

Discontinuity-preserving Normal Integration with Auxiliary Edges

Hyomin Kim*

Yucheol Jung*

Seungyong Lee

POSTECH

{min00001, ycjung, leesy}@postech.ac.kr

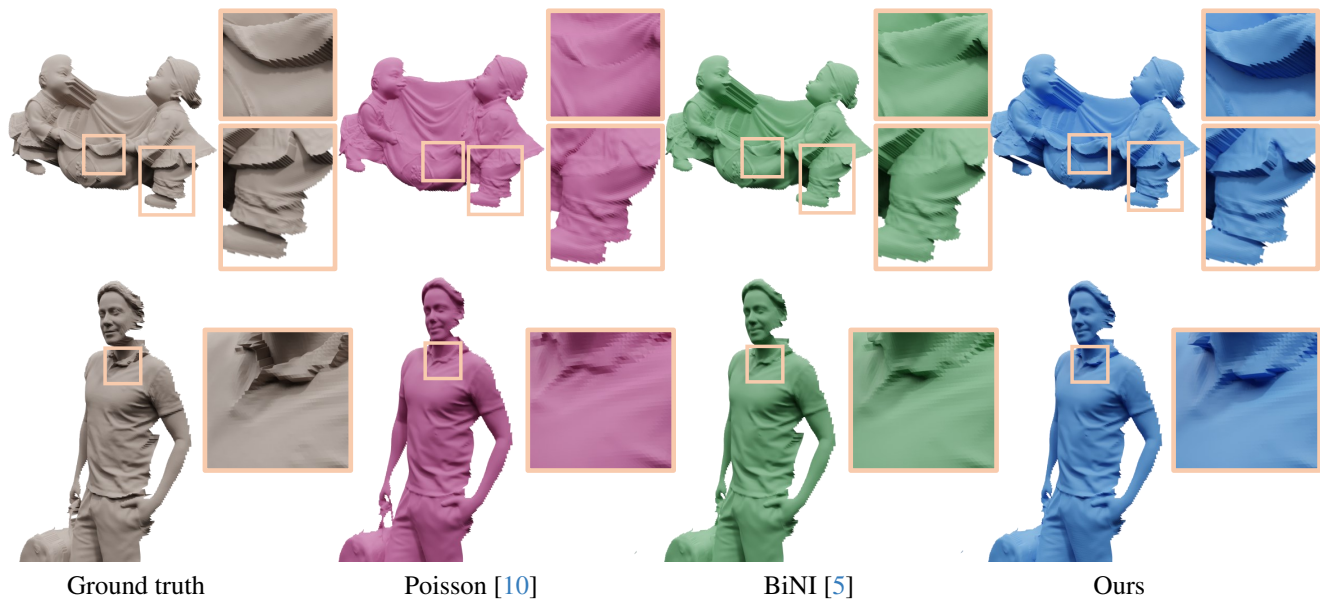


Figure 1. Our normal integration recovers a depth map from an input normal map while preserving surface discontinuities. Compared to the state-of-the-art discontinuity-preserving normal integration method, BiNI [5], our method recovers small discontinuities accurately.

Abstract

Many surface reconstruction methods incorporate normal integration, which is a process to obtain a depth map from surface gradients. In this process, the input may represent a surface with discontinuities, e.g., due to self-occlusion. To reconstruct an accurate depth map from the input normal map, hidden surface gradients occurring from the jumps must be handled. To model these jumps correctly, we design a novel discretization scheme for the domain of normal integration. Our key idea is to introduce auxiliary edges, which bridge between piecewise-smooth patches in the domain so that the magnitude of hidden jumps can be explicitly expressed. Using the auxiliary edges, we design a novel algorithm to optimize the discontinuity and the depth map from the input normal map. Our method optimizes discontinuities by using a combination of iterative re-weighted

least squares and iterative filtering of the jump magnitudes on auxiliary edges to provide strong sparsity regularization. Compared to previous discontinuity-preserving normal integration methods, which model the magnitudes of jumps only implicitly, our method reconstructs subtle discontinuities accurately thanks to our explicit representation of jumps allowing for strong sparsity regularization.

1. Introduction

Normal integration is a fundamental problem in 3D shape reconstruction. Most photometric stereo frameworks estimate the surface normal of objects, and normal integration recovers the depth of the object from the estimated normal map [2]. Normal integration is often used in single-image 3D reconstruction to reconstruct a detailed surface because a high-quality normal map is relatively easy to infer from an image compared to a fully detailed 3D surface [19].

