mattes” separating pixels of the canopy from those in the background.

The isotropic assumption is valid if the orientation of leaves is considered random within the canopy. As the synthetic results show, a single scattering approximation is qualitatively similar to the behaviour of indirect lighting in a tree canopy.

As we shall see later (Section 4) our relighting does not depend on reflectance; we thus use a constant monochromatic phase function with unit value.

We show a rendering of our volumetric lighting model (Section 4.1) in Figure 4(a). We also show the effect of the sky and sun separately in (b) and (c). If we compare Figure 3(d) (all indirect) and 4 (c) (sunlight component), outlined in red, we see similar overall distribution of light. As we can see, first order scattering in the volumetric model provides an approximation of indirect light. This provides a strong indication that it is in effect unnecessary to consider multiple scattering.

In addition, Figure 3(d) and Figure 4(c) (highlighted in red) show that our model captures the overall behavior of skylight.

The above comparisons are a qualitative indication that the participating media approximation we propose provides a good approximation for the overall distribution of lighting within a tree. We make no claims about the accuracy of the absolute illumination levels, nor physical fidelity from our model; however in the context of our algorithm for tree canopy relighting, these approximations appear to work well (see synthetic validation and results on real images in Figures 11-21, as well as the accompanying video.)

4 A PARTICIPATING MEDIA APPROACH FOR TREE RELIGHTING

In this section, we first discuss our relighting approach, and then present our participating media lighting model.

To achieve tree canopy relighting, we use a method based on image ratios ([22], [23], [24], [25]). Assume that we have two images of the same tree canopy, at two different times of day I_{in} and I_{targ} (see Figure 5(a) and (b)). We can write: $I_{in} = \rho E_{in}$ and $I_{targ} = \rho E_{targ}$. The above leads to the well known property governing all diffuse materials, which we call ratio I_r :

$$I_r = \frac{\rho E_{targ}}{\rho E_{in}} = \frac{E_{targ}}{E_{in}} \quad (3)$$

An example image I_r is given in Figure 5(c). Clearly, if we only have I_{in} as input, we only need the ratio I_r to compute I_{targ} , since

$$I_{targ} = I_r \cdot I_{in} \quad (4)$$

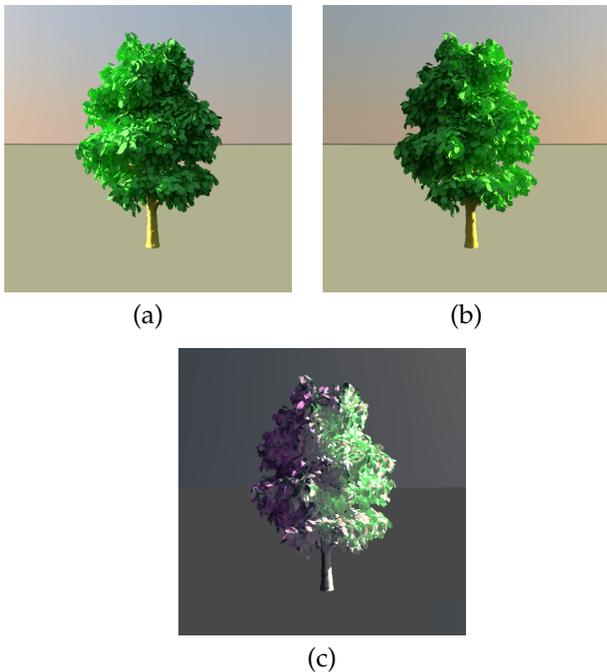


Fig. 5. (a) The input image I_{in} . (b) The target lighting condition I_{targ} . (c) The ratio image I_r of (b)/(a).

There are two interesting observations to be made from Eqs. 3 and 4. First, we can compute relighting *without estimating reflectance*, since only irradiance values are required to compute I_r . Second, only *relative* values are required, since only the ratio of the input and target irradiance values is required, not their absolute values. As a consequence, all we need is to compute appropriate approximations for E_{in} and E_{targ} so that their *ratio* is accurate.

The analysis presented in the previous section indicates that a good way to compute such approximations \tilde{E}_{in} and \tilde{E}_{targ} would be with a participating media lighting model. We describe our volumetric lighting solution in the following section.

4.1 Participating Media Lighting Model

Our goal is to compute values \tilde{E}_{in} and \tilde{E}_{targ} . This will give us an approximation of E_{in} and E_{targ} and ultimately a good approximation of I_r . Once we have I_r we can relight the canopy using Eq. 4.

We start by constructing an approximate voxel representation of the tree; we use the approach of [19] to achieve this. We briefly recall the process here: a set of photographs (around 12-15) is taken around the tree, the cameras are calibrated, and a set of mattes created to separate the canopy from the background.

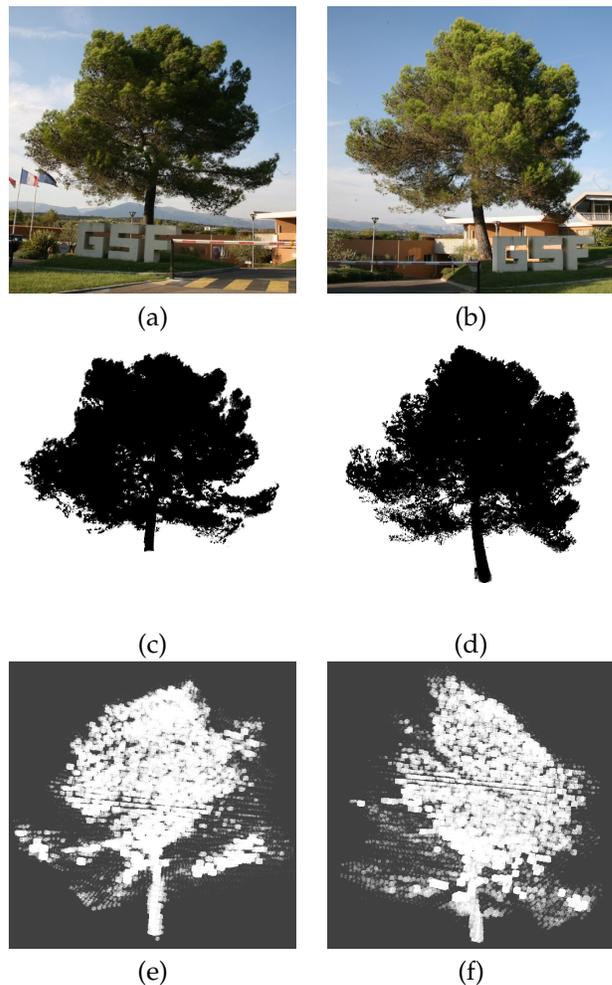


Fig. 6. Two input photos of the same tree ((a) and (b)), the corresponding mattes ((c) and (d)) and two views of the resulting volume ((e) and (f)).

The quality of the mattes plays an important role since our method for relighting heavily depends on the quality of the reconstructed volume. When possible, blue screens were utilized in order to isolate the subject tree properly from other vegetation behind it (see Figure 7). This process improves the quality of the resulting mattes.



Fig. 7. When conditions allowed, a bluescreen was placed behind the tree (*left*) to aid matte extraction (*right*).

Once the input images with calibrated cameras and

the mattes are in place, an optimization process is run to estimate opacities in each voxel and to “carve out” the shape of the tree. We show an example of tree photos and the resulting reconstructed volume in Figure 6.

As mentioned in the introduction, the basic intuition for the model we adopted is that the leaves of the tree canopy can be seen as a participating media volume in terms of lighting; such a model thus appears suitable for the computations of \tilde{E}_{in} and \tilde{E}_{target} . We thus adopt a standard participating media rendering approach, based on Beer’s law [16], [27], [28]. Radiance at point x is given as follows:

$$L(x) = L(x_0)e^{-\int_{x_0}^x k_t(u)du} \quad (5)$$

where k_t is the extinction coefficient of the medium and light travels from x_0 to x . We write:

$$\tau(x_0, x) = e^{-\int_{x_0}^x k_t(u)du} \quad (6)$$

$\tau(x_0, x)$ is the *transmittance* from x_0 to x . We will also write $\tau(x_0, \omega)$ which denotes the transmittance along a direction ω from the entry point of the volume to the point x_0 .

