

What Is a Good Day for Outdoor Photometric Stereo?

Yannick Hold-Geoffroy¹ Jinsong Zhang^{*2} Paulo F. U. Gotardo³ Jean-François Lalonde^{†1}

¹Laval University, Quebec City ²Beihang University, Beijing ³Disney Research, Pittsburgh

Abstract

Photometric Stereo has been explored extensively in laboratory conditions since its inception. Recently, attempts have been made at applying this technique under natural outdoor lighting. Outdoor photometric stereo presents additional challenges as one does not have control over illumination anymore. In this paper, we explore the stability of surface normals reconstructed outdoors. We present a data-driven analysis based on a large database of outdoor HDR environment maps. Given a sequence of object images and corresponding sky maps captured in a single day, we investigate natural factors that impact the uncertainty in the estimated surface normals. Quantitative evidence reveals strong dependencies between expected accuracy and the normal orientation, cloud coverage, and sun elevation. In particular, we show that partially cloudy days yield greater accuracy than sunny days with clear skies; furthermore, high sun elevation—recommended in previous work—is in fact not necessarily optimal when taking more elaborate illumination models into account.

1. Introduction

Photometric Stereo (PS) is a well-known shape reconstruction technique which has been studied extensively since its inception in the early 80s [23]. The original idea—that shape can be recovered from several images of the same object lit from different directions—has since matured to provide robust algorithms able to deal with increasingly complex materials and lighting conditions [4–6, 11, 12].

The vast majority of the work in PS has so far focused on images captured indoors, in laboratory conditions. Recently however, new algorithms have targeted the case of images captured outdoors [1, 2, 19, 20, 24]. Clearly, this new scenario with uncontrolled, natural illumination differs sig-

nificantly from laboratory setups: we are left at the mercy of the sun, its motion through the sky, and the complex effects created by sunlight interacting with the atmosphere.

To tackle this new challenge, a natural first strategy has been to experiment with Lambertian reflectance and to model the sun as a point light source, to match a well-studied lab condition. Unfortunately, approaches based on this model have practical limitations caused by the movement of the sun in the sky for a given day. Depending on the latitude and time of year, its trajectory may lie too close to a 3D plane [19], yielding an under-constrained, two-source PS problem [9]. Possible solutions include waiting for a day when the sun trajectory is non-planar [19], or capturing several months of data [1, 2] to ensure good conditioning.

While these techniques can lead to well-defined solutions, they are not always practical in many scenarios with strict temporal or geographical constraints. A second strategy has therefore been to combine PS with other techniques such as multi-view stereo [13, 20], or use reference objects as in [15] or example-based PS [3, 10]. But can we accurately reconstruct surface geometry simply based on the photometric cue in an outdoor setting, without overly restrictive temporal and geographical constraints?

A promising approach to answer this question is to use more elaborate models of illumination—high dynamic range (HDR) environment maps [18]—as input to outdoor PS. Promising results have been reported in [24] for outdoor images taken within an interval of just eight hours (in a single day). However, the quality of outdoor results is reported to be inferior to that obtained in indoor environments, the decline being attributed to modest variation in sunlight. This observation leads to an interesting, yet unanswered question: had the sun path and atmosphere conditions been different on that day, could the quality of their results have been better? Or, in other words: what makes it a good day for outdoor Photometric Stereo?

As one might expect, the answer to the question above is intrinsically tied to the orientation of a particular surface patch, the associated hemisphere of lighting directions observed by the patch, and the variation in lighting intensity

^{*}This work was performed while Jinsong Zhang was an intern at Laval University.

[†]Contact author: jflalonde@gel.ulaval.ca




Figure 1. Examples from our dataset of HDR outdoor illumination conditions. In all, our dataset contains 3,800 different illumination conditions, captured from 10:30 until 16:30, during 23 days, spread over ten months and at two geographical locations. Each image is stored in the 32-bit floating point EXR format, and shown tone mapped here for display (with $\gamma = 1.6$). The companion video² shows time-lapse sequences for these sky environment maps.

in that hemisphere over the course of a day. So far, this question has only been explored in laboratory conditions or with simple directional illumination, where optimal lighting configurations can be theoretically derived [7, 16, 19]. No attempt has been made at answering this question with more realistic illumination models in an outdoor setup, where lighting cannot be controlled and atmospheric effects are difficult to predict.

In this paper, we present a systematic analysis of the expected performance of photometric stereo algorithms in outdoor settings. To explore the influence of these effects, we exploit a large database of outdoor HDR environment maps, which provides a rich sampling of the variability in outdoor illumination. Through a theoretical analysis (inspired by [22]) and supported by an extensive empirical evaluation as well as a preliminary quantitative validation, we derive confidence intervals that allow us to predict when surface normals can be reconstructed more accurately and solely from the photometric cue. Instead of focusing exclusively on sun position [19], our analysis incorporates the influence of all components of natural lighting: sun, sky, clouds, etc., as well as noise and surface albedo.