*Equal contribution.

A challenge in normal integration is the fact that the normal map may represent a depth map with discontinuities. Discontinuities in a depth map are universally observed in captures of real-world objects due to self-occlusions and the hidden portions of surfaces. Not handling the discontinuity and assuming the surface is continuous everywhere introduces global distortions in normal integration results. In this work, we design a discontinuity-preserving normal integration method to automatically detect and handle surface discontinuity.

Previous discontinuity-preserving normal integration methods model such jumps by locating and weakening, or removing the effect of edges located across the discontinuity in the image grid. In this work, we design a new approach that models the jumps at discontinuity by *adding* edges corresponding to the jumps, instead of *removing* edges from the original image grid. We implement this design by introducing *auxiliary* edges that connect neighboring pixels in the original grid.

Our intuition behind this additive design is that the surface discontinuities can be viewed as additional hidden gradients that are not addressed by the input normal map. Unlike previous approaches [4, 5, 17] that models the jumps implicitly without considering what the values of the jumps across discontinuities should be, our approach enables explicit control over the locations and the degrees of the jumps. Benefiting from this explicit control for the jumps, we design an iterative optimization scheme that recovers sparse surface discontinuity from the input normal map by combining iterative re-weighted least squares (IRLS) [6] and iterative filtering of the jump magnitudes on the auxiliary edges. Compared to the state-of-the-art method [5], our approach accurately detects challenging subtle discontinuities thanks to our sparsity regularization (Fig. 1).

In summary, our technical contributions are as follows:

- We propose a novel normal integration framework based on a discrete graph with auxiliary edges that explicitly models the locations and magnitudes of jumps across surface discontinuities.
- We design a novel sparsity-regularized iterative optimization method for discontinuity and the depth map by combining IRLS and iterative filtering of jump magnitudes.

2. Problem Statement

Given a single-view normal map \mathbf{n} of an object, normal integration is a process to recover the depth map \mathbf{d} from \mathbf{n} . First, the depth gradient $\hat{\mathbf{g}}$ is calculated from the normal map \mathbf{n} . Assuming normal maps captured using orthographic cameras, the gradient map is

$$\hat{\mathbf{g}} = \frac{\mathbf{n}_{xy}}{n_z} \quad (1)$$

with the xy component $\mathbf{n}_{xy} = (n_x, n_y)$ and the z component n_z . Then, the depth map \mathbf{d} is reconstructed using the relation

$$\nabla \mathbf{d} = \hat{\mathbf{g}}, \quad (2)$$

where $\nabla \mathbf{d}$ is the gradients of \mathbf{d} . In the case of using a perspective camera, a similar relation can be formed [5]. The perspective case only requires two changes: 1) n_z is scaled according to the image coordinates and 2) $\log \mathbf{d}$ is used instead of \mathbf{d} in the optimization; then the final depth is recovered by applying exponential to the log values. The details for the perspective case can be found in the supplementary document. In this paper, for notational convenience, we describe our algorithm using n_z and \mathbf{d} .

Ideally, if the vector field $\hat{\mathbf{g}}$ is the true gradient map of the depth map to be reconstructed, the depth value at a position can be obtained by accumulating $\hat{\mathbf{g}}$ along any continuous path to the position from a position with a known depth [13]. However, for real-world data, $\hat{\mathbf{g}}$ is rarely a true gradient field for the depth map, causing different depths for different integration paths. This error can be caused by multiple factors [15], including sensor noise, quantization error, and surface discontinuity. As a result, variational methods have been proposed [14] to implicitly edit the input gradient $\hat{\mathbf{g}}$ using an editing \mathbf{g}' so that the vector field becomes a true depth gradient field \mathbf{g} ;

$$\mathbf{g} = \hat{\mathbf{g}} + \mathbf{g}'. \quad (3)$$

Multiple design choices can be considered to build the gradient editing \mathbf{g}' . Among them, modeling surface discontinuity is a critical one that should be considered to avoid large errors in reconstructions of real-world objects. For example, when the target surface contains large jumps along the view direction, e.g., due to self-occlusions, the jump cannot be effectively described using only the surface normal captured in an image. As a result, an integration path on the image grid crossing the occlusion boundary and another path not crossing the boundary produce largely different results. If the editing \mathbf{g}' correctly adds the jump at the discontinuous boundary, the two integration paths will produce the same results. In this work, given the input normal map \mathbf{n} , we aim to optimize \mathbf{g}' to obtain a correct depth map as a result of normal integration.