In our case, x will typically be the eyepoint, or equivalently the pixel, and x_0 will be the furthestmost point of the volume along a ray emanating from the eye. We illustrate these quantities, and those of Eq. 7 and 8, in Figure 8.

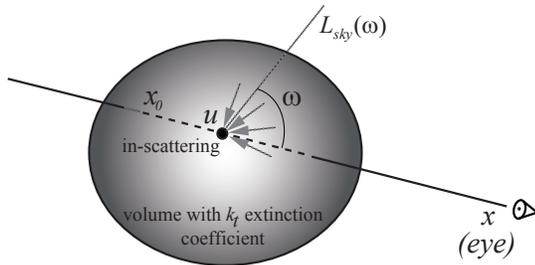


Fig. 8. The geometry of our single scattering participating media model, illustrating the quantities of Eq. 5.

Given these assumptions, and expanding $L(x_0)$ in Eq. 5, radiance at L is given as [16]:

$$L(x) = \int_{x_0}^x \tau(u, x)k_t(u)L_{ss}(u)du \quad (7)$$

where L_{ss} is the single-scattering radiance arriving at u . In our case we write:

$$L_{ss}(u) = L_{sun}\tau(u, \omega_{sun}) + \int_{\Omega} L_{sky}(\omega)\tau(u, \omega)d\omega \quad (8)$$

where Ω is the positive hemisphere of directions ω and $\tau(u, \omega_{sun})$ is the transmittance through the volume along the sun direction ω_{sun} to the point u .

The integrals of Eqs. 6 and 7 can be approximated with a standard method, similar to [28], using ray traversal through the discrete voxels of the grid and numerical integration.

The participating media model we use computes radiance at a given pixel. However, since our phase function is equal to one in all frequencies, the final quantity we compute is actually irradiance. We thus write $\tilde{E}(x)$ for the value computed by Eq. 7.

To compute the values \tilde{E}_{in} and \tilde{E}_{target} , we use the discrete volume, a k_t coefficient for each voxel (both computed using [19]), and values for L_{sun} and $L_{sky}(\omega)$ using the Preetham model [11]. At each pixel we evaluate Eq. 7, with appropriate values for the sun and sky illumination depending on the time of day of the input and target images.

4.2 Sunlight

The contribution of the sun at the eyepoint x is given as follows:

$$L^{sun}(x) = \int_{x_0}^x \tau(u, x)k_t(u)L_{sun}\tau(u, \omega_{sun})du \quad (9)$$

L_{sun} is attenuated according to its respective zenith angle position, as described in [13]:

$$L^{sun} = \alpha e^{(-\beta m(\omega_{sun}))} N_{max} \quad (10)$$

where $m(\omega_{sun})$ is given as in [13], and is the relative optical path length through Earth’s atmosphere [29], N_{max} is the maximum sun brightness value, α is a scale factor and β is a scattering coefficient. We used $N = 1 \times 10^6$ and α, β as described in [13]. The final contribution of sunlight is then:

$$L_{sun}(x) = \int_{x_0}^x \tau(u, x)k_t(u)\alpha e^{(-\beta m(\omega_{sun}))} N_{max}\tau(u, \omega_{sun})du \quad (11)$$

5 EFFICIENT RELIGHTING ALGORITHM

To approximate Eqs. 7 and 8 at each pixel we need to compute sun and skylight along the corresponding viewing ray. We separate the contributions of the sun and sky in Eq. 7 for clarity of presentation, and without loss of generality we write:

$$L(x) = L_{vol}^{sun}(x) + L_{vol}^{sky}(x) \quad (12)$$

The contribution $L_{vol}^{sun}(x)$ of sunlight is computed by directly estimating the value of $\tau(u, \omega_{sun})$ at a sample point u at the center of each voxel along the viewing ray, providing a discrete approximation of Eq. 9.

Evaluating sky illumination in Eq. 8 is the most expensive part of our computation. Using the naive approach discussed previously, to relight each pixel we need to sample in the order of 1024 directions in the sky, stepping through the voxelization for each sample direction. We do this for *each* voxel visited along a viewing ray. Clearly, this results in prohibitively high relighting costs. We will next show how to use spherical harmonics to precompute the accumulated effect of transmittance through the

volume, both in each sky sample direction and along the viewing ray for a given pixel.

Radiance L_{vol}^{sky} due to the sky in our volumetric model is given as follows:

$$L_{vol}^{sky}(x) = \int_{x_0}^x \tau(u, x) k_t(u) \int_{\Omega} L_{sky}(\omega) \tau(u, \omega) d\omega du \quad (13)$$

where the quantities τ , ω , Ω and u are the same as in Eq. 6-8.

Since $\tau(u, x)$ does not depend on ω we can write:

$$L_{vol}^{sky}(x) = \int_{x_0}^x \int_{\Omega} \tau(u, x) k_t(u) L_{sky}(\omega) \tau(u, \omega) d\omega du \quad (14)$$

We can then invert the order of integration and factor out L_{sky} :

$$L_{vol}^{sky}(x) = \int_{\Omega} L_{sky}(\omega) \left(\int_{x_0}^x \tau(u, \omega) \tau(u, x) k_t(u) du \right) d\omega \quad (15)$$

We can write the inner integral as $g(\omega)$:

$$g(\omega, x) = \int_{x_0}^x \tau(u, \omega) \tau(u, x) k_t(u) du \quad (16)$$

One way to reason about the above equation is that, for a given point u along a viewing ray (in the direction $\vec{x}\vec{u}$), this expression simultaneously captures the contribution in skylight sample direction ω at the point u , and the attenuation from point u to the eye-point x (see Figure 8).

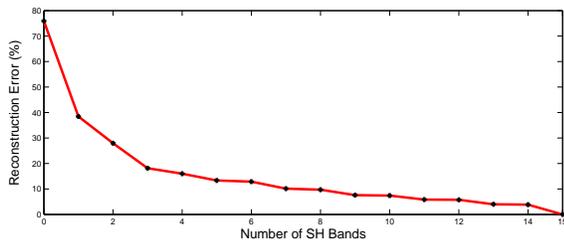


Fig. 9. Graph plotting the reconstruction error (Y) for a given number of bands.

We precompute a discretized version of $g(\omega, x)$ and project it onto a spherical harmonic basis $\tilde{g}(\omega, x)$ using 6 bands.

Spherical Harmonics were first introduced by Sloan et al. [30] in the context of pre-computed radiance transfer. More recently, Jansen et al. [31] introduced Fourier opacity mapping, in which absorption is projected along one ray onto a 1D Fourier basis to speed up the integral computation. In our case however, we project the directional (2D) distribution of transmissivities onto a sphere to allow for faster integration. This precomputation is done only once for each reconstructed tree, which allows faster relighting.

The reconstruction error using 6 bands is 12.86% in L_2 norm. See Figure 9 for a graph which plots the errors in reconstruction using different number of bands. At

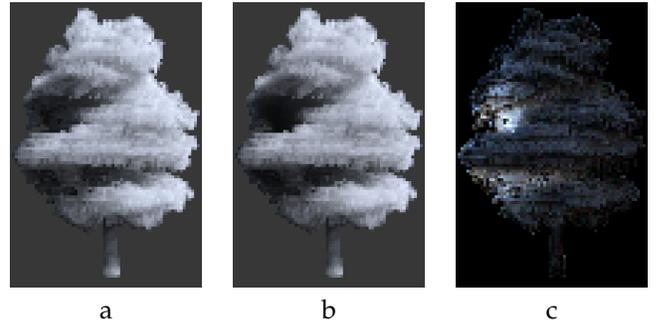


Fig. 10. (a) Image computed by explicit sampling of the sun dome at each sample point. (b) The result of our spherical harmonics approximation. (c) Difference image ($\times 10$).

runtime, we can then simply compute a dot product of the skylight projected on-the-fly onto spherical harmonics (we could alternatively use an analytical formulation [32]), and the precomputed combined transmittance:

$$L_{vol}^{sky}(x) = \int_{x_0}^x L_{sky}(\omega) \tilde{g}(\omega, x) \quad (17)$$

This allows us to compute \tilde{E}_{in} and \tilde{E}_{target} very efficiently (i.e., in the order of 3 seconds each for a 512×512 image¹). Precomputation for a 128^3 volume required 20 minutes. However, computing the sky contribution on a coarsened 64^3 volume does not result in significant degradation, but reduces the precomputation time to 5 minutes.