This paper does *not* introduce a novel photometric stereo algorithm. Rather, we explore the conditions in which PS may or may not work, when faced with the challenging case of uncontrollable outdoor illumination. Our main goal is to provide guidance for future research by assessing the *reconstructibility* of surface patches as a function of their orientation and the illumination conditions, given a set of HDR environment maps captured throughout a single day.

To make this novel analysis tractable, we make the following assumptions. We consider Lambertian surface patches, and assume that at most one day of data can be

used. We also assume that the dominant part of the light comes from the sky hemisphere (sun, sky, clouds, etc.) and surface patches are independent (cast shadows and inter-reflections are not modeled).

2. Dataset of HDR outdoor lighting conditions

To analyze the influence of outdoor lighting on photometric stereo, we rely on a rich dataset of high dynamic range images of the sky, captured under a wide variety of conditions. We use the environment map database of [17], which contains HDR images of the sky captured using the approach described in [21]. We augment the dataset of [17] with an additional set of images captured using a similar setup, but at a different geographical location. In all, the dataset we used for our analysis totals 3,800 illumination conditions, captured over 23 different days. To ensure the data is properly aligned temporally, the HDR sky photos were captured during a continuous 6 hour time interval on each of these days, from 10:30 until 16:30. Fig. 1 shows examples of the sky environment maps used in our analysis. Note that while the examples have been tone mapped for display, the actual sky images have extremely high dynamic range, and span the full 22 stops required to properly capture outdoor lighting [21]. In addition, all the images are converted to grayscale before the analysis is performed.

3. Sensitivity analysis for outdoor PS

This section derives the image formation model and the probabilistic method used to assess the uncertainty in the solution of outdoor PS. Although previous work have presented sensitivity analyses for standard PS with point light sources (e.g., [14, 16, 19, 22]), so far as we know, the case of outdoor PS with more complex illumination models has received little attention.

²Available at project webpage: <http://vision.gel.ulaval.ca/~jflalonde/projects/outdoorPS/>




Figure 2. A normal \mathbf{n} defines an integration domain $\Omega_{\mathbf{n}}$ equivalent to a hemisphere on the entire spherical environment map. Only light emanating from this hemisphere contribute to the shading on that patch. Therefore, patches with different normals are lit differently even if the environment map is the same.

3.1. Image formation model

Consider a small, Lambertian surface patch with normal vector \mathbf{n} and albedo ρ (w.l.o.g., assume albedo is monochromatic). At time t , this surface patch is observed under natural, outdoor illumination represented by the environment map $L_t(\omega)$ (e.g., fig. 1), with ω denoting a direction in the unit sphere. With an orthographic camera model, this patch is depicted as an image pixel with intensity

$$b_t = \frac{\rho}{\pi} \int_{\Omega_{\mathbf{n}}} L_t(\omega) \langle \omega, \mathbf{n} \rangle d\omega, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product. Integration is carried out over the hemisphere of incoming light, $\Omega_{\mathbf{n}}$, defined by the local orientation \mathbf{n} of the surface, fig. 2. This hemisphere corresponds to an occlusion (or attached shadow) mask; only half of the pixels in the environment map contribute to the illumination of the surface patch. To make the analysis tractable and independent on object geometry, this paper focuses on the simpler case without cast shadows.

This image formation model is then discretized as,

$$b_t = \frac{\rho}{\pi} \sum_{\omega_j \in \Omega_{\mathbf{n}}} \hat{L}_t(\omega_j) \langle \omega_j, \mathbf{n} \rangle, \quad (2)$$

with $\hat{L}_t(\omega_j) = L_t(\omega_j) \Delta\omega_j$ representing the environment map weighted by the solid angle $\Delta\omega_j$ spanned by pixel j ($\Delta\omega_j, \forall j$, are normalized as to sum to 2π). Eq. (2) can be further summarized into the equivalent form

$$b_t = \bar{\mathbf{l}}_t^T \mathbf{x} \quad (3)$$

where $\mathbf{x} = \rho \mathbf{n}$ is the albedo-scaled normal vector and

$$\bar{\mathbf{l}}_t = \frac{1}{\pi} \sum_{\omega_j \in \Omega_{\mathbf{n}}} \hat{L}_t(\omega_j) \omega_j \in \mathbb{R}^3 \quad (4)$$

is interpreted as the *mean light vector* for the environment map at a time t .

Given multiple images taken at times $t \in \{1, 2, \dots, T\}$, we collect all photometric constraints for patch \mathbf{x} to obtain:

$$\mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_T \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{l}}_1^T \\ \bar{\mathbf{l}}_2^T \\ \vdots \\ \bar{\mathbf{l}}_T^T \end{bmatrix} \mathbf{x} = \mathbf{L} \mathbf{x}. \quad (5)$$

With (5), this model of natural environmental illumination becomes quite similar to a model with a distant point light source, the well-known case in PS. However, note that each $\bar{\mathbf{l}}_t$ in \mathbf{L} is a function of $\Omega_{\mathbf{n}}$ and, thus, of \mathbf{n} .

