

Normal Integration via Inverse Plane Fitting with Minimum Point-to-Plane Distance Supplementary Material

Xu Cao¹ Boxin Shi^{2,3} Fumio Okura¹ Yasuyuki Matsushita¹
¹Osaka University ²Peking University ³Peng Cheng Laboratory

This supplementary material details the analytically computed height maps and normal maps for reproducibility (Sec. 1), and displays additional comparison results from orthographic and perspective normal integration methods (Sec. 2).

1. Experimental Settings

This section details analytically computed height maps and their corresponding normal maps under orthographic and perspective projection, for the sake of reproducibility.

1.1. Orthographic Surfaces

Figure 1 shows the coordinate system. We sample a $H \times W$ height map $z(u, v)$, defined over the uv -coordinates on the image plane, from a 2D function $z(x, y)$ in the object coordinates. We then perform normal integration in the camera coordinates. The 2D function $z(x, y)$ is defined over a continuous domain $[x_-, x_+] \times [y_-, y_+]$ in the object coordinates, where x_-, x_+, y_-, y_+ is the lower and upper limit along x and y axis, respectively. We sample the height values $z(u, v)$ over a regular grid $\{x_i \mid x_i = x_- + (i-1)w, i = 1, \dots, W\} \times \{y_i \mid y_i = y_- + (i-1)h, i = 1, \dots, H\}$, where

$$w = \frac{x_+ - x_-}{H - 1}, \quad h = \frac{y_+ - y_-}{W - 1} \quad (1)$$

are the step-sizes along x and y direction, respectively.

Once the analytic form of $z(x, y)$ is known, we can analytically compute its partial derivatives $z_x(x, y)$ and $z_y(x, y)$, and sample $z_x(u, v)$ and $z_y(u, v)$ on the same uv -coordinates. Let $\mathbf{n}^o(u, v)$ be a normal vector in the object coordinates, indicated by the superscript o . We can compute $\mathbf{n}^o(u, v)$ from sampled partial derivative as

$$\mathbf{n}^o(u, v) = \frac{[-z_x(u, v), -z_y(u, v), 1]^\top}{\|[-z_x(u, v), -z_y(u, v), 1]\|_2} \equiv [n_x^o, n_y^o, n_z^o]^\top. \quad (2)$$

As we perform normal integration in the camera coordinates, we transfer the normal map in the object coordinates

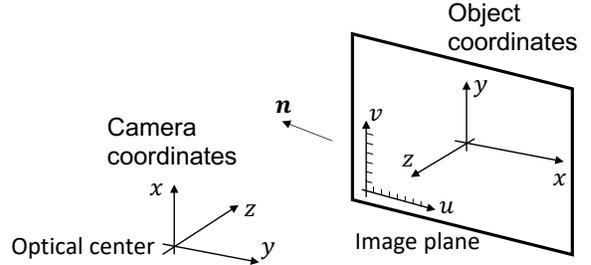


Figure 1. We compute analytic height or normal maps in the object coordinates and perform normal integration in the camera coordinate.

to the camera coordinate as

$$\mathbf{n}^c(u, v) = [n_y^o, n_x^o, -n_z^o]^\top, \quad (3)$$

where the superscript c indicates the camera coordinates.

The followings are the 2D functions $z(x, y)$ we used in our experiments.

SPHERE

$$z(x, y) = \sqrt{1 - (x^2 + y^2)} \quad \text{for } x^2 + y^2 < 1 \quad (4)$$

VASE

$$z(x, y) = \sqrt{\hat{y}^2 - x^2}, \quad \text{where} \\ \hat{y} = -138.4\tilde{y}^6 + 92.16\tilde{y}^5 + 84.48\tilde{y}^4 - 48.64\tilde{y}^3 \\ - 17.60\tilde{y}^2 + 6.4\tilde{y} + 3.2, \quad \text{and} \quad (5) \\ \tilde{y} = y/12.8 \\ \text{for } (x, y) \in [-6.4, 6.4] \times [-6.4, 6.4]$$

ANISOTROPIC GAUSSIAN The anisotropic Gaussian surface used in [2] is given by the sum of anisotropic Gaussian probability density functions as

$$z(x, y) = \sum_{k=1}^n a_k \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \mathbf{p}_k \right)^\top \Lambda_k^{-1} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \mathbf{p}_k \right)\right).$$