To validate the quality of this approximation, we compare the result of Eq. 17 to the ground truth reference in which we approximate the integral of Eq. 13. Ground truth is computed by explicitly sampling 1024 sky directions at each sample point of the line integral along the viewing ray. In Figure 10, we see that the difference is very low (note that the difference image is multiplied by $\times 10$).

6 VALIDATION RESULTS ON SYNTHETIC TREES

We first apply our model on several synthetic trees; the advantage of these tests is that it is much easier to generate reference solutions and obtain a variety of data for analysis, and we can vary the location and dates of the images freely. Note that we only use the *images* generated by PBRT as input. The only difference with the real photographs is that no calibration is required for the camera, since they are provided as input. All other steps (volumetric construction etc.) are identical to the process for real photographs.

We also investigate the effect of varying reflectance parameters (specularity, translucency), of the leaves, and the resulting effect on the performance of our algorithm.

1. Timings on 4-core 2.3Ghz Xeon; only relighting is multi-threaded.

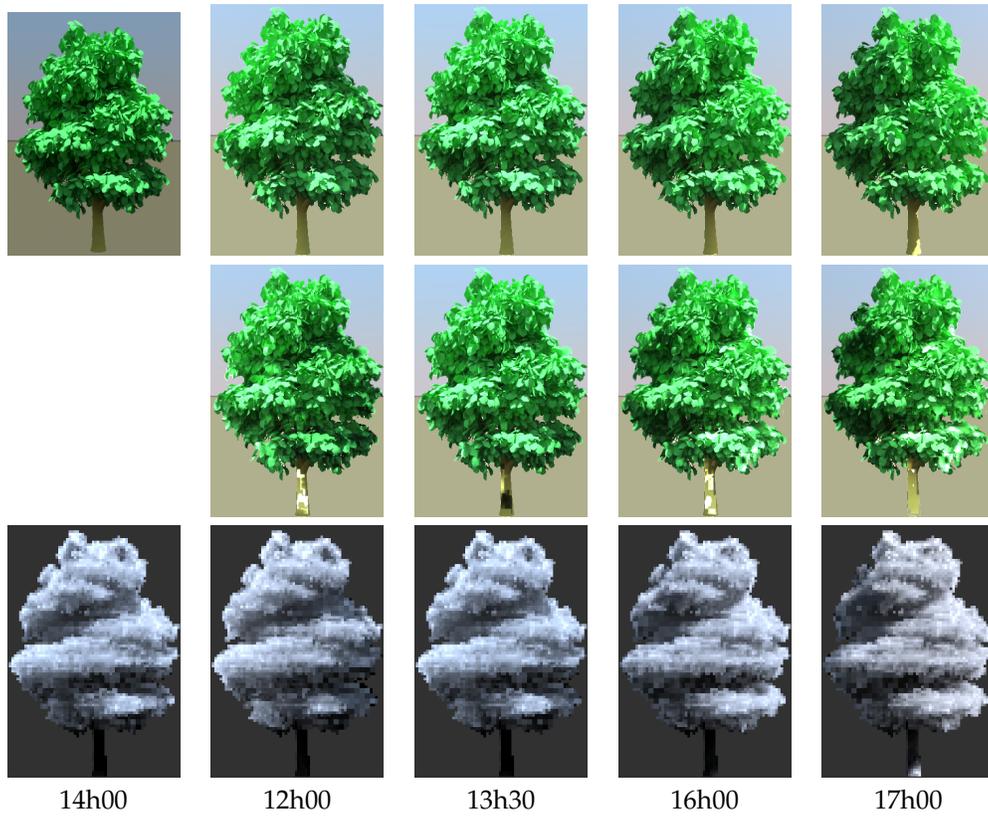


Fig. 11. **Horse Chestnut** Top row: input images and 4 target ground truth images with corresponding times of day. Middle row: 4 resulting relit images using our approach. Bottom row: \tilde{E}_{in} and the four \tilde{E}_{targ} images.

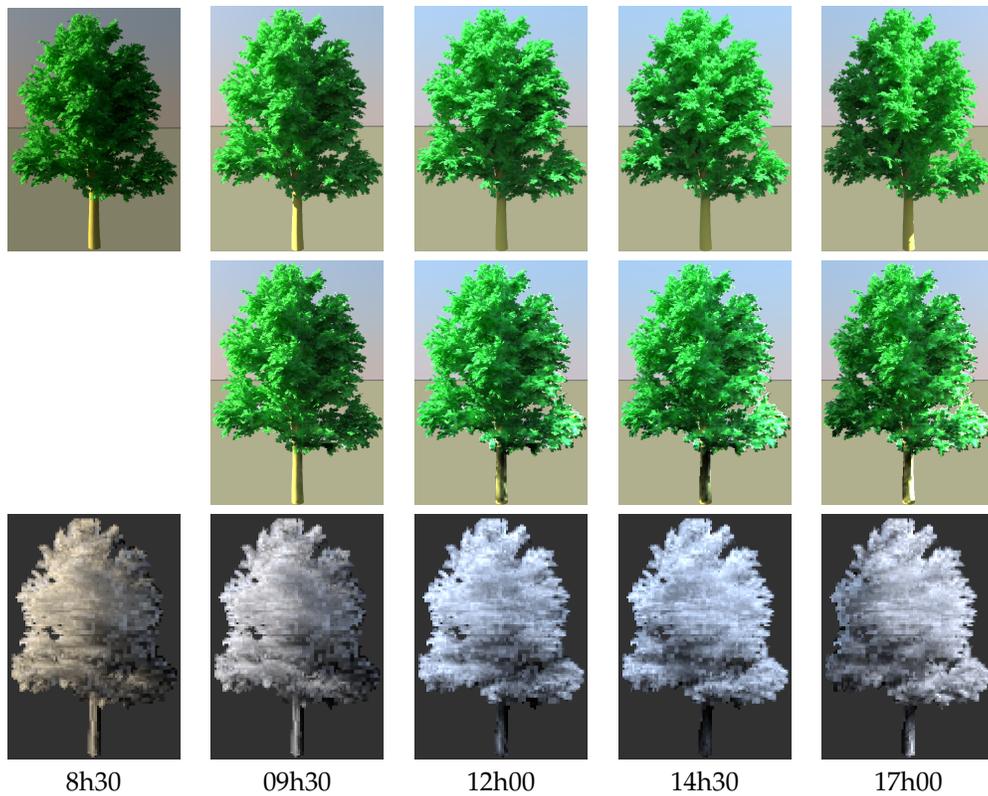


Fig. 12. **European Beech** Top row: input images and 4 target ground truth images with corresponding times of day. Middle row: 4 resulting relit images using our approach. Bottom row: \tilde{E}_{in} and the four \tilde{E}_{targ} images.

6.1 Synthetic Validation Results

We show two examples taken from the *Xfrog* (<http://www.xfrog.com>) European tree database. Two more examples are provided in the additional material, as well as the entire time-lapses for all trees. The date used is July 22nd, and the location is East Coast US.

For each tree on the top row, we show the input image and 4 target images at different times of day. In the middle row we show the relit images using our method and in the lower row we show the \bar{E} images (Figure 11 and 12 are respectively called *Horse Chestnut* and *European Beech*). Please also see the accompanying video, where the movement of the sun and the corresponding illumination is much easier to comprehend.

