



Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views

Pierre-Yves Laffont, Adrien Bousseau, George Drettakis

► To cite this version:

Pierre-Yves Laffont, Adrien Bousseau, George Drettakis. Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views. IEEE Transactions on Visualization and Computer Graphics, 2013, 19 (2). hal-00761121

HAL Id: hal-00761121

<https://inria.hal.science/hal-00761121>

Submitted on 5 Dec 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views

Pierre-Yves Laffont, Adrien Bousseau and George Drettakis

Abstract—Intrinsic images aim at separating an image into its reflectance and illumination components to facilitate further analysis or manipulation. This separation is severely ill-posed and the most successful methods rely on user indications or precise geometry to resolve the ambiguities inherent to this problem. In this paper we propose a method to estimate intrinsic images from multiple views of an outdoor scene without the need for precise geometry and with a few manual steps to calibrate the input. We use multiview stereo to automatically reconstruct a 3D point cloud of the scene. Although this point cloud is sparse and incomplete, we show that it provides the necessary information to compute plausible sky and indirect illumination at each 3D point. We then introduce an optimization method to estimate sun visibility over the point cloud. This algorithm compensates for the lack of accurate geometry and allows the extraction of precise shadows in the final image. We finally propagate the information computed over the sparse point cloud to every pixel in the photograph using image-guided propagation. Our propagation not only separates reflectance from illumination, but also decomposes the illumination into a sun, sky and indirect layer. This rich decomposition allows novel image manipulations as demonstrated by our results.

Index Terms—Intrinsic images, image-guided propagation, multi-view stereo, mean-shift algorithm

1 INTRODUCTION

Editing materials and lighting is a common image manipulation task that requires significant expertise to achieve consistent results. The main difficulty in such manipulations resides in the fact that a pixel color aggregates the effect of both material and lighting, so that standard color manipulations are likely to affect both components.

A typical usage scenario would be to take a photograph and manipulate its content *after* the capture, e.g., to change the color of an object, soften the shadows, or add in virtual objects. In this paper we present a method to achieve this just by taking a few extra photographs, which represents a minimal additional cost. We focus on outdoor scenes and separate photographs into a material layer (also called *reflectance*) and several *illumination* layers that describe the contributions of sun, sky and indirect lighting (Fig. 1e-h). This *intrinsic images* decomposition [1] allows easy editing of each component separately, and subsequent compositing with consistent lighting. We illustrate applications of our decomposition in Fig. 1. We first alter the floor material with a graffiti while maintaining consistent shadows (b); we then add a virtual object in the scene with consistent lighting (c); and we finally change the lighting color and blur the shadows to simulate sunset (d). Please also see the video which illustrates the image manipulation process.

Estimating intrinsic images from photographs is an ill-posed problem that has been traditionally addressed with restrictive assumptions on reflectance and illumination variations [2], significant user intervention [3], or multiple images under varying lighting [4] captured over several hours

for outdoor scenes. In addition, most existing methods only generate a single illumination layer, grouping the effects of sun, sky and indirect lighting. Inverse global illumination methods [5], [6] also estimate the material properties of a scene but require very accurate geometric models, either built by hand or acquired with complex laser scanners.

We address the above shortcomings and generate intrinsic image layers for sun, sky and indirect illumination. Our method combines sparse geometric reconstruction [7], [8] with image-guided propagation [9], thus leveraging their respective strengths. We exploit the automatically reconstructed 3D information to compute lighting information for a subset of pixels, and use image-guided propagation to decompose the photographs into intrinsic images. While our method requires users to capture additional images, it alleviates the need for user intervention apart from some initial calibration steps.

Our algorithm takes as input a small number of photographs of the scene captured *at a single time of the day*, along with an environment map that represents the illumination coming from the sky and distant environment in all directions. From this lightweight capture we use recent computer vision algorithms to reconstruct a sparse 3D point cloud of the scene. The point cloud only provides an imprecise and incomplete representation of the scene. However, we show that this is sufficient to compute plausible sky and indirect illumination at each reconstructed 3D point, without complex inverse global illumination computations. The coarse geometry is however unreliable for sun illumination that typically contains high-frequency features such as cast shadows. We introduce a new parameterization of reflectance with respect to sun visibility that we integrate in an optimization algorithm to robustly identify the 3D points that are in shadow. We developed an optimization inspired




Fig. 1: Starting from multiple views of the scene (a), our method decomposes photographs into four intrinsic components — the reflectance (e), the illumination due to sun (f), the illumination due to sky (g) and the indirect illumination (h). Each intrinsic component can then be manipulated independently for advanced image editing applications (b-d).

by mean shift [10] where we use asymmetric regions of influence and constrain the evolution of the estimates.

Image-guided propagation algorithms are typically used to propagate user scribbles [3], [9]; we show how to use these algorithms to propagate the illumination information computed at 3D points over the image pixels. Our approach generates intrinsic images of similar quality as scribble-based approaches, with only a small amount of user intervention for capture and calibration, notably to find the sun position, orient the environment, and calibrate the sphere used to capture the environment map. In addition, our ability to separate sun, sky and indirect illumination opens the door for advanced image manipulations, as demonstrated in Fig. 1b-d.

In summary, this paper makes the following contributions:

- We show how to compute sky, indirect, and sun (ignoring cast shadows) illumination at automatically reconstructed 3D points, using incomplete and imprecise geometry and a small set of input images.
- We introduce an algorithm to reliably identify points in shadow based on a new parameterization of the reflectance with respect to sun visibility. Our algorithm compensates for the lack of accurately reconstructed and complete 3D information.
- We show how to propagate reflectance, sun, sky and indirect illumination to all pixels in an image, without user intervention or involved inverse global illumination computation. We achieve this by using the illumination values computed at 3D points as constraints for image propagation algorithms.

After a presentation of previous work, the definition of our image formation model and a description of capture, the structure of our paper follows our three contributions.

2 RELATED WORK

Intrinsic images. Several methods have been proposed to estimate reflectance and illumination from a single image. This decomposition is severely ill-posed and can only be solved with additional knowledge or assumptions about the scene content. The Retinex algorithm [2], [11] assumes smooth illumination and piece-wise constant reflectance, while Shen and Yeo [12] assume that neighboring pixels with similar chromaticity have the same reflectance and that the image is composed of a small set of reflectances. Tappen et al. [13] train a classifier to discriminate image derivatives due to reflectance and illumination, and Shen et al. [14] introduce texture constraints to ensure that pixels with similar texture share the same reflectance. These various approaches produce encouraging decompositions on isolated objects, as evaluated by the ground-truth dataset of Grosse et al. [15]. However, the automatic decomposition of complex outdoor images remains an open challenge, in part because most existing methods assume monochrome lighting while outdoor scenes are often lit by a mixture of colored sun and sky light.

Weiss [4] demonstrates how multiple images of a scene under different illuminations can be factored into a reflectance image and a set of illumination images. Sunkavalli et al. [16] decompose similar image sequences of outdoor scenes into a shadow mask and images illuminated only by skylight or sunlight. By also capturing an environment map at multiple times of the day, Matusik et al. [17] estimate the reflectance field of an outdoor scene that can then be used for relighting. These methods assume a fixed viewpoint and varying illumination (i.e., *timelapse* sequences), while our method relies on images captured under different viewpoints and fixed illumination. The main advantage of our capture approach is to reduce the acquisition time to a few minutes while timelapses typically require at least several hours of capture to cover as many lighting directions as

possible. Most related to our capture strategy is the system of Melendez et al. [18] that relights buildings reconstructed from multiple photographs. However their method necessitates the additional acquisition of flash / no-flash image pairs to capture a *material exemplar* for every material of a building. Users then need to associate a material exemplar to each texture region of the reconstructed building.

Bousseau et al. [3] and Shen et al. [19] rely on user scribbles to disambiguate reflectance from illumination while Okabe et al. [20] propagate a sparse set of user-specified normals over the image. Given an intrinsic image decomposition, Carroll et al. [21] propose a user-assisted decomposition to isolate the indirect contribution of each colored material in the scene. In work developed concurrently, Karsch et al. [22] annotate a single photograph to extract a model of the scene that is suitable for inserting virtual objects. While our method shares similarities with these user-assisted approaches, we ask users to capture a few additional pictures instead of providing scribbles or annotations. Putting more effort on the capture allows us to automate the computation of illumination constraints and to extract multiple illumination components (sun, sky, indirect) that would be hard to disambiguate by users.

Inverse rendering. Inverse rendering methods [5], [6], [23], [24] recover the reflectance and illumination of a scene by inverting the rendering equation [25]. These methods require an accurate 3D model of the scene that is either modeled manually or acquired with expensive laser scanners and then solve an often costly global illumination system. In contrast our method is robust to incomplete geometry and handles sparse point clouds automatically reconstructed from a few photographs of the scene. In addition, we are able to estimate indirect illumination at every pixel without the need to solve inverse global illumination. Nonetheless, inverse rendering methods allow applications that are beyond the scope of this paper, such as free viewpoint navigation and dynamic relighting.

