Lighting Estimation in Outdoor Image Collections
Supplementary Material

Jean-François Lalonde
Laval University
jflalonde@gel.ulaval.ca

Iain Matthews
Disney Research, Pittsburgh
iainm@disneyresearch.com

1. Image database

Dataset	Hamburg	Smith	UC1	UC2	Arts1	Arts2	MM1	MM2
#imgs SfM	283	353	260	469	392	320	291	465
#imgs light	65	75	43	67	77	75	63	85
Mesh								

Dataset	D1	D2	D3	Health	Rodef	Byzantine	Ascension1	Ascension2
#imgs SfM	316	431	231	436	313	524	164	521
#imgs light	71	63	66	91	95	143	61	150
Mesh								

Dataset		Hammer	CFA	Apts1	Apts2	Apts3	Resnick	
#imgs SfM		742	270	373	291	461	458	
#imgs light		165	78	84	76	72	80	
Mesh								

Figure 1. Statistics of our novel light and image collection database. For each model in our novel light and image collections database, we show (row-wise) its name, the number of images used in the SfM+PMVS reconstruction, the number of images with corresponding HDR light probes, and a rendering of the mesh obtained with Poisson reconstruction. The mesh textures are obtained by projecting the colored PMVS points onto the mesh, and are used here for display purposes only.

2. Angular profiles of sun models

Fig. 2 shows additional justification for why our double-exponential model is a better sun model than the well-known von-Mises Fisher (vMF) model by comparing their behavior on angular sun intensity profiles. While both behave similarly for very small angles $\theta < 1^{\circ}$, the *EPD* model falls sharply afterwards and does not accurately capture the data. The *exp* model, on the other hand, is a better predictor of the intensity variation due to the sun. Note that the y axis is in log space to account for the extremely high dynamic range of the sun.

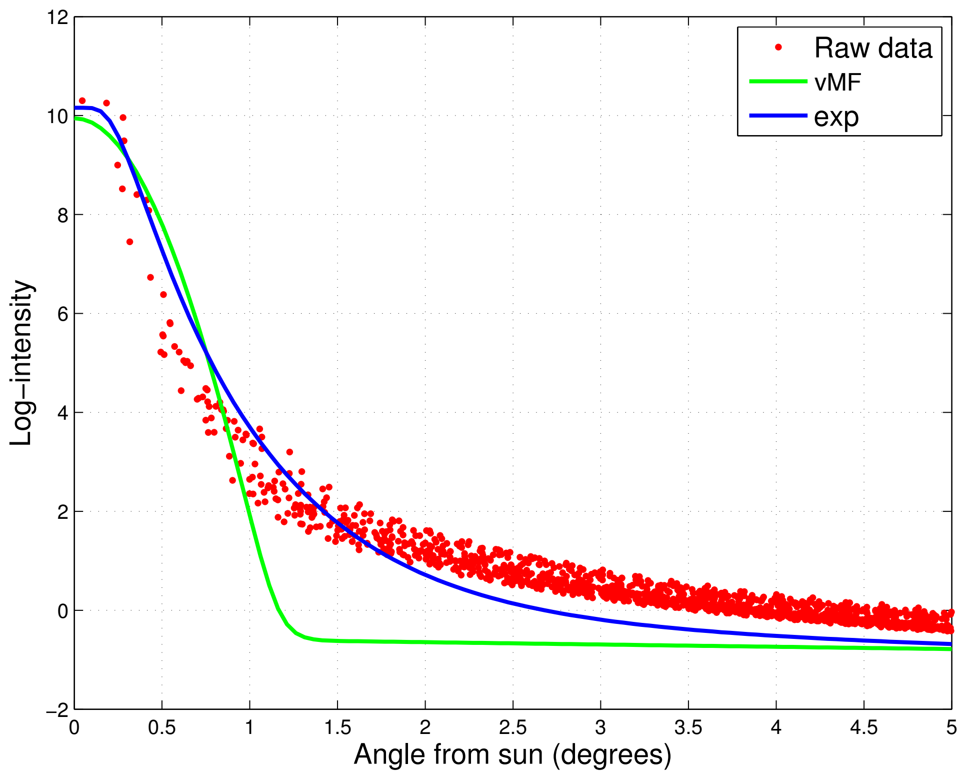

Figure 2. Comparison of angular profiles for the *vMF* and our sun models, on log-intensity data. Notice how our model better captures the sun falloff at $\theta > 1^{\circ}$ than the alternative, which falls short too abruptly.