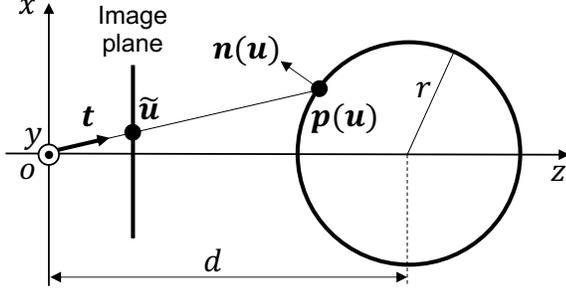


Figure 2. The geometry of computing the height map and the normal map of a sphere captured by a perspective pinhole camera.

Following [2], we sample a 150×150 height map over the domain $(x, y) \in [-1, 10] \times [-1, 10]$, with $n = 5$ peaks, and

$$\begin{aligned}
 a_1 &= 2.5, & \mathbf{p}_1 &= \begin{bmatrix} 1 \\ 2 \end{bmatrix}, & \Lambda_1 &= \begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix}, \\
 a_2 &= 3, & \mathbf{p}_2 &= \begin{bmatrix} 7 \\ 4 \end{bmatrix}, & \Lambda_2 &= \begin{bmatrix} 2 & -1 \\ -1 & 4 \end{bmatrix}, \\
 a_3 &= -5, & \mathbf{p}_3 &= \begin{bmatrix} 5 \\ 5 \end{bmatrix}, & \Lambda_3 &= \begin{bmatrix} 2 & 1 \\ 1 & 5 \end{bmatrix}, \\
 a_4 &= -2, & \mathbf{p}_4 &= \begin{bmatrix} 2 \\ 8 \end{bmatrix}, & \Lambda_4 &= \begin{bmatrix} 5 & 1 \\ 1 & 3 \end{bmatrix}, \\
 a_5 &= 5, & \mathbf{p}_5 &= \begin{bmatrix} 6 \\ 8 \end{bmatrix}, & \Lambda_5 &= \begin{bmatrix} 4 & -1 \\ -1 & 1 \end{bmatrix}.
 \end{aligned}$$

1.2. Perspective Sphere

In main paper's Fig. 7, we applied normal integration methods on the normal map of a sphere captured by a perspective pinhole camera. To this end, a sphere with radius r at $[0, 0, d]$ is first centered in the camera coordinate, as shown in Fig. 2. A point $\mathbf{p} = [x, y, z]^\top$ on the sphere surface should thus satisfy

$$x^2 + y^2 + (z - d)^2 = r^2. \quad (6)$$

The intersections of camera rays and the sphere surface is then computed. Let $\mathbf{K} = \begin{bmatrix} f & 0 & o_x \\ 0 & f & o_y \\ 0 & 0 & 1 \end{bmatrix}$ be the intrinsic matrix of a perspective pinhole camera, where f is the focal length and $[o_x, o_y]^\top$ is the principle point in the pixel coordinates. A camera ray $\mathbf{t} \in \mathbb{R}^3$ passing through a pixel $\mathbf{u} = [u, v]^\top$ is $\mathbf{t} = \mathbf{K}^{-1}\tilde{\mathbf{u}} \equiv [t_x, t_y, 1]^\top$, where $\tilde{\mathbf{u}} = [u, v, 1]^\top$ is \mathbf{u} in homogeneous coordinates. The third component of \mathbf{t} is 1, and the point $\mathbf{p}(\mathbf{u})$ on this ray is

$$\mathbf{p}(\mathbf{u}) = z(\mathbf{u})\mathbf{t} = [z(\mathbf{u})t_x, z(\mathbf{u})t_y, z(\mathbf{u})]^\top. \quad (7)$$

Plugging Eq. (7) into Eq. (6) yields a quadratic equation

$$(z(\mathbf{u})t_x)^2 + (z(\mathbf{u})t_y)^2 + (z(\mathbf{u}) - d)^2 = r^2. \quad (8)$$

Two solutions exist, and we take the smaller one as the height value $z(\mathbf{u})$, since the intersection point is the one close to the image plane. The corresponding normal vector $\mathbf{n}(\mathbf{u})$ is given by

$$\mathbf{n}(\mathbf{u}) = \frac{[z(\mathbf{u})t_x, z(\mathbf{u})t_y, z(\mathbf{u}) - d]^\top}{\|[z(\mathbf{u})t_x, z(\mathbf{u})t_y, z(\mathbf{u}) - d]\|_2}. \quad (9)$$

In our experiment, we sample the height values and normal vectors on a 128×128 image with $f = 600$, $d = 10$, and $r = 1$. The principle point is located at the center of the image, *i.e.*, $o_x = o_y = 63.5$.

