An Adaptive Screen-Space Meshing Approach for Normal Integration

Moritz Heep[®] and Eduard Zell[®]

University of Bonn, Bonn, Germany moritzheep.github.io/adaptive-screen-meshing

Abstract. Reconstructing surfaces from normals is a key component of photometric stereo. This work introduces an adaptive surface triangulation in the image domain and afterwards performs the normal integration on a triangle mesh. Our key insight is that surface curvature can be computed from normals. Based on the curvature, we identify flat areas and aggregate pixels into triangles. The approximation quality is controlled by a single user parameter facilitating a seamless generation of low- to high-resolution meshes. Compared to pixel grids, our triangle meshes adapt locally to surface details and allow for a sparser representation. Our new mesh-based formulation of the normal integration problem is strictly derived from discrete differential geometry and leads to wellconditioned linear systems. Results on real and synthetic data show that 10 to 100 times less vertices are required than pixels. Experiments suggest that this sparsity translates into a sublinear runtime in the number of pixels. For 64 MP normal maps, our meshing-first approach generates and integrates meshes in minutes while pixel-based approaches require hours just for the integration.

Keywords: Surface Reconstruction \cdot Normal Integration \cdot Meshing



Fig. 1: Our screen-space remeshing pipeline decimates smooth, featureless areas efficiently before the normal integration while preserving high-frequency details. Depicted results illustrate high, mid and low-resolution triangulations. Wireframes are rendered as vector graphics for closer examination.

1 Introduction

Normal integration is renowned for recovering fine and delicate surface details as a subsequent step to photometric stereo [28] or shape from shading. Increasing the resolution of the normal map improves the accuracy of fine structures. At the same time, smooth, featureless regions will reveal little additional information but increase computational costs. Real-world objects often consist of both fine structures and featureless regions and one must choose between sufficient image resolution and reasonable computational cost. In general, to resolve structures half the size, *both* the image height and width must be doubled. This quadratic growth in the number of variables has a significant impact on the runtime of the normal integration and makes regular pixel grids an increasingly inefficient geometric representation.

In practice, many 3D reconstruction pipelines operate at high resolutions and switch to a sparser triangle mesh in the last processing step [18, 29, 30]. However, this ignores the underlying problem: all steps up to the final mesh representation are still performed at full resolution, with obvious negative effects on the computational performance. In contrast to previous work [15, 17], we do not aim to speed up computations in the pixel domain. Instead, we introduce a flexible and locally adaptive triangle mesh *before* the normal integration and solve the problem at its origin, see Fig. 1.

To the best of our knowledge and despite the rich literature on normal integration over the last 30 years, we appear to be the first to propose normal integration on general triangle meshes. Our main insight is that curvature can be extracted from normal maps and that curvature is sufficient to refine triangle meshes locally and in screen space. In addition, we present a novel formulation of the normal integration for triangle meshes. It is strictly derived from discrete differential geometry and is state-of-the-art in that it avoids the checkerboard artefacts and the Gibbs phenomenon described in [7]. Finally, our meshing-first approach can be easily integrated into most existing photometric stereo pipelines and is well-suited for high-quality reconstructions where we achieve significantly sparser representations.

2 Related Work

In this section we will focus on dedicated work on normal integration and (re)meshing. We refer the reader to [1] for a starting point on photometric stereo and how to obtain the normal maps, required by our method.

Most state-of-the-art normal integration methods are so-called *variational* methods, which find a depth map by minimizing an L^2 functional containing the difference between the actual depth-map gradients and the observed gradients, e.g. from photometric stereo. Broadly, there exist two types of variational methods: Either, functional analysis is used to derive a Poisson equation which is then discretized and solved [21] or the functional itself is discretized [13]. Both cases lead to a linear system of equations. Methods using the more robust L^1





results in only 5 minutes.

norm have been proposed but are computationally more involved [11]. More recently, authors have raised concerns about checkerboard artefacts [38] and the Gibbs phenomenon [7] (Fig. 3) occurring in the discretized functional setting. Still, these problems are not innate to the variational approach: [6] and [20] independently proposed a functional using normals over gradients in the pixelbased integration setting. This functional avoids the artefacts mentioned above without resorting to larger stencils for the partial derivatives [37] or introducing additional variables [7]. Our method starts with the same functional but discretizes it for general triangle meshes. An overview of variational methods can be found in [28].

Regarding non-variational approaches, Xie *et al.* [32,34] consider the regular pixel grid as a quad mesh. They alternate between tilting these quads to align them with the normal directions and glueing adjacent quads into a continuous or even discontinuous [33] surface. Although superficially mesh-based, this approach



Fig. 3: Normal integration results with (left) and without (right) the Gibbs phenomenon.

ultimately operates on a pixel-by-pixel basis. Despite being non-variational, Xie *et al.*'s approach still suffers from the Gibbs phenomenon around sharp corners. Cao *et al.* [7] fix the Gibbs phenomenon by optimizing pixel-wise depth values to form a set of interlocking planes defined by the normals.

