Supplemental Material
Whitened Expectation Propagation: Non-Lambertian Shape from Shading and Shadow

Brian Potetz
Google Inc.
potetz@google.com

Mohammadreza Hajiarbabi
University of Kansas
mhajiarb@ittc.ku.edu

1. Additional Results

In this section, we show the output on the algorithm on additional images. In figure 1, we show results for the canonical Stanford Bunny under various lighting arrangements and surface reflectances. In figure 2, we show surface estimates for renderings of elevation data of the Puget Sound obtained from the United States Geological Survey (USGS). In this figure, the bottom two rows show the Puget Sound under oblique lighting both with cast shadows (bottom row) and without (next-to-bottom row). On one hand, cast shadows contribute to the ambiguity of SfS because the surface normal within shadowed regions may have any orientation. On the other hand, cast shadows also provide valuable depth cues. Here, when cast shadows are simulated by the renderer, the algorithm is able to infer the height of mountain peaks with much greater accuracy. By enabling inference within large cliques, whitened EP permits depth inference in scenes with cast shadows, which allows SfS to be applied to more general and more realistic scenes, and also allows SfS to benefit from an additional depth cue.

2. Quantitative Surface Error Measurements

For all estimated 3D surfaces, we quantify image error by rendering the estimated surface using the same illumination arrangement and surface BRDF as the input image, and then computing the mean squared error between the input image and the rerendered surface estimate. We will refer to this measure as *image error*. This form of error measurement is useful because it quantifies the how faithfully the inference algorithm was able to adhere to the constraints imposed by the image. If image error is sufficiently low, then whitened EP was successfully able to exploit all available shading cues in the image and to fully optimize the data likelihood component of the probabilistic model. Further improvements to the surface reconstruction would require making improvements to the MRF model, such as the use of more accurate spatial priors $P(Z)$.

Other error measurements are possible, and many systems of quantifying error have been proposed for SfS [1, 5, 6, 10, 4]. Because the paper prioritizes the evaluation of the whitened EP algorithm, and not the evaluation of the MRF model, we avoided error measures that quantify the accuracy of the inferred surface itself. Instead, we report the image error, which measures the fidelity of the reconstructed surface to the constraints imposed by the image. Because the inferred 3D surfaces render to match the input image almost exactly, even for non-Lambertian surfaces and multiple lighting directions, the optimization algorithm has performed well. Satisfying this constraint does not necessarily imply that the reconstructed 3D surface should match the ground-truth 3D shape. SfS is known to be an underconstrained problem: for any input image there are large families of 3D surfaces which all satisfy the reflectance constraints exactly [3]. Selecting the most plausible shape estimate among these families requires a strong 3D spatial prior. Improving the quality of the output of SfS will therefore require improvements to the MRF model and the 3D spatial priors it employs.

In the paper, we show how whitened EP enables several improvements to the MRF model. Some improvements affect the data likelihood, such as the introduction of shadow cues, and the generalization to non-Lambertian reflectance. These improvements entail relaxing the constraints of classical SfS, and allowing depth inference to occur in more general scenarios, and to exploit additional depth cues. In the main paper, we evaluate these improvements by demonstrating that whitened EP was able to successfully infer depth from images that had these properties. In supplemental section 3, we show how performance is severely degraded when algorithms that assume Lambertian reflectance are applied to images with non-Lambertian surfaces. We hope that in the future, tthe flexibility of whitened EP will facilitate the inference of depth under substantially more general scenarios.