As we can see, our relighting approach captures the overall behavior of lighting well. Evidently, the limitation of the volumetric reconstruction compared to the actual detailed geometry results in minor differences in levels of illumination. In the additional material and video we can also see that our approach works quite well even for cases of relatively sparse trees, and for trees of different canopy shapes. For the case of the European Mountain Ash, we see slight “banding” artifacts, which are more visible in the video (see Figure 13). We believe this is due to the close-to-spherical nature of this particular canopy, and the consequent inability of the volumetric reconstruction to capture fine geometric details. This renders the grid structure more visible in the final results.



Fig. 13. The European mountain Ash. (Left) Ground Truth; (Right) Relit result. The almost spherical nature of the canopy results in lack of detail from the volumetric reconstruction. Slight banding artifacts can be seen (better seen in the video).

6.2 Varying Reflectance Parameters

To study the effect of high specularity on the performance of our algorithm we rendered a synthetic tree in PBRT with a leaf material that contains both diffuse

and glossy specular reflections. The material uses Blinn BRDF with a microfacet model for surface roughness (see [26] for details). Glossy reflectivity was set up to be $5 \times$ stronger than the diffuse reflectivity to exaggerate the results. Surface roughness was set up to 0.1 which indicates smaller but highly specular highlights (see Figure 16). From our test, we clearly see that specular highlights are not well captured by our relighting method since we do not have geometry (and normal) information.

An additional reason why modeling specularly may be of low importance in our context, could be the low-frequency nature of sky and secondary illumination.

Despite the above test, we believe that our model works well in practice for photographs because real tree leaves are in reality translucent. To investigate this hypothesis we created a synthetic tree with very translucent leaves. We rendered synthetic images of a tree using a translucent material with glossy/diffuse transmissivity for leaves - values for glossy transmissivity are 0.15 and diffuse/glossy reflection are 0.85 – see PBRT [26] for more details (see Figure 17). As we can see, the results are significantly improved compared to the specular only test in Figure 16.

7 RESULTS ON PHOTOGRAPHS

We next present our results on real photographs. We first present issues related to the procedure and implementation, then present and discuss our results.

7.1 Procedure and Implementation

As mentioned previously, we use the method of [19] to construct the volume. We currently use [33] for the mattes, and ImageModeler (<http://usa.autodesk.com>) for calibration. Automatic camera calibration using e.g., [34] can also be used, simply requiring a larger number of photographs. Mattes would still only be required on 10-12 photos however.

We applied two modifications to the initialization described in [19]. Due to inaccurate camera calibration and the fact that we were unable to use blue screens everywhere, we observed that the algorithm culls voxels too aggressively, resulting in too sparse volumes and relighting artifacts. The first modification involves keeping voxels even if they are not present in 2-3 mattes/photos. The second modification involves artificially “densifying” the voxel reconstruction in very sparse areas. We first find voxels with less than 8 non-empty neighbors. For each such voxel, we collect its 6 axial neighbors, and we give those that are empty an extinction coefficient value corresponding to the average of the coefficients of the non-empty neighbors. The denser voxel grid significantly improves the results of relighting. Better camera calibration and matting algorithms would render this step unnecessary.

We tested a Mulberry, an Oak and a Pine tree (shown respectively in Figs. 14, 15, 21.) For the volume reconstruction, we used 11, 12 and 11 images respectively for

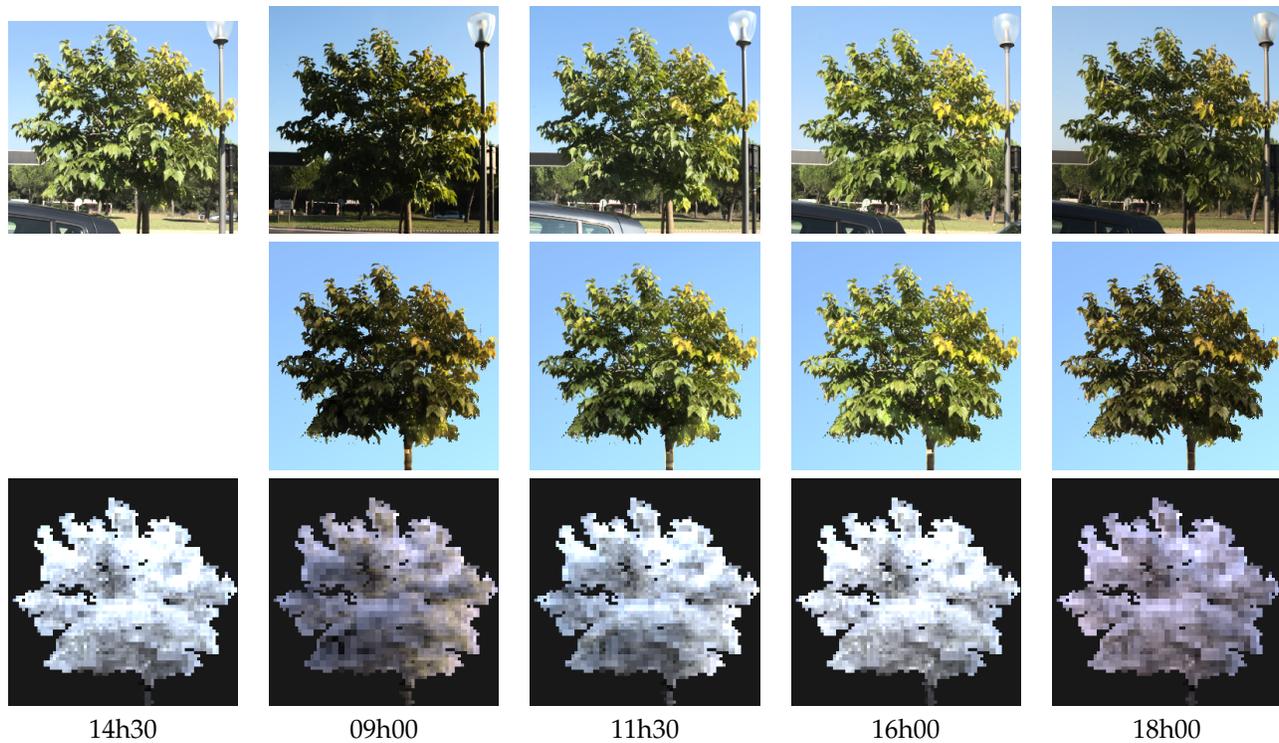


Fig. 14. **Mulberry tree** Top row: input images and 4 target images with corresponding times of day. Middle row: 4 resulting relit images using our approach. Bottom row: \tilde{E}_{in} and the four \tilde{E}_{targ} images.