In the next section, we review previous normal integration frameworks and their implications in terms of gradient editing and handling of surface discontinuities.

3. Related Work

3.1. Normal integration

Normal integration is an extensively studied topic in computer vision. Refer to [13] and [14] for a survey of methods that aim to achieve various desirable properties for normal

integration, such as fast computation or robustness to noisy input. In this work, we focus on handling discontinuity on the surface. We first review a few notable methods for normal integration in general, then review various previous approaches for discontinuity-preserving normal integration.

A standard approach to reconstruct the depth satisfying Eq. (1), despite $\hat{\mathbf{g}}$ not being a true gradient map, is to find a least squares solution minimizing the difference between the gradient of the optimized depth and $\hat{\mathbf{g}}$. This approach of Horn and Brooks [10] is equivalent to solving a Poisson equation. As a result, the depth obtained from this approach implicitly applies minimal editing to the input gradient map by minimizing the l_2 norm of \mathbf{g}' . However, the minimization of l_2 norm produces dense \mathbf{g}' , whereas discontinuities should be located sparsely at occlusion boundaries. Consequently, this approach cannot correctly handle the discontinuities.

Xie *et al.* [18] proposed a normal integration method based on discrete geometry processing (DGP) where a quadrilateral is constructed for each pixel in the normal map, and normal integration is performed using the constructed mesh. Yet, this DGP method does not handle surface discontinuities. Our normal integration method constructs a similar discrete grid using a quadrilateral per pixel. In contrast to the previous DGP method, we introduce auxiliary edges that bridge between pixels to model discontinuities on the surface.

3.2. Discontinuity-preserving normal integration

Surface discontinuity in normal integration has been handled in various ways. An early approach [11] proposed automatic localization of discontinuities using the error residuals between the input gradients and the gradients of the surface reconstructed assuming no discontinuity. Agrawal and Raskar [3] considered the integration domain as a graph and cut the graph based on the curl of the input gradients. Fraile [7] proposed construction of an integration tree by computing a minimum spanning tree of a weighted graph with weights computed from the gradient map, e.g., using the principal curvature direction. However, such approaches based on a single analysis of the error residuals or the input gradients may not localize the discontinuity effectively; often, the residuals or the curls highlight only a part of the discontinuity.

To overcome the limitations of single-analysis approaches, iterative optimization methods have been employed for discontinuity-preserving normal integration, where the depth and the location of discontinuities are optimized alternately. Agrawal *et al.* [4] proposed a framework to build a spanning tree for integration by iteratively cutting the graph based on the residuals computed from intermediate depth values. Wu *et al.* [17] proposed an expectation-maximization algorithm alternating between

the discontinuity map and the parameters used to estimate the discontinuity map. Quéau *et al.* [14] proposed several approaches to optimize sparse discontinuity using l_1 gradient minimization, sparse gradient minimization using robust estimators, integration by anisotropic diffusion, and the minimization of Mumford-Shah functional.

Recently, Cao *et al.* [5] proposed a bilaterally weighted functional to perform discontinuity-preserving normal integration and demonstrated a large improvement over the methods of Quéau *et al.* [14] in terms of depth reconstruction error. The bilateral weights are iteratively computed based on intermediate depth reconstruction results so that the gradients from the pairs of depth pixels crossing the occlusion boundary weakly influence the normal integration. However, as the bilateral weights provide only a limited level of sparsity using smooth Heaviside function, small discontinuities may not be detected.

Our method is also based on the alternating optimization of the depth and discontinuity. Unlike previous approaches handling the gradient at discontinuity only implicitly via graph cuts or weight computation for functionals, we explicitly model the gradients across the surface discontinuity using auxiliary edges. Our approach detects small discontinuities accurately via sparsity regularization based on the explicit filtering of gradients at auxiliary edges, allowing more accurate reconstruction of the depth map.