1.3. Noise and Outliers

To obtain a normal map with noise, we add Gaussian noise with standard deviation σ to the gradient field $[z_x, z_y]^\top$ as

$$z_x^{(\text{noise})} = z_x + \mathcal{N}(0, \sigma), \quad z_y^{(\text{noise})} = z_y + \mathcal{N}(0, \sigma).$$

To simulate outliers, we scale the gradient field by 5 as

$$z_x^{(\text{outlier})} = 5z_x, \quad z_y^{(\text{outlier})} = 5z_y.$$

We then recover the normal map with noise or/and outliers according to Eq. (2).

2. Additional Comparison

This section shows more results in addition to the ones displayed in the main paper.

Analytic Orthographic Surfaces Figures 3 and 4 display more comparison results between orthographic normal integration methods discussed in main paper's Sec. 5. For each normal map, we used its three variants: the noise-free one, the one with Gaussian noise $\mathcal{N}(0, 0.1)$, and the one with 1% outliers. The checkerboard artifact exists in the surface estimated by discrete functional [2], and discrete functional performs the worst in all cases. Discrete Poisson's equation [4] and DGP [6] show similar error patterns. Both discrete Poisson's equation and DGP are sensitive to natural boundary or outliers. Our five-point plane fitting method estimates the height map on the same domain Ω_n as discrete functional or discrete Poisson's equation; our four-point plane fitting method is on the same domain Ω_z as DGP. Both our methods perform robustly to natural boundary or outliers. By comparing the evaluation scores between discrete Poisson's equation and DGP as well as our five-point and four-point version, we can realize that estimating the height map on Ω_z is more robust to noise than that on Ω_n . But for outliers, our five-point version is more robust than the four-point version.

Figure 5 compares the robustness of discrete Poisson’s equation [4], DGP [6], discrete functional [2], and our methods using a sphere’s normal map under orthographic projection. On the left, we gradually increase the percentages of randomly selected outliers. The RMSE of our method increases much slower than the other three methods, verifying that our method is more robust. Further, DGP has an overlapped curve as discrete Poisson’s equation (red and green), validating that there is no crucial difference between the two methods. On Fig. 5 right, we add Gaussian noise to all normal vectors and gradually increase the standard deviation. Our method performs comparably to DGP or discrete Poisson’s equation.

Real-world Perspective Surfaces Figures 6 and 7 display more results from perspective normal integration methods on real-world normal maps, in addition to HARVEST displayed in main paper’s Fig. 7. We obtained the normal maps by applying the state-of-the-art photometric stereo method [1] on DiLiGenT benchmark [5]. In terms of evaluation scores, no method achieves the best performance in most objects due to complex discontinuities in the surfaces. However, each method exhibits its feature. For discrete functional [2] and discrete Poisson’s equation [4], spikes are more likely to occur. Besides, the checkerboard artifact occurs in discrete functional [2] and Zhu and Smith’s method [7]. Our method performs stably; there is no checkerboard artifact, and spikes are not likely to occur.

References

- [1] Guanying Chen, Michael Waechter, Boxin Shi, Kwan-Yee K Wong, and Yasuyuki Matsushita. What is learned in deep uncalibrated photometric stereo? In *Proc. of European Conference on Computer Vision (ECCV)*, 2020.
- [2] Matthew Harker and Paul O’leary. Regularized reconstruction of a surface from its measured gradient field. *Journal of Mathematical Imaging and Vision*, 2015.
- [3] Yvain Quéau, Jean-Denis Durou, and Jean-François Aujol. Normal integration: a survey. *Journal of Mathematical Imaging and Vision*, 2018.
- [4] Yvain Quéau, Jean-Denis Durou, and Jean-François Aujol. Variational methods for normal integration. *Journal of Mathematical Imaging and Vision*, 2018.
- [5] Boxin Shi, Zhipeng Mo, Zhe Wu, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2019.
- [6] Wuyuan Xie, Yunbo Zhang, Charlie CL Wang, and Ronald C-K Chung. Surface-from-gradients: An approach based on discrete geometry processing. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [7] Dizhong Zhu and William A. P. Smith. Least squares surface reconstruction on arbitrary domains. In *Proc. of European Conference on Computer Vision (ECCV)*, 2020.