Whitened EP also enabled improvements to the spatial

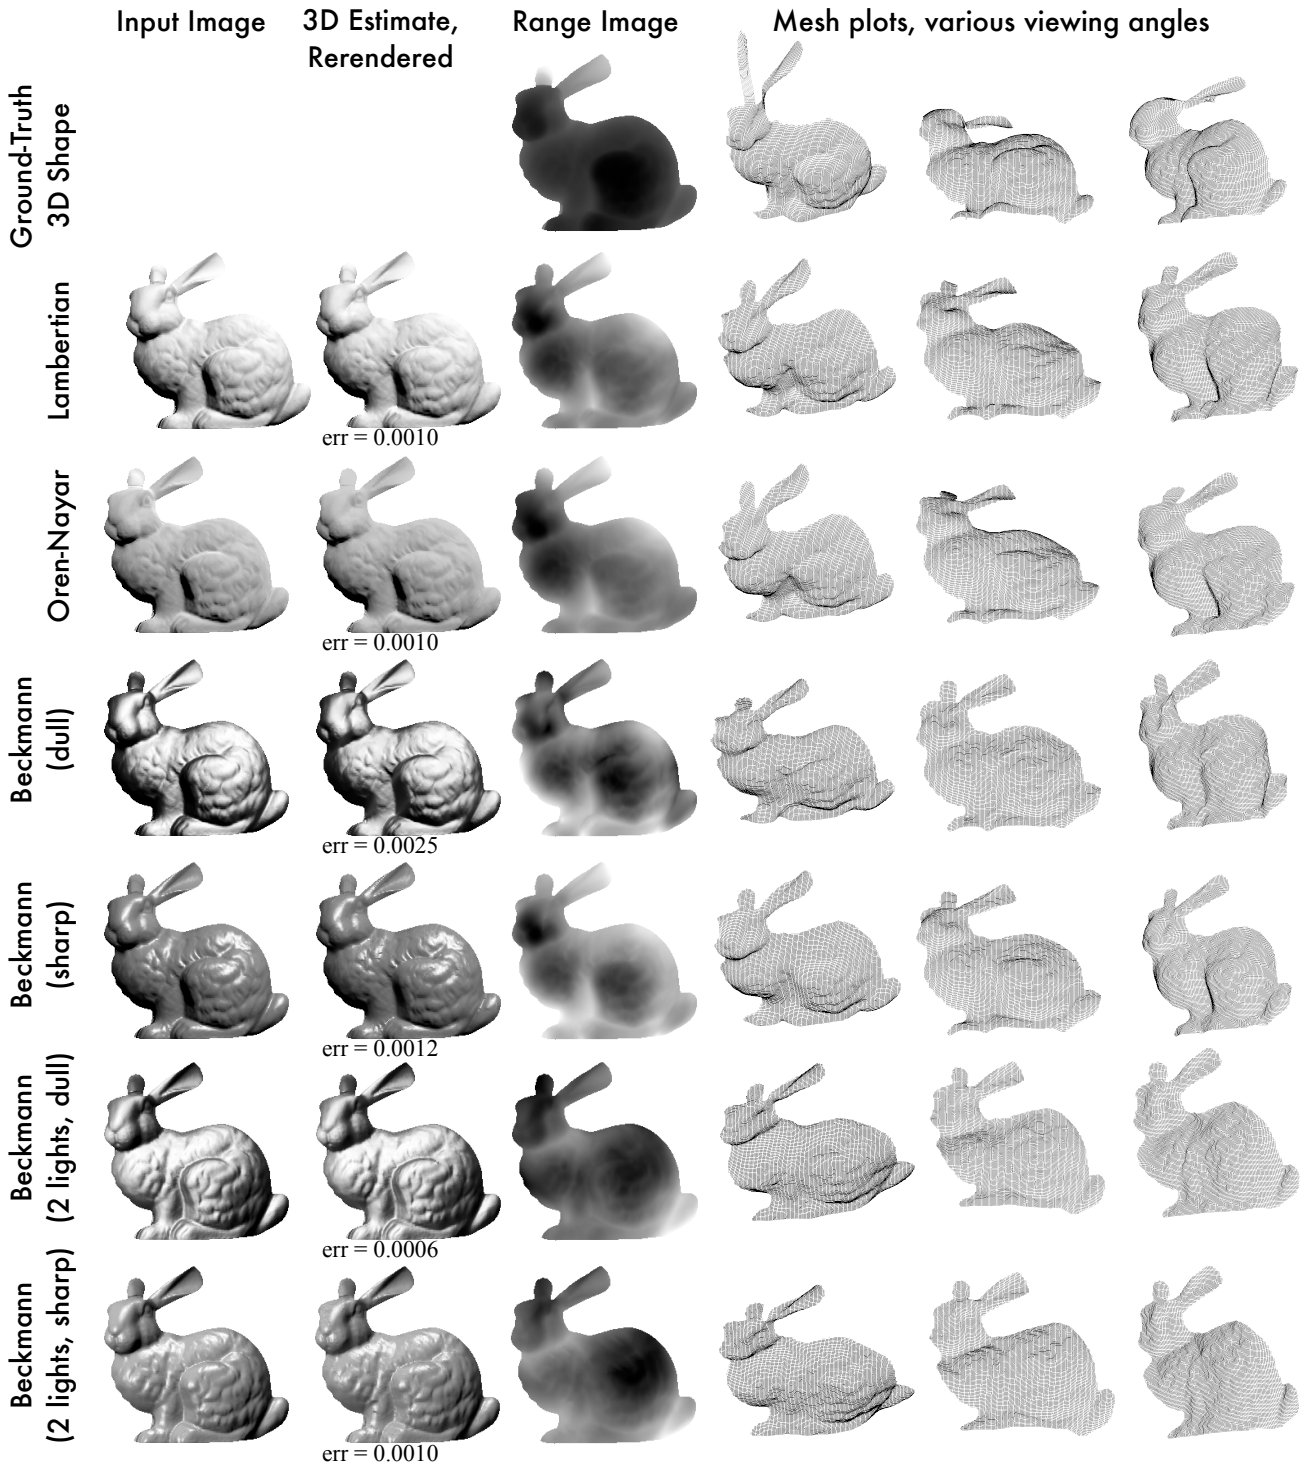

Figure 1. Results of whitened EP on the Stanford Bunny, under several reflectances and lighting conditions. The top row shows ground-truth surface plots and range image (distant points are shown as being brighter). The left column shows the rendering of the ground-truth 3D shape, and the second column shows the inferred shape rendered under equal conditions.

