Diffuse Mirrors: 3D Reconstruction from Diffuse Indirect Illumination Using Inexpensive Time-of-Flight Sensors (Supplementary)

Felix Heide[1] Lei Xiao[1] Wolfgang Heidrich[1,3] Matthias B. Hullin[2]
1) The University of British Columbia
2) University of Bonn
3) King Abdullah University of Science and Technology

1. Additional Experiments

Here we show additional experiments that illustrate the performance of our method for non-Lambertian surfaces, and examine the impact of the sparsity inducing priors.

1.1. Non-Lambertian Surfaces

Figures 1 to 4 show a progression of results for surfaces that violate the assumption of diffuse Lambertian reflectors. In these images, we show both the color-coded depth map as well as a cross-section through the volume densities $v(\mathbf{x})$ for one scanline.

Figure 1 shows a single foam board. This material is quite diffuse, and as a result the volume cross-section shows a good localization of the depth value, revealing the board was slightly tilted in depth.

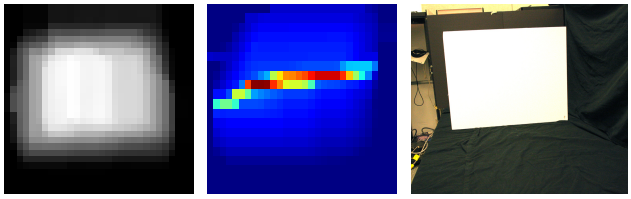

Figure 1. Reconstruction of scene image with diffuse foamboard (right). Smooth depth map, occluder-probability weighted depth along z-coordinate visualized in jet color map (left). Slice through the reconstruction volume along y-coordinate (middle).

Figure 2 shows results for two small whiteboards, which are slightly more shiny than the foam boards, but still very diffuse. We see that the sharpness of depth values suffers a bit from the glossiness of the scene, although isolating the strongest peak for each pixel still yields a good depth map.

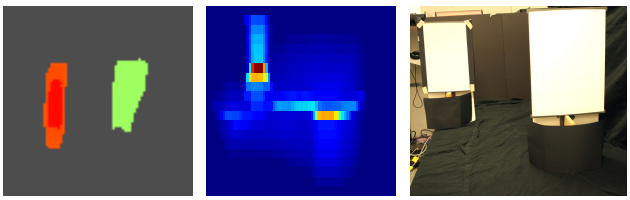

Figure 2. Reconstruction of scene image with two whiteboards, that are more reflective than the foam board (right). Color-coded depth map of strongest peak along z-coordinate visualized (left) (left). Slice through the reconstruction volume along y-coordinate (middle).

Figure 3 shows an extreme example with two mirrors. The geometry of the two surfaces can be recovered well, since the reflections from the mirrors do not interfere with each other on the diffuse wall - they create two spatially separated caustics. Our optimization procedure was adapted to handle this special case, where everything stays as before exept for the bounce in the volume which is then made spectular.

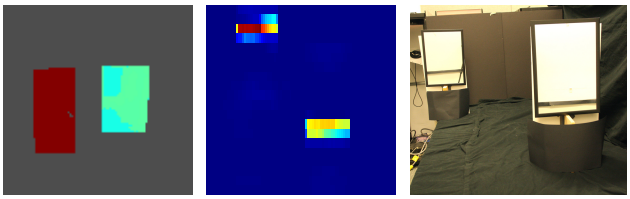

Figure 3. Reconstruction of scene image with two mirrors, that are the most reflective objects tested (right). Color-coded depth map of strongest peak along z-coordinate visualized (left). Slice through the reconstruction volume along y-coordinate (middle).

Finally, Figure 4 shows the result of mixing Lambertian and non-Lambertian surfaces. The F-shape is made out of cardboard material painted with white fairly diffuse paint, while the large rectangular surface is made out of brushed metal. The reflection from the brushed metal surface is extremely strong and glossy, and thus overlaps with the

one from the “F” in both space and time. This causes our method to fail to reconstruct the dimmer “F”-shape.

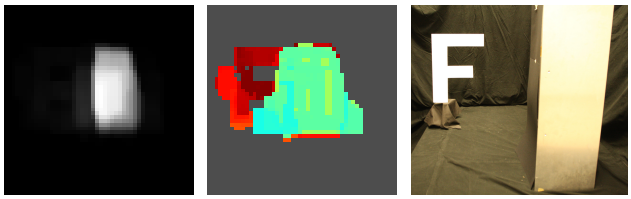

Figure 4. Reconstruction of challenging scene with a strong quite diffuse reflector made from brushed metal and a letter made from painted cardboard (right). Smooth depth map, occluder-probability weighted depth (left). Color-coded depth map of strongest peak along z-coordinate visualized (middle).

1.2. Sparsity Analysis

Figure 5 shows the effect of the sparsity inducing terms in the opimization on the mirror scene from above. We show slices through the recoverd volume without the sparsity terms (left) and with the full objective function (center). We can see that our method nicely prefers a very sparse solution and therefore finds good solutions that represent our assumption of height field geometry well.