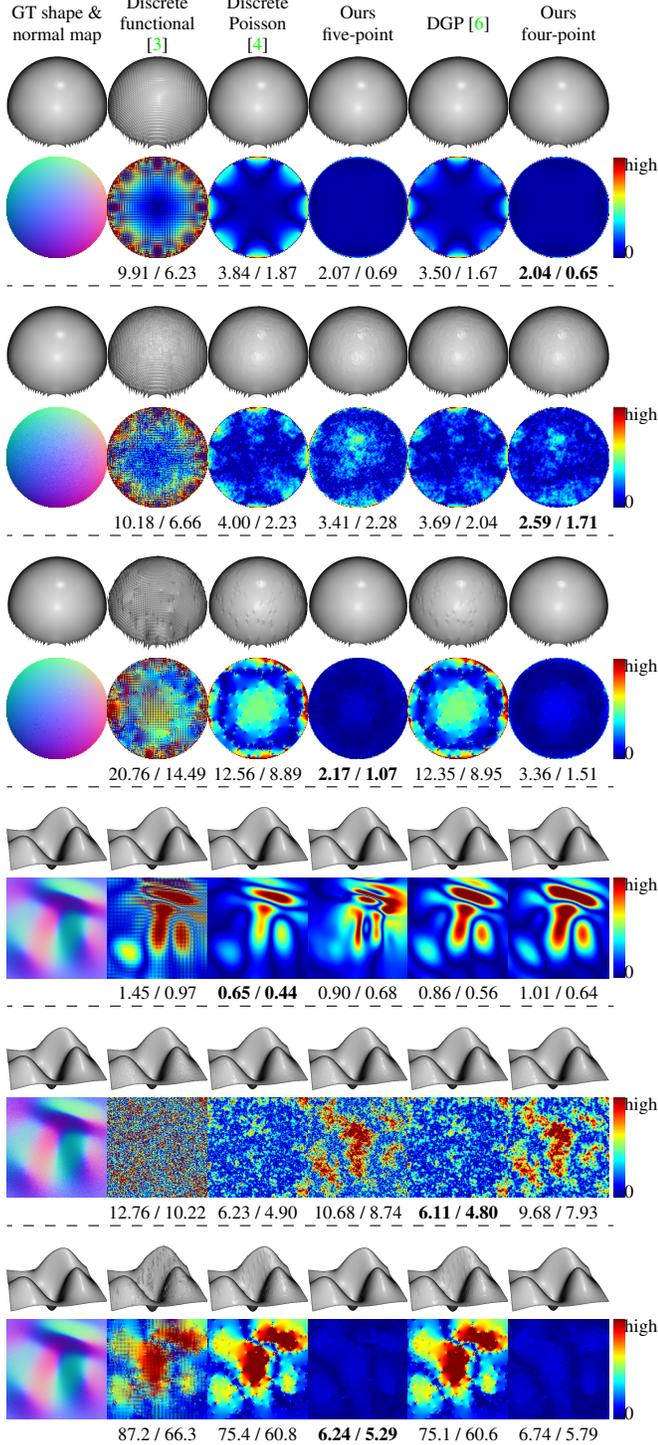


Figure 3. Comparison of orthographic normal integration methods on analytically computed SPHERE and ANISOTROPIC GAUSSIAN. The numbers underneath the absolute error maps are RMSE / MAE [$\times 10^{-3}$].