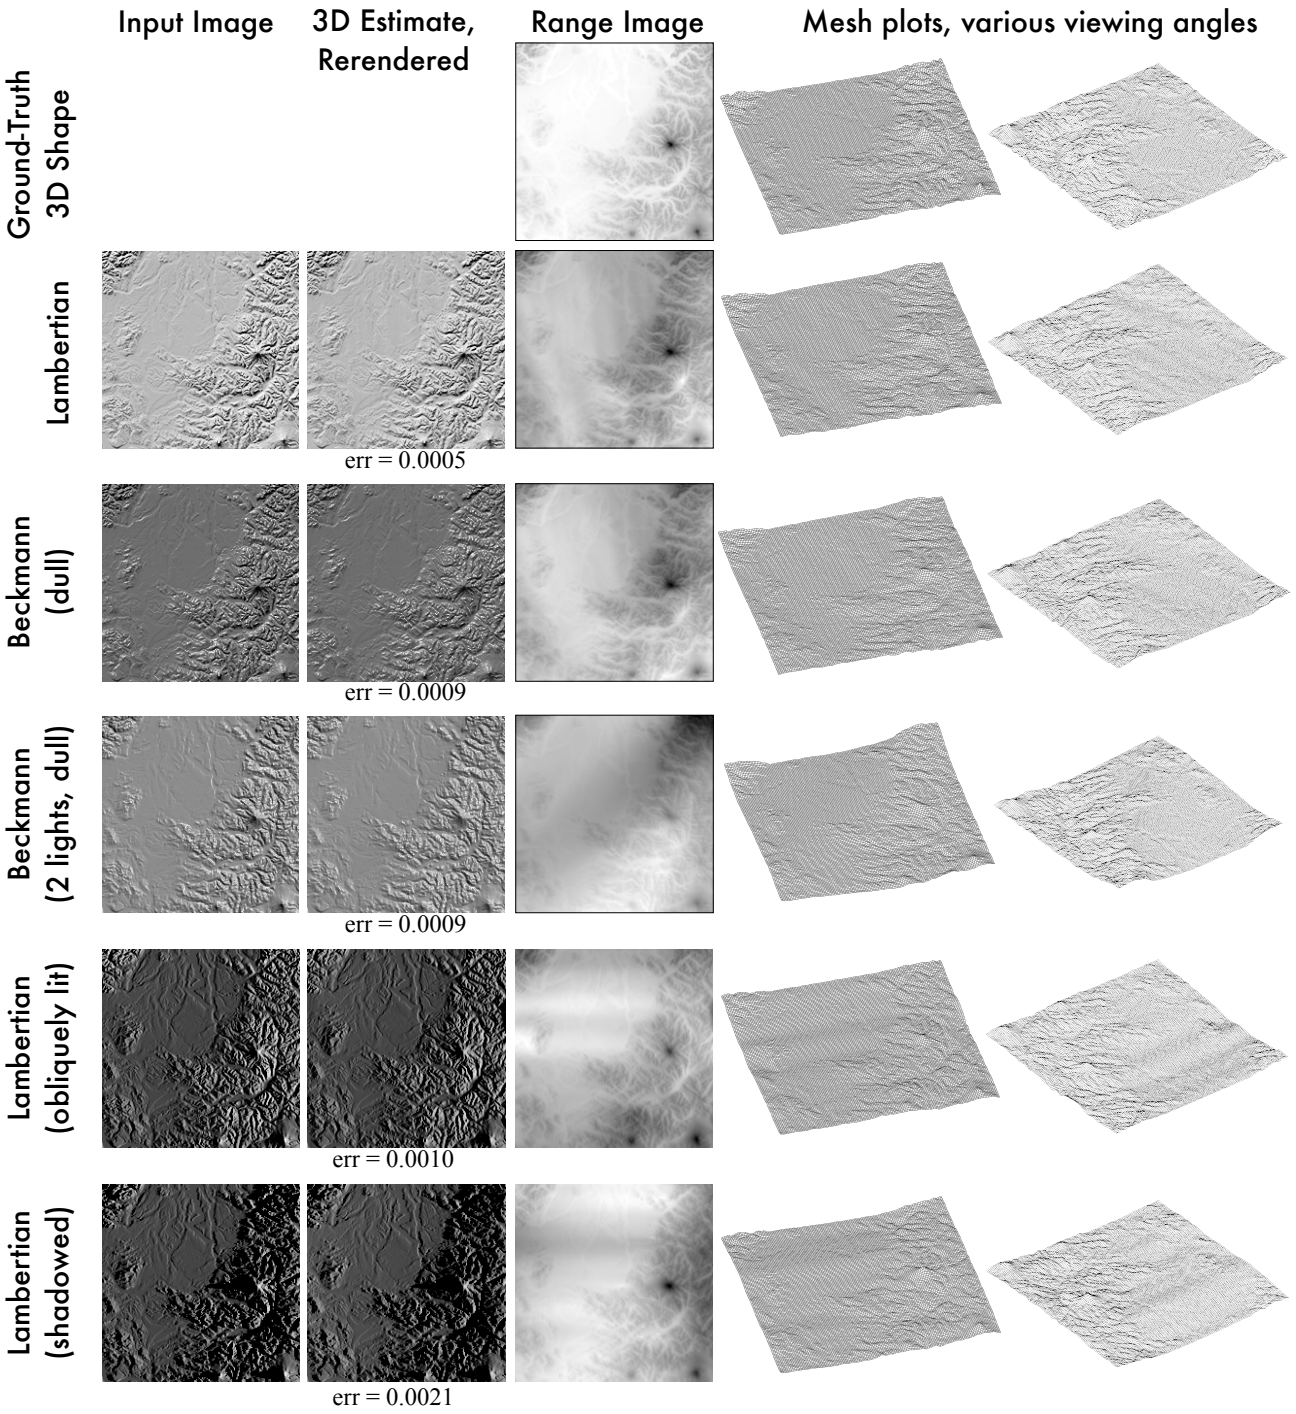

Figure 2. Results of whitened EP on the Stanford Bunny, under several reflectances and lighting conditions. The top row shows ground-truth surface plots and range image (distant points are shown as being brighter). The left column shows the rendering of the ground-truth 3D shape, and the second column shows the inferred shape rendered under equal conditions.

Table 1. For each object and reflectance, the percentage of surface normals that lie within the given angular separation of the ground-truth surface normal.

Reflectance	Object	Method	1°	2°	3°	4°	5°	10°	15°	20°	25°
Lambertian	Mozart	Worth. & Hand. [10]	2.4	5.8	8.0	10.4	13.4	25.0	33.4	40.5	46.8
Lambertian	Mozart	Haines & Wilson [5]	0.2	0.8	2.1	4.5	7.9	21.9	33.3	43.1	50.4
Lambertian	Mozart	Khan et al. [6]	0.6	2.0	4.2	7.0	10.3	26.2	41.3	54.5	65.9
Lambertian	Mozart	Whitened EP	4.2	9.9	15.1	20.0	24.8	49.1	68.1	80.3	86.4
OrenNayar	Mozart	Whitened EP	3.7	8.3	12.9	18.0	22.8	42.9	62.9	77.2	87.4
Beckmann	Mozart	Whitened EP	2.7	6.9	12.1	16.8	21.5	43.4	57.6	67.9	75.2
Doubly Lit Beck.	Mozart	Whitened EP	4.5	13.0	20.4	26.4	31.4	49.6	61.5	71.0	77.9