In the depth-from-stereo context, image pyramids [17] or iteratively refined grids of deformation points [15] have been proposed to exploit locally similar depth values in smooth regions to speed up reconstruction. While these approaches progressively refine their grid, the grid remains regular and refinement is applied equally over the entire image. In contrast, we adaptively coarsen our grid to allocate computational resources where they are most needed.

The triangulation and remeshing of 2D and 3D surfaces is an intensively studied area [22] with varying objectives targeting different applications. For rendering and computer-aided design, the overall shape should be preserved and consist of a minimal amount of vertices [16, 35, 36]. However, for numerical simulations like mechanical stress, heat transfer, etc., these very sparse meshes are usually inadequate and more regular triangulations are preferred. Here, Centroidal Voronoi Tesselations (CVT) [10] as well as Optimal Delaunay Triangulations (ODT) [8] are widely used. Both approaches generate uniform or isotropic meshes where each vertex is the (weighted) centroid of some local neighbourhood. The isotropic variants are typically controlled by an application-dependent density. Densities based on maximum absolute curvature are popular to balance geometric faithfulness and regularity of the mesh [2, 9]. An argument for this choice can be found in Dunyach et al. [12] where a relation between curvature and approximation error is derived. All these approaches have in common that they operate on an existing 3D geometry. In contrast, our approach generates a geometry-adapted triangle mesh from the normal maps alone. We extend the approach of Dunyach et al. [12] because it directly controls the approximation error. Furthermore, the resulting isotropic meshes increase the numerical stability of our mesh-based normal integration.

Besides surface reconstruction, applications of triangle meshes in 2D include image vectorization [25], although conversion to parametric curves is often preferred in this context [14, 26]. The key difference is that image vectorization is based on colours, while our approach approximates the triangulation of a 3D object in 2D screen space.

3 Definitions, Normals and Curvature

In this section, we will briefly summarize some fundamental concepts in differential geometry and the relation between surfaces, normals and curvatures. This knowledge will be required to derive our solutions for screen-space meshing in Sec. 4 and mesh-based integration in Sec. 5.

For a single view, the 3D reconstruction problem can be formulated as finding a projection

$$\phi: \Omega \to \mathbb{R}^3 \tag{1}$$



Fig. 4: Left: The projection ϕ maps points on the screen into three-dimensional space. Distances on the screen appear shorter than their three-dimensional counterparts (fore-shortening). This is captured in the Jacobian J of ϕ . Right: Illustration on the relationship between the edge lengths and the user-defined approximation error ϵ as derived in [12].

from the image foreground $\Omega \subset \mathbb{R}^2$ into 3D space. We write $\boldsymbol{u} = (u, v)^{\mathrm{t}}$ for screen coordinates and $\boldsymbol{x} = (x, y, z)^{\mathrm{t}}$ for object coordinates. The projection is usually parametrized by a depth map $h : \Omega \to \mathbb{R}$ as

$$\boldsymbol{x} := \phi(\boldsymbol{u}) = \left(u, v, h(\boldsymbol{u})\right)^{\mathrm{t}} \tag{2}$$

in the orthographic case and

$$\boldsymbol{x} := \phi(\boldsymbol{u}) = \boldsymbol{r}(\boldsymbol{u}) \cdot h(\boldsymbol{u}) \tag{3}$$

in the projective case where $\mathbf{r} := \mathbf{C}^{-1} \cdot (u, v, 1)$ is the ray given by the camera matrix \mathbf{C} . As h = z for both projections, we will use z synonymously.

3.1 Normal Integration

Given a normal map $n : \Omega \to S^2$, the normal integration problem is finding the depth map z that minimizes

$$E_{\text{Int}} = \int_{\Omega} \langle \boldsymbol{n}, \partial_u \boldsymbol{x} \rangle^2 + \langle \boldsymbol{n}, \partial_v \boldsymbol{x} \rangle^2 \, du \, dv \,. \tag{4}$$

This modified Dirichlet energy has only recently been proposed in the context of discontinuity preserving normal integration [6] and depth-from-stereo [20]. Using normals instead of depth map gradients overcomes the problem of distortions near sharp corners. This phenomenon, also known as the Gibbs phenomenon [7], Fig. 3, is present in traditional variational approaches and is discussed in more depth within the supplementary material. Relying on discrete differential geometry, we derive our novel mesh-based discretization of the modified Dirichlet energy in Sec. 5.