Image-based propagation. Image guided interpolation methods have been introduced by Levin and Lischinski to propagate colors [9] and tonal adjustments [26]. In this paper we use the propagation algorithm of Bousseau et al. [3] that was originally designed to propagate user indications for intrinsic image decompositions. This algorithm is inspired by the *matting Laplacian* of Levin and colleagues [27] that has been used to decompose an image into a foreground and background layer, and to recover white balanced images under mixed lighting [28].

Shadow removal. Our work is also related to shadow removal methods [29], [30], [31], [32], [33] that aim at removing cast shadows in an image, either automatically or with user assistance. While our method is also able to identify cast shadows, our main goal is to extract a reflectance image as well as smooth illumination variations. We also separate the contribution of sun, sky and indirect illumination, which enables novel image manipulations.

3 IMAGE FORMATION MODEL

We assume Lambertian surfaces and model the image values at each pixel as the product between the incident illumination and the object reflectance \mathbf{R} . Formally, the radiance \mathbf{I} towards the camera at each non-emissive, visible point corresponding to a pixel is given by the equation:

$$\mathbf{I} = \mathbf{R} * \int_{\Omega} \cos \theta_{\omega} \mathbf{L}(\omega) d\omega \quad (1)$$

where we integrate over the hemisphere Ω centered on the normal at the visible point, $\mathbf{L}(\omega)$ is the incoming radiance in direction ω , θ_{ω} is the angle between the normal at the visible point and direction ω . Capital bold letters represent RGB color values and $*$ denotes per-channel multiplication.

For our purposes, we will separate out the incoming radiance into three components: the radiance due to the sun, that due to the sky and that due to indirect lighting. To simplify notation, we define two subsets of the hemisphere: Ω_{sky} , i.e., the subset of directions in which the visible point sees the sky, and Ω_{ind} the subset of directions in which another object is visible, and thus contributes to indirect lighting. We however explicitly represent the sun visibility v_{sun} , first because precise computation of v_{sun} is necessary to capture sharp shadows, and second because estimating v_{sun} robustly is one of our main contributions.

We can now re-write Equation 1:

$$\mathbf{I} = \mathbf{R} * \left(v_{sun} \max(0, \cos \theta_{sun}) \mathbf{L}_{sun} + \int_{\Omega_{sky}} \cos \theta_{\omega} \mathbf{L}_{sky}(\omega) d\omega + \int_{\Omega_{ind}} \cos \theta_{\omega} \mathbf{L}_{ind}(\omega) d\omega \right)$$

where \mathbf{L}_{sun} , \mathbf{L}_{sky} and \mathbf{L}_{ind} are radiance from the sun, the sky and indirect lighting respectively, θ_{sun} is the angle between the normal at the visible point and the sun modeled as a directional light source, and θ_{ω} is the angle between the normal and the direction of integration ω over the hemisphere. The scalar $v_{sun} \in [0, 1]$ models the visibility of the sun (0 for completely hidden, 1 for fully visible).

We next define simplified quantities at each pixel:

$$\mathbf{S}_{sun} = v_{sun} \max(0, \cos \theta_{sun}) \mathbf{L}_{sun} = v_{sun} \hat{\mathbf{S}}_{sun} \quad (2)$$

$$\mathbf{S}_{sky} = \int_{\Omega_{sky}} \cos \theta_{\omega} \mathbf{L}_{sky}(\omega) d\omega \quad (3)$$

$$\mathbf{S}_{ind} = \int_{\Omega_{ind}} \cos \theta_{\omega} \mathbf{L}_{ind}(\omega) d\omega. \quad (4)$$

where $\hat{\mathbf{S}}_{sun}$ corresponds to the sun illumination when cast shadows are ignored. We define from these quantities a simplified image formation model:

$$\mathbf{I} = \mathbf{R} * (\mathbf{S}_{sun} + \mathbf{S}_{sky} + \mathbf{S}_{ind}) \quad (5)$$

$$= \mathbf{R} * \mathbf{S}_{total} \quad (6)$$

where \mathbf{R} is the object RGB reflectance. \mathbf{S}_{sun} , \mathbf{S}_{sky} and \mathbf{S}_{ind} are the RGB incident illumination (or shading) from the sun, sky and indirect lighting respectively.




Fig. 2: Overview of our approach. Users capture a small set of pictures of the scene, along with an environment map and two pictures of a gray card in sun light and in shadow (a). We illustrate our intrinsic image decomposition with the picture highlighted in orange. We use multi-view stereo to reconstruct a point cloud of the scene and a coarse proxy geometry (b). Users align the environment map and the sun with this point cloud and use the gray card to calibrate their intensity. Once this calibration is performed, all the remaining steps are automatic. We use the reconstructed 3D geometry to compute sun, sky and indirect lighting over the point cloud (c). We also compute an initial guess of the sun visibility using the coarse proxy. These lighting values give us the necessary information to compute a set of candidate reflectances for each 3D point. The candidate reflectances form curves in color space parametrized by the sun visibility (d). We introduce an iterative optimization that identifies the reflectance of each 3D point from these candidates, along with a precise estimation of the sun visibility. The final step of our method consists in propagating the illumination values computed at 3D points to every pixel in the image (e). We decompose the propagated illumination into the sun, sky and indirect lighting components.

Our first goal is to extract the reflectance \mathbf{R} and illumination $\mathbf{S}_{\text{total}}$ from this image formation model. We demonstrate how to make this problem tractable by leveraging the sparse geometric information generated by multiview stereo algorithms to compute $\hat{\mathbf{S}}_{\text{sun}}$, \mathbf{S}_{sky} and \mathbf{S}_{ind} at each reconstructed 3D point (Section 5). We then introduce a new algorithm to estimate the sun visibility v_{sun} precisely despite the approximate available geometry (Section 6), which will allow us to obtain all illumination components using image-guided propagation (Section 7). Fig. 2 illustrates the main steps of our approach in the form of a block diagram.

4 CAPTURE AND RECONSTRUCTION



Our method relies on a lightweight capture setup composed of a digital camera (preferably on a tripod), a photographer’s gray card and a simple reflective sphere¹ to capture the

1. The reflective sphere does not need to be as accurate as described in [6]; in practice we used an inexpensive *pétanque* ball.

environment map (see inset). No other special capture or measurement hardware is required.

4.1 Photography

We first capture a few ordinary low-dynamic range (LDR) photographs which we use to perform approximate geometric reconstruction of the scene. This set of photographs should have a good coverage of the scene from different viewpoints and sufficient overlap between neighboring viewpoints to facilitate multiview stereo reconstruction. The number of photographs required to obtain an acceptable reconstruction depends on the complexity of the scene and the presence of image features. We captured between 10 and 31 LDR photographs for the scenes in this paper.

We then capture two linear, high-dynamic range (HDR), images of the front and side of the reflective sphere, placed in the scene, to obtain an angular environment map of the scene. We do this using the standard HDR assembly technique of Debevec et al. [34].




Fig. 3: Multiview geometry provides us with a sparse oriented point cloud (b, red pixels). Compared to the original image (a), the initial guess of the sun shadows computed using the proxy reconstructed from the point cloud is very inaccurate due to geometric errors, and incomplete (c). In particular, a bush casting a large shadow on the floor (d) is not reconstructed.

We finally capture linear HDR images of the viewpoints that we want to decompose. Recall that we capture all images in one session, at a single time of day.

4.2 Geometry and Illuminant Calibration

We apply structure-from-motion using Bundler [7] and the patch-based multi-view stereo (PMVS) algorithm [8] on the set of LDR+HDR photographs – using the publicly available implementations². The result of this process is an oriented point cloud (3D positions and normals), and calibrated cameras (extrinsic and intrinsic parameters). The process also returns whether each point is visible from each camera. We rectify the images for radial distortion, and we consider only visible PMVS points in each image.

We found the PMVS normals to be too noisy for illumination computation. We instead estimate normals at each 3D point by fitting a plane on the local point cloud using PCA [35]. We discard 3D points for which the neighborhood is too sparse or degenerate. We consider a neighborhood as degenerate when the first singular value of the PCA is twice greater than the second one.

We recover a geometric proxy from the oriented point cloud using the scale space meshing method of Digne et al. [36]. This automatic approach produces detailed accurate meshes in regions where the point cloud is dense, while leaving holes in areas where the point cloud is irregularly sampled. We use the automatic hole filling tool available

in MeshLab [37] to further improve the reconstruction in those areas. We also experimented with the Poisson reconstruction of Kazhdan et al. [38]. Although this algorithm generates a closed surface with no holes, we found that it tends to produce bumpy surfaces due to the irregular sampling of the point cloud. We nevertheless used the Poisson reconstruction on the “Rocks” scene for which the point cloud is dense and uniform. Note that recent approaches could improve the quality of reconstructions in specific scenarios, for instance by detecting planar regions (e.g., [39], [40], [41]) in urban scenes.