Lambertian	Penny	Potetz [8]	1.4	6.2	13.7	21.6	29.3	55.8	73.5	83.6	89.0
Lambertian	Penny	Whitened EP	2.8	8.1	13.9	19.6	24.9	51.1	67.7	80.6	87.7
OrenNayar	Penny	Whitened EP	3.3	9.0	15.5	21.5	27.3	53.0	71.6	83.4	90.3
Beckmann	Penny	Whitened EP	2.8	7.2	12.3	17.4	22.4	42.8	56.7	65.1	71.9
Doubly Lit Beck.	Penny	Whitened EP	1.6	4.6	8.0	11.1	14.3	29.4	44.3	56.5	65.7
Lambertian	Bunny	Whitened EP	2.7	7.3	12.1	17.2	22.6	45.8	63.5	73.4	79.6
OrenNayar	Bunny	Whitened EP	1.3	3.6	6.8	11.0	16.3	39.5	56.4	68.9	79.0
Beckmann (dull)	Bunny	Whitened EP	2.6	7.5	12.1	16.8	21.2	39.6	52.1	61.7	69.5
Doubly Lit Beck.	Bunny	Whitened EP	2.4	6.0	9.4	12.9	16.4	33.3	46.5	57.1	67.6
Beckmann (sharp)	Bunny	Whitened EP	4.2	11.3	18.6	25.0	30.7	50.6	61.8	69.2	75.6
Doubly Lit Beck.	Bunny	Whitened EP	1.5	4.5	7.2	10.0	12.8	25.8	39.2	54.2	64.9
Lambertian	Puget	Whitened EP	15.6	35.4	50.9	62.2	70.4	87.9	94.9	97.9	99.2
Beckmann	Puget	Whitened EP	18.8	38.8	52.4	61.6	67.6	81.3	88.8	93.8	96.8
Doubly Lit Beck.	Puget	Whitened EP	7.0	21.2	34.4	45.5	54.2	75.5	84.8	90.7	94.4
Oblique Lambertian	Puget	Whitened EP	9.7	21.7	34.8	46.1	55.3	77.5	87.9	93.7	97.0
Shadowed	Puget	Whitened EP	2.5	12.2	24.5	36.6	47.0	71.9	83.3	90.5	94.7