3.2 Curvature

Strong changes in the normals indicate geometric details. Mathematically speaking, this is described by curvature κ . Just as normals provide a planar approximation of the surface, so does curvature provide a spherical approximation. The radius of this approximating sphere is exactly $|\kappa|^{-1}$. Curvatures are calculated from the first and second fundamental forms:

$$\boldsymbol{I}_{ij} = \partial_i \boldsymbol{x} \cdot \partial_j \boldsymbol{x}, \qquad \qquad \boldsymbol{I}_{ij} = -\partial_i \boldsymbol{x} \cdot \partial_j \boldsymbol{n}. \tag{5}$$

where $i, j \in \{u, v\}$. The first fundamental form quantifies foreshortening and measures the 3D distance between two points on the 2D screen. The second fundamental form is virtually the gradient of the normal map. However, the second fundamental form itself is subject to foreshortening. Instead, we consider the two eigenvalues κ_i and eigenvectors v_i of the generalized eigenvalue problem

$$\kappa_i \cdot \boldsymbol{I} \, \boldsymbol{v}_i = \boldsymbol{I} \boldsymbol{I} \, \boldsymbol{v}_i \,. \tag{6}$$

They are called the principal curvatures and principal directions respectively. The addition of the first fundamental form in the generalized eigenvalue problem exactly compensates for foreshortening.

Given a set of surface normals, we can calculate tangents since $\mathbf{n} \cdot \partial_i \mathbf{x} = 0$. In the orthographic case, this gives

$$\partial_u \boldsymbol{x} = \boldsymbol{e}_x - \frac{\boldsymbol{n}_x}{\boldsymbol{n}_z} \cdot \boldsymbol{e}_z \qquad \qquad \partial_v \boldsymbol{x} = \boldsymbol{e}_y - \frac{\boldsymbol{n}_y}{\boldsymbol{n}_z} \cdot \boldsymbol{e}_z \tag{7}$$

where e_x, e_y, e_z are the unit vectors in the respective coordinate axis directions. As no knowledge about the depth map z is required, one can perform all these calculations using normal maps alone. This is the foundation for our screen-space triangulation.

In the projective case, we get

$$\partial_u \boldsymbol{x} = \left(\partial_u \boldsymbol{r} - \frac{\boldsymbol{n} \cdot \partial_u \boldsymbol{r}}{\boldsymbol{n} \cdot \boldsymbol{r}} \cdot \boldsymbol{r}\right) \cdot \boldsymbol{z} \qquad \partial_v \boldsymbol{x} = \left(\partial_v \boldsymbol{r} - \frac{\boldsymbol{n} \cdot \partial_v \boldsymbol{r}}{\boldsymbol{n} \cdot \boldsymbol{r}} \cdot \boldsymbol{r}\right) \cdot \boldsymbol{z} \qquad (8)$$

where the ray r is again given by the camera matrix. Unlike the orthographic case, the tangents depend on the depth map z. However, if distances within the object are small compared to the camera-to-object distance, we can adopt the weak perspective setting and set z to the average camera-to-object distance. In the supplementary material, we go into detail on how to handle objects very close to the camera.

4 Screen-Space Meshing

To create a feature-adaptive mesh before the actual surface integration, we must reliably predict detailed surface areas and adjust the density of the triangle mesh



Fig. 5: The meshing pipeline: We start by calculating curvature from the normal maps. This curvature is used to derive an optimal edge length. Then, an initial mesh is iteratively refined to move its edge lengths closer to the calculated optimum. The result is a sparse, isotropic mesh.

accordingly. While mesh simplification algorithms like [16] achieve an impressive reduction in vertex number, their sole focus on geometric faithfulness often leads to skinny triangles. As we will see in 5, small angles decrease the numerical stability of the mesh-based integration. In contrast, remeshing algorithms also take mesh quality into account and are much more suited for our use-case. That being said, existing remeshing methods operate on 3D surfaces and are not directly applicable to our case. Searching for a fast and robust method with only a few parameters, we identified the work by Dunyach *et al.* [12] as the most suitable candidate. In their algorithm, all edge lengths are optimized towards an optimal edge length

$$L = \sqrt{\frac{6\epsilon}{|\kappa|} - \epsilon^2} \tag{9}$$

that is based on curvature and a user-defined approximation error ϵ , see Fig. 4b. However, there are two obstacles: The first issue is the edge length; lengths on screen are subject to foreshortening, while the optimal edge length is not. They cannot be simply compared without compensation (Sec. 4.1). The second issue is that we require a vertex-based curvature measure to obtain the optimal lengths. In Sec. 3.2, we discussed how to compute pixel-wise curvatures from normal maps and we will present a solution for computing vertex-based curvatures in screen space in Sec. 4.2. Finally, we will discuss the optimization of vertex positions (Sec. 4.3). For all these steps, we assume that the screen is already covered by a triangle mesh. In Sec. 4.4, we show how we obtain the initial triangulation of the screen and summarize our screen-space algorithm to obtain the final triangulation. An overview of our approach is given in Fig. 5.