Some manual interaction is required to calibrate the sun and sky illumination in our current system. The user specifies the sun position, using cues from the shadows in the image and reconstructed geometry. She labels sky pixel in the environment map, and rotates it until it aligns with the specified sun position and the scene horizon. Finally, she takes two photographs of a gray card placed in the scene, in sunlight and in shadow, in order to estimate the color transfer function of the reflective sphere and the radiance L_{sun} of the sun. The details of this process are described in the Appendix.

The output of our capture is: a moderately dense point cloud reconstruction of the scene which we will refer to as the *PMVS points* (Fig. 3b), a very approximate geometric *proxy*, which often contains significant geometric errors (Fig. 3c), the direction and radiance of the sun and a correctly aligned and scaled HDR environment map containing the sky and distant indirect radiance (Fig. 5c).

5 GEOMETRY-BASED COMPUTATION

We describe here how to compute sun, sky and indirect illumination values for each PMVS point. These points have been generated using multiview stereo and also have normals (Section 4.2).

We first compute sun illumination \hat{S}_{sun} ignoring cast shadows, i.e., unoccluded sunlight. We already know the required quantities for this computation, i.e., the normal at the point and the sun radiance and direction. Treating sun visibility requires much higher precision than the one provided by the proxy, and is treated separately in Section 6.

We then compute sky and indirect illumination at each PMVS point. Fig. 4 illustrates this computation that we detail below. In a nutshell, we use the HDR environment map to compute both sky illumination and *distant* indirect illumination, while we compute the *near-field* indirect illumination from the proxy geometry. Note however that we do not need to know the reflectance of the proxy for this step, we instead use the captured photographs to recover the necessary outgoing radiance over the geometry.

Assigning radiance to the proxy geometry. We first assign radiance to each PMVS point by looking up its pixel values in each HDR image where it appears (Fig. 4). We assign the average value as the outgoing radiance of

². Available at <http://phototour.cs.washington.edu/bundler/> and <http://grail.cs.washington.edu/software/pmvs/>, as of 2011/10/10




Fig. 4: Evaluation of sky and indirect illumination at each PMVS point. We look up the radiance value of the PMVS points in each image and average the values over the images where they are visible (a). We then project and interpolate this radiance over the proxy geometry (b). Finally we gather the sky and indirect illumination at each PMVS point by shooting rays which sample the hemisphere (c, dotted lines). Rays that reach sky pixels in the environment map (c, blue lines) contribute to sky illumination, while other rays contribute to indirect illumination (c, dark lines).




Fig. 5: We estimate sky and indirect illumination at PMVS points using (b) a colored proxy, obtained by projecting and interpolating the radiance of PMVS points on the reconstructed geometry, and (c) an environment map, shown in a latitude-longitude parameterization where the red curve separates pixels contributing to sky illumination and those contributing to distant indirect illumination.

the point, which is assumed constant in all directions. We then assign the radiance of the closest PMVS point to each vertex of the proxy, and interpolate these values over the faces of the mesh. This step yields a colored proxy that

stores the outgoing radiance of the scene as captured in the photographs (Fig. 5b).

Sky and indirect illumination. We separate the environment map into two regions (see Fig. 5c), one for the sky and the other one for distant objects. At every PMVS point we cast a set of rays towards the hemisphere. If a ray reaches a sky pixel of the environment map, the resulting radiance contributes to sky illumination \mathbf{S}_{sky} . Otherwise the ray contributes to indirect illumination \mathbf{S}_{ind} , either from an intersection with the proxy or from a non-sky pixel of the environment map. If the ray hits the proxy, it uses the radiance assigned to the geometry as described in the previous paragraph.

We rewrite Equations 3 and 4 to express this computation, changing the integration domain to the entire hemisphere Ω for both integrals, and introducing a visibility term $v_{\text{sky}}(\omega)$ which is 1 when the ray in direction ω reaches the sky, and 0 otherwise (it contributes to indirect illumination):

$$\mathbf{S}_{\text{sky}} = \int_{\Omega} v_{\text{sky}}(\omega) \cos \theta_{\omega} \mathbf{L}_{\text{sky}}(\omega) d\omega \quad (7)$$

$$\mathbf{S}_{\text{ind}} = \int_{\Omega} (1 - v_{\text{sky}}(\omega)) \cos \theta_{\omega} \mathbf{L}_{\text{ind}}(\omega) d\omega. \quad (8)$$

The computation of Equations 7 and 8 is robust to the coarse geometry of the proxy since the integration over Ω averages the values over the entire hemisphere. We implement this computation with a custom renderer in the PBRT stochastic raytracer [42].

Our approach shares similarities with the techniques of Yu and Malik [23], but while their method was designed for accurate geometry constructed manually, we handle sparse incomplete geometry reconstructed automatically.

Finally we compute an initial guess of the sun visibility $v_{\text{sun}}^{\text{init}}$ at each PMVS point by intersecting a ray with the proxy geometry in the direction of the sun. Note however that this visibility test is very sensitive to errors in the reconstructed proxy, as illustrated in Fig. 3c. This fact underlines the importance of accurately estimating v_{sun} and we next show how to refine this initial estimate.

6 ESTIMATING SUN VISIBILITY AT 3D POINTS

One key contribution of our approach is a novel algorithm for identifying visibility v_{sun} with respect to the sun for each PMVS point. From our image formation model in Equations 2 and 5 we express the reflectance at each PMVS point as a function of the visibility term:

$$\mathbf{R}(v_{\text{sun}}) = \frac{\mathbf{I}}{(v_{\text{sun}} \hat{\mathbf{S}}_{\text{sun}} + \mathbf{S}_{\text{sky}} + \mathbf{S}_{\text{ind}})} \quad (9)$$

where \mathbf{I} is the RGB pixel value in the image we wish to process, \mathbf{S}_{sky} , \mathbf{S}_{ind} and $\hat{\mathbf{S}}_{\text{sun}}$ are the illumination values of the corresponding PMVS point computed in Section 5, and the division is per-channel. With this parameterization, varying




Fig. 6: Multiple PMVS points sharing the same reflectance will generate intersecting curves in color space. (a) We selected two PMVS points with the same reflectance but different illuminations (red and blue squares). (b) The corresponding candidate reflectance curves (nearly) intersect at one end, which corresponds to the reflectance of both points. Diamond markers on the curves correspond to the initial guess for visibility (randomly set in this example). (c) Each curve affects a region of the surrounding color space, according to Equation 10. Regions of influence at initialization are illustrated as isosurfaces with varying opacity. Regions closer to the curve and the current visibility estimate are more affected. (d) After one iteration, the visibility estimates have moved towards the intersecting end of the curves, increasing the overlap between the regions of influence (i.e., the energy E_{total} in Equation 12).

v_{sun} in $[0, 1]$ generates a curve of candidate reflectances in RGB space. Our goal is to find for each PMVS point, the position on its *candidate curve* corresponding to its reflectance \mathbf{R} (or correspondingly visibility v_{sun}).

The intuition of our approach is that multiple PMVS points sharing the same reflectance will generate intersecting curves in color space; their (shared) reflectance will be the color where the candidate curves intersect (Fig. 6a-b). By finding these intersections we can deduce the value of v_{sun} for the PMVS point corresponding to each curve.

However, imprecision in the capture process and in the geometry-based computation prevents the curves from perfectly intersecting in color space. In addition, multiple intersections can occur along a curve, giving multiple candidates for the visibility. We address these issues with a robust iterative procedure inspired by the *mean shift* algorithm.

Overview. Mean shift [10], [43] is a non-parametric mode-seeking algorithm that aims to locate the maxima of a density function, given a set of *data points*. First, a *kernel* (or window) is placed at each data point; it represents the region of influence of this point. In an iterative process, each kernel is then moved in a direction that increases the local density, computed as the weighted average of nearby data points. The process stops when all kernels have reached a stationary point (or mode).

In our approach, we define an asymmetric region of influence for each candidate curve. We use *mean shift iterations* to maximize an energy that measures the overlap among pairs of curves, and iteratively update the estimated reflectances while constraining them to lie on their candidate curves. After convergence, for each curve we

obtain the reflectance (and corresponding visibility) that tends to maximize the number of PMVS points sharing a similar reflectance, and correspondingly tends to minimize the number of reflectances in the scene.

This algorithm assumes that the scene is composed of a sparse set of reflectances shared by multiple points, which is a common assumption recently used in image segmentation [44] and white balance algorithms [28].

Region of influence. Equation 9 defines the candidate reflectances of a PMVS point as a rational curve that is parameterized non-uniformly by v_{sun} . In order to obtain a uniform parametrization we first approximate each curve c as a piecewise linear curve in CIE $L^*a^*b^*$ space, and parameterize it by arc length from the shadowed end of the curve: t in $[0, 1]$ so that $\mathbf{R}(v_{sun} = 0) = \mathbf{R}(t = 0)$. We use this uniform parametrization to compute distances along the curve. We chose to work in CIE $L^*a^*b^*$ space because it defines a perceptually uniform distance metric.