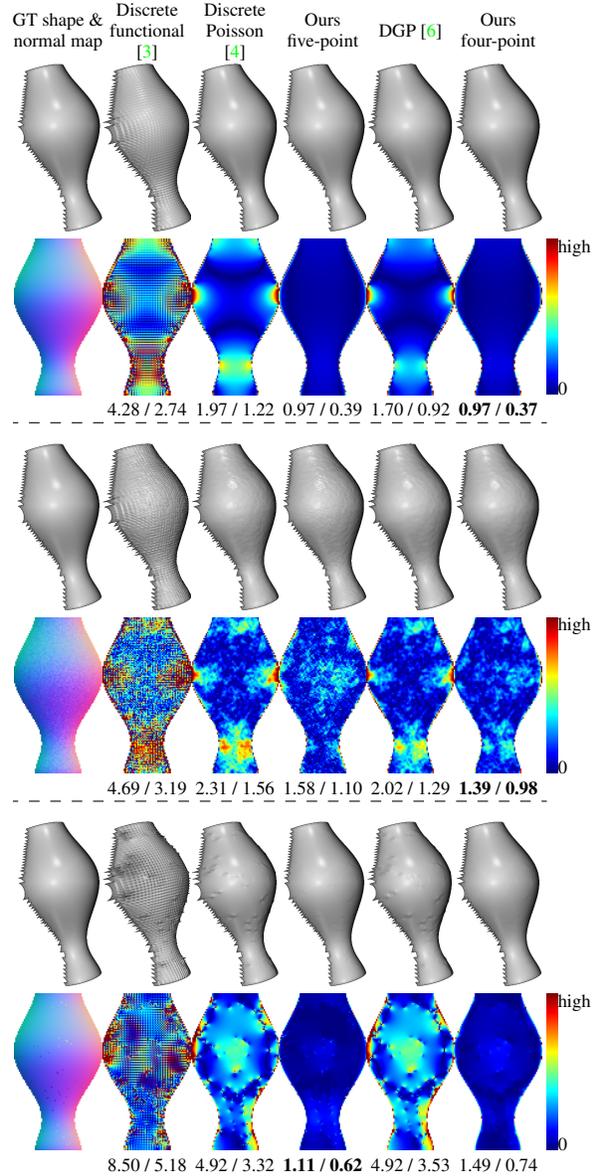


Figure 4. Comparison of orthographic normal integration methods on analytically computed VASE. The numbers underneath the absolute error maps are RMSE / MAE [$\times 10^{-2}$].

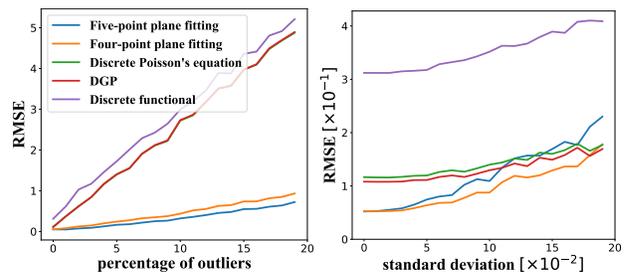


Figure 5. Comparison of robustness against (Left) outliers and (Right) Gaussian noise. Note that in the left graph the red and green curves are almost overlapped.