4.1 Measuring Lengths

Edges in screen-space are subject to foreshortening, *i.e.* appear shorter on screen than in 3D, see Fig. 4a. This foreshortening is exactly described by the first fundamental form I, Eq. (5). We have shown in Sec. 3.2, how to derive the first fundamental form from normals. To apply this to the mesh setting, we use the

pixel-wise normal map to obtain face normals

$$\boldsymbol{n}_f := \text{normalize}\left(\sum_{p \in \mathcal{P}_f} \boldsymbol{n}_p\right)$$
 (10)

where \mathcal{P}_f contains all pixels in the image that are covered by triangle f. As large triangles occur only in regions with low curvature, where normals vary slowly, the variance of n_f is small. Plugging this normal into Eqs. (7) and (8) respectively, we obtain fundamental forms for each face. With these fundamental forms, the length of an edge $(i, j) \in \mathcal{E}$ is

$$l_{ij} = \sqrt{\frac{1}{2}} \cdot \left((\boldsymbol{u}_i - \boldsymbol{u}_j)^t \cdot \boldsymbol{I}_f \cdot (\boldsymbol{u}_i - \boldsymbol{u}_j) + (\boldsymbol{u}_i - \boldsymbol{u}_j)^t \cdot \boldsymbol{I}_{f'} \cdot (\boldsymbol{u}_i - \boldsymbol{u}_j) \right)$$
(11)

where f, f' are the two faces adjacent to the edge and u_i, u_j are the *screen* coordinates of the edge's endpoints. For boundary edges, there is only one adjacent face and we omit the factor 1/2.

4.2 Calculating Curvatures

To obtain vertex curvatures, we start by calculating maximum absolute curvature $\kappa_p^{(\text{max})}$ for each pixel by solving the eigenvalue problem in Eq. (6). For this, we require the first and second fundamental forms, Eq. (5). Depending on the projections, we use Eqs. (7) and (8) respectively for the surface gradients and finite difference for gradients of the normal maps. We lift the pixel-wise curvatures to the mesh setting via

$$\kappa_v := \max\left(\left\{\kappa_p^{(\max)} \middle| p \in \mathcal{P}_{\mathcal{S}_v}\right\}\right).$$
(12)

where S_v is the star of v, i.e. the union of all triangles touching v. Taking the maximum ensures small triangles at sharp corners. Sticking to the work of Dunyach *et al.*, we use Eq. (9) to obtain optimal lengths L_v which are nominally vertex properties. The optimal length for an edge is simply the mean of these vertex properties.

4.3 Tangential Smoothing

To obtain an optimal Delaunay triangulation [8] and avoid triangles with small angles, we move each vertex to the centroid of its star, *i.e.* the union of its adjacent faces. In this case, we have to address foreshortening twice. First, the centroid is affected by the triangle size which is typically smaller after projection. Second, the centroid of the star is a weighted sum over the centroids of the faces in the star. The position of the central vertex relative to these face centroids is foreshortened too. Considering both aspects yields a linear system

$$\left(\sum_{f\in\mathcal{F}_v}\frac{A_f\cdot\sqrt{\det\mathbf{I}_f}}{L_f^2}\cdot\mathbf{I}_f\right)\cdot\mathbf{u}_v^{(k+1)} = \left(\sum_{f\in\mathcal{F}_v}\frac{A_f\cdot\sqrt{\det\mathbf{I}_f}}{L_f^2}\cdot\mathbf{I}_f\,\mathbf{c}_f^{(k)}\right) \quad (13)$$

where the areas A_f and centroids c_f are measured as they appear on-screen and foreshortening is compensated by the inclusion of the face's first fundamental forms I_f . L_f is simply the average of optimal lengths for the three vertices in f.

4.4 Iterative Mesh Refinement

After deriving all measurements for curvature and edge lengths in screen-space, we formulate our screen-based remeshing. We start with a uniform triangulation of the pixel grid, where each foreground pixel is split into two triangles. Following previous remeshing methods [3, 12], which are visualized in Fig. 5, the mesh is iteratively refined by performing the following steps:

- 1. Collapse short edges and split long edges.
- 2. Flip non-Delaunay edges.
- 3. Shift vertices (tangential smoothing).

To determine whether an edge is too long or too short, we use the optimal edge lengths and the heuristic constants derived in [3].

5 Mesh-Based Integration

For normal integration, previous authors [6, 28] have noted that a unified treatment of the orthographic and perspective case can be achieved by using the functional

$$E_{\text{Int}} = \int_{\Omega} \left(\boldsymbol{n} \cdot \boldsymbol{r} \, \partial_{u} z + n_{x} \right)^{2} + \left(\boldsymbol{n} \cdot \boldsymbol{r} \, \partial_{v} z + n_{y} \right)^{2} \, du \, dv \tag{14}$$

where $\mathbf{r} = \mathbf{e}_z$ and z the depth map in the orthographic case. In the perspective case, z is the logarithmic depth and \mathbf{r} is the camera ray given by the intrinsics. Using this unified formulation, we do not differentiate between the orthographic and perspective cases for the rest of this section. The full derivation of the mesh-based normal integration contains a lot of details and the interested reader is referred to the supplementary material. In this section, we will summarize only the main results.