We then define the influence of curve c on a point \mathbf{x} of color space as:

$$A_{\mathbf{x}c} = h_{\perp} \left(\frac{d_{\perp}^2(\mathbf{x}, c)}{\sigma_{\perp}^2} \right) h_{\parallel} \left(\frac{d_{\parallel}^2(\mathbf{x}, c)}{\sigma_{\parallel}^2} \right) \quad (10)$$

where

- $d_{\perp}^2(\mathbf{x}, c) = \|\mathbf{x} - proj_{3D}(\mathbf{x}, c)\|^2$ is the squared distance in color space between \mathbf{x} and its projection on curve c (i.e., the distance perpendicular to the curve).
- $d_{\parallel}^2(\mathbf{x}, c) = (t_c - proj_t(\mathbf{x}, c))^2$ is the squared difference between the position of the current reflectance estimate along the curve t_c and the arc length $proj_t(\mathbf{x}, c)$

corresponding to the projection of \mathbf{x} on the curve (i.e., the distance along the curve).

- h_{\perp} and h_{\parallel} are Gaussian kernel profiles with the form $h(x) = e^{-x}$ controlled by the standard deviations σ_{\perp}^2 and σ_{\parallel}^2 (larger σ values correspond to a wider region of influence).

The first term in Equation 10 compensates for curves that do not exactly intersect, by defining a region of influence around the curve with a Gaussian falloff orthogonal to the curve. The second factor makes our algorithm robust to “false intersections”, i.e., intersections of curves that in fact do not share the same reflectance. Due to the Gaussian kernel h_{\parallel} , each curve only influences regions close to its current reflectance estimate. As the reflectance estimate converges toward the most likely reflectance, the value of h_{\parallel} at intersections lying at other positions along the curve will decrease.

The regions of influence of two curves are illustrated in Fig. 6c.

Energy Formulation. We then define an energy that measures the overlap between the regions of influence of pairs of curves:

$$E = \int_V \left(\sum_{c \in C} \sum_{c' \neq c} A_{\mathbf{x}c} A_{\mathbf{x}c'} \right) d\mathbf{x} \quad (11)$$

where we integrate over the entire color space V .

We evenly discretize the 3D color space into a set S of samples, and rewrite this energy as a discrete sum:

$$E_{\text{total}} = \sum_{\mathbf{s} \in S} E_{\text{sample}}(\mathbf{s}). \quad (12)$$

where the energy of a sample $E_{\text{sample}}(\mathbf{s})$ accumulates the contribution of each pair of curves intersecting nearby:

$$E_{\text{sample}}(\mathbf{s}) = \sum_{c \in C} \sum_{c' \neq c} A_{\mathbf{s}c} A_{\mathbf{s}c'}$$

With this formulation, two curves will contribute to the energy of a sample only if they (almost) intersect near this sample.

Derivation of the energy gradient. We seek the positions of the reflectance estimates along the curves t_c for $c \in C$ that maximize E_{total} :

$$\operatorname{argmax}_{t_c} E_{\text{total}} \quad (13)$$

i.e. that are located at the zeros of the gradient function.

The derivative of E_{total} with respect to t_c is given by:

$$\frac{\partial E_{\text{total}}}{\partial t_c} = \frac{4}{\sigma_{\parallel}^2} \sum_{\mathbf{s} \in S} \sum_{c' \neq c} (\operatorname{proj}_t(\mathbf{s}, c) - t_c) A_{\mathbf{s}c} A_{\mathbf{s}c'} \quad (14)$$

Setting Equation 14 to 0 for all $c \in C$ gives:

$$\begin{aligned} \frac{\partial E_{\text{total}}}{\partial t_c} = 0 \quad \Leftrightarrow \\ \left(\sum_{\mathbf{s} \in S} \sum_{c' \neq c} A_{\mathbf{s}c} A_{\mathbf{s}c'} \right) \\ \left(t_c - \frac{\sum_{\mathbf{s} \in S} \operatorname{proj}_t(\mathbf{s}, c) \left(\sum_{c' \neq c} A_{\mathbf{s}c} A_{\mathbf{s}c'} \right)}{\sum_{\mathbf{s} \in S} \sum_{c' \neq c} A_{\mathbf{s}c} A_{\mathbf{s}c'}} \right) = 0 \end{aligned}$$

which is analogous to the form obtained in the mean-shift algorithm [10].

Iterative process. We define our iterative procedure recursively by computing at iteration i the weights $A_{\mathbf{s}c}^i$ (i.e., the regions of influence) using the estimates t_c^i for all curves $c \in C$, then updating the estimates t_c^{i+1} for each curve c using a weighted average of the projection of nearby samples on c :

$$t_c^{i+1} = \frac{\sum_{\mathbf{s} \in S} \operatorname{proj}_t(\mathbf{s}, c) \left(\sum_{c' \neq c} A_{\mathbf{s}c}^i A_{\mathbf{s}c'}^i \right)}{\sum_{\mathbf{s} \in S} \sum_{c' \neq c} A_{\mathbf{s}c}^i A_{\mathbf{s}c'}^i} \quad (15)$$

We initialize t_c^0 using the initial guess of the sun visibility $v_{\text{sun}}^{\text{init}}$, obtained in Section 5 by casting the shadow of the geometric proxy. We ensure that this iterative process converges to a maximizer by checking that the energy does not decrease after each iteration [45]; if $E_{\text{total}}^{i+1} < E_{\text{total}}^i$ we set $t_c^{i+1} = (t_c^{i+1} + t_c^i)/2$ for all c and iterate. The algorithm stops when $|t_c^{i+1} - t_c^i|$ is small enough for all c .

As a result of this process, the reflectance estimates converge to regions of space where many curves intersect, while being constrained to lie on their candidate curves due to our parameterization. This is illustrated in Fig. 6d.

Orientation separation. We observed that false intersections sometimes occur when two curves of different reflectances have different orientations, sky or indirect illumination. In order to further reduce the presence of such intersections, we separate PMVS points into three groups based on the orientation of their normals. We first apply mean-shift clustering [10] where the 3D feature vectors contain the normals of the points. The first two clusters correspond to the two dominant groups of normal orientations in the scene (Fig. 7a, blue and green), while the remaining points form a third group (Fig. 7a, red). We then run our optimization independently on each group.

Curve clustering. Even within one group of orientations, the candidate curves corresponding to the numerous PMVS points (up to dozens of thousands, in our scenes) tend to intersect each other at several positions along the curves (Fig. 7c), since many similar (but not equal) reflectances might be present in the scene. We enforce sparsity in our algorithm by grouping all the curves corresponding to a similar reflectance and illumination (Fig. 7b). Note that this grouping is only used to identify the sun visibility, we do not enforce sparsity in the final reflectance image.

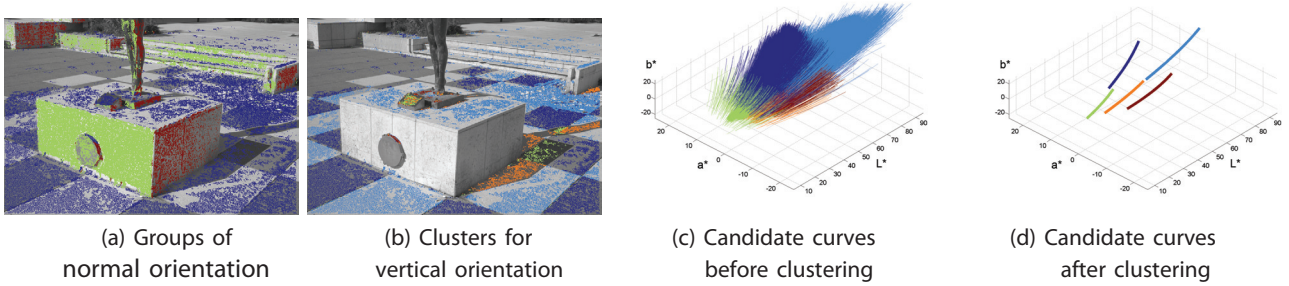


Fig. 7: Orientation separation and curve clustering. (a) PMVS points are separated into three groups according to the orientation of their normals. Each group is then processed independently. (b-c) Within one group, PMVS points are clustered according to the endpoints of their candidate curves. (d) After clustering, pairs of curves intersect as expected (orange and light blue, green and dark blue), since each pair represents PMVS points sharing the same reflectance.

We use mean-shift clustering to perform this grouping, representing each curve as a 6D feature vector that contains the $L^*a^*b^*$ reflectances of its endpoints. Clusters that contain few curves (less than 1% of the largest cluster) are discarded, and the corresponding PMVS points are ignored in the remaining of the algorithm. This leads to the clustering shown in Fig. 7b-c.