Most importantly, in outdoor PS, a well-defined solution \mathbf{x} may exist even if the relative sun motion is nearly planar during a certain time interval. Instead of relying solely on sun direction, now, the solution requires non-coplanar mean light vectors $\bar{\mathbf{l}}_t$, which are determined by a comprehensive set of natural illumination factors.

3.2. Modeling uncertainty

From (5), the least-squares solution $\mathbf{x} = (\mathbf{L}^T \mathbf{L})^{-1} \mathbf{L}^T \mathbf{b}$ of outdoor PS is clearly affected by the condition number of \mathbf{L} . Thus, we next characterize how well the solution \mathbf{x} is constrained by natural, outdoor illumination within a given time interval (e.g., one day)—which is encoded by the set of mean light vectors $\bar{\mathbf{l}}_t$ in \mathbf{L} or, equivalently, the set of environment maps $L_t(\cdot)$.

To assess the reliability of a solution \mathbf{x} , we follow standard practice in PS [16, 22] and consider image measurements corrupted by zero-mean Gaussian noise with equal variance σ^2 (as least squares estimation is only optimal for this practical, most common noise model). Thus, \mathbf{b} in (5) follows a normal distribution:

$$\mathbf{b} \sim \mathcal{N}(\boldsymbol{\mu}_b, \sigma^2 \mathbf{I}), \quad (6)$$

where $\boldsymbol{\mu}_b$ has the (unknown) uncorrupted pixel values.

Since the desired least-squares solution for the albedo-scaled normal, $\mathbf{x} = (\mathbf{L}^T \mathbf{L})^{-1} \mathbf{L}^T \mathbf{b}$, is a linear transformation of a Gaussian random vector, it is easy to show that

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_x, \sigma^2 (\mathbf{L}^T \mathbf{L})^{-1}), \quad (7)$$

where $\boldsymbol{\mu}_x = (\mathbf{L}^T \mathbf{L})^{-1} \mathbf{L}^T \boldsymbol{\mu}_b$ is the expected value of \mathbf{x} . Once the albedo of a surface patch is known, we analyze its contribution to the uncertainty in the estimated normal vector, $\mathbf{n} = \rho^{-1} \mathbf{x}$, using a similar distribution,

$$\mathbf{n} \sim \mathcal{N}\left(\frac{\boldsymbol{\mu}_x}{\rho}, \frac{\sigma^2}{\rho^2} (\mathbf{L}^T \mathbf{L})^{-1}\right). \quad (8)$$

The marginal distributions in (8) allow us to derive confidence intervals that indicate the uncertainty in each component of the least squares estimate $\hat{\mathbf{n}} = [\hat{n}_x \ \hat{n}_y \ \hat{n}_z]^T$ of

$\mathbf{n} = [n_x \ n_y \ n_z]^T$. The corresponding 95% confidence interval [8] is given by

$$\hat{\mathbf{n}} \pm \delta, \quad \text{with } \delta_k = 1.96 \frac{\sigma \lambda_k}{\rho}, \quad (9)$$

where λ_k is the square root of the k th element on the diagonal of $(\mathbf{L}^T \mathbf{L})^{-1}$. As expected, the sensor-dependent noise level σ is not the only factor that determines uncertainty. The gain factor λ_k/ρ in (9) reveals how outdoor illumination (the conditioning of \mathbf{L}) and albedo can amplify the effect of noise on the solution $\hat{\mathbf{n}}$. The lower the albedo ρ is, the larger is the variance in the obtained estimate $\hat{\mathbf{n}}$ (as less light is reflected towards the camera). Our goal is then to answer the remaining question: how do natural changes in outdoor illumination affect this gain factor (λ_k) and, therefore, uncertainty?

3.3. An intuitive measure of uncertainty

To provide a measure of uncertainty that is more intuitive than (9), we consider angular distances in degrees,

$$\theta^\pm = \cos^{-1}(\mathbf{n}^T \hat{\mathbf{n}}^\pm), \quad \text{where } \hat{\mathbf{n}}^\pm = \frac{\hat{\mathbf{n}} \pm \delta}{\|\hat{\mathbf{n}} \pm \delta\|}. \quad (10)$$

The uncertainty in the estimate of \mathbf{n} is then summarized in a single confidence interval, in degrees,

$$\mathcal{C}_{\mathbf{n}} = [0, \max(\theta^\pm)], \quad (11)$$

which indicates the expected accuracy of the estimated surface orientation $\hat{\mathbf{n}}$.