In the mesh-based normal integration, we have one unknown depth value per vertex. Depth values within the triangle are obtained through linear interpolation using barycentric coordinates. By linearity, the derivatives in Eq. (14) are constant within each face [4] and so are the normals. The integral can then be split into a sum of integrals over faces which can be carried out explicitly. The result is a discretized energy that is quadratic in the vertex depths. Taking the derivative, the optimality condition is

$$\sum_{i \in \mathcal{V}_j} \sum_{f \in \mathcal{F}_{ij}} \cot(\alpha_{f,ij}) \cdot \left(2m_f \cdot (z_j - z_i) + b_f \cdot \binom{n_x}{n_y} \cdot (\boldsymbol{u}_j - \boldsymbol{u}_i) \right) = 0 \quad (15)$$

Table 1: Compression rates at three different quality settings for the five objects in the DiLiGenT-MV dataset. Compression rates are averaged over all 20 views for each object. Results are computed from ground truth normals as well as normals obtained from photometric stereo using [24] and [20].

	GT	Norm	nals		Li [24]		Heep [20]			
Dataset	Low	Mid	High	Low	Mid	High	Low	Mid	High	
Bear	98.8%	97.5%	94.4%	99.0%	97.7%	94.7%	99.0%	97.8%	95.1%	
Buddha	96.1%	92.4%	82.4%	97.2%	94.1%	86.2%	97.0%	93.7%	85.4%	
Cow	98.9%	97.6%	94.3%	99.0%	97.7%	94.5%	97.9%	95.4%	89.3%	
Pot2	98.1%	96.1%	91.2%	98.4%	96.6%	92.1%	98.2%	96.2%	91.6%	
Reading	98.2%	96.0%	90.7%	98.4%	96.4%	91.4%	95.8%	91.2%	79.7%	

for each vertex *i*. \mathcal{V}_i are the vertices adjacent to *i* and $\alpha_{f,ij}$ is the angle in face *f* opposite to the edge (i, j) which is calculated in screen-space without foreshortening using the first fundamental form. The two constants are

$$m_f := \frac{1}{12} \sum_{\substack{i,j \in \mathcal{V}_f \\ i < j}} (\boldsymbol{n}_f \cdot \boldsymbol{r}_i) (\boldsymbol{n}_f \cdot \boldsymbol{r}_j) \qquad b_f := \frac{1}{3} \sum_{i \in \mathcal{V}_f} (\boldsymbol{n}_f \cdot \boldsymbol{r}_i)$$
(16)

where \mathcal{V}_f are the three vertices of a face f. Analogous to the pixel-based case, the linear system is underdetermined as Eq. (15) only contains depth differences, *i.e.* adding a constant offset does not change E_{Int} .

Readers familiar with the topic of mesh processing will notice the similarity to the seminal cotangent weights for the vertex Laplacian [27]. These weights ensure that our integration energy is independent of the triangulation. Generally, the cotan may diverge but as angles in isotropic meshes are close to 60° , the cotans are well-behaved and yield numerically stable linear systems [4]. To our surprise and to the best of our knowledge, this discretized version of Eq. (14) and its derivation for triangle meshes, has not been published yet.

6 Implementation Details

Our meshing and integration algorithms rely on the SURFACEMESH [31] data structure and are written in C++. To translate between pixels and triangles we rely on the NVDIFFRAST renderer [23]. Before calculating pixel-curvatures, the normal maps are low-pass filtered with a Gaussian blur ($\sigma = \sqrt{2}$) in OPENCV [5]. The sparse linear systems of the normal integration part are solved with EIGEN [19]. In all tests, we run ten iterations of mesh refinement with five iterations of tangential smoothing. The optimal edge lengths are clamped to lie within 1 and 100 pixels. A reference implementation can be found under moritzheep.github.io/adaptive-screen-meshing.

7 Evaluation

Unlike previous work, our method is not bound to the image resolution. Instead, the mesh resolution is controlled by a user parameter ϵ . Therefore, we conduct all experiments at three quality settings: 'high' (0.1 mm), 'medium' (0.3 mm) and 'low' (1 mm), see Figs. 1 and 8. With the 'high' setting, we aim for accurate reconstructions, ideally matching pixel-level methods. The 'low' setting focuses on high compression rates, while the 'medium' setting balances quality and compression.

Our evaluation consists of three parts: First, we analyze the compression rates and runtime advantages of our meshing-first approach over pixel-based methods in Sec. 7.1. We continue by comparing the reconstruction error on samples of the DiLiGenT-MV dataset [24] in Sec. 7.2. Finally, we discuss various design decisions of our algorithm in an ablation study and investigate the reliability of our single user-parameter, see Sec. 7.3. The DiLiGenT-MV dataset contains 5 objects with 20 different views, leading to 100 normal maps of real data for the evaluation. It is an extensive and recent dataset but lacks high-resolution normal maps. We compensate for that by testing rendered normal maps of virtual objects with resolutions of up to 8192^2 . Additional results are found in the supplemental material.