Finally, we replace all the curves belonging to each cluster by one *representative curve*: the curve closest to the median of this cluster’s feature vectors. The initial guess of the sun visibility $v_{\text{sun}}^{\text{init}}$ is assigned to the median visibility of each cluster. This process greatly simplifies and cleans up the set of curves, and lowers the computational cost of our algorithm. As shown in Fig. 7d, the curves corresponding to groups of PMVS points with similar reflectances (nearly) intersect after clustering.

After clustering, we run the optimization to maximize E_{total} . Upon convergence, for each cluster we obtain the position of the estimated reflectance along its representative curve t^{final} . This position value is assigned to all curves belonging to this cluster, from which we can deduce the final estimated reflectance of each PMVS point corresponding to these curves.

We found that clustering candidate curves largely improves the stability of the algorithm. The clustering groups candidate curves corresponding to PMVS points with same reflectance *and* same illumination, while our subsequent optimization identifies the reflectances shared between clusters with different illumination. We show more examples of candidate curves after clustering in the supplemental materials.

At the end of this process, we obtain a list of PMVS points for which the position along the curve t , the sun visibility v_{sun} and the reflectance \mathbf{R} have been estimated.

Implementation. In practice we use a truncated kernel profile h_{\perp} (Equation 10) so that $h_{\perp}(x) = 0$ when $x > \lambda$ (we use $\lambda = 3$). This means that each curve will only influence a limited number of samples; for each curve c we can precompute the indices of these samples, as well as

their orthogonal distance $\|\mathbf{s} - \text{proj}_{3D}(\mathbf{s}, c)\|$ and position of their projection along the curve $\text{proj}_t(\mathbf{s}, c)$.

We evenly discretized the CIE $L^*a^*b^*$ color space into $60 \times 36 \times 36$ samples, with L^* in $[5, 95]$, a^* in $[-25, 25]$, b^* in $[-25, 25]$. We used a fixed bandwidth for the curve clustering using 6D mean-shift clustering, as well as for the 3D mean-shift clustering for orientation separation.

The clustering is the most costly part of the algorithm and takes from 25 seconds to a few minutes with our Matlab implementation, depending on the number of PMVS points. Once the clustering has been performed, the iterative optimization takes around 10 seconds, which allowed us to test many parameters for σ_{\perp}^2 and σ_{\parallel}^2 .

We found that the algorithm produces good results for a wide range of parameters, and that the best values were scene-dependent. In our experiments we found that large values of σ_{\perp}^2 can compensate for calibration errors that prevent reflectance curves from perfectly intersecting, while large values of σ_{\parallel}^2 can compensate for the erroneous initialization of $v_{\text{sun}}^{\text{init}}$ provided by the approximate proxy. However, σ values should remain small enough to prevent curves of different materials from influencing one another. Table 1 summarizes the parameter values used for the examples in this paper.

	Statue	Rocks	Ramp	Stairs
$(\sigma_{\parallel}^2, \sigma_{\perp}^2)$	(0.1, 1)	(0.01, 1)	(0.3, 1)	(0.01, 20)
	(0.1, 5)	(0.01, 5)	(0.3, 5)	

TABLE 1: Sets of region of influence parameters used on the four scenes.

7 ESTIMATING ILLUMINATION AT EACH PIXEL

In previous steps, we have used multiview stereo methods to generate a sparse set of 3D points on which we compute the illumination values $\hat{\mathbf{S}}_{\text{sun}}$, \mathbf{S}_{sky} and \mathbf{S}_{ind} (Section 5) along with the visibility of the sun v_{sun} (Section 6). We next show how to leverage image-guided propagation methods




Fig. 8: *Separation of Reflectance and Total Illumination. The optimization procedure described in Equation 17 enforces illumination constraints at PMVS points (b) and propagates them to all pixels, in order to separate an input photograph (a) into total illumination (c) and reflectance (d).*

to assign reflectance and illumination values to all pixels in the input photographs. We first show how to propagate the total illumination $\mathbf{S}_{\text{total}}$ and then describe a method to subsequently separate the contribution of each lighting component (i.e., sun, sky and indirect).

7.1 Image Guided Propagation

We use the intrinsic images algorithm of Bousseau et al. [3] that was designed to propagate user indications for separating reflectance and illumination in a single image. This algorithm makes the intrinsic image decomposition tractable by assuming that the reflectance values in a pixel neighborhood lie in a plane in color space. This planar reflectance assumption translates to a set of linear equations so that the illumination image is expressed as the minimizer of a least-square energy

$$\underset{\bar{\mathbf{S}}_{\text{total}}}{\operatorname{argmin}} \bar{\mathbf{S}}_{\text{total}}^T \mathbf{M} \bar{\mathbf{S}}_{\text{total}} \quad (16)$$

where the vector $\bar{\mathbf{S}}_{\text{total}}$ stacks the pixels of the estimated illumination image and the matrix \mathbf{M} encodes the planar reflectance assumption (see the paper by Bousseau et al. [3] for the complete derivation). For colored illumination the optimization is solved for each color channel separately.

In their original paper, Bousseau et al. constrain the least-square system with user indications. Users can specify the value of $\bar{\mathbf{S}}_{\text{total}}$ over a few pixels or indicate that several pixels share the same illumination or reflectance. In our approach we use instead the illumination and visibility estimated at PMVS points to constrain the optimization. We express these constraints as an additional quadratic penalty

$$\underset{\bar{\mathbf{S}}_{\text{total}}}{\operatorname{argmin}} \bar{\mathbf{S}}_{\text{total}}^T \mathbf{M} \bar{\mathbf{S}}_{\text{total}} + w \sum_{p \in \mathcal{P}} (\bar{\mathbf{S}}_{\text{total}}^p - \mathbf{S}_{\text{total}}^p)^2 \quad (17)$$

where \mathcal{P} is the set of pixels covered by PMVS points and $\mathbf{S}_{\text{total}}^p = v_{\text{sun}}^p \hat{\mathbf{S}}_{\text{sun}}^p + \mathbf{S}_{\text{sky}}^p + \mathbf{S}_{\text{ind}}^p$ their illumination values computed in Sections 5 and 6. The weight w controls the importance of the constraints, we use $w = 1$ to give equal importance to the constraints and the propagation model.

Fig. 8 shows the results of the image guided propagation. This result shows the power of our approach. We exploit the information provided by the sparse and imprecise geometry

to automatically generate constraints for the algorithm of Bousseau et al. [3], thus eliminating the need for user scribbles. A visualization of the constraints for all scenes can be found in the supplemental material. While our method supports user scribbles to guide the decomposition in regions that could not be reconstructed (e.g., trees, specular objects), all our results have been generated using only the automatically computed constraints.

7.2 Light Source Separation

Given the estimated illumination image $\bar{\mathbf{S}}_{\text{total}}$, we wish to separate the contribution of each illumination component $\bar{\mathbf{S}}_{\text{sun}}$, $\bar{\mathbf{S}}_{\text{sky}}$ and $\bar{\mathbf{S}}_{\text{ind}}$. Inspired by previous work on white balance under mixed lighting [28], we express our light source separation as two successive matting problems. We first decompose the illumination into a sun component and a *diffuse* component $\bar{\mathbf{S}}_{\text{diff}}$ that includes the contribution of both sky and indirect lighting. In a second step we decompose the diffuse component into its two terms.

We first express each illumination term $\bar{\mathbf{S}}$ as the product between a scalar intensity $s = \|\bar{\mathbf{S}}\|$ and a chromaticity $\mathbf{C} = \bar{\mathbf{S}}/\|\bar{\mathbf{S}}\|$:

$$\begin{aligned} \bar{\mathbf{S}}_{\text{total}} &= s_{\text{sun}} \mathbf{C}_{\text{sun}} + s_{\text{sky}} \mathbf{C}_{\text{sky}} + s_{\text{ind}} \mathbf{C}_{\text{ind}} \\ &= s_{\text{sun}} \mathbf{C}_{\text{sun}} + s_{\text{diff}} \mathbf{C}_{\text{diff}}. \end{aligned} \quad (18)$$

Denoting $\alpha = s_{\text{sun}}/(s_{\text{sun}} + s_{\text{diff}})$ we express the illumination image values at each pixel as a mixture between two values weighted by α :

$$\begin{aligned} \bar{\mathbf{S}}_{\text{total}} &= \alpha(s_{\text{sun}} + s_{\text{diff}}) \mathbf{C}_{\text{sun}} \\ &\quad + (1 - \alpha)(s_{\text{sun}} + s_{\text{diff}}) \mathbf{C}_{\text{diff}} \end{aligned} \quad (19)$$

$$= \alpha \mathbf{F} + (1 - \alpha) \mathbf{B} \quad (20)$$

We can now recover $\bar{\mathbf{S}}_{\text{sun}} = \alpha \mathbf{F}$ and $\bar{\mathbf{S}}_{\text{diff}} = (1 - \alpha) \mathbf{B}$ by solving a standard matting problem. We compute α at each PMVS point from the illumination values estimated in Sections 5 and 6. We then propagate α over the image $\bar{\mathbf{S}}_{\text{total}}$ using the *matting Laplacian* algorithm of Levin et al. [27]. Finally, given α at every pixel and the known \mathbf{S}_{sun} and \mathbf{S}_{diff} at each PMVS point, we solve for the sun and diffuse illumination images with the following least-square




