

Rich Intrinsic Image Decomposition of Outdoor Scenes from Multiple Views

Pierre-Yves Laffont, Adrien Bousseau, George Drettakis
REVES / Inria Sophia-Antipolis*




Figure 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).

Introduction

Intrinsic images aim at separating an image into reflectance and illumination layers to facilitate analysis or manipulation. Most successful methods rely on user indications [Bousseau et al. 2009], precise geometry, or need multiple images from the same viewpoint and varying lighting to solve this severely ill-posed problem. We propose a method to estimate intrinsic images from multiple views of an outdoor scene at a single time of day without the need for precise geometry and with only a simple manual calibration step.

We use multi-view stereo to automatically reconstruct a 3D point cloud of the scene. Although this point cloud is sparse and incomplete, we show 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 extraction of precise cast shadows. We finally propagate the information computed over the sparse point cloud to every pixel in the photograph using image-guided propagation.

Our method not only **separates reflectance and illumination**, but also **decomposes the illumination into sun, sky and indirect layers**. This rich decomposition allows novel image manipulations.

Our Approach



Our method relies on a relatively lightweight capture setup composed of a digital camera, a photographer’s gray card for calibration, and a simple reflective sphere to capture an environment map. We capture 10-30 pictures from different viewpoints in addition to the images to decompose.

Geometry-based computation We use structure-from-motion and multi-view stereo to reconstruct a cloud of oriented 3D points, and surface reconstruction to obtain a *proxy* geometric model of the scene. The user specifies the orientation and color of sun and sky through a simple calibration step.

The geometric proxy is approximate and incomplete, and cannot be directly used to estimate the illumination at each pixel. In particu-

lar, it produces inaccurate or even missing cast shadows. However, it can give a reasonable approximation for low-frequency lighting components: we estimate sky and indirect illumination *at reconstructed 3D points* by casting rays towards all directions. Rays that intersect the proxy contribute to indirect lighting (we use the captured photographs to lookup the outgoing radiance at intersected points), while rays that hit the environment map above the horizon contribute to sky lighting. We also use the proxy to compute an initial estimate of sun visibility (cast shadows), which we later refine.

Sun visibility estimation We introduce an algorithm to reliably identify points in shadow based on a new parameterization of reflectance with respect to sun visibility. Our algorithm compensates for the lack of accurately reconstructed and complete geometry. We show that the reflectance of each 3D point lies on a *candidate curve* in color space, once sky and indirect illuminations are estimated. Multiple points which share a similar reflectance generate intersecting candidate curves. We use an iterative optimization method that reliably estimates the reflectance and sun visibility at reconstructed points, by finding regions where multiple candidate curves intersect. We illustrate this process in the supplementary document.

Image-based propagation and lighting separation At this stage, the total illumination has been estimated at sparse reconstructed 3D points. We propagate it to all pixels of the image to decompose by using an image-guided propagation method [Bousseau et al. 2009], which yields a decomposition into reflectance and total illumination. We further decompose the total illumination into sun, sky and indirect illumination, by casting this as a Matting problem and enforcing the illumination values estimated at 3D points.

Results Our algorithm decomposes an image into four layers, which can be then modified independently. Fig. 1 shows examples of changes made possible by our approach, while the video demonstrates how to use the decomposition in image editing software. More results and comparisons with single-image approaches are shown in the supplementary document.

References

BOUSSEAU, A., PARIS, S., AND DURAND, F. 2009. User-assisted intrinsic images. *ACM Trans. Graph.* 28, 5.

* e-mail: {pierre-yves.laffont, adrien.bousseau, george.drettakis}@inria.fr