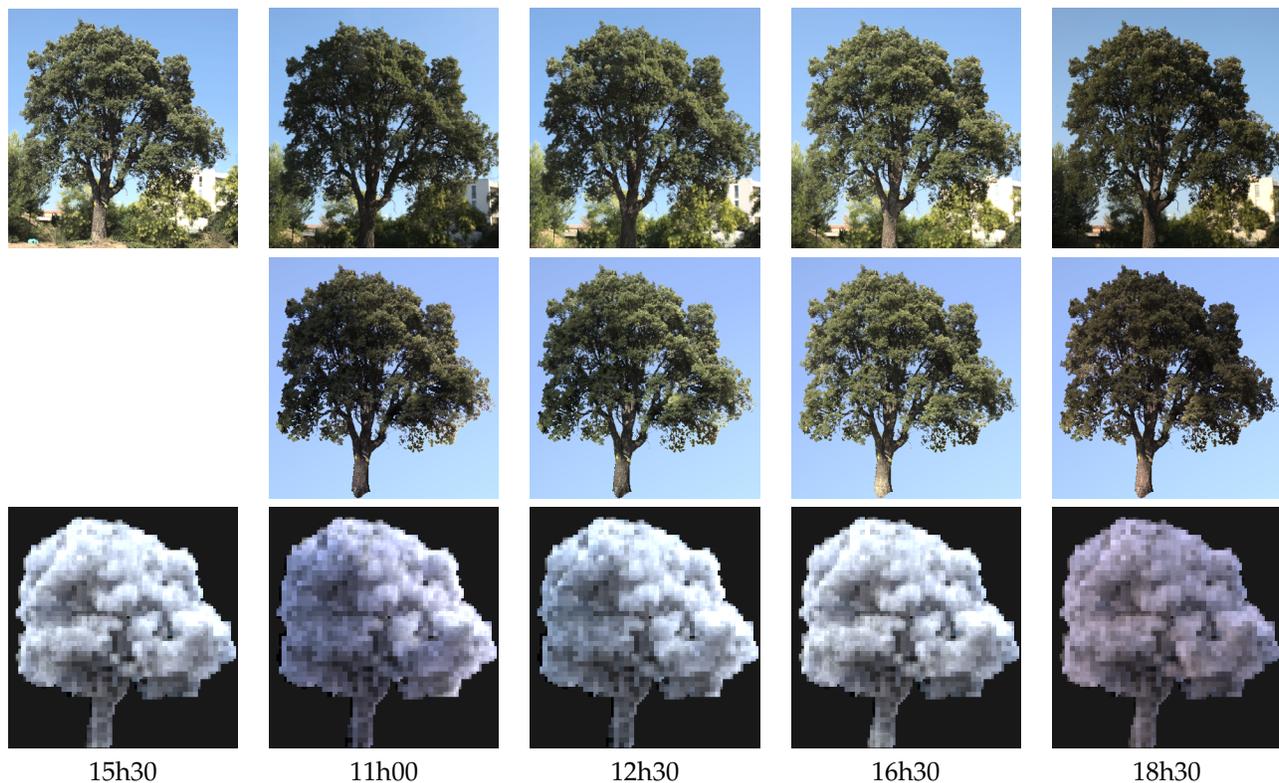


Fig. 15. **Oak tree** Top row: input images and 4 target images with corresponding times of day. Middle row: 4 resulting relit images using our approach. Bottom row: \tilde{E}_{in} and the four \tilde{E}_{targ} images.

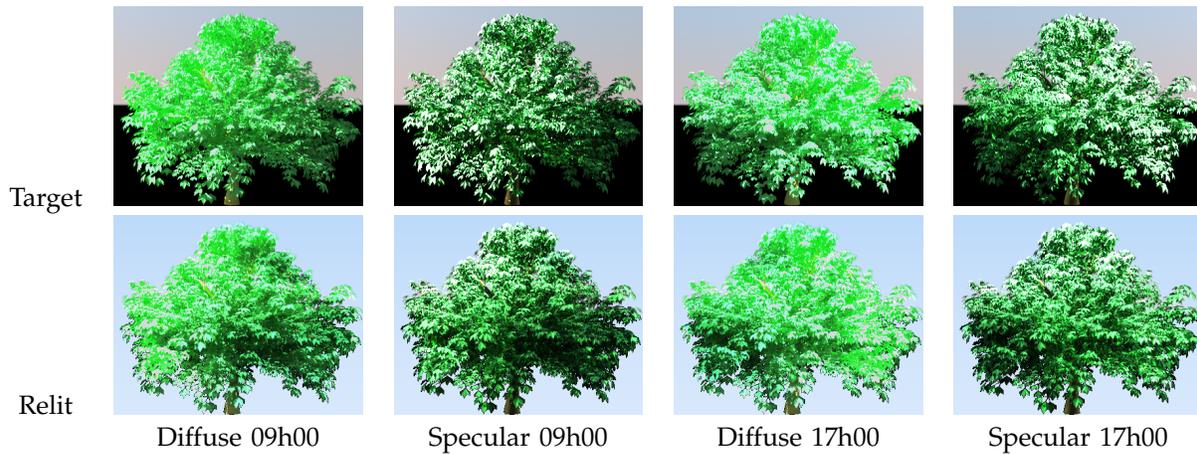


Fig. 16. **Comparison with Specular/Non-Specular leaves** Top row: target images with corresponding times of day for the same tree rendered with diffuse (columns 1 and 3) and specular (columns 2 and 4) leaves. Bottom row: the four corresponding relit images.

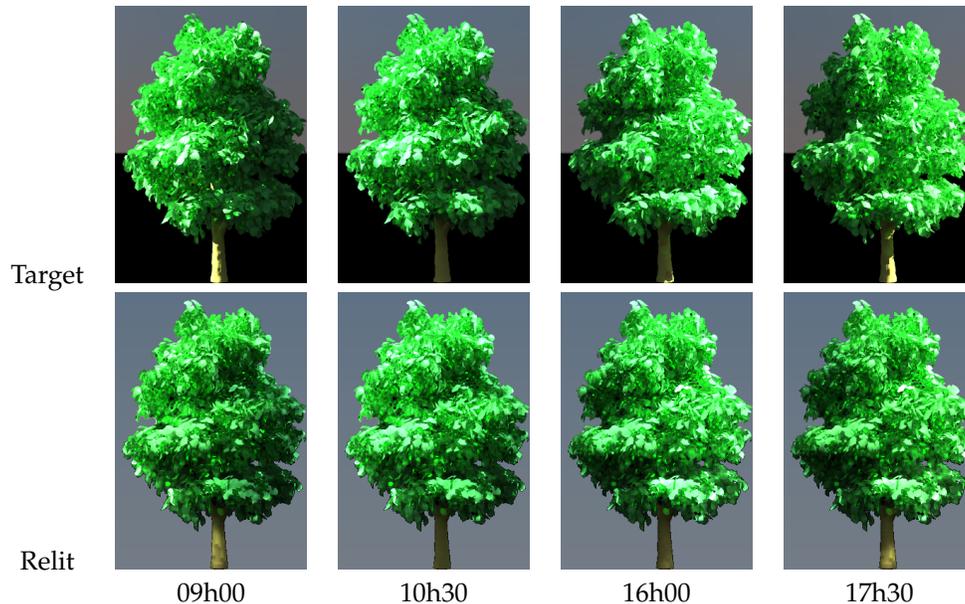


Fig. 17. **A tree canopy with a diffuse/glossy transmissive material for leaves; note that the result of our algorithm is significantly better than the specular only case.**

each tree. The Mulberry tree was the only one where a blue screen was used to aid matte extraction. The Oak and Pine tree locations and sizes prevented us from using a blue screen. The values of k_t computed using [19] can be modulated by a global scaling factor s . For the three examples used, we chose s so that the E_{in} image best represents the input lighting (see first column of Figures 15-21). The values for s we used were 0.95, 0.90 and 0.85 respectively, and account for the density of leaves in each tree. In some cases the resulting masks have small imperfections, and additional manual editing is required. In all the cases shown here, manual editing of masks required less than 5 minutes. We used a Canon EOS 5D camera, and performed all processing on linearized, 12-bit “.RAW” images. HDR images of trees

(composed using different exposures) are hard to capture because of the inherent motion of leaves due to wind. We found that the 12-bit images contain sufficient dynamic range for our method. Care has to be taken to ensure relatively high-quality camera calibration, otherwise the volume reconstruction is unsatisfactory and will not give good relighting results.

Photos are always captured at a single time of day. The camera is positioned around the tree, with the tree canopy at the focal point. The viewpoints around the tree (10 – 12) were chosen in a way to minimize occlusion and interference of nearby objects. When possible, blue screens were utilized to provide a better matte from vegetation behind (see Figure 7). For each shot a compass was used to determine the viewpoint direction, and a

standard GPS provided coordinates to determine the spatial location of the tree. These data were then used to compute the sky and sun models.

For the Preetham sky model, we use a turbidity parameter of 4.2 which appears to work well in all cases. All photos were taken on September 23rd and 24th.

Since our method is based on the Preetham model, sky luminance values early in the morning or late in the evening can be inaccurate, as mentioned in [35]. We alleviate this problem by making the sun brightness decay towards sunset and sunrise stronger. To do this we change the relative optical path $m(\omega_{sun})$ to $m(\omega_{sun})^k$, where $k = 2.5$ for the red channel and $k = 2.0$ for the green and blue channels. This heuristic yields good results in all experiments, significantly improving the quality of the results. An example is shown in Figure 18.