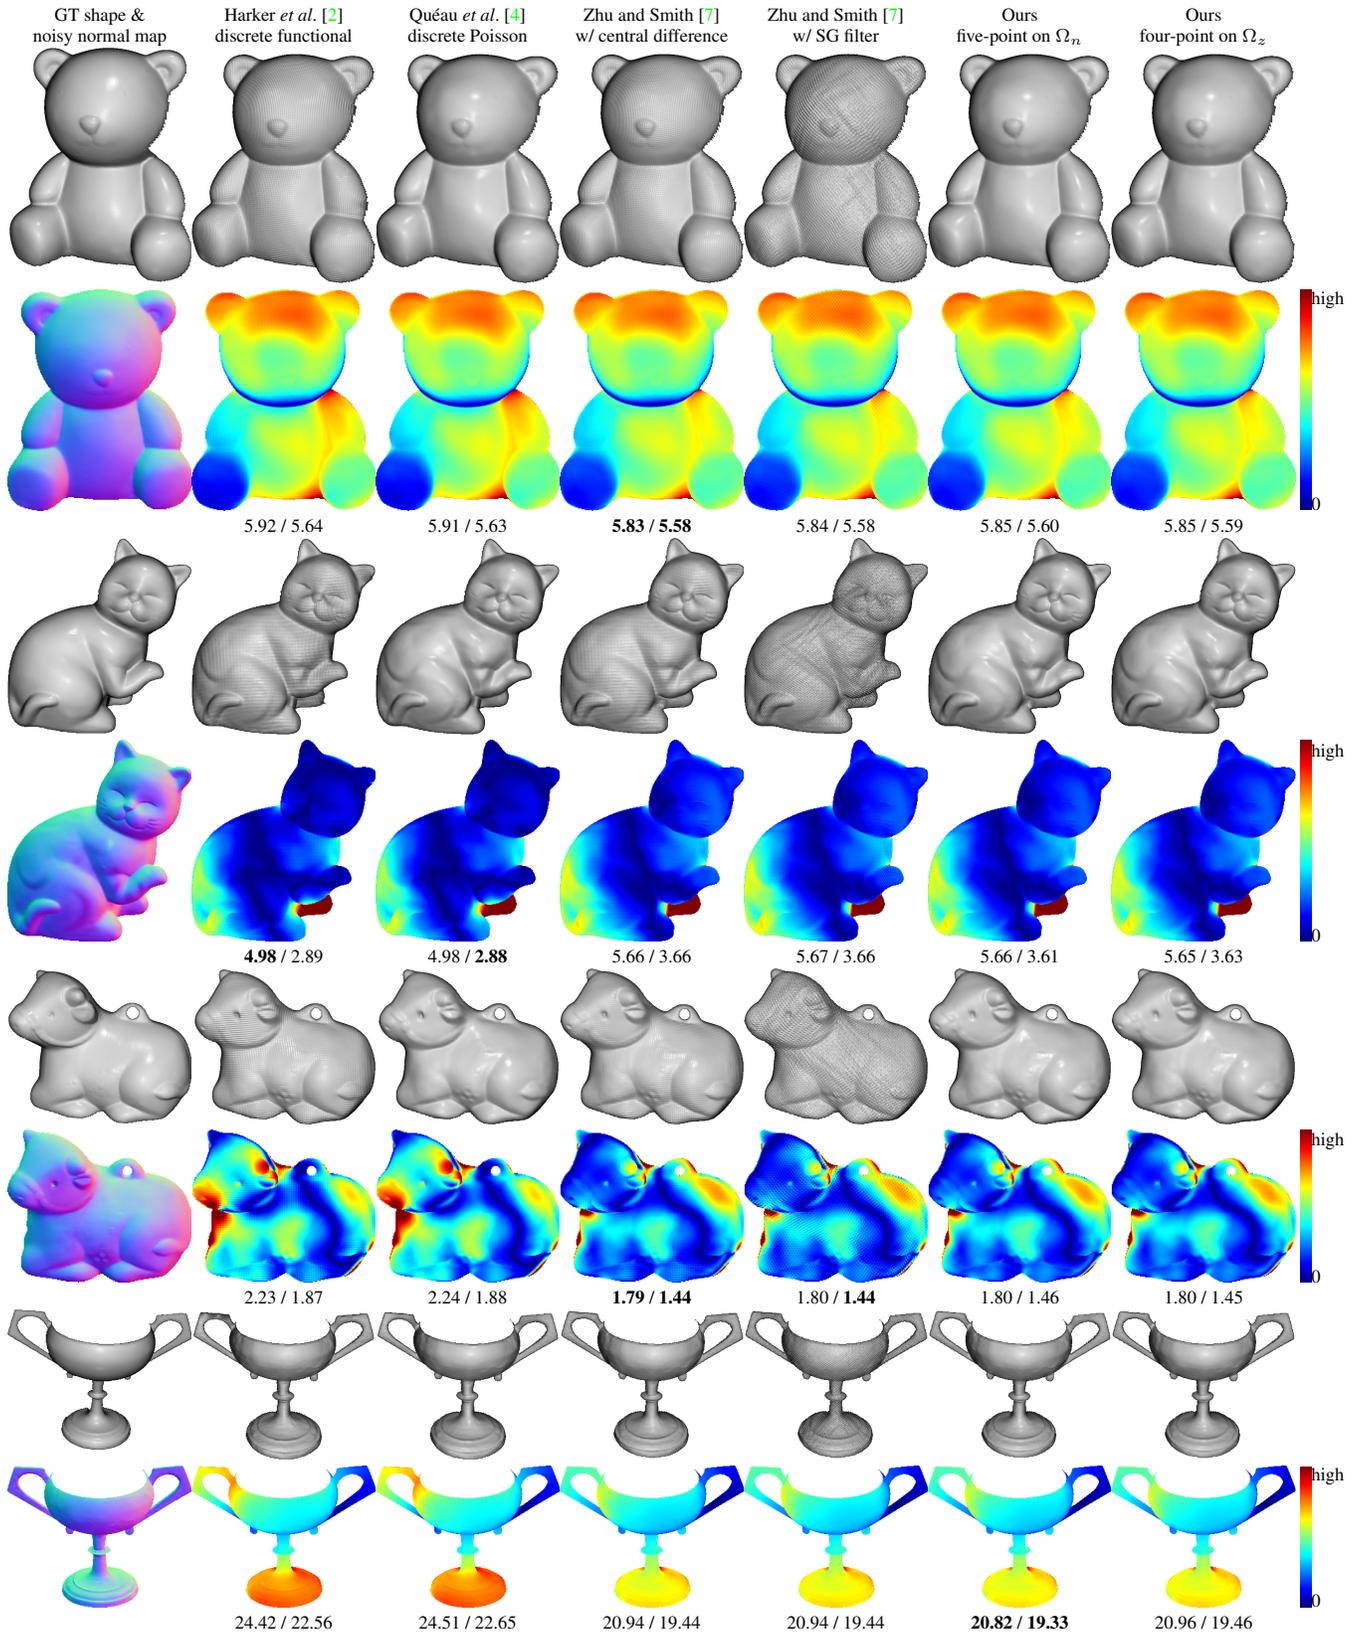


Figure 6. Comparison of perspective normal integration methods on real-world normal maps. The numbers underneath the absolute error maps are RMSE / MAE $\times 1$.