prior. These included the ability to exploit a Gaussian spatial prior of clique size D (where D is the number of pixels in the image), and the ability to infer depth directly instead of surface normals p and q. While the performance of whitened EP and image error is the primary concern of this paper, in this section, we quantify the quality of the inferred 3D surface of Whitened EP. One measure of surface accuracy is to compute the angle between the true surface normal and the inferred surface normal for each pixel within an image [5, 6]. In table 1, we provide the percentage of pixels where the surface normal lies within certain thresholds of the true normal. The surface normals estimated by whitened EP outperform several past methods [10, 5, 6], and are competitive with LBP [8].

These changes to the spatial prior demonstrate some of the novel capabilities of whitened EP. However, they are only minor improvements; substantial improvement to the classical Lambertian SfS problem will likely require the development of spatial priors that are capable of exploiting several more sophisticated statistics trends. For example:

- **Occlusion:** For images of single objects, the silhouette provides important cues to its shape. Additionally, interior lines within the image can provide evidence of self-occlusion, cusps, or corners on the surface of the object [7, 9, 11].

- **Latent variables over local surface texture:** Regions of a scene that are rough are likely to lie near other regions that are also rough. A strong spatial prior will recognize global or regional trends within a surface and exploit these trends to improve inference. These latent variables may correlate with external factors that participated in the creation of the surface, such as manufacturing processes for man-made objects. Cues that stem from non-local similarity within a scene have been applied very successfully towards image denoising [2].

- **The use of symmetry and other gestalt cues:** Humans are able to better estimate the 3D shape from images such as the Mozart or Penny image by recognizing and exploiting the symmetry present in those objects. Strong 3D spatial priors will also incorporate these cues.

We hope that whitened EP will facilitate new research in these directions by providing new ways to efficiently perform inference over MRFs with large cliques.