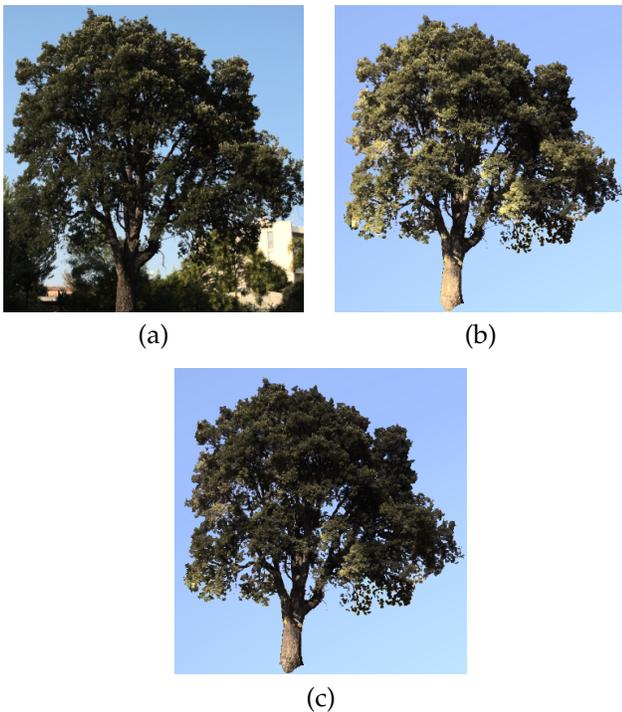


Fig. 18. (a) Target image at 17:30. (b) the relit image - the color shift is due to the inaccuracies of luminance values in the Preetham sky model. (c) A simple factor correction brings the result much closer to the target.

The quality of our results is degraded if the input data used lies within this period of the day, as can be seen in Figure 19. Our experiments indicate that to obtain optimal results with our method, it is best to take input images between 12am and 3pm.

7.2 Results

We took sparse “time-lapse” sequences of the three trees to provide ground truth references, from 9am to 6pm over one hour intervals. In Figures 14, 15 and 21 we show the results in the top row, and the ground truth photographs in the second row. We show four times

here (indicated in the figure); the entire sequences are provided in the additional material and in the accompanying video. As in the case of the synthetic validation, watching the video provides a better sense of the moving sun.

As we can see the quality of the relighting results is satisfactory. Our diffuse reflectance assumption, and the consequent volumetric approximation appears to work well on the examples tested, despite its apparent simplicity. From the results we see that the quality for the real photographs is on a par with that of the synthetic trees, despite the inaccuracies in camera calibration and the non-diffuse nature of the real tree leaves. In addition, even for the case of a relatively sparse tree (the Mulberry example, but also the Oak to a lesser extent), the results are of high quality.

The quality of the volumetric reconstruction does affect the results. A blue screen was used on the Mulberry tree and we can see that the results are slightly better. Adverse capture conditions such as the impossibility to use a blue screen to aid matte extraction due to size and the difficulty of calibrating cameras using scene features (see Figure 20 for the capture conditions of the Pine tree) produces somewhat lower quality relighting results as can be seen in Figure 21.

We believe that improvements in vision and image processing algorithms, which are beyond the scope of this paper, will allow the creation of better mattes and higher quality camera calibration, thus improving the results of our approach, and removing the need for the heuristics for the volumetric reconstruction.

Qualitative comparisons between the real trees and the synthetic results show that our approach works slightly better on the real trees photographs. We believe that this is due to the quality of the geometric model used for the synthetic renderings. Although each synthetic tree is on the order of $300K$ triangles, individual leaves were flat, containing only a few faces. As a result a lot of the geometric subtlety is missing, which allows the volumetric reconstruction method to produce good results.



Fig. 20. Capture conditions for the Pine Tree (Figure 11).

7.3 Limitations

Our method is intended for distant cameras only, such that the entire canopy is present in the image. We thus

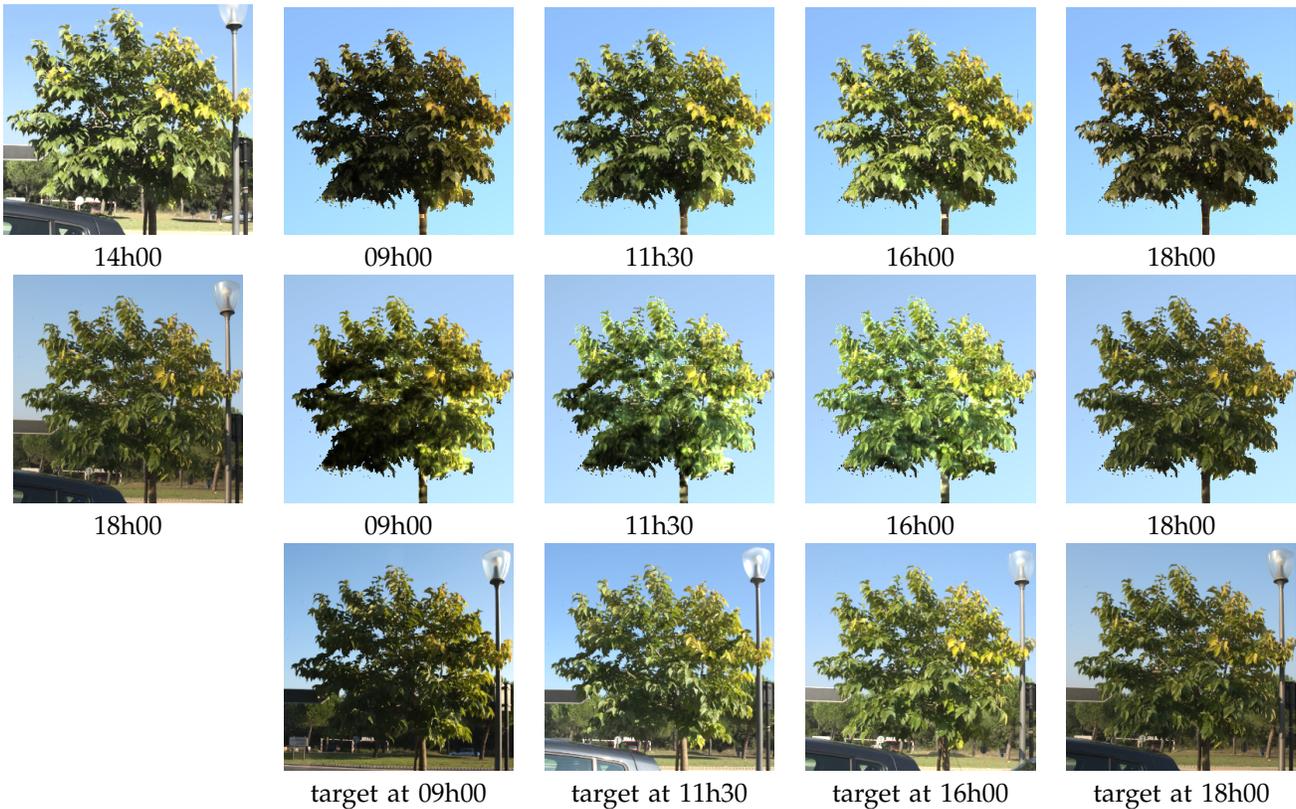


Fig. 19. **Mulberry tree** In contrast to Figure 14, in this example we show relit results when the input data photographs were taken at different times. Top row: relit images using input data taken at 14h00. Middle row: relit images using input data taken at 18h00. Bottom row: target images with corresponding times of day.

propose a solution to the specific case when no 3D information can be reliably captured in a pixel accurate way, and robustly enough to handle small variations in the geometry (presence of slight wind). If the camera gets close enough to the tree to be able to perceive large regions of individual leaves and branches, our volumetric assumption would not hold.