Fig. 9: *The reflectance, sky illumination (here, close up) and indirect illumination estimated by our algorithm can contain residual variations along hard shadow boundaries (a). We use inpainting to remove these artifacts (b).*

optimization:

$$\begin{aligned} \operatorname{argmin}_{\mathbf{F}, \mathbf{B}} \sum_{i \in \mathcal{I}} & \left(\left(\bar{\mathbf{S}}_{\text{total}}^i - (\alpha^i \mathbf{F}^i + (1 - \alpha^i) \mathbf{B}^i) \right)^2 \right. \\ & + \lambda \left((\mathbf{F}_x^i)^2 + (\mathbf{F}_y^i)^2 + (\mathbf{B}_x^i)^2 + (\mathbf{B}_y^i)^2 \right) \\ & \left. + w \sum_{p \in \mathcal{P}} (\alpha^p \mathbf{F}^p - \mathbf{S}_{\text{sun}}^p)^2 + ((1 - \alpha^p) \mathbf{B}^p - \mathbf{S}_{\text{diff}}^p)^2 \right) \end{aligned}$$

where \mathcal{I} is the image domain, \mathcal{P} is the set of pixels covered by PMVS points and \mathbf{F}_x , \mathbf{F}_y , \mathbf{B}_x and \mathbf{B}_y are the x and y derivatives of \mathbf{F} and \mathbf{B} computed with finite differences.

The first term of this functional ensures that the decomposition explains the input illumination $\bar{\mathbf{S}}_{\text{total}}$ and follows the estimated ratio α . The second term adds a smoothness regularization on each component while the third term constrains the solution to agree with the illumination values computed at PMVS points. We used $\lambda = 0.1$ and $w = 0.01$ for all our results.

As a second step we apply the same matting approach to further separate the diffuse illumination $\bar{\mathbf{S}}_{\text{diff}}$ as the sum of the sky illumination $\bar{\mathbf{S}}_{\text{sky}}$ and the indirect illumination $\bar{\mathbf{S}}_{\text{ind}}$. We show in Fig. 10 and supplemental materials the results of this decomposition. The overall decomposition takes around 90 seconds to compute for a 3.2 megapixel image, where 30 seconds are necessary to compute $\bar{\mathbf{S}}_{\text{total}}$ and 60 seconds to perform the two subsequent separations.

Inpainting This overall process gives a satisfactory decomposition in the scenes we have tested. However, a small residual border can remain around the hard shadow boundary (see Fig. 9 (left)), a common artifact of shadow removal [30], [32]. We identify these shadow boundaries as pixels located on sun illumination discontinuities but not on normal discontinuities. We first propagate normals from the PMVS points over the image using the method of Okabe et al. [20]. We then run an edge detector over the sun illumination image and label edge pixels that do not correspond to edges in the normal image. We remove the labeled pixels and their immediate neighbors from the reflectance, sky and indirect illumination images and use inpainting to fill in the holes (see Fig. 9 (right)). This post-process takes 60 seconds on average.

8 RESULTS AND DISCUSSION

In Fig. 10 we show results on four different scenes. For all results, we show reflectance, sun, sky and indirect illumination layers. The number of photographs used for each scene is shown in Table 2. The estimated illumination at PMVS points and additional views for each scene are shown in the supplemental material.

The first scene shows the base of a statue on a square; we took HDR images for three different views and we show results for all three. The reflectance layer is plausible in all views. The sun and sky layers have been successfully separated in all cases, however the third view shows a slight color shift of the reflectance in the shadow area; this is due to the insufficient number of PMVS points on the ground. Please refer to the supplemental material for an illustration of the distribution of 3D points. The indirect layer clearly shows the indirect light bouncing off the front of the base (first and second views) which is in direct sunlight. The sides of the base do not receive the same amount of sky and indirect illumination due to the configuration of the scene (see Fig. 5c).

The second scene is challenging, since the reconstruction process is unable to capture details of the railings and does not reconstruct the vegetation casting the main shadows (see also Fig. 3d). In this view, despite the intricate geometric configuration, shadows are successfully removed from the reflectance layer, and the three other layers show good results. In the third scene, a staircase is shown. There are some residual artifacts at the shadow boundaries because the vegetation moved in the breeze during HDR acquisition. In the fourth scene, an umbrella casts a shadow onto a rock wall. Notice how the indirect layer well represents lighting in the cracks between rocks where neither sun nor skylight is present. The third and fourth scenes also show that our method succeeds in creating plausible reflectance, sun, sky and indirect layers.

Statue	Rocks	Ramp	Stairs
30	11	31	10

TABLE 2: *Number of photographs captured for the geometric reconstruction, for each scene in this paper.*

Next we compare our approach to three state-of-the-art methods in Fig. 12. All these methods take a single image as input. The user-assisted approach of Bousseau et al. [3] produces results of a similar quality to ours, but requires a significant number of user indications (between 25 and 105 scribbles). The automatic method of Shen et al. [14] is able to extract most of the illumination variations but colored shadows remain in the reflectance image. These colored residuals are due to the variation in color between sun and sky illumination, which violates the gray illumination assumption of this method. Residual shadows are also present in the reflectance estimated with the automatic method of Shen et al. [19], as well as reflectance residuals in the illumination image (tiled floor in the statue scene).