Note that the condition number, determinant, and trace of matrix $(\mathbf{L}^T \mathbf{L})^{-1}$ can also be used as measures of total variance in the estimated solutions—as done in [22]—to find the optimal location of point light sources in PS. These measures are closely related to the rank of matrix \mathbf{L} , which must be three for a solution to exist; that is, $\mathbf{L}^T \mathbf{L}$ must be nonsingular. In practice, this matrix is always full-rank, although it is often poorly conditioned [19]. In the remaining sections, we consider confidence intervals $\mathcal{C}_{\mathbf{n}}$ in degrees, as they provide a more intuitive measure of uncertainty in the obtained solutions. Finally, in (10), the normalization of $\hat{\mathbf{n}}^\pm$ to unit length reflects the fact that part of the estimation error is propagated to the estimated albedo ρ (*i.e.*, the length of the computed least-squares solution vector \mathbf{x}). In this paper, we focus on analyzing our ability to recover geometry and will assume that the albedo is constant and known.

4. Analyzing the outdoor lighting dataset

Now that we are equipped with a tool to characterize the stability of surface normal estimation in outdoor PS, we proceed to apply it to all 23 days from our dataset (sec. 2) in order to determine the conditions in which surface normals

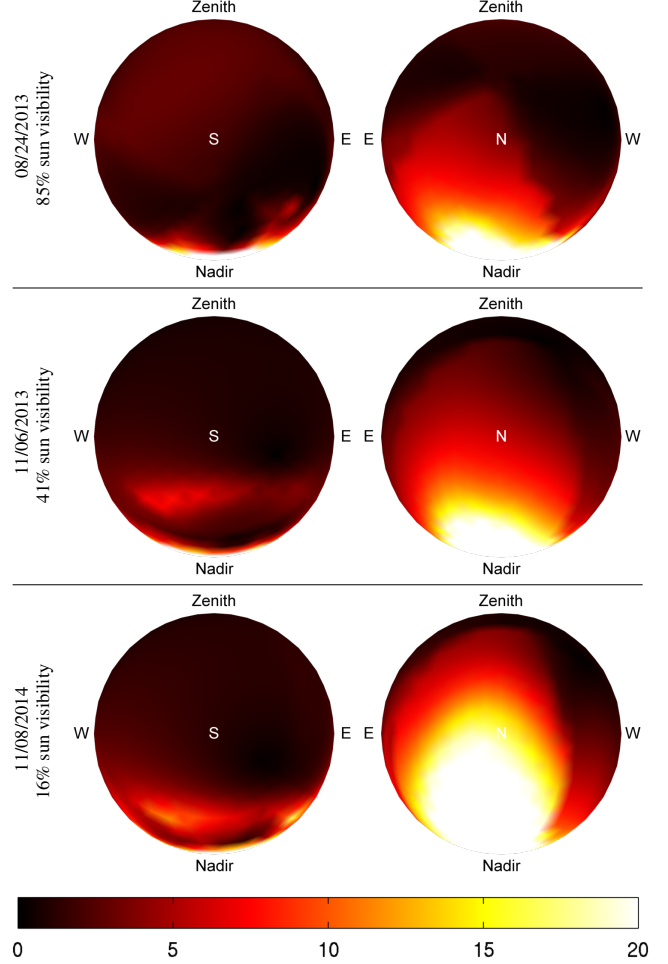


Figure 3. Uncertainty in normal estimation with $\sigma = 1\%$ is indicated by 95% confidence intervals (in degrees), as a function of ground-truth surface normal. Results are shown for three different days in our dataset (same days as in fig. 1). The plots show the full sphere of normals from two different viewpoints: South (left), and North (right). Cardinal directions are shown for reference. The color-coding is indicated below the plots. See companion video for an animated version of these plots.

can accurately be reconstructed. We first describe how the confidence intervals $\mathcal{C}_{\mathbf{n}}$ are computed and visualized, then analyze the effect of two important characteristics of outdoor lighting: the degree of cloud coverage in sec. 4.2, and the sun elevation in sec. 4.3.

4.1. Computing the stability of outdoor PS

We aim to apply the stability analysis of sec. 3 on all 23 days from our dataset, and on all possible normal directions. To do so, we consider each day independently, and first sample directions on the sphere by subdividing an icosahedron three times, yielding a set of 642 potential orientations $\mathcal{N} = \{\mathbf{n}_1, \dots, \mathbf{n}_{642}\}$. Then, for each $\mathbf{n}_j \in \mathcal{N}$, we




Figure 4. Median confidence interval of normal estimates (red line) as a function of mean sun visibility over the course of the day for various values of σ . Our analysis predicts that normal reconstruction errors will likely be high if the sky is completely overcast (low sun visibility), or completely clear (high sun visibility). Good results can thus be expected in partially cloudy conditions, as shown in fig. 5. The lower (upper) edge of each blue box indicates the 25th (75th) percentile. Statistics are computed only on normals pointing upwards to lessen ground effects.

build the illumination matrix \mathbf{L} in (5), given all 6 hours of data for that day. Finally, the 95% confidence interval $\mathcal{C}_{\mathbf{n}_j}$ is computed using (11), indicating the uncertainty (possible reconstruction error) for that normal—the higher the interval, the less stable the solution.

To compute $\mathcal{C}_{\mathbf{n}_j}$, values for the noise level σ and the surface albedo ρ from (9) must be chosen. For the albedo, we will consider the best case with $\rho = 1$. To set σ , we first render noise-less pixel values \mathbf{b}_j using the (ground-truth) normal \mathbf{n}_j and (5), with \mathbf{L} including our entire collection of environment maps. All figures have been generated with σ set at 1% of the 95th percentile value of the resulting \mathbf{b}_j , unless otherwise stated. This particular value is chosen to yield a small, yet non-negligible, level of noise that is similar to that in previous work [16]. The values of ρ and σ are kept constant in order to focus solely on the influence of lighting conditions, but could be set to match particular experimental conditions if needed.

Fig. 3 shows the result of such an analysis on the three days of fig. 1. For each day, two sides of the sphere of normal directions are shown: seen from the South (left), and from the North (right). The spheres are color-coded according to the confidence interval $\mathcal{C}_{\mathbf{n}}$ for each normal direction. Note that vertex interpolation is used to display full spheres, but valid data is available only at vertices (thus the staircase effect in some plots).

At first glance, we notice that normals pointing down (towards Nadir) consistently have high confidence intervals, irrespective of the illumination conditions. The stability of outdoor PS on these normals is thus expected to be low. This behavior concords with expectation: normals pointing down define integration domains $\Omega_{\mathbf{n}}$ (see fig. 2) which mostly include the ground, whose intensity does not vary

spatially throughout the day. Another interesting observation from fig. 3 is that the same normal exhibits different confidence intervals depending on the day. This raises the question: what is the relation between outdoor illumination conditions and uncertainty in the recovered surface normal?

4.2. Influence of cloud cover

It is already apparent from fig. 3 that cloud coverage has an effect on the uncertainty of normal reconstruction, since an overcast day (last row) clearly does not behave the same way as a day with light clouds (top row). Here, we present a more systematic analysis of the influence of cloud coverage. To control for the effect of the sun elevation (which will be explored in sec. 4.3), the analysis is performed on days with similar sun elevations by keeping only the skies captured in October and November.

We approximate cloud coverage by computing the fraction of time that the sun is visible, *i.e.*, that it fully shines on the scene, for a given day. To do so, we simply find the brightest spot in a sky image, and determine that the sun shines on the scene if the intensity of the brightest pixel is greater than 20% of the maximum sun intensity—we determined empirically that this is the point at which the sun is bright enough to start creating cast shadows. Cloud coverage is represented by computing the mean sun visibility for a given day. A value of less than 10–15% would indicate a mostly overcast sky, while skies are mostly clear with values above 85–90%.

The relation between sun visibility and the confidence interval $\mathcal{C}_{\mathbf{n}}$ is shown in fig. 4. Normal reconstruction errors will likely be quite high in two situations: when the sky is completely overcast (low sun visibility), or completely




Figure 5. Influence of cloud cover on the 95% confidence intervals with $\sigma = 1\%$. Each row shows a different type of sky, based on sun visibility. For example, the top row shows confidence intervals averaged over all days with direct sun visibility in the range 85%–100%. The averaged days presented similar sun elevations. As with fig. 3, the plots show the full sphere of normals from two different viewpoints: South (left), and North (right).

clear (high sun visibility). Interestingly, good reconstruction results are expected for a wide range of cloud coverage conditions, ranging from 10–90% mean sun visibility.

These results are corroborated by fig. 5, which shows the confidence intervals themselves on a plot similar to fig. 3.

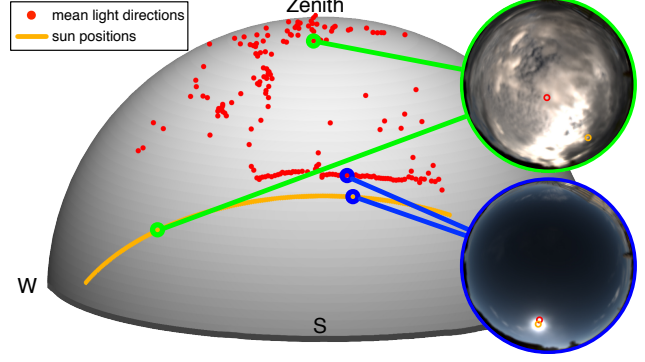


Figure 6. Cloud effect on the mean light direction over one day: while the sun path (orange) yields nearly co-planar directions of illumination, the mean light directions (red dots) provide a much more varied set (data from 11/06/2013, 2nd row of fig. 1).

Results there are presented by averaging the intervals of skies belonging to four groups: overcast (0–15%), mixed overcast (15–50%), mixed clear (50–85%), and clear (85–100%) days. Again, high uncertainty results are visible for the two extreme cases of fully overcast and fully clear days, while the remainder indicate more stable solutions.

The improved conditioning in mixed skies is explained by the following key observation: cloud cover shifts the mean light vectors $\bar{\mathbf{l}}_t$ towards zenith and away from sun trajectory in the sky. Therefore, even when the sun moves along a trajectory that nearly lies on a 3D plane, as shown in fig. 6, cloud cover effectively causes an *out-of-plane* shift of the mean light vectors, making reconstruction possible.

This key observation also demonstrates the advantages of adopting more elaborate illumination models (e.g., [24]). For instance, the simpler point light model was used in [19] to study the conditioning of outdoor PS. Because the atmospheric component is not modeled, the conclusion was that single-day reconstruction breaks down in two cases of nearly coplanar sun directions: closer to the poles near the winter solstice, and worldwide near an equinox. Our results suggest that more attention should be placed on the illumination model, without focusing exclusively on the sun.

4.3. Influence of sun elevation

We also hypothesized the position of the sun to have an important effect on the ability to recover surface normals outdoors. Since our dataset is already aligned in terms of sun azimuths (see sec. 2), we now analyze the influence of sun elevation. For this purpose, we retain only the days for which the sun visibility was between 15–85% of the time. We then compute the relation between sun elevation and expected reconstruction confidence, as illustrated in fig. 7.

Sun elevation does appear to have an influence, with higher sun elevations resulting in slightly increased uncer-




Figure 7. Median confidence interval of normal estimates (red line) as a function of maximum sun elevation over the course of the day for various values of σ . Although the effect is not as significant as with cloud cover, we anticipate a small decrease in performance at higher sun elevations. The lower (upper) edge of each blue box indicates the 25th (75th) percentile. Statistics are computed only on normals pointing upwards to lessen ground effects.

tainty than when the sun is lower in the sky. However, the impact is less drastic than that of cloud cover from sec. 4.2. These results can be explained, at least in part, by the smaller sun-zenith shift introduced by cloud cover when the sun is already high in the sky; therefore, smaller improvements in conditioning are obtained.

Our data-driven analysis shows that higher sun elevations—predicted as preferable by the analysis in [19]—are in fact not necessarily optimal when taking more elaborate illumination models into account.

4.4. Analyzing full objects

So far, our analysis has focused solely on one patch (one normal direction) at a time. But can we also say something about an entire object? Clearly, a full answer to this question would require analyzing non-local effects such as occlusions, inter-reflections, cast shadows, etc. However, we hypothesize that a simpler analysis ignoring these effects can still provide useful results. We therefore predict the performance of an outdoor PS algorithm by computing the expected value of the confidence interval, $\mathbb{E}_n[C_n]$, over an entire object; expectation over the normal vector is computed using a prior distribution taken from a simpler convex shape (*e.g.*, a sphere). Fig. 8 shows the expected reconstruction performance for two objects: a bottle, and the bunny. Overall, surface reconstruction performance is expected to be best when objects face south (*i.e.*, the sun).

5. Validation on real object images

The analysis performed on our dataset indicates that surface patches may be better reconstructed in certain conditions, dependent upon cloud coverage, sun elevation, and the orientation of the patch itself. At the time of this writing, a thorough validation on real outdoor images of different objects is currently under development. This section reports preliminary results on real outdoor images of an owl

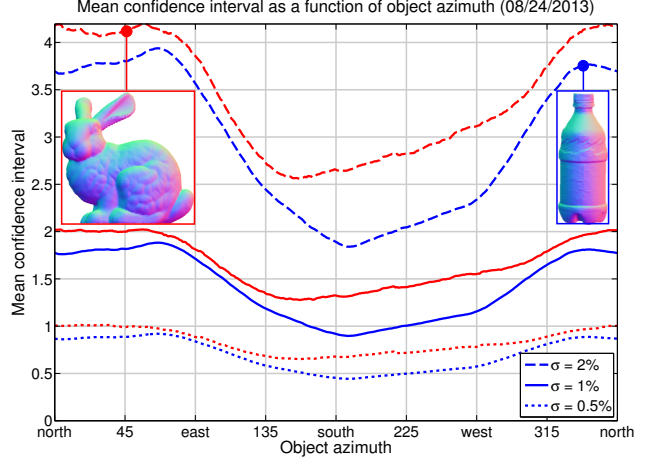


Figure 8. Predicting object reconstruction performance for the 08/24/2013 dataset. These curves show the mean surface reconstruction variance for two objects: a bottle (blue) and the bunny (red). Irrespective of the noise level, surface reconstruction performance is expected to be best when the objects face south.

statuette (fig. 9), which was scanned with a Creaform MetraSCAN at a resolution of 0.05mm to provide a reference of true surface orientation.

We oriented this owl statuette towards south and took 66 HDR captures using a Canon EOS Rebel SL1 between 10h30 and 16h30, local time, in Quebec City. These captures were synchronized with the HDR environment maps described in sec. 1, providing high-fidelity estimates of the illumination conditions for each image, as shown in fig. 9. The laser scan was then manually aligned to these images.

A simple outdoor PS algorithm was run on evenly distributed patches on the owl images. For each patch, the surface normal is obtained as to minimize the rendering error given by the model in (5), in a least-squares sense. Since this estimation problem is nonlinear on the normal \mathbf{n} , and we consider a small number of surface patches, a solution is computed via simple exhaustive search on a set of 321 candidate orientations for the light hemisphere Ω_n . Given Ω_n , a candidate solution \mathbf{x} is obtained via linear least squares. For the best solution, an estimation error, in degrees, is computed using the reference normal from the laser scan. The uncertainty measure (confidence interval C_n) derived in sec. 3.3 is also computed for each patch.

Results from this quantitative evaluation are given in fig. 10. As shown on the left, normals pointing upwards are generally recovered more accurately than normals pointing downwards. This is in concordance with the predicted confidence intervals shown on the right in fig. 10. The behavior of the reconstruction error is, in general, well-predicted by the confidence intervals.




Figure 9. Real outdoor HDR images of owl statuette and corresponding HDR environment maps (top row) providing synchronized, high-fidelity estimates of illumination conditions. All images were acquired on 10/11/2014 and tone-mapped for display only (with $\gamma = 1.6$).

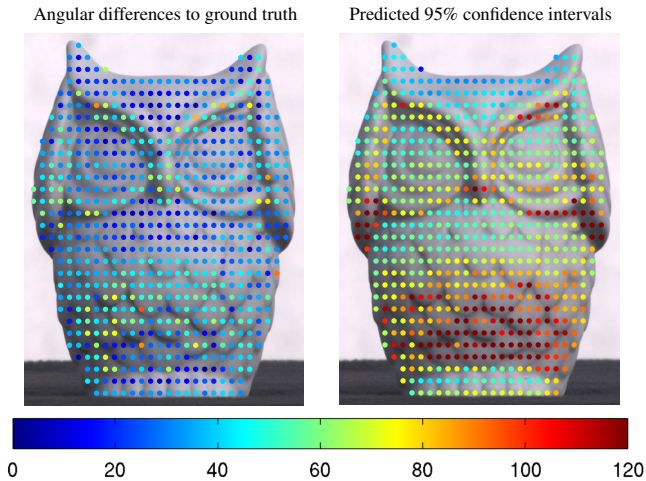


Figure 10. Quantitative validation on images of a real object. Evenly distributed patches of the owl images were used to perform a simple calibrated outdoor PS algorithm. Error estimates of the surface normals are reported on the left using the reference surface orientation in the laser scan. The predicted 95% confidence intervals C_n from (11) are shown on the right. A few large estimation errors on the left image may be caused by imperfect alignment between the scanned model and the images.

6. Discussion and future work

This paper has presented the first data-driven analysis of the expected behavior of PS under outdoor illumination. In this scenario, we have no control over the illumination conditions themselves, so existing methods to determine an optimal illumination setup [7, 16] cannot be applied. Our goal is, therefore, to reveal natural factors that distinguish good and unfavorable daylight conditions for outdoor PS.

The recent work of Shen et al. [19] presents an insightful, theoretical analysis of the conditioning of outdoor PS using the standard illumination model with a point light source. Their work suggests that PS reconstruction becomes unstable in two cases of nearly planar sun motion: near the poles during the winter solstices, and worldwide near the

equinoxes. However, these conclusions are drawn based on a simple illumination model that focuses exclusively on direct sun light, without considering other atmospheric components. Our approach is fundamentally different in that we adopt a more realistic illumination model and then follow a data-driven approach to investigate the conditioning of outdoor PS. We identify not only solar, but also atmospheric and object-intrinsic factors that contribute to stability, or uncertainty, in the photometric reconstruction.

To achieve our goal, we exploit a large database of outdoor HDR environment maps that provide a rich sampling of the natural variability in outdoor illumination. To make comprehensive use of this dataset, our photometric model considers not only direct sun illumination, but also the full atmospheric (sky) component and an additional ground effect. We then derive a theoretical analysis to investigate the stability of surface normal reconstruction, as measured by an intuitive angular confidence interval. Our empirical analysis reveals how reconstruction is affected by surface orientation, cloud coverage, and sun elevation. Our preliminary quantitative experiments corroborate the predictions with actual reconstruction results.

In short, our analysis revealed the following important observations of direct practical value. We show that better stability can be achieved when:

1. surface patches are oriented South and above the horizon: since our data was captured in the Northern Hemisphere, these are the patches (normals) that “see” the sun while it moves during the day;
2. the sky is partially cloudy throughout the day: this resulted in improved stability than either completely sunny, or completely overcast days;
3. the sun is lower in the sky: this resulted, on average, in slightly better performance than with days where the sun is higher in the sky.

While these observations might seem intuitive to an experienced practitioner, as far as we know, our work is the first attempt to confirm or contradict such intuition in outdoor PS. For instance, a common belief has been that clear days were preferable [1, 2, 13, 24], instead of mixed weather.

Another goal of this paper has been to inform future research in outdoor PS. Clearly, a better understanding of the sources of uncertainty is useful in developing better reconstruction algorithms (*e.g.*, better informed regularization). Despite the interesting observations above, this work still has some limitations that also open the way to interesting future work. First, it would be interesting to explore the correlation between the predicted reconstruction performance (*e.g.*, fig. 8) and the performance of an actual outdoor photometric stereo algorithm such as [24]. Second, our analysis assumes white Gaussian noise, for practical purposes; more realistic noise models for HDR photography should be investigated. Third, our current analysis does not model random perturbations in the light probes themselves (*i.e.*, \mathbf{L}). These are interesting extensions that increase the complexity of our derivations above and, thus, are left as future work. Finally, some seasons are not well represented in our dataset yet. We plan on continuing data capture to further enrich this database of outdoor environment maps.

Acknowledgements

We thank Mathieu Garon, Marc-André Gardner, and Martin Gagnon for their help with data capture. This work was partially supported by the REPARTI Strategic Network, NSERC Discovery Grant RGPIN-2014-05314, and a new faculty grant from the ECE department at Laval University.

References

- [1] A. Abrams, C. Hawley, and R. Pless. Heliometric stereo: Shape from sun position. In *European Conference on Computer Vision*, 2012.
- [2] J. Ackermann, F. Langguth, S. Fuhrmann, and M. Goesele. Photometric stereo for outdoor webcams. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 262–269, 2012.
- [3] J. Ackermann, F. Langguth, S. Fuhrmann, A. Kuijper, and M. Goesele. Multi-view photometric stereo by example. In *International Conference on 3D Vision*, 2014.
- [4] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [5] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1239–1252, 2003.
- [6] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007.
- [7] O. Drbohlav and M. Chantler. On optimal light configurations in photometric stereo. In *IEEE International Conference on Computer Vision*, pages 1707–1712, 2005.
- [8] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, New York, NY, 2009.
- [9] C. Hernández, G. Vogiatzis, and R. Cipolla. Overcoming shadows in 3-source photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(2):419–426, 2011.
- [10] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, 2005.
- [11] S. Ikehata and K. Aizawa. Photometric stereo using constrained bivariate regression for general isotropic surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- [12] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [13] K. Inose, S. Shimizu, R. Kawakami, Y. Mukaigawa, and K. Ikeuchi. Refining outdoor photometric stereo based on sky model. *IPSJ Transactions on Computer Vision and Applications*, 5(0):104–108, 2013.
- [14] X. Jiang and H. Bunke. On error analysis for surface normals determined by photometric stereo. *Signal Processing*, 23(3):221 – 226, 1991.
- [15] M. K. Johnson and E. H. Adelson. Shape estimation in natural illumination. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [16] M. Kludiny and A. Hilton. Error analysis of photometric stereo with colour lights. *Pattern Recognition Letters*, 48(0):81 – 92, 2014.
- [17] J.-F. Lalonde and I. Matthews. Lighting estimation in outdoor image collections. In *IEEE International Conference on 3D Vision*, 2014.
- [18] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec. *High dynamic range imaging*. Morgan Kaufman, 2005.
- [19] F. Shen, K. Sunkavalli, N. Bonneel, S. Rusinkiewicz, H. Pfister, and X. Tong. Time-lapse photometric stereo and applications. *Computer Graphics Forum (Proc. Pacific Graphics)*, 33(7), Oct. 2014.
- [20] B. Shi, K. Inose, Y. Matsushita, P. Tan, S.-K. Yeung, and K. Ikeuchi. Photometric stereo using internet images. *International Conference on 3D Vision*, 2014.
- [21] J. Stumpfel, A. Jones, A. Wenger, C. Tchou, T. Hawkins, and P. Debevec. Direct HDR capture of the sun and sky. In *Proceedings of AFRIGRAPH*, 2004.
- [22] J. Sun, M. Smith, L. Smith, and A. Farooq. Examining the uncertainty of the recovered surface normal in three light photometric stereo. *Image and Vision Computing*, 25(7):1073–1079, 2007.
- [23] R. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Eng.*, 19(1):139–144, 1980.
- [24] L.-F. Yu, S.-K. Yeung, Y.-W. Tai, D. Terzopoulos, and T. F. Chan. Outdoor photometric stereo. In *IEEE International Conference on Computational Photography*, 2013.