7.1 Compression and Runtime

The core motivation for our meshing-first approach is to substantially reduce the number of free variables *prior* to the integration step. We test all three quality settings of our screen-space meshing algorithm on all 20 views of the DiLiGenT-MV dataset (Tab. 1) with ground truth normals as well as estimated normals from the photometric stereo methods in [24] and [20]. Despite the relatively small image size of the DiLiGenT-MV dataset, our mesh representation requires only 6-18% of free variables for high-resolution meshes. For some low-resolution meshes, our method requires only 1% the number of vertices compared to foreground pixels. The results are consistent across all test cases and robust even



Fig. 6: Runtime and compression comparison between our meshing-first and pixelbased integration for increasing image resolutions. Left: The number of free variables (pixels or vertices) as a function of foreground pixels. Right: Pixel-based integration compared to mesh first, *i.e.* meshing and integration. Please note the log-log scale.

Table 2: Average RMSE over all 20 views of the DiLiGenT-MV dataset for pixel-based integration [7], [34], [20], a combination of pixel-based integration [7] and subsequent remeshing [12] as well as mesh integration of uniform and our adaptive meshes. Errors are reported in mm.

	Pixel-l			sed	Combined $[7, 12]$			Uniform			Ours		
		[7]	[34]	[20]	Low	Mid	High	Low	Mid	High	Low	Mid	High
Ortho	Bear	2.97	3.36	n/a	3.12	2.99	2.98	4.39	4.07	3.74	3.95	3.65	3.37
	Buddha	6.74	6.85	n/a	6.78	6.76	6.75	7.79	7.62	7.41	7.74	7.54	7.33
	Cow	2.45	2.51	n/a	2.67	2.55	2.53	3.59	3.39	3.10	3.42	3.12	2.96
	Pot2	5.15	5.02	n/a	5.18	5.15	5.15	5.99	5.84	5.73	5.89	5.77	5.65
	Reading	<u>6.34</u>	6.05	n/a	6.37	6.35	6.34	7.18	7.00	6.87	7.08	6.93	6.83
Persp	Bear	2.91	n/a	3.05	3.09	2.94	<u>2.92</u>	4.62	4.17	3.84	3.94	3.72	3.47
	Buddha	6.75	n/a	6.60	6.80	6.77	6.76	7.85	7.64	7.43	7.74	7.53	7.40
	Cow	2.35	n/a	2.51	2.59	2.46	2.43	3.71	3.47	3.21	3.49	3.24	3.07
	Pot2	4.99	n/a	5.23	5.04	4.99	4.99	6.08	5.97	5.81	6.04	5.86	5.76
	Reading	6.28	n/a	<u>6.29</u>	6.33	6.29	6.29	7.20	7.03	6.89	7.19	6.94	6.85

for less accurate normals. The image resolutions of the DiLiGenT-MV dataset are comparatively low (less than 1 MP). To obtain resolutions ranging from 512^2 to 8192^2 pixels, we render normal maps from three highly detailed meshes in Blender. In this application, we only use the high-resolution mesh setting. The resulting plots in Fig. 6 suggest that the number of vertices grows sublinearly with the number of pixels, meaning that our compression factor increases at higher resolutions. For the female face, we obtain an impressive reduction by about 200 at 8192^2 pixels and maintain runtimes of a few minutes instead of hours.

7.2 Benchmark Evaluation

We compare the reconstruction error of our adaptive mesh integration against different pixel-based approaches [7, 20, 34]. Even at the highest mesh resolution, we require on average less than 10% of free variables. Given this level of compression, it is not surprising that pixel-based approaches perform slightly better, see Tab. 2. Still, discontinuities remain the main source of integration errors.

7.3 Ablation Study

Our mesh refinement algorithm consists of three iterative steps, all described and evaluated in previous work or textbooks [3,4]. The remaining design choice of the algorithm is our locally adaptive over a uniform triangulation, see Fig. 7. We create uniform meshes as a baseline by directly setting a constant optimal edge length such that the resulting mesh roughly matches the vertex count of its adaptive counterpart. To exclude possible side-effects of the surface integration,



Fig. 7: Comparison between uniform (top) and adaptive screen-space meshing (bot-tom). Despite up to 5% more vertices, the RMSE is always higher for the uniform meshes. Wireframes are rendered as vector graphics.

the 2D meshes are non-rigidly fitted to ground truth depth maps. Thus, the remaining RMSE quantifies the error introduced by the mesh representation alone. Although our locally adaptive triangulations always contain equal or fewer vertices than the uniform triangulations, they still achieve a lower alignment error in all cases, see Fig. 7.