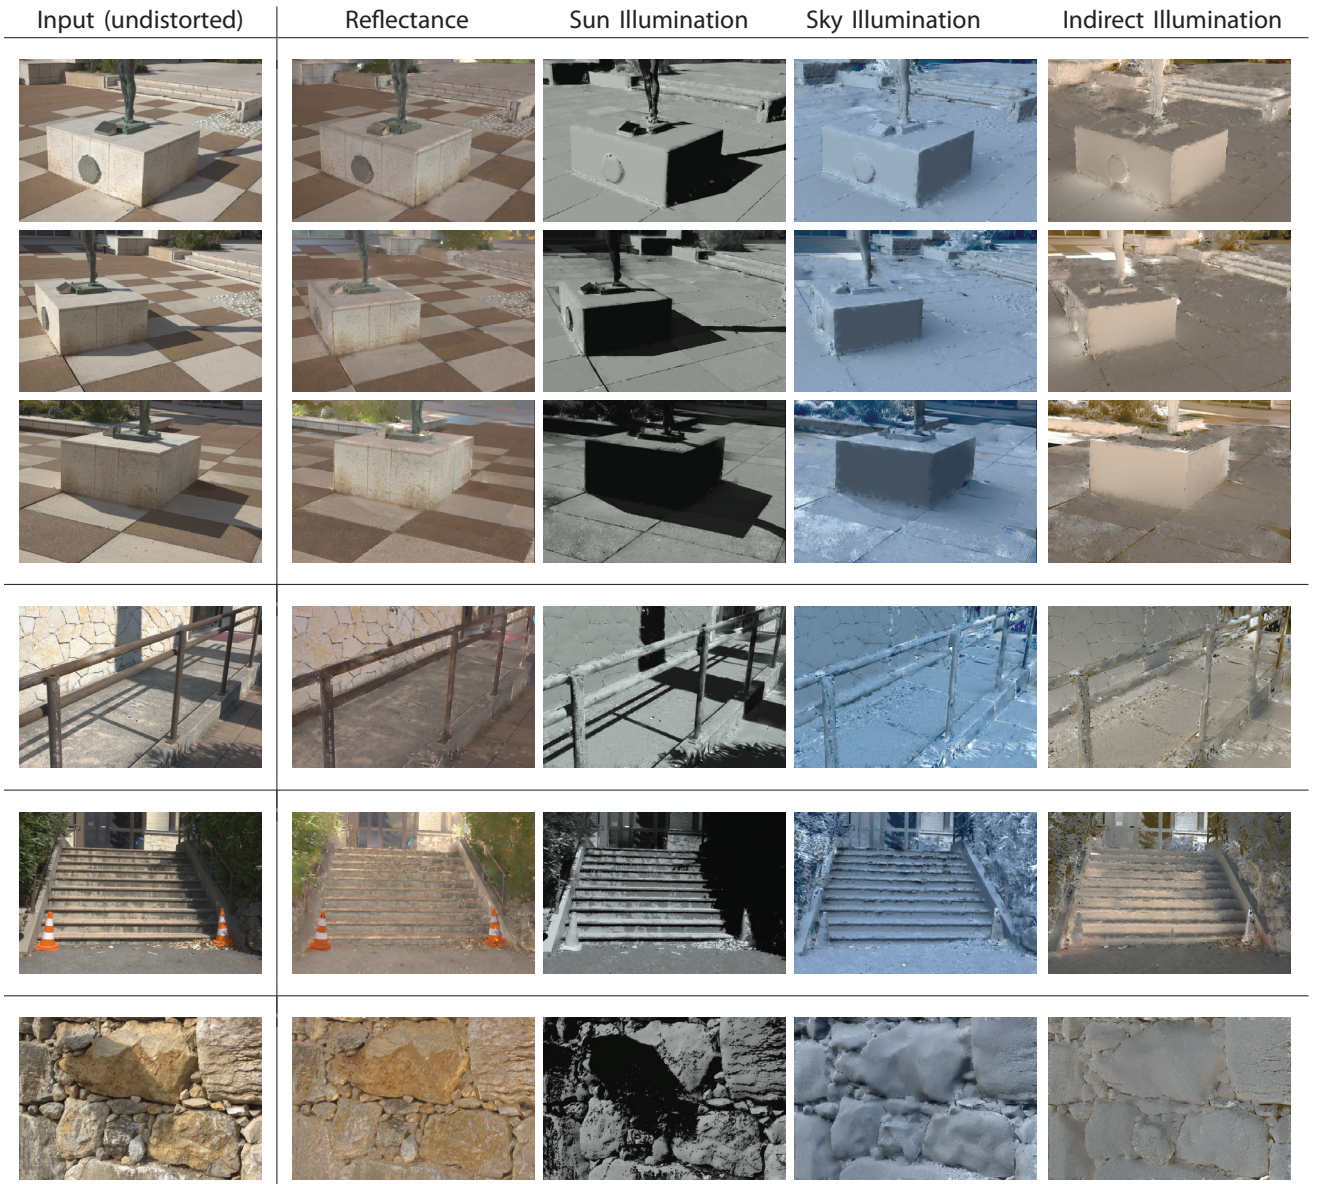


Fig. 10: Results of our decomposition. We adjusted the brightness of images for illustration purposes (the scaling factors used can be found in the supplementary materials). For each scene, the sun illumination is usually much more intense than sky illumination on average, and the sky illumination is more intense than indirect illumination on average.

Although this method can support user scribbles, the authors reported that scribbles did not improve the result significantly in these examples. All the results of these comparisons have been kindly provided by the respective authors. While these approaches can be applied to single images from various sources, to our knowledge ours is the first to provide a richer decomposition by separating the illumination into different components.

In Fig. 11 we show a comparison to the result obtained simply by using the proxy to compute sun, sky and indirect illumination using PBRT, and then inverting Equation 5 to obtain reflectance. As we can see, for a given image, there are many holes due to the incomplete proxy. In addition, the sun illumination is completely erroneous, due to the lack of reconstruction of the surrounding objects. In contrast, our

method correctly captures these sun shadows, and removes them in the reflectance layer as well.

We illustrate applications of our decomposition in Fig. 1. We first alter the reflectance of the ground to insert a graffiti while maintaining consistent shadows (b). We then add a virtual object in the scene with consistent lighting and shadows (c). We used PBRT to render the dinosaur surrounded by the captured environment map, and its shadow cast on a horizontal plane. Finally we simulate a sunset by changing the color and intensity of each illumination component separately; in addition, our decomposition allows us to blur shadows without affecting the reflectance of the scene (d). All these manipulations can be performed easily in image-editing software with layer support; the supplementary video shows how we created the images in Fig. 1b-d with




Fig. 11: Comparison between the decomposition estimated directly from the geometric proxy (first row), and our results (second row). We obtain the proxy reflectance by dividing the input image by the sum of the illumination components. Holes and inaccuracies in the proxy translate to artifacts and residual shadows in the reflectance.

Adobe Photoshop.

Discussion A drawback of our current method is the need for the reflective sphere to capture the environment map. We could investigate fitting a sky model to sky pixels visible in the other input photographs, similar to [23], and extracting information about non-sky surroundings in a similar manner. We will investigate ways to simplify the calibration process of the sphere so that it can be done once and subsequently used for any scene.

Since we were interested in separating sun lighting from other sources, we have not shown overcast scenes. There is no fundamental reason that our method would not work with overcast scenes; the decomposition would simply rely on the S_{sky} and S_{ind} values and ignore sun illumination.

Our method tolerates holes in the proxy geometry if the hole is sufficiently distant from the 3D points where sky and indirect illumination are computed. In such configurations, rays that do not hit the proxy will hit the environment map and provide a plausible color. However, in Fig. 10 (fifth row), the near bush over the stairs was not reconstructed while being too close to be properly captured in the environment map. As a result, rays emitted from the stairs reach the sky instead of the bush and yield a residual shadow in the indirect lighting.

Our method can fail if the initial guess for sun visibility is completely wrong. In the supplemental material we show a case where a spurious object appeared in a single view, and was thus not reconstructed at all, resulting in this type of failure. Similarly, if objects have very dark reflectance, or specularities, the PMVS reconstruction procedure does not provide a sufficient number of points, resulting in errors. Similarly to other intrinsic image methods, we have assumed that the scene is Lambertian. Incorporating a specular layer in the intrinsic image formulation is a challenging avenue for future work.

Our optimization for sun visibility exploits the redundant

information provided by points of the same reflectance that are in sun light *and* shadow. Nevertheless our method can also handle reflectances that are only in light *or* shadow. Most often in such cases the proxy initialization will result in the correct answer. The only case which could potentially cause errors is when other reflectance curves incorrectly intersect with or influence the curve of this material. We have carefully designed our algorithm to avoid this situation, by making each curve vote only for the region in which it is confident, by processing PMVS points with drastically different orientations separately, and by clustering similar candidate reflectance curves. The results were satisfactory for our examples and the values shown in Table 1.

Finally, the candidate curves in our sun visibility method may not intersect if the illuminant calibration is too imprecise. We are investigating ways to improve calibration and reduce the amount of user intervention required for capture and calibration.

9 CONCLUSION

We have presented a method to estimate rich intrinsic images for outdoor scenes. In addition to reflectance, our algorithm generates a separate image for the sun, sky and indirect illumination. Our method relies on a lightweight capture (10-31 photographs in the scenes shown here) to estimate a coarse geometric representation of the scene. This geometric information allows us to estimate illumination terms over a sparse 3D sampling of the scene. We then introduce an optimization algorithm to refine the inaccurate estimations of sun visibility. While incomplete, we demonstrate that this sparse information provides the necessary constraints for an image-guided propagation algorithm that recovers the reflectance and illumination components at each pixel of the input photographs.

Our intrinsic image decomposition allows users of image manipulation tools to perform consistent editing of material

Input	Bousseau et al. [3]	Shen et al. [14]	Shen et al. [19]	Our approach
  82 user scribbles for [3]	 	 	 	 
  105 user scribbles for [3]	 	 	 	 
  72 user scribbles for [3]	 	 	 	 
  26 user scribbles for [3]	 	 	 	 

Fig. 12: Comparison of our method with a user-assisted method [3] and two automatic algorithms [14], [19]. In the first column, the user scribbles used for the method of Bousseau et al. [3] are shown under the input image. In the next columns, the first row contains the estimated reflectance while the second one corresponds to the total illumination. Our results are shown in the last column. Our multiview approach outperforms single-image automatic algorithms and achieves results of comparable quality to the user-assisted approach with significantly less user intervention. White arrows point to residual reflectance or shading variations. The brightness has been adjusted for comparison.

and lighting in photographs. An interesting direction of future work is to adapt the method of [21] to provide an alternative way of computing the illumination components. It will be interesting to investigate the tradeoffs between complexity, speed and quality obtained from such an adaptation compared to our approach.

Our current method works independently for each view. Enforcing consistent intrinsic image properties between views is an interesting future research direction. With such consistency, our method will open the way for dynamically relightable environments and free-viewpoint navigation for Image Based Rendering systems [46], in addition to the applications demonstrated in this paper. An important step for this goal is the generation of plausible shadow motion in the sun illumination layer in order to simulate moving light sources.


ACKNOWLEDGEMENTS

We thank the reviewers for their comments and suggestions. Thanks also to P. Vangorp, G. Chaurasia and C. Bosch for their help with capture. J. Digne and P. Alliez provided valuable help for the geometry reconstruction, while T. Grenier provided helpful advice on mean-shift algorithms. We thank the authors of the other methods of kindly providing comparison images included here. We acknowledge the generous support of Autodesk (software donation) and Adobe (research donation). This work was partly supported by the EU FP7 IP project VERVE (<http://www.verveconsortium.eu/>).