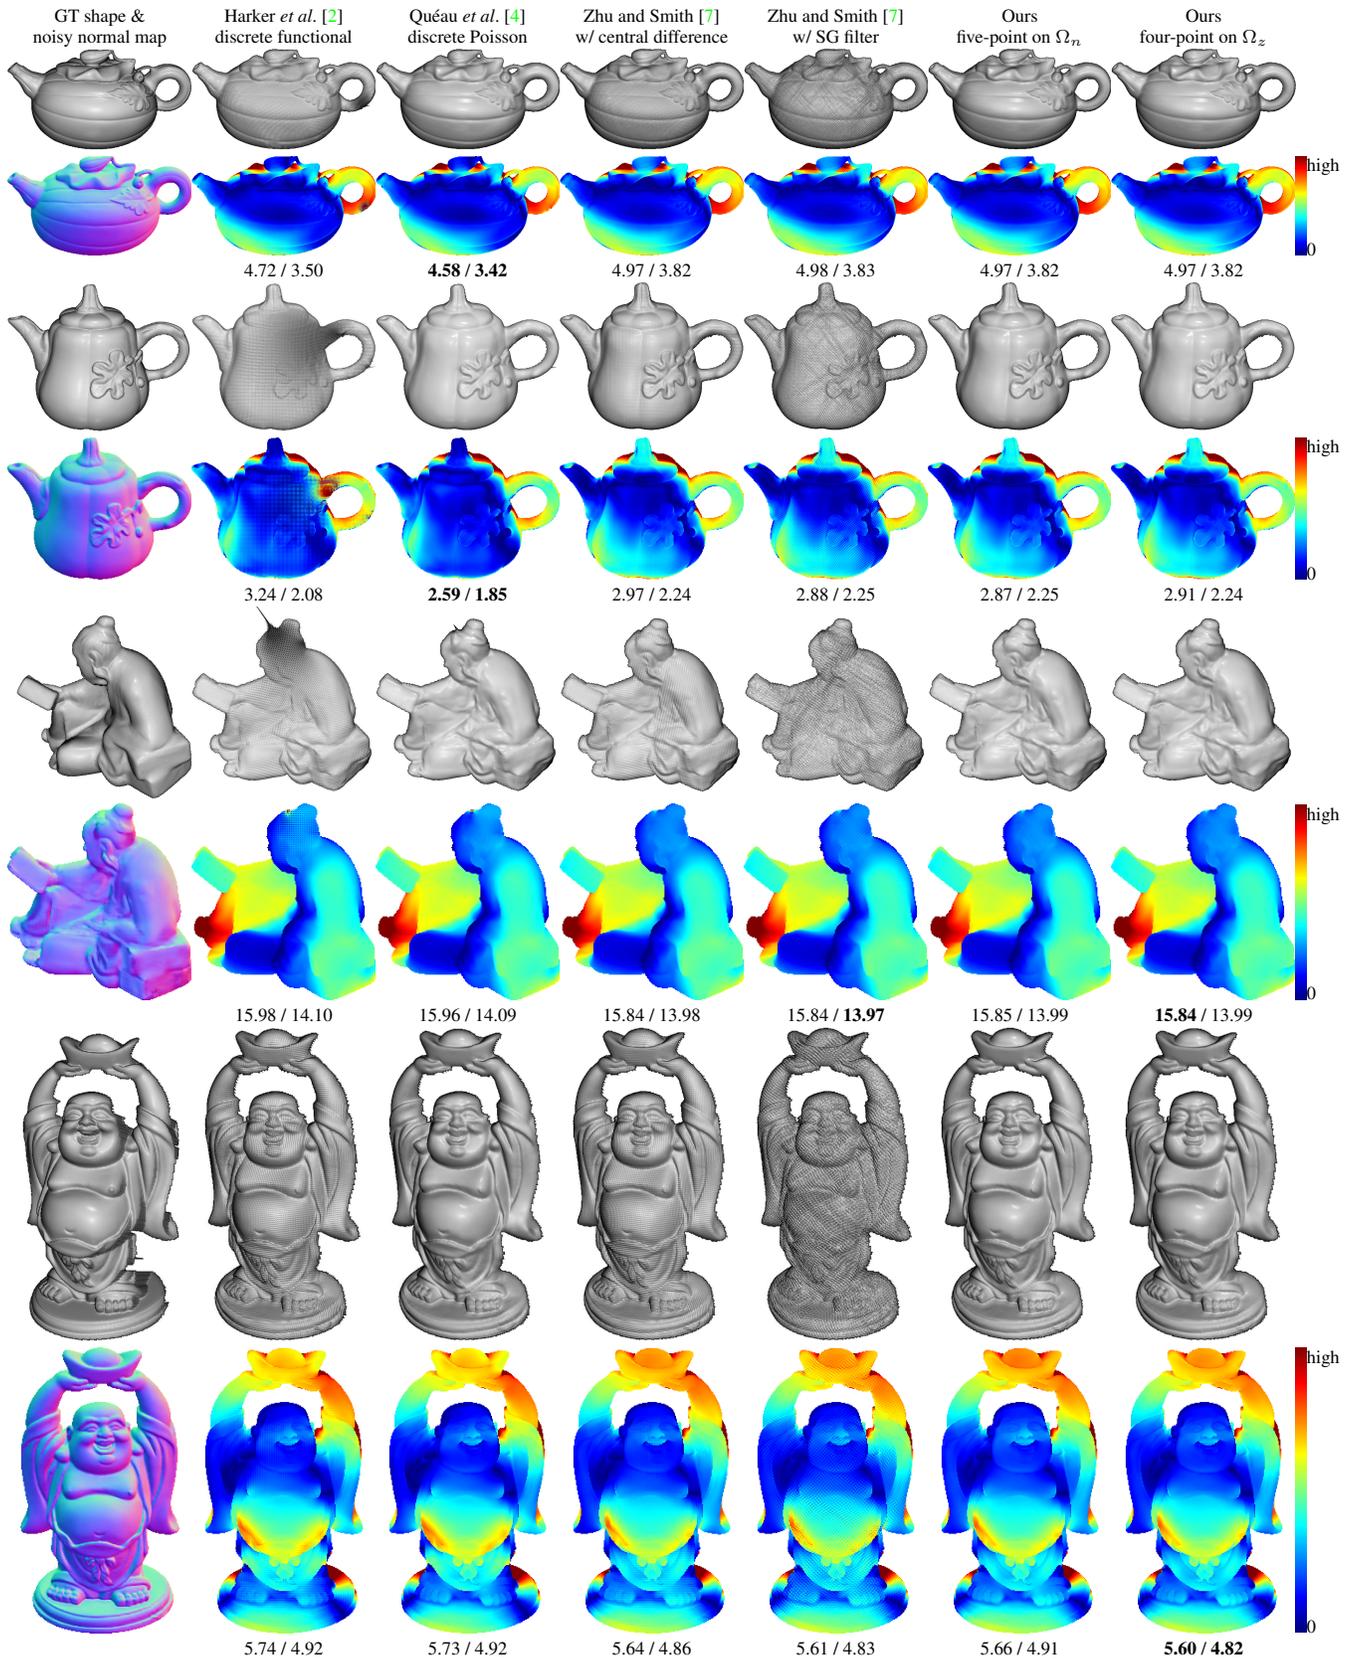


Figure 7. Comparison of perspective normal integration methods on real-world normal maps. The numbers underneath the absolute error maps are RMSE / MAE $[\times 1]$.