Lastly, our method relies on a single user parameter to control the density of the mesh, namely ϵ , the permitted error of the surface. Throughout our evaluations, we consistently reported the results for three settings the parameter. In Fig. 8, we set $\epsilon \in [0.1 \text{ mm}, 3 \text{ mm}]$ and non-rigidly fit the 2D meshes to ground truth depth maps. The results show a strong correlation between the intended and achieved approximation error. After a dataset-dependent offset, the achieved error grows approximately linearly with ϵ . Additional results in the supplementary material suggest that this linear relation is still present after integration.

8 Conclusion & Limitations

We proposed a novel normal integration method, where normal maps are converted into triangle meshes *before* the integration. A single user parameter controls the mesh resolution. We showed that our data structure is substantially



Fig. 8: Impact analysis of the user-parameter. Left: Obtained high (0.1 mm), medium (0.3 mm) and low (1 mm) resolution meshes (as vector graphics) with respective error maps. Right: Average RMSE over all 20 views as a function of the user-parameter ϵ .

sparser, but well-suited to obtain detailed reconstructions. Our results show that the mesh representation is especially advantageous for high-resolution normal maps, where it facilitates enormous speed-ups. As with all work on normal integration, the resulting surfaces are only unique up to a global scale and absolute depth remains ill-posed. Furthermore, we do not handle depth discontinuities implicitly at the moment. For triangle meshes, we would need to control the mesh topology and align the triangle edges with the discontinuities along the surface. Developing such a method for our mesh-based integration is not straightforward and is an interesting research direction. Alternatively, depth discontinuities could be precomputed and introduced as small holes into the foreground masks. Anisotropic meshes to improve the handling of sharp corners are another interesting direction for future work. Here, skinny triangles might be a drawback as they may cause numerical instabilities due to diverging cotangents. Finally, our method can be easily integrated into most existing photometric stereo pipelines.

Acknowledgements

The "David Head" by 1d_inc and the "Football Medal2 - PhotoCatch" by Moshe Caine were licensed under CC BY 4.0. This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 - 390732324 (PhenoRob).

References

- Ackermann, J., Goesele, M.: A Survey of Photometric Stereo Techniques. Foundations and Trends in Computer Graphics and Vision 9(3-4), 149-254 (2015). https://doi.org/10.1561/0600000065
- Alliez, P., de Verdière, É.C., Devillers, O., Isenburg, M.: Centroidal Voronoi diagrams for isotropic surface remeshing. Graphical Models 67(3), 204–231 (May 2005). https://doi.org/10.1016/j.gmod.2004.06.007

An Adaptive Screen-Space Meshing Approach for Normal Integration

- Botsch, M., Kobbelt, L.: A remeshing approach to multiresolution modeling. In: Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing. pp. 185–192 (2004)
- Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., Lévy, B.: Polygon Mesh Processing. CRC press (2010)
- 5. Bradski, G.: The OpenCV library. Dr. Dobb's Journal of Software Tools (2000)
- Cao, X., Santo, H., Shi, B., Okura, F., Matsushita, Y.: Bilateral Normal Integration. In: Computer Vision – ECCV 2022, vol. 13661, pp. 552–567. Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-19769-7_32
- Cao, X., Shi, B., Fumio, O., Matsushita, Y.: Normal Integration via Inverse Plane Fitting With Minimum Point-to-Plane Distance. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2382–2391 (Jun 2021)
- Chen, L., Xu, J.c.: Optimal delaunay triangulations. Journal of Computational Mathematics pp. 299–308 (2004)
- Chen, Z., Cao, J., Wang, W.: Isotropic Surface Remeshing Using Constrained Centroidal Delaunay Mesh. Computer Graphics Forum **31**(7), 2077–2085 (Sep 2012). https://doi.org/10.1111/j.1467-8659.2012.03200.x
- Du, Q., Faber, V., Gunzburger, M.: Centroidal Voronoi Tessellations: Applications and Algorithms. SIAM Review 41(4), 637–676 (Jan 1999). https://doi.org/10. 1137/S0036144599352836
- Du, Z., Robles-Kelly, A., Lu, F.: Robust Surface Reconstruction from Gradient Field Using the L1 Norm. In: 9th Biennial Conference of the Australian Pattern Recognition Society on Digital Image Computing Techniques and Applications (DICTA 2007). pp. 203–209 (Dec 2007). https://doi.org/10.1109/DICTA.2007. 4426797
- Dunyach, M., Vanderhaeghe, D., Barthe, L., Botsch, M.: Adaptive remeshing for real-time mesh deformation. In: Eurographics 2013 - Short Papers. The Eurographics Association (2013)
- Durou, J.D., Courteille, F.: Integration of a Normal Field without Boundary Condition. In: Proceedings of the First International Workshop on Photometric Analysis For Computer Vision-PACV 2007. pp. 8–p. INRIA (2007)
- 14. Dziuba, M., Jarsky, I., Efimova, V., Filchenkov, A.: Image Vectorization: A Review (Jun 2023). https://doi.org/10.48550/arXiv.2306.06441
- Gan, J., Wilbert, A., Thormählen, T., Drescher, P., Hagens, R.: Multi-view photometric stereo using surface deformation. The Visual Computer 34(11), 1551–1561 (Nov 2018). https://doi.org/10.1007/s00371-017-1430-5
- Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '97. pp. 209–216. ACM Press, Not Known (1997). https: //doi.org/10.1145/258734.258849
- Gotardo, P.F.U., Simon, T., Sheikh, Y., Matthews, I.: Photogeometric Scene Flow for High-Detail Dynamic 3D Reconstruction. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 846–854. IEEE, Santiago, Chile (Dec 2015). https://doi.org/10.1109/ICCV.2015.103
- Griwodz, C., Gasparini, S., Calvet, L., Gurdjos, P., Castan, F., Maujean, B., Lillo, G.D., Lanthony, Y.: AliceVision Meshroom: An open-source 3D reconstruction pipeline. In: Proceedings of the 12th ACM Multimedia Systems Conference - MM-Sys '21. ACM Press (2021). https://doi.org/10.1145/3458305.3478443
- 19. Guennebaud, G., Jacob, B., et al.: Eigen v3 (2010)