REFERENCES

- [1] H. Barrow and J. Tenenbaum, "Recovering intrinsic scene characteristics from images," *Computer Vision Systems*, 1978.
- [2] E. H. Land and J. J. McCann, "Lightness and retinex theory," *Journal of the optical society of America*, vol. 61, no. 1, 1971.
- [3] A. Bousseau, S. Paris, and F. Durand, "User-assisted intrinsic images," *ACM Trans. Graph.*, vol. 28, no. 5, 2009.
- [4] Y. Weiss, "Deriving intrinsic images from image sequences," in *ICCV*, 2001, pp. 68–75.
- [5] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: recovering reflectance models of real scenes from photographs," in *SIGGRAPH*, 1999, pp. 215–224.
- [6] P. Debevec, C. Tchou, A. Gardner, T. Hawkins, C. Poullis, J. Stumpfel, A. Jones, N. Yun, P. Einarsson, T. Lundgren, M. Fajardo, and P. Martinez, "Estimating surface reflectance properties of a complex scene under captured natural illumination," USC Institute for Creative Technologies, Tech. Rep., 2004.
- [7] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: exploring photo collections in 3D," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 835–846, 2006.
- [8] Y. Furukawa and J. Ponce, "Accurate, dense, and robust multi-view stereopsis," *IEEE Trans. PAMI*, vol. 32, no. 8, pp. 1362–1376, 2009.
- [9] A. Levin, D. Lischinski, and Y. Weiss, "Colorization using optimization," *ACM TOG (proc. of SIGGRAPH)*, vol. 23, no. 3, pp. 689 – 694, 2004.
- [10] D. Comaniciu and P. Meer, "Mean shift: a robust approach toward feature space analysis," *IEEE Trans. PAMI*, vol. 24, no. 5, pp. 603–619, 2002.
- [11] B. K. Horn, *Robot Vision*, 1st ed. McGraw-Hill Higher Education, 1986.
- [12] L. Shen and C. Yeo, "Intrinsic image decomposition using a local and global sparse representation of reflectance," in *CVPR*, 2011.
- [13] M. F. Tappen, W. T. Freeman, and E. H. Adelson, "Recovering intrinsic images from a single image," *IEEE Trans. PAMI*, vol. 27, no. 9, 2005.
- [14] L. Shen, P. Tan, and S. Lin, "Intrinsic image decomposition with non-local texture cues," in *CVPR*, 2008.
- [15] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman, "Ground-truth dataset and baseline evaluations for intrinsic image algorithms," in *ICCV*, 2009.
- [16] K. Sunkavalli, W. Matusik, H. Pfister, and S. Rusinkiewicz, "Factored time-lapse video," *ACM Trans. Graph.*, vol. 26, no. 3, 2007.
- [17] W. Matusik, M. Loper, and H. Pfister, "Progressively-refined reflectance functions from natural illumination," in *Eurographics Symposium on Rendering*, 2004, pp. 299–308.
- [18] F. Melendez, M. Glencross, G. J. Ward, and R. J. Hubbard, "Relightable Buildings from Images," in *Eurographics: Special Area on Cultural Heritage*, 2011, pp. 33–40.
- [19] J. Shen, X. Yang, Y. Jia, and X. Li, "Intrinsic images using optimization," in *CVPR*, 2011.
- [20] M. Okabe, G. Zeng, Y. Matsushita, T. Igarashi, L. Quan, and H.-Y. Shum, "Single-view relighting with normal map painting," in *Pacific Graphics*, 2006, pp. 27–34.
- [21] R. Carroll, R. Ramamoorthi, and M. Agrawala, "Illumination decomposition for material recoloring with consistent interreflections," *ACM TOG (Proc. of SIGGRAPH)*, vol. 30, no. 4, 2011.
- [22] K. Karsch, V. Hedau, D. Forsyth, and D. Hoiem, "Rendering synthetic objects into legacy photographs," *ACM TOG (Proc. of SIGGRAPH Asia)*, vol. 30, no. 6, 2011.
- [23] Y. Yu and J. Malik, "Recovering photometric properties of architectural scenes from photographs," in *SIGGRAPH*, 1998, pp. 207–217.
- [24] C. Loscos, G. Drettakis, and L. Robert, "Interactive virtual relighting of real scenes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 6, no. 4, pp. 289–305, 2000.
- [25] J. T. Kajiya, "The rendering equation," *Computer Graphics (proc. of SIGGRAPH)*, vol. 20, no. 4, pp. 143–150, 1986.
- [26] D. Lischinski, Z. Farbman, M. Uyttendaele, and R. Szeliski, "Interactive local adjustment of tonal values," *ACM TOG (proc. of SIGGRAPH)*, vol. 25, no. 3, pp. 646–653, 2006.
- [27] A. Levin, D. Lischinski, and Y. Weiss, "A closed-form solution to natural image matting," *IEEE Trans. PAMI*, 2008.
- [28] E. Hsu, T. Mertens, S. Paris, S. Avidan, and F. Durand, "Light mixture estimation for spatially varying white balance," *ACM TOG (proc. of SIGGRAPH)*, vol. 27, no. 3, p. 70, 2008.
- [29] G. D. Finlayson, S. D. Hordley, and M. S. Drew, "Removing shadows from images," in *ECCV*, 2002.
- [30] G. D. Finlayson, M. S. Drew, and C. Lu, "Intrinsic images by entropy minimization," in *ECCV*, 2004, pp. 582–595.
- [31] A. Mohan, J. Tumblin, and P. Choudhury, "Editing soft shadows in a digital photograph," *IEEE Computer Graphics and Applications*, vol. 27, no. 2, pp. 23–31, 2007.
- [32] T.-P. Wu, C.-K. Tang, M. S. Brown, and H.-Y. Shum, "Natural shadow matting," *ACM TOG*, vol. 26, no. 2, p. 8, 2007.
- [33] Y. Shor and D. Lischinski, "The shadow meets the mask: Pyramid-based shadow removal," *Computer Graphics Forum*, vol. 27, no. 2, pp. 577–586, 2008.

- [34] P. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *SIGGRAPH*, 1997, pp. 369–378.
- [35] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," *Computer Graphics (proc. of SIGGRAPH)*, vol. 26, no. 2, pp. 71–78, 1992.
- [36] J. Digne, J.-M. Morel, C.-M. Souzani, and C. Lartigue, "Scale space meshing of raw data point sets," *Computer Graphics Forum*, vol. 6, no. 4, pp. 1630–1642, 2011.
- [37] P. Cignoni, M. Corsini, and G. Ranzuglia, "Meshlab: an open-source 3d mesh processing system," *ERCIM News*, no. 73, pp. 45–46, 2008.
- [38] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Eurographics Symposium on Geometry Processing*, 2006, pp. 61–70.
- [39] S. Sinha, D. Steedly, and R. Szeliski, "Piecewise planar stereo for image-based rendering," in *ICCV*, 2009, pp. 1881–1888.
- [40] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski, "Manhattan-world stereo," in *CVPR*, 2009, pp. 1422–1429.
- [41] D. Gallup, J.-M. Frahm, and M. Pollefeys, "Piecewise planar and non-planar stereo for urban scene reconstruction," in *CVPR*, 2010, pp. 1418–1425.
- [42] M. Pharr and G. Humphreys, *Physically Based Rendering: From Theory to Implementation, second edition*. Morgan Kaufmann Publishers Inc., 2010.
- [43] K. Fukunaga and L. Hostetler, "The estimation of the gradient of a density function, with applications in pattern recognition," *IEEE Transactions on Information Theory*, vol. 21, no. 1, pp. 32–40, 1975.
- [44] I. Omer and M. Werman, "Color lines: Image specific color representation," in *CVPR*, 2004, pp. 946–953.
- [45] D. Comaniciu, V. Ramesh, and P. Meer, "Real-time tracking of non-rigid objects using mean shift," in *CVPR*, 2000, pp. 142–149.
- [46] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen, "Unstructured lumigraph rendering," in *SIGGRAPH*, 2001, pp. 425–432.



Pierre-Yves Laffont is a PhD candidate at INRIA Sophia-Antipolis in the REVES group, under the supervision of George Drettakis and Adrien Bousseau. His PhD focuses on intrinsic images, image editing, image-based rendering and relighting, with geometric cues from multi-view reconstruction. Before joining INRIA, he studied at INSA Lyon (France) and KAIST (South Korea). He visited UC Berkeley and MIT CSAIL during his PhD.



Adrien Bousseau is a junior researcher at REVES / INRIA Sophia Antipolis. His current research focuses on image creation and manipulation for artists and designers. Adrien Bousseau did his PhD at INRIA Rhône Alpes under the supervision of Joëlle Thollot and François X. Sillion, and his postdoc at UC Berkeley with Maneesh Agrawala and Ravi Ramamoorthi. During his PhD he also spent a few months at Adobe Advanced Technology Labs and MIT CSAIL as an intern

and visiting student. Adrien Bousseau received one of the three Eurographics 2011 PhD award for his research on expressive image manipulations.



Dr. George Drettakis (PhD, Toronto 1994; Habilitation, Grenoble, 1999) 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 all leading international conferences and journals. His recent work is on image-based rendering and textures, perceptual rendering, and in particular audio-visual cross-modal rendering, interactive illumination, shadows,

relighting and interactive rendering in general. He has coordinated or participated in several EU and French national projects. He has been a member of program committees of many international conferences, was programme co-chair of the Eurographics Workshop on Rendering 98, and the Eurographics 2002 and 2008 Conferences. He was technical papers chair of ACM SIGGRAPH Asia 2010. He won the Eurographics Outstanding Technical Achievement Award in 2007 and is a Eurographics Fellow.