4. Normal Integration with Auxiliary Edges

Given the gradient field $\hat{\mathbf{g}}$ in Eq. (1) calculated from the input normal map \mathbf{n} , our goal is to optimize the gradient editing \mathbf{g}' in Eq. (3) to create correct jumps across surface discontinuities when the gradient field is integrated to obtain the depth map \mathbf{d} . As the jumps for the discontinuities would occur only at the occlusion boundary, we assume \mathbf{g}' to be zero almost everywhere. In other words, the l_0 norm of \mathbf{g}' should be minimized while allowing the resulting depth \mathbf{d} to match the gradient field $\hat{\mathbf{g}}$ as much as possible:

$$\mathbf{g}', \mathbf{d} = \arg \min_{\mathbf{g}', \mathbf{d}} E_{data}(\hat{\mathbf{g}} + \mathbf{g}', \mathbf{d}) + \|\mathbf{g}'\|_0, \quad (4)$$

where E_{data} refers to the *data* term that measures the discrepancy between the gradient map $\hat{\mathbf{g}} + \mathbf{g}'$ and the normal integration result \mathbf{d} . The data term requires \mathbf{g}' not to be zero everywhere: the jumps for the discontinuity must be properly added to minimize the discrepancy.

l_0 -regularized optimization using strategies based on auxiliary variables has been studied in a wide range of applications, including image smoothing [20] and 3D shape filtering [8]. The minimization of l_0 norm is NP-hard, and different formulations to obtain a local optimum have been proposed depending on the problem domain. In this work, we design a novel l_0 -regularized optimization scheme specialized for normal integration using a discrete grid with




Figure 2. Construction of our graph for normal integration. For each pixel in the input gradient map, we construct a quadrilateral with four vertices for depth and their edges \mathcal{E}_v for directional derivatives. The quadrilaterals are bridged with auxiliary edges \mathcal{E}_a that model gradients across the discontinuity.




Figure 3. Construction of depth map from the normal integration results on the graph. The depth values on the four vertices of each quadrilateral are averaged and assigned to a depth pixel.

auxiliary edges.

Our key idea is to model the gradients $\hat{\mathbf{g}}$ and the editing \mathbf{g}' on separate finite elements in a directional graph that defines the discrete domain for normal integration. In this graph, each vertex is assigned a depth, and an edge is assigned a directional derivative. We first construct a quadrilateral for each pixel in the input normal map. Then, we bridge each quadrilateral with auxiliary edges (Fig. 2). The set of edges \mathcal{E}_v belonging to the quadrilaterals are assigned the projection of gradient $\hat{\mathbf{g}}$ to the edges, and the set of auxiliary edges \mathcal{E}_a are assigned the directional gradients \mathbf{g}' across quadrilaterals. After performing the normal integration on this graph, we construct the depth map on a pixel grid by averaging the depth values of vertices for each quadrilateral (Fig. 3).

On this graph with auxiliary edges, the data term E_{data} measuring the discrepancy between the gradient map and the depth can be defined as a sum of two terms, assuming an orthographic camera:

$$E_{data}(\hat{\mathbf{g}} + \mathbf{g}', \mathbf{d}) = E_v(\hat{\mathbf{g}}, \mathbf{d}) + E_a(\mathbf{g}', \mathbf{d}) \quad (5)$$

with

$$E_v(\hat{\mathbf{g}}, \mathbf{d}) = \sum_{e \in \mathcal{E}_v} \|n_z(e)(D_e \mathbf{d} - \hat{\mathbf{g}}(e))\|^2, \quad (6)$$

$$E_a(\mathbf{g}', \mathbf{d}) = \sum_{e \in \mathcal{E}_a} \|n'_z(e)(D_e \mathbf{d} - \mathbf{g}'(e))\|^2, \quad (7)$$

where $n_z(e)$ is the z component of the normal in the pixel containing the edge e , $D_e \mathbf{d}$ is the directional derivative of \mathbf{d} at e along the edge direction, and $\hat{\mathbf{g}}(e)$, $\mathbf{g}'(e)$ are the gradient values at edge e . The residuals for E_v are scaled with n_z for numerical stability; occasionally $n_z \approx 0$ and then $\hat{\mathbf{g}}$ obtained from Eq. (1) may diverge to infinity. We scale the residuals for E_a using a constant $n'_z = 1$ as the 3D surface normals corresponding to auxiliary edges are not clearly defined. Formulations of Eqs. (6) and (7) in the perspective case can be found in the supplementary document.

We model surface discontinuities as sparse gradient editing \mathbf{g}' , which are contained in the auxiliary edges. We incorporate the sparsity constