Supplemental Material: Implementation Details for Multi-View Intrinsic Images of Outdoors Scenes with an Application to Relighting

SYLVAIN DUCHENE, CLEMENT RIANT, GAURAV CHAURASIA, JORGE LOPEZ MORENO, PIERRE-YVES LAFFONT, STEFAN POPOV, ADRIEN BOUSSEAU, GEORGE DRETTAKIS
Inria

1. INTRODUCTION

In this supplemental material we present the implementation details for our algorithm. Specifically, we present the details for the indirect light compensation (Sec. 4 in the main text), the refinement for S_{env} and visibility (Sec. 7), and an additional comparison for the Toys scene.

2. COMPENSATING FOR SUPERFLUOUS INDIRECT LIGHT

The outdoors scenes we target contain perpendicular and horizontal surfaces (walls, floors, etc.). The reconstruction of such corners is often inaccurate, with geometry being added to the proxy. We often observe such geometry at grazing angles in the photographs, resulting in a high median value. When gathering indirect light at a given point x this can result in a higher contribution from such points. Finding the correct attenuation factor would require complete geometry and BRDF data, so we can only provide an approximate scale factor. Consider such a point x at which we gather light, and a point y on another surface contributing to x. The incoming angle θ_i is the angle between the direction $y - x$ and the normal n_y at y. We attenuate incoming lighting by $\cos\theta_i$, thus reducing the contribution at grazing angles, which is amplified by the incorrect reconstruction. This is a coarse approximation, but is well adapted to the case of perpendicular surfaces such as walls and ground which are predominant in outdoor scenes. This approach improves the result in all scenes we tested, in particular in regions containing evidently non-diffuse surfaces.

3. IMPLEMENTATION DETAILS OF S_{ENV} REFINEMENT

To refine the estimation of S_{env} we first find a set of light/shadow pairs, we then compute the offset values x_{sl} and propagate the refined S^n_{env} values over the image. The implementation has two main steps: finding pairs and offset values and smooth propagation.

Pairs and Offset Values. We find pairs by traversing shadow boundaries, pairs, in a manner similar to the L_{sun} estimation process (Sec.5 in the main text). We keep pairs with same reflectance, which we identify by a small D_{ij} value, since the visibility labels i and j are mostly correct. We also only keep pairs that satisfy the chromatic alignment of shadow/light pairs used in [Guo et al. 2011]; we thus avoid creating pairs on incorrectly classified boundaries.

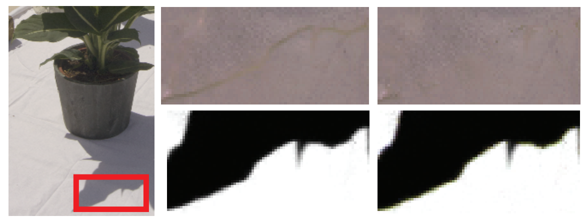

Fig. 1: The reflectance contains halo artifacts in penumbra regions due to errors in the visibility (top middle). We re-estimate the visibility (bottom, mid and right) to remove these artifacts (top right). The differences in visibility are very subtle, please zoom into the pdf to see them.

For each pair, we add an offset x_{sl} to S_{env} to make the two reflectances equal:

$$R_s = R_l \Rightarrow \tag{1}$$

$$\frac{I_s}{v^s_{\text{sun}} S^s_{\text{sun}} + S^s_{\text{env}} + x_{sl}} = \frac{I_l}{v^l_{\text{sun}} S^l_{\text{sun}} + S^l_{\text{env}} + x_{sl}} \tag{2}$$

Re-arranging the terms gives the offset value:

$$x_{sl} = \frac{I_s(v^l_{\text{sun}} S^l_{\text{sun}} + S^l_{\text{env}}) - I_l(v^s_{\text{sun}} S^s_{\text{sun}} + S^s{}_{\text{env}})}{I_l - I_s} \tag{3}$$

Smooth propagation. The pairs of light/shadow pixels provide us with the values of $S^n_{\text{env}} = S_{\text{env}} + x_{sl}$ along the shadow boundaries. We propagate this information to all pixels by solving for the S^n_{env} image that minimizes

$$\underset{S^n_{\text{env}}}{\operatorname{argmin}} \sum_{\partial S} ||S_{\text{env}} + x_{sl} - S^n_{\text{env}}||^2 + \sum_{\mathcal{P}} ||\nabla S_{\text{env}} - \nabla S^n_{\text{env}}||^2 + w \sum_{\mathcal{P}} ||S_{\text{env}} - S^n_{\text{env}}||^2 \tag{4}$$

where ∂S is the set of constrained pixels along the shadow boundaries and $\mathcal{P}$ is the set of all image pixels. The first term encourages the constraint satisfaction, the second term preserves the variations of the original S_{env}, and the last term is a weak regularization that encourages the solution to remain close to S_{env} away from the shadow boundaries, using a small weight $w = 0.01$. This optimization can be solved using any standard least squares solver (we use the `backslash` operator in matlab).

Since x_{sl} can be negative, we can obtain negative values of S^n_{env} for a very small number of pixels. This can occur for example in regions which are poorly reconstructed as cavities, resulting in S_{env} values close to zero. We iterate by adding constraints for such points, setting $x_{sl} = 0$ such that S^n_{env} is equal to S_{env}. In all our experiments a single iteration was required to remove all negative values, which were always less than 1% of the pixels in the image.

Correcting Penumbra. The re-estimation of S_{env} described above ensures that both sides of a hard shadow boundary receive the same reflectance. However, errors also occur in the penumbra regions due to approximate continuous visibility, yielding halo artifacts in these regions (Fig. 1(mid top)). We correct these visibility values by associating each penumbra pixel to its closest pair of same reflectance light/shadow pixels as detected above. We then deduce the value of v_{sun} that makes the pixel receive the same reflectance. Fig. 1(right top) shows the final corrected reflectance. The effects are overall quite subtle, but this step does improve the result overall.

4. COMPARISON FOR TOYS SCENE

In Fig. 2 we present a comparison with other intrinsic image methods for the Toys scene. The single image methods [Chen and Koltun 2013; Barron and Malik 2013] both have residues in the reflectance. The method of [Laffont et al. 2013] has similar results with ours for this scene: ours has slightly less residue in reflectance, but does miss-classify some of the checkerboard colors as shadow. In addition, that method overestimates indirect light in corners with inaccurate reconstruction, which we attenuate with the cosine factor. It is important to recall again that the method of [Laffont et al. 2013] is not fully automatic, requiring several manual steps described in the main text.

REFERENCES

BARRON, J. T. AND MALIK, J. 2013. Intrinsic scene properties from a single RGB-D image. *CVPR*.

CHEN, Q. AND KOLTUN, V. 2013. A simple model for intrinsic image decomposition with depth cues. In *ICCV*. IEEE.

GUO, R., DAI, Q., AND HOIEM, D. 2011. Single-image shadow detection and removal using paired regions. In *CVPR, 2011*. IEEE, 2033–2040.

LAFFONT, P.-Y., BOUSSEAU, A., AND DRETTAKIS, G. 2013. Rich intrinsic image decomposition of outdoor scenes from multiple views. *IEEE Trans. on Visualization and Computer Graphics 19,* 2, 210–224.

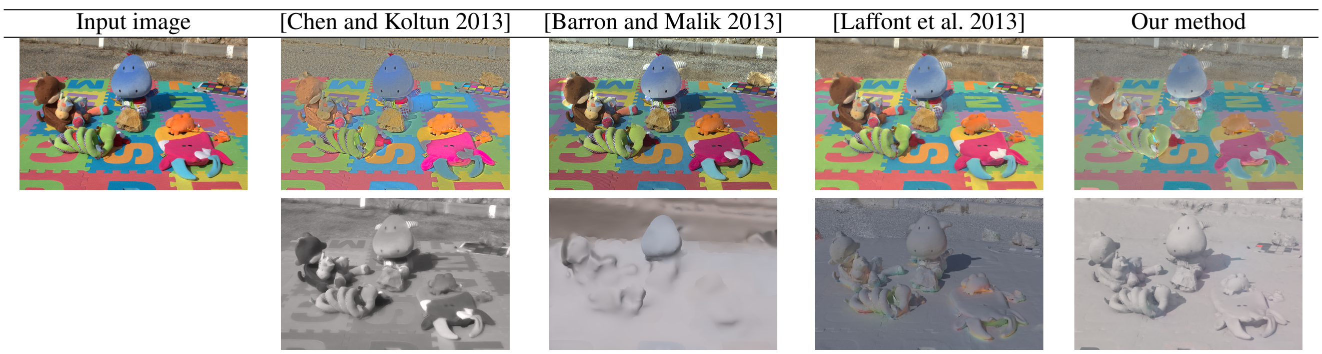

Fig. 2: Reflectance and shading respectively top and bottom row. Results are shown with scale factor and gamma-correction.