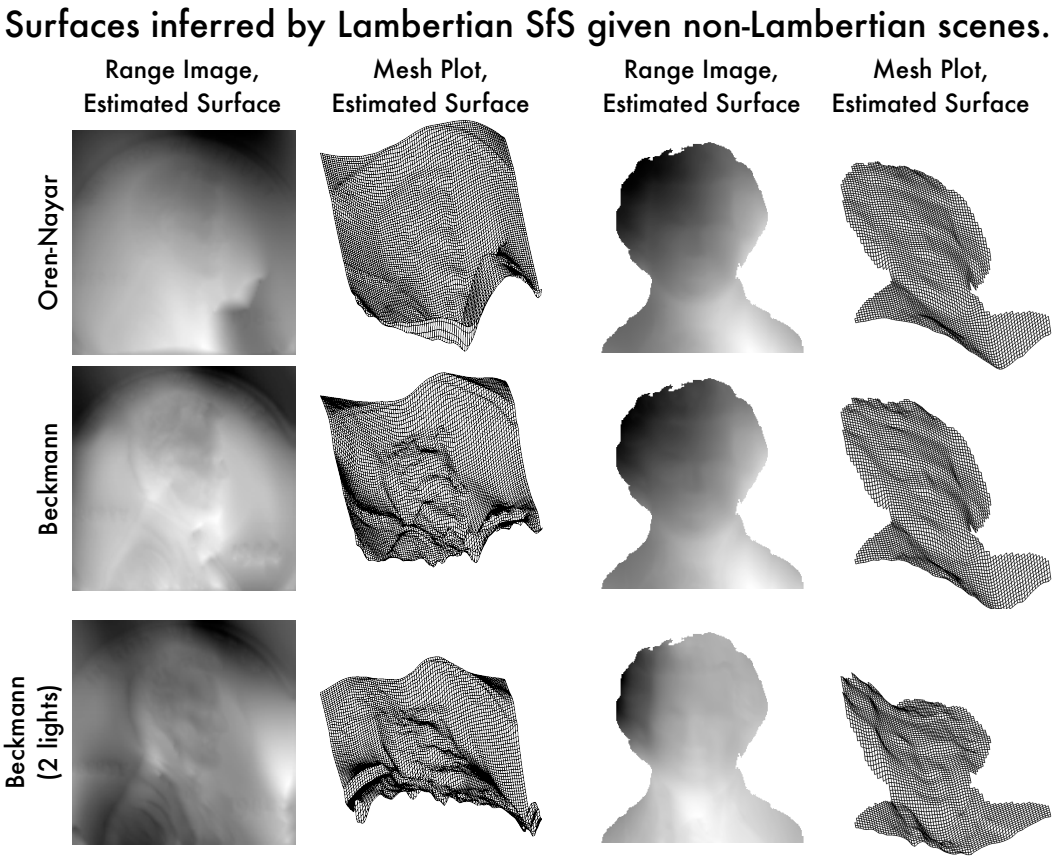

Figure 3. Inferred 3D surfaces when the algorithm assumes Lambertian reflectance despite the surface having non-Lambertian reflectance. In general, the reconstructed surface is poor when the assumptions of the algorithm are violated.

3. Effectiveness of the Lambertian Assumption for Non-Lambertian Surfaces

It is reasonable to ask whether non-Lambertian inference is necessary. In particular, it may not be immediately evident as to how adversely performance would be affected if a Lambertian SfS algorithm were naively applied to the image of a non-Lambertian surface. In figure 3 we show the results of the algorithm when Lambertian reflectance is assumed in spite of non-Lambertian reflectance of the object. The reconstructed 3D shapes are typically poor in comparison to the results when the correct reflectance is used by the algorithm.

References

[1] J. T. Barron and J. Malik. High-frequency shape and albedo from shading using natural image statistics. In *CVPR*, 2011. 1

[2] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, pages 60–65, Washington, DC, USA, 2005. IEEE Computer Society. 4

[3] A. Ecker and A. D. Jepson. Polynomial shape from shading. In *CVPR*, pages 145–152, June 2010. 1

[4] D. A. Forsyth. Variable-source shading analysis. *Int. J. of Computer Vision*, 91(3):280–302, Feb. 2011. 1

[5] T. S. F. Haines and R. C. Wilson. Belief propagation with directional statistics for solving the shape-from-shading problem. In *ECCV*, pages 780–791, 2008. 1, 4

[6] N. Khan, L. Tram, and M. Tappen. Training many-parameter shape-from-shading models using a surface database. In *3DIM Workshop at ICCV*, pages 1433–1440, 2009. 1, 4

[7] H.-S. Ng, T.-P. Wu, and C.-K. Tang. Surface-from-gradients with incomplete data for single view modeling. *ICCV*, 0:1–8, 2007. 4

[8] B. Potetz. Efficient belief propagation for vision using linear constraint nodes. In *CVPR*. 2007. 4

[9] M. Prasad, A. Zisserman, and A. Fitzgibbon. Single view reconstruction of curved surfaces. In *CVPR*, volume 2, pages 1345–1354, Los Alamitos, CA, USA, 2006. IEEE Computer Society. 4

[10] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(12):1250–1267, 1999. 1, 4

[11] L. Zhang, G. Dugas-Phocion, J.-S. Samson, and S. M. Seitz. Single view modeling of free-form scenes. In *CVPR*, volume 1, pages 990–997, 2001. 4