

Reasoning about Photo Collections using Models of Outdoor Illumination

Daniel Hauagge
 hauagge@cs.cornell.edu
 Scott Wehrwein
 swehrwein@cs.cornell.edu
 Paul Upchurch
 paulup@cs.cornell.edu
 Kavita Bala
 kb@cs.cornell.edu
 Noah Snavely
 snavely@cs.cornell.edu

Cornell University
 Ithaca, NY, USA

Natural illumination from the sun and sky plays a significant role in the appearance of outdoor scenes. We propose the use of sophisticated outdoor illumination models, developed in the computer graphics community, for estimating appearance and timestamps from a large set of uncalibrated images of an outdoor scene. We first present an analysis of the relationship between these illumination models and the geolocation, time, surface orientation, and local visibility at a scene point. We then use this relationship to devise a data-driven method for estimating per-point albedo and local visibility information from a set of Internet photos taken under varying, unknown illuminations. Our approach significantly extends prior work on appearance estimation to work with sun-sky models, and enables new applications, such as computing timestamps for individual photos using shading information.

1 Modeling illumination in outdoor scenes

The illumination arriving at a point in an outdoor scene depends on several key factors, including:

- geographic location
- time and date
- surface orientation
- local visibility

Our model describes the irradiance incident at an outdoor scene point on a clear day as a function $L(\phi, \lambda, t, \alpha, \vec{n})$ where ϕ, λ are latitude and longitude, t is the time and date, \vec{n} is the normal, and α is the *local visibility angle*. This angle α is a parameterization of local visibility based on a model of ambient occlusion proposed by Hauagge et al. [1], which models local geometry around a point as a cylindrical hole with angle α from the normal to the opening. Figure 1 shows examples of L , in the form of spheres rendered under predicted outdoor illumination at various times and α angles, at a given location on Earth.

2 Method

A georegistered 3D point cloud built using SfM and MVS provides geographic location (ϕ, λ) , surface normals (\vec{n}) , and a set of observed pixel values for each point (I_x) . We first estimate the albedo of each point, then use the albedo to estimate lighting and capture time for each photo.

Estimating Albedo. We adopt a simple Lambertian image formation model $I_x = \rho L_x$ where I_x is the observed color of a point x in a given image I , ρ_x is the (assumed constant) albedo at that point, and L_x is the irradiance as defined above. Given many observations of a point I_x , we derive the albedo ρ_x by dividing the average observed color $\mathcal{E}[I_x]$ by an estimate of the average illumination $\mathcal{E}[L_x]$.

Our key insight is that we can use a sun/sky model to predict illumination for a given condition, or indeed the *average* illumination for a given scene. For a given location, time, and visibility angle, we compute a physically-based environment map (we use the model of Hosek and Wilkie [2]) and, for each normal, integrate over the visible portion of the environment map to produce a database of spheres giving values for L at each normal direction, as illustrated in Figure 1(a-b). We then estimate expected illumination $\bar{L}(\vec{n}, \alpha)$ as a function of normal and visibility angle by taking the average over a set of times sampled throughout the year.

For each point x , we have a surface normal estimate \vec{n}_x from the 3D reconstruction; however we also need the visibility angle α_x to look up

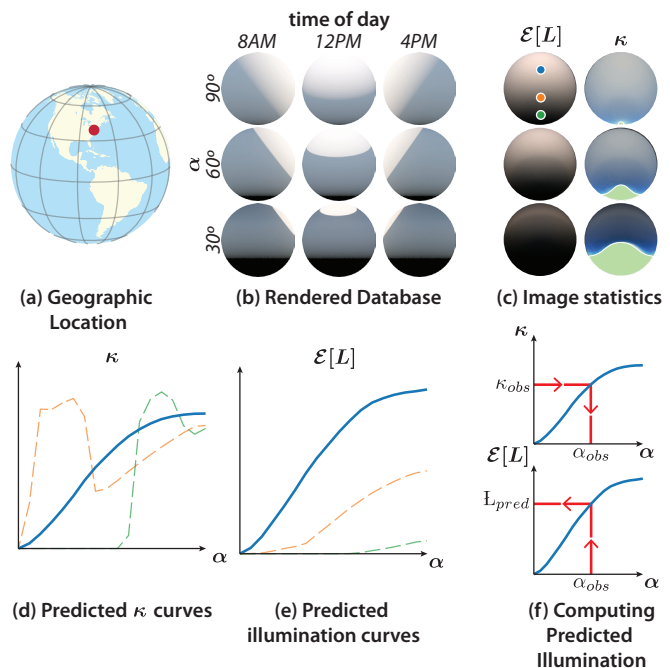


Figure 1: For a given geographic location (a), we render a database of spheres (b) covering all possible times (over a full year) and visibility angles. (c) We compute $\mathcal{E}[L]$ and κ for each α . Green regions correspond to combinations of normal direction and crevice for which we cannot reliably recover albedo.

the appropriate expected illumination $\bar{L}(\vec{n}_x, \alpha_x)$. Under a simpler lighting model, Hauagge et al. showed that α can be determined analytically as a function of an albedo-invariant image statistic $\kappa_x = \mathcal{E}[I_x]^2 / \mathcal{E}[I_x^2]$. Under our more complex illumination model, we instead relate κ to α numerically by computing $\kappa(\vec{n}, \alpha)$ over the predicted illumination values provided by the sun/sky model, as shown in Figure 1(c). We let α_x be the alpha for which $\kappa(\vec{n}_x, \alpha)$ most closely matches the observed κ_x .

Estimating Time of Day. With albedos in hand, we can estimate illumination for an image by dividing each visible point’s observed color value by the estimated albedo $L_x = \frac{I_x}{\rho_x}$. To estimate the time for that image, we can compare this estimated per-point illumination to the illumination predicted by the sun/sky model at a set of times candidate times t (potentially sampled over the entire year). The predicted time t^* is then the time for which the observed and predicted illumination are most similar.

Results. Our technique recovers the albedo of outdoor scenes more accurately than Hauagge et al. [1] and successfully identifies and discards points whose albedo cannot be recovered. The timestamp estimates using our albedo have median error under one hour (about 15 degrees of sun position) on our test datasets, which significantly outperforms random chance and a state-of-the-art single image method. Please see the full paper for more details and complete results.

[1] Daniel Hauagge, Scott Wehrwein, Kavita Bala, and Noah Snavely. Photometric ambient occlusion. *CVPR*, 2013.
 [2] Lukas Hosek and Alexander Wilkie. An analytic model for full spectral sky-dome radiance. *ACM Transactions on Graphics*, 2012.