3. Evaluating light models

To evaluate how well our illumination model approximates the original light probes, we compute the ratio r_{ill} of total light irradiance received by a flat ground plane (with normal pointing up) when lit by our estimated light probe, compared to when lit by the captured ground truth light probes. $r_{ill} = 1$ corresponds to a perfect match. Fig. 3 shows the ratio obtained by computing r_{ill} over the full sky dome (left plot), over a $5°$-diameter region centered around the sun position only (second plot), and by masking that same region for clear and overcast skies (third and fourth plots respectively), so that only skies are compared. Our model faithfully captures the sun and sky intensities, with a minimal number of parameters (11 for 3 color channels).

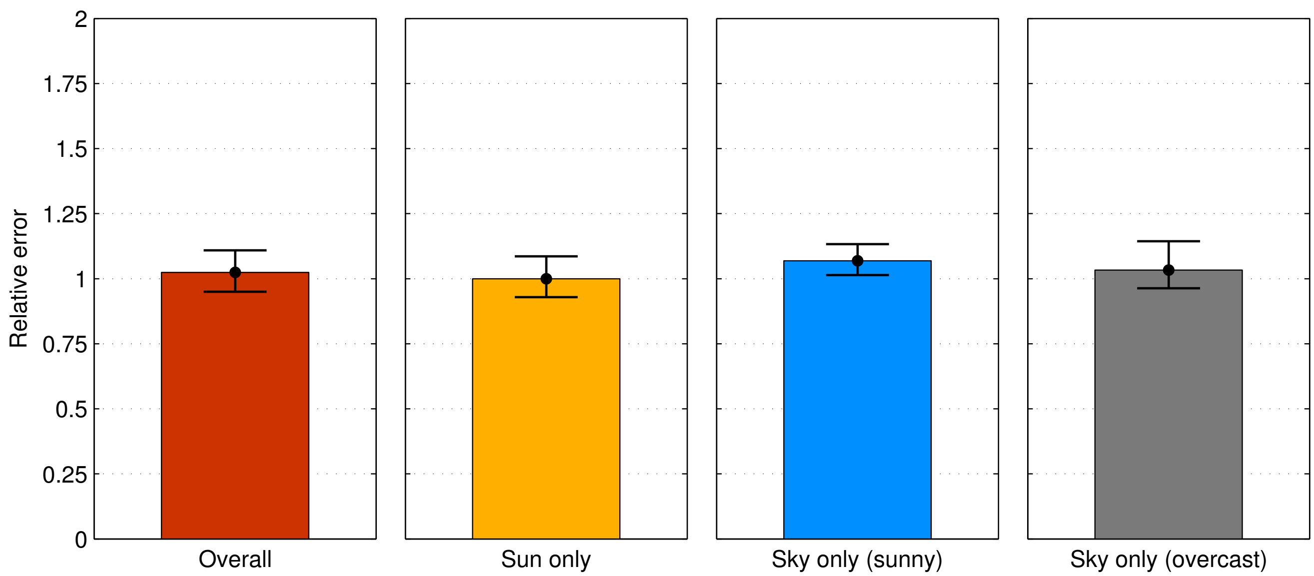

Figure 3. Quantitative error of our illumination model, compared to the captured ground truth light probes from sec. 3 in the main paper. The error is defined as the ratio of total light irradiances received by a flat ground plane, see text for more details. A perfect match should result in $r_{ill} = 1$. The first plot shows the ratio computed over the entire hemisphere. The second one shows the ratio computed within a $5°$-diameter angular window around the ground truth sun position. The third and fourth plots show ratios of the sky model on clear and overcast skies respectively, where the sun has been masked out. Error bars show the 25th and 75th percentiles.

4. Consistent relighting

Our approach can be used to automatically relight virtual objects consistently in every image in the collection, as is demonstrated in fig. 4. To do so, the user must first position the virtual objects using one image from the collection. Using the camera parameters obtained from the structure-from-motion algorithm, the objects can automatically be placed at the correct position in the other images. Our approach provides the high dynamic range illumination conditions to light the objects in a physically-plausible manner. Thanks to our lighting estimates, the objects appear realistic in their new surroundings, irrespective of the illumination conditions of the images.

Figure 4. Two virtual statues are inserted in all the images in the collection. The middle image in the bottom row was used as a reference to manually position the two virtual objects. Objects are lit by using our environment map estimates as the sole light sources, and results are generated with the LuxRender physically-based rendering engine within the Blender modeling software.