It would be possible to use a high-dynamic range (HDR) light-probe to capture the input sky. However, the advantage of using the Preetham model is that it is available for both the input and the target times. Such an additional capture overhead is thus unnecessary for our approach, and would defeat the purpose of our method which is to have no lighting information at the target time. Additionally, using different models for the input and target skies will inevitably lead to inconsistencies and probably give worse results.

We assume clear skies in our method since these are well simulated by the Preetham model [11]. To our knowledge, overcast skies cannot be accurately and consistently represented to allow our ratio-based approach to work.

Due to the small number of SH bands used for reconstruction of natural skies created by the Preetham model, our method is expected to produce lower quality results for higher frequency lighting such as those encountered in some high frequency captured HDR environment

maps. Additionally, our method requires consistent *input* and *target* lighting models to work. Since our method focuses on minimizing the capture procedure to a *single* time of day, there is no captured HDR available for the target time.

Although our method can relight a canopy for different lighting conditions, such as winter or summer sun positions, our method cannot simulate changes in the leaf structure of the tree canopy that occurs between two extremely different seasons - this would require explicit geometric reconstruction of leaves and branches, and appropriate modification.

Finally, our method depends heavily on the quality of the volumetric reconstruction. Irregularly shaped trees with complex isolated branch structures pose a difficult problem for reconstruction, and thus our relighting results are of somewhat lower quality. This can be seen in Figure 12 at 14h30 and 17h00.

8 DISCUSSION AND CONCLUSIONS

Other approaches could be envisaged to solve the tree canopy relighting problem. For example, texture synthesis could be used rather than a ratio, in which we would search for similar luminance pixels in the input. Our experiments with such an approach showed that the visual quality was not as good as that presented

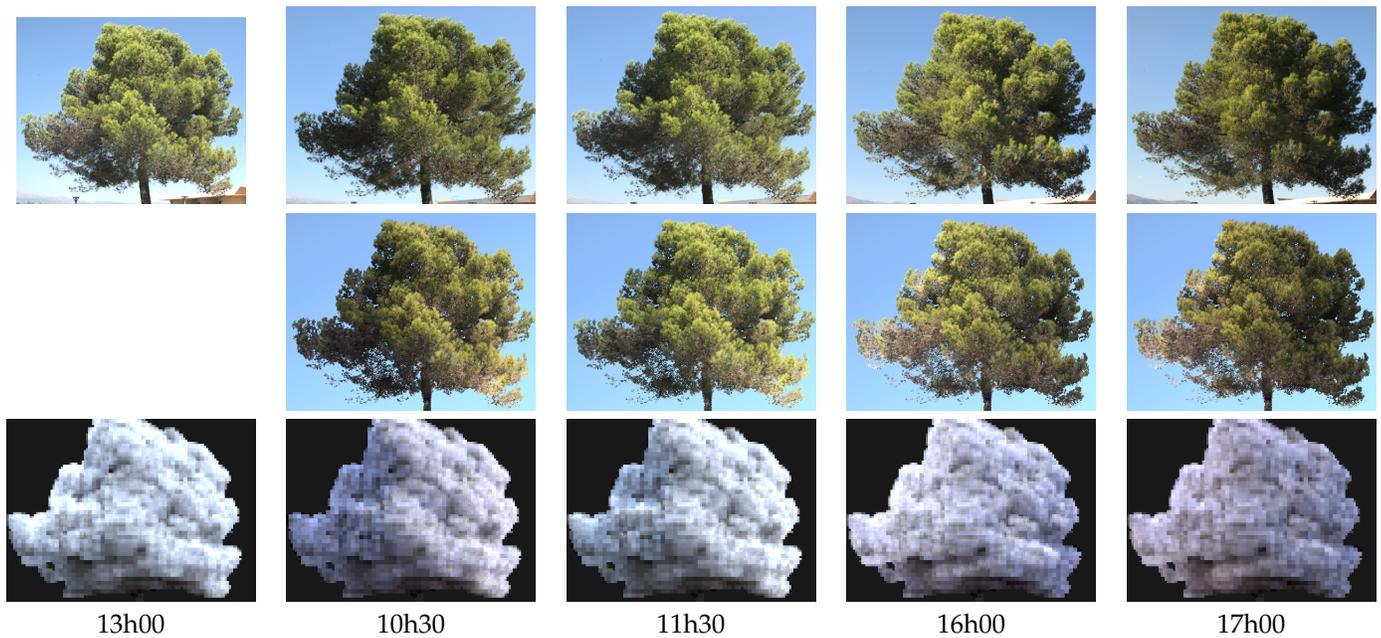


Fig. 21. **Pine tree** Top row: input images and 4 target images with corresponding times of day. Middle row: 4 resulting relit images using our approach. Bottom row: \tilde{E}_{in} and the four \tilde{E}_{targ} images.

here, since the synthesis stage alters the image and thus degrades the overall quality of the result. A hybrid volumetric/geometric approach can also be considered, in which 3D leaf positions of pixels are estimated to perform relighting. We also experimented with this idea, and despite promising initial results, extensive tests showed that the pixels positions could not be reliably estimated in the general case.

Concerning the application in Figure 1(a)-(b), it could be argued that simple histogram transfer would suffice. As we can see in Figure 22, colors are correctly reproduced, but lighting is incorrect: the right side of the tree should be in shadow. This is correctly reproduced by our approach (Figure 1(c)).

While beyond the scope of this paper, it would of course be highly desirable to be able to relight the entire environment including the tree canopies. We consider this to be important future work, starting for example with shadow removal (including the cast shadows from the tree canopies and trunks) and then treating general relighting. We also expect our method to have somewhat reduced performance on very sparse trees, or trees where leaves have a very strong preferred orientation.

The volumetric model should be applicable to all photographs of materials which are either truly volumetric (e.g., clouds or smoke), or have behavior which is similar to a volume (e.g., large collections of small objects or grass). While the model presented here is not directly transposable on all geometries and all scales, we do believe that some of the ideas presented here could well generalize to non-volumetric materials. In addition, we believe that our approach could fit well with more traditional approaches such as [6], resulting in a general

and complete relighting method.

ACKNOWLEDGMENTS

We would like to thank Sylvain Paris, Frédo Durand, Ares Lagae and the reviewers of an earlier version for their input and suggestions. We also thank David Grelaud, Gaurav Chaurasia, Matthaus Chajdas and Fernanda Andrade Cabral for their help with the tree photography sessions. Finally we thank our colleagues from REVES/INRIA Sophia-Antipolis for fruitful discussions throughout the development of this project.

REFERENCES

- [1] Y. Yu and J. Malik, "Recovering photometric properties of architectural scenes from photographs," in *SIGGRAPH '98*. New York, NY, USA: ACM, 1998, pp. 207–217.
- [2] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: recovering reflectance models of real scenes from photographs," in *SIGGRAPH '99*, 1999.
- [3] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems, "Relighting with 4d incident light fields," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 613–620, 2003.
- [4] S. Boivin and A. Gagalowicz, "Image-based rendering of diffuse, specular and glossy surfaces from a single image," in *SIGGRAPH '01*, 2001, pp. 107–116.
- [5] C. Loscos, G. Drettakis, and L. Robert, "Interactive virtual relighting of real scenes," *IEEE Trans. on Visualization and Computer Graphics*, vol. 6, no. 3, July–September 2000. [Online]. Available: <http://www-sop.inria.fr/revs/Basilic/2000/LDR00>
- [6] P. Debevec, "Image-based lighting," *IEEE Computer Graphics and Applications*, vol. 22, no. 2, pp. 26–34, 2002.
- [7] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *SIGGRAPH '00*, 2000, pp. 145–156.
- [8] P. Peers, N. Tamura, W. Matusik, and P. Debevec, "Post-production facial performance relighting using reflectance transfer," *ACM Trans. Graph.*, vol. 26, no. 3, p. 52, 2007.



Fig. 22. (Left): The input image has been pasted to the background after an histogram transfer. Although the colors are correctly reproduced, the overall lighting direction, and corresponding shadows, are not well captured compared to our method (Right).

- [9] B. Neubert, T. Franken, and O. Deussen, "Approximate image-based tree-modeling using particle flows," *ACM Trans. on Graphics (SIGGRAPH 2007)*, vol. 26, no. 3, p. 88, 2007.
- [10] P. Tan, T. Fang, J. Xiao, P. Zhao, and L. Quan, "Single image tree modeling," *ACM Trans. Graph. (SIGGRAPH Asia)*, vol. 27, no. 5, pp. 1–7, 2008.
- [11] A. J. Preetham, P. Shirley, and B. Smits, "A practical analytic model for daylight," in *SIGGRAPH*, 1999.
- [12] I. Sato, Y. Sato, and K. Ikeuchi, "Acquiring a radiance distribution to superimpose virtual objects onto a real scene," *IEEE Trans. on Visualization and Computer Graphics*, vol. 5, no. 1, pp. 1–12, January 1999. [Online]. Available: <http://www.hci.iis.u-tokyo.ac.jp/~imarik/publications/TVCG99.pdf>
- [13] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan, "Webcam clip art: Appearance and illuminant transfer from time-lapse sequences," *ACM Trans. on Graphics (SIGGRAPH Asia 2009)*, vol. 28, no. 5, December 2009.
- [14] C. Soler, F. X. Sillion, F. Blaise, and P. Dereffye, "An efficient instantiation algorithm for simulating radiant energy transfer in plant models," *ACM Trans. Graph.*, vol. 22, no. 2, pp. 204–233, 2003.
- [15] F. de Castro and N. Fetcher, "Three dimensional model of the interception of light by a canopy," *Agricultural and Forest Meteorology*, vol. 90, no. 3, pp. 215 – 233, 1998.
- [16] E. Cerezo, F. Perez-Cazorla, X. Pueyo, F. Seron, and F. Sillion, "A survey on participating media rendering techniques," *the Visual Computer*, 2005.
- [17] K. Boulanger, K. Bouatouch, and S. N. Pattanaik, "Rendering trees with indirect lighting in real time," *Comput. Graph. Forum*, vol. 27, no. 4, pp. 1189–1198, 2008.
- [18] I. Shlyakhter, M. Rozenoer, J. Dorsey, and S. J. Teller, "Reconstructing 3d tree models from instrumented photographs," *IEEE Comp. Graphics and App.*, vol. 21, no. 3, 2001.
- [19] A. Reche, I. Martin, and G. Drettakis, "Volumetric reconstruction and interactive rendering of trees from photographs," *ACM Trans. on Graphics (SIGGRAPH Conference Proceedings)*, vol. 23, no. 3, July 2004. [Online]. Available: <http://www-sop.inria.fr/reves/Basilic/2004/RMD04>
- [20] P. Debevec, "Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography," in *SIGGRAPH '98*, 1998, pp. 189–198.
- [21] H. Rahman, B. Pinty, and M. Verstraete, "Coupled surface-atmosphere reflectance (csar) model 2. semiempirical surface model usable with noaa advanced very high resolution radiometer data," *J. Geophys. Res.*, vol. 98(D11), 1993.
- [22] S. R. Marschner and D. P. Greenberg, "Inverse lighting for photography," in *Proceedings of the Fifth Color Imaging Conference, Society for Imaging Science and Technology*, 1997, pp. 262–265.
- [23] Z. Liu, Y. Shan, and Z. Zhang, "Expressive expression mapping with ratio images," in *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 2001, pp. 271–276.
- [24] A. Shashua and T. Riklin-Raviv, "The quotient image: Class-based re-rendering and recognition with varying illuminations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 2, pp. 129–139, 2001.
- [25] A. Stoschek, "Image-based re-rendering of faces for continuous pose and illumination directions," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 1, p. 1582, 2000.
- [26] M. Pharr and G. Humphreys, *Physically Based Rendering: From Theory to Implementation*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2004.
- [27] F. Perez-Cazorla, X. Pueyo, and F. Sillion, "Global illumination techniques for the simulation of participating media," in *8th Eurographics Wkshp on Rendering*, June 1997.
- [28] J. T. Kajiya and B. P. Von Herzen, "Ray tracing volume densities," *SIGGRAPH Comput. Graph.*, vol. 18, no. 3, pp. 165–174, 1984.
- [29] F. Kasten and A. Young, "Revised optical air mass tables and approximation formula," in *Applied Optics*, vol. 28, 1989, pp. 4735–4738.
- [30] P.-P. Sloan, J. Kautz, and J. Snyder, "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments," in *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 2002, pp. 527–536.
- [31] J. Jansen and L. Bavoil, "Fourier opacity mapping," in *I3D '10: Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*. New York, NY, USA: ACM, 2010, pp. 165–172.
- [32] R. Habel, B. Mustata, and M. Wimmer, "Efficient spherical harmonics lighting with the preetham skylight model," in *Eurographics 2008 - Short Papers*, Apr. 2008. [Online]. Available: http://www.cg.tuwien.ac.at/research/publications/2008/Habel_08_SSH/
- [33] A. Levin, D. Lischinski, and Y. Weiss, "A closed-form solution to natural image matting," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 2, pp. 228–242, 2008.
- [34] N. Snavely, S. M. Seitz, and R. Szeliski, "Modeling the world from Internet photo collections," *International Journal of Computer Vision*, vol. 80, no. 2, pp. 189–210, November 2008. [Online]. Available: <http://phototour.cs.washington.edu/>
- [35] G. Zotti, A. Wilkie, and W. Purgathofer, "A critical review of the preetham skylight model," in *Journal of WSCG 2007*, V. Skala, Ed., 15, 2007.



Marcio Cabral is a PhD candidate at the REVES group of INRIA Sophia Antipolis, under the supervision of George Drettakis. His PhD focuses on user-aided content creation, using techniques such as image-based relighting and interactive mesh editing with constraint-based systems. His current research is on 3D immersive interaction for content creation. Before joining INRIA, he completed his undergraduate and Master studies at University of São Paulo, Brazil, where he participated in the installation and development of applications for the first VR CAVETM system in Latin America. In 2005, he co-authored the ACM SIGGRAPH Emerging Technologies exhibit *Virtual hang-gliding over Rio de Janeiro*.



Nicolas Bonneel received a PhD degree from INRIA-REVES in Sophia Antipolis, France, under the supervision of George Drettakis, in 2009. He participated in the European project CROSSMOD on audio-visual perceptual rendering. His research focused on audio and visual rendering, reflectance acquisition and texture synthesis. He is now a postdoctoral researcher at the University of British Columbia (UBC) in Vancouver, Canada, under the supervision of Michiel van de Panne, working on interpolation techniques.



Sylvain Lefebvre is a researcher at INRIA, France. He completed his PhD in 2004, on the topic of texturing and procedural texture generation using GPUs. After graduating he joined Microsoft Research (Redmond, USA) as a postdoctoral researcher during the year 2005. His current research focuses are in procedural content generation, end-user content manipulation and compact data structures for interactive applications and games. Sylvain received the EUROGRAPHICS Young Researcher Award in 2010 for his work on texture synthesis and GPU data-structures. More details about his research can be found at <http://www.loria.fr/~slefebvr>



Dr. George Drettakis (PhD, Toronto 1994) is head of the research group REVES, specializing on Rendering for Virtual Environments with Sound, at INRIA Sophia-Antipolis. He has published many papers in graphics and 3D sound, in particular on global illumination, visibility, relighting and augmented reality and interactive rendering at all leading international conferences and journals. He has coordinated or participated in several EU (SIMULGEN, ARCADE, CREATE and CROSSMOD) and French national projects (ARC ARCHEOS, ARC NIEVE). He has been part of program committees of many international conferences, has co-chaired EGWR 98 and the Eurographics 2002 and 2008 Conferences, and is papers chair for ACM SIGGRAPH Asia 2010. He won the Eurographics Outstanding Technical Achievement Award in 2007 and is a Eurographics Fellow. His current research interests include image-based rendering and relighting; procedural textures, texture synthesis and weathering; fast illumination/shadow algorithms as well as perceptual aspects of audio and visual rendering.