- 16 M. Heep and E. Zell
- Heep, M., Zell, E.: ShadowPatch: Shadow Based Segmentation for Reliable Depth Discontinuities in Photometric Stereo. Computer Graphics Forum 41(7), 635–646 (Oct 2022). https://doi.org/10.1111/cgf.14707
- Horn, B.K., Brooks, M.J.: The variational approach to shape from shading. Computer Vision, Graphics, and Image Processing 33(2), 174–208 (1986)
- Khan, D., Plopski, A., Fujimoto, Y., Kanbara, M., Jabeen, G., Zhang, Y.J., Zhang, X., Kato, H.: Surface remeshing: A systematic literature review of methods and research directions. IEEE Transactions on Visualization and Computer Graphics 28(3), 1680–1713 (2020)
- Laine, S., Hellsten, J., Karras, T., Seol, Y., Lehtinen, J., Aila, T.: Modular primitives for high-performance differentiable rendering. ACM Transactions on Graphics (TOG) 39(6), 1–14 (2020)
- 24. Li, M., Zhou, Z., Wu, Z., Shi, B., Diao, C., Tan, P.: Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. IEEE Transactions on Image Processing 29, 4159–4173 (2020)
- Liao, Z., Hoppe, H., Forsyth, D., Yu, Y.: A subdivision-based representation for vector image editing. IEEE Transactions on Visualization and Computer Graphics 18(11), 1858–1867 (2012)
- Ma, X., Zhou, Y., Xu, X., Sun, B., Filev, V., Orlov, N., Fu, Y., Shi, H.: Towards layer-wise image vectorization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16314–16323 (2022)
- Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. Experimental mathematics 2(1), 15–36 (1993)
- Quéau, Y., Durou, J.D., Aujol, J.F.: Normal Integration: A Survey. Journal of Mathematical Imaging and Vision 60(4), 576–593 (2018)
- Schönberger, J.L., Frahm, J.M.: Structure-from-Motion Revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise View Selection for Unstructured Multi-View Stereo. In: European Conference on Computer Vision (ECCV) (2016)
- 31. Sieger, D., Botsch, M.: The polygon mesh processing library (2019)
- 32. Xie, W., Dai, C., Wang, C.C.: Photometric stereo with near point lighting: A solution by mesh deformation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4585–4593 (2015)
- 33. Xie, W., Wang, M., Wei, M., Jiang, J., Qin, J.: Surface Reconstruction from Normals: A Robust DGP-based Discontinuity Preservation Approach. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5328–5336 (2019)
- Xie, W., Zhang, Y., Wang, C.C., Chung, R.C.K.: Surface-from-gradients: An approach based on discrete geometry processing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2195–2202 (2014)
- Yi, R., Liu, Y.J., He, Y.: Delaunay mesh simplification with differential evolution. ACM Transactions on Graphics (TOG) 37(6), 1–12 (2018)
- Zhao, T., Busé, L., Cohen-Steiner, D., Boubekeur, T., Thiery, J.M., Alliez, P.: Variational shape reconstruction via quadric error metrics. In: ACM SIGGRAPH 2023 Conference Proceedings. pp. 1–10 (2023)
- 37. Zhou, M., Ding, Y., Ji, Y., Young, S.S., Yu, J., Ye, J.: Shape and Reflectance Reconstruction using Concentric Multi-Spectral Light Field. IEEE Transactions on Pattern Analysis and Machine Intelligence 42(7), 1594–1605 (2020)

An Adaptive Screen-Space Meshing Approach for Normal Integration 17

Zhu, D., Smith, W.A.P.: Least Squares Surface Reconstruction on Arbitrary Domains. In: Computer Vision – ECCV 2020, vol. 12367, pp. 530–545. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-58542-6_32