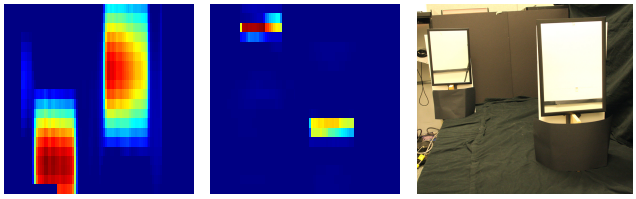

Figure 5. Example showing the effect of our sparse optimization (all the sparsity terms explained so far): Scene with two planar mirror surfaces shown on the right. Slice through the reconstruction volume computed *with simple backprojection* along y-coordinate (middle) shown on the left. Slice through the reconstruction volume computed *our whole sparsity preferring optimization* shown in the middle.

2. Optimization

This section provides a derivation of the optimization method used in our paper.

2.1. Objective Function

We had formulated the following optimization problem in the main paper

$$\mathbf{v}_{\text{opt}} = \underset{\mathbf{v}}{\text{argmin}} \ \frac{1}{2} \left\| \mathbf{CPv} - \mathbf{h} \right\|_2^2 + \Gamma(\mathbf{v}), \tag{1}$$

where the regularizer is:

$$\Gamma(\mathbf{v}) = \lambda \sum_z \left\| \nabla_{x,y} \mathbf{v}_z \right\|_1 + \theta \left\| \mathbf{Wv} \right\|_1 + \omega \sum_{x,y} \text{ind}_{\mathcal{C}}(\mathbf{v}_{x,y}) \tag{2}$$

The individual components of the objective function have been described in detail in the main paper.

Now, in order to solve the optimization problem from Equation 1, we split the regularization term into a linear operator $\mathbf{K}$ and a function $F(.)$: $\Gamma(\mathbf{v}) = F(\mathbf{Kv})$, with

$$\mathbf{K} = \left[\mathbf{D}_x^T, \mathbf{D}_y^T, \mathbf{W}\mathbb{I}^T, \mathbb{I}^T \right]^T, \tag{3}$$

where $\mathbf{D}_x, \mathbf{D}_y$ are derivative operators for the x, y dimensions for all z coordinates (stacked ontop of each other) and $\mathbb{I}$ is the identity matrix. We note that the minimum of $\Gamma(\mathbf{v})$ is obtained by independently minimizing F for each component of $\mathbf{Kv}$.

The particular definition of $\mathbf{K}$ allows us now in the next section to map our optimization problem to one that can be efficiently solved with the alternate direction method of multipliers method (ADMM) with a modified linearization step.

2.2. Solving the Optimization Problem

We first rewrite the Eq. (1) as a constrained optimizeation problem by introducing the variable $\mathbf{j}$:

$$\begin{aligned} \mathbf{v}_{\text{opt}} = &\underset{\mathbf{v}}{\text{argmin}} \ G(\mathbf{v}) + F(\mathbf{j}) \\ &\text{subject to} \quad \mathbf{Kv} = \mathbf{j} \end{aligned}. \tag{4}$$

We can then form the augmented Lagrangian

$$\begin{aligned} \mathcal{L}_\rho(\mathbf{v}, \mathbf{j}, \lambda) = &G(\mathbf{v}) + F(\mathbf{j}) + \\ &\lambda^T (\mathbf{Kv} - \mathbf{j}) + \frac{\rho}{2} \left\| \mathbf{Kv} - \mathbf{j} \right\|_2^2, \end{aligned} \tag{5}$$

where λ is a dual variable associated with the consensus constraint. ADMM now minimizes $\mathcal{L}_\rho(\mathbf{v}, \mathbf{j}, \lambda)$ w.r.t. one variable at a time while fixing the remaining variables. The dual variable is then the scaled sum of the consensus constraint error. For more details see, for example [2].

For our problem, the minimization of the augmented Lagrangian in each step leads to the following algorithm.

The individual steps of this algorithm are as follows:

Algorithm 1 ADMM algorithm

1: $\mathbf{v}^{k+1} := \underset{\mathbf{v}}{\text{argmin}} \ \mathcal{L}_\rho(\mathbf{v}, \mathbf{j}^k, \lambda^k)$ // **v-step**
2: $\mathbf{j}^{k+1} := \underset{\mathbf{j}}{\text{argmin}} \ \mathcal{L}_\rho(\mathbf{v}^{k+1}, \mathbf{j}, \lambda^k)$ // **j-step**
3: $\lambda^{k+1} := \lambda^k + \rho(\mathbf{Kv}^{k+1} - \mathbf{j}^{k+1})$ // λ-**step**

v-step The update of the volume $\mathbf{v}$ proceeds as follows:

$$\begin{aligned} \mathbf{v}^{k+1} = &\underset{\mathbf{v}}{\text{argmin}} \ \mathcal{L}_\rho(\mathbf{v}, \mathbf{j}^k, \lambda^k) \\ = &\underset{\mathbf{v}}{\text{argmin}} \ \frac{1}{2} \left\| \mathbf{CPv} - \mathbf{h} \right\|_2^2 + (\lambda^k)^T (\mathbf{Kv} - \mathbf{j}^k) + \\ &\frac{\rho}{2} \left\| \mathbf{Kv} - \mathbf{j}^k \right\|_2^2 \\ \approx &\underset{\mathbf{v}}{\text{argmin}} \ \frac{1}{2} \left\| \mathbf{CPv} - \mathbf{h} \right\|_2^2 + (\lambda^k)^T (\mathbf{Kv} - \mathbf{j}^k) + \\ &\rho (\mathbf{K}^T \mathbf{Kv}^k - \mathbf{K}^T \mathbf{j}^k)^T \mathbf{v} + \frac{\mu}{2} \left\| \mathbf{v} - \mathbf{v}^k \right\|_2^2 \\ = &\left(\mathbf{P}^T \mathbf{C}^T \mathbf{CP} + \mu \mathbb{I} \right)^{-1} \left(\mathbf{P}^T \mathbf{C}^T \mathbf{h} + \mu \mathbf{v}^k - \right. \\ &\left. \rho (\mathbf{K}^T \mathbf{Kv}^k - \mathbf{K}^T \mathbf{j}^k) + \mathbf{K}^T \lambda^k \right) \end{aligned} \tag{6}$$

Note that in the third step we have made an approximation that linearizes the quadratic term from the second line in the proximity of the previous solution $\mathbf{v}^k$. This linearization approach is known under several different names, including Linearized ADMM or inexact Uzawa method (e.g. [3, 1]). The additional parameter μ satisfies the relationship $0 < \mu \leq 1/\left(\rho \|\mathbf{K}\|_2^2\right)$.

j-step The slack variable $\mathbf{j}$ is updated as follows:

$$\begin{aligned} \mathbf{j}^{k+1} = &\underset{\mathbf{j}}{\text{argmin}} \ \mathcal{L}_\rho(\mathbf{v}^{k+1}, \mathbf{j}, \lambda^k) \\ = &\underset{\mathbf{j}}{\text{argmin}} \ F(\mathbf{j}) + \\ &(\lambda^k)^T (\mathbf{Kv}^{k+1} - \mathbf{j}) + \frac{\rho}{2} \left\| \mathbf{Kv}^{k+1} - \mathbf{j} \right\|_2^2 \\ = &\underset{\mathbf{j}}{\text{argmin}} \ F(\mathbf{j}) + \frac{\rho}{2} \left\| \left(\mathbf{Kv}^{k+1} - \frac{\lambda^k}{\rho} \right) - \mathbf{j} \right\|_2^2 \end{aligned} \tag{7}$$

Both $F(.)$ and the least square term can be minimized independently for each component in $\mathbf{j}$. Using the slack variable $\mathbf{j}$, the minimization involving the difficult function *F has now been turned into a sequence of much simpler problems in just a few variables*.

To derive the specific solutions to these problems, we note that the last line in Equation 7 can be interpreted as a proximal operator:

$$\mathbf{j}^{k+1} = \mathbf{prox}_{(1/\rho)F} \left(\mathbf{Kv}^{k+1} - \frac{\lambda^k}{\rho} \right) \tag{8}$$

using the standard definition [2]:

$$\mathbf{prox}_{\gamma F}(\nu) = \underset{\mathbf{j}}{\operatorname{argmin}} \left(F(\mathbf{j}) + \frac{1}{2\gamma} \|\mathbf{j} - \nu\|_2^2 \right). \quad (9)$$

Proximal operators are well-known in optimization and have been derived for many terms. Please see [2] for a detailed overview and applications. For our problem, we require the proximal operators for the ℓ_1 norm and for the indicator set. These are given as

$$\begin{aligned} \mathbf{prox}_{\gamma|\cdot|}(a) =& (a-\gamma)_+ - (-a-\gamma)_+ \\ \mathbf{prox}_{\gamma\, \mathrm{ind}_{\mathcal{C}}(\cdot)}(a) =& \Pi_{\mathcal{C}}(a) \end{aligned} \quad (10)$$

The first term is the well-known point-wise shrinkage and the second is the projection on the set $\mathcal{C}$.

λ-step The final step of the ADMM algorithm is to update the Lagrange multiplier by adding the (scaled) error:

$$\lambda^{k+1} := \lambda^k + \rho \left(\mathbf{K}\mathbf{v}^{k+1} - \mathbf{j}^{k+1} \right) \quad (11)$$

References

[1] E. Esser, X. Zhang, and T. Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM J. Imag. Sci*, 3(4):1015–1046, 2010. 3

[2] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, pages 1–96, 2013. 3, 4

[3] X. Zhang, M. Burger, and S. Osher. A unified primal-dual algorithm framework based on bregman iteration. *J. Sci. Comp.*, 46(1):20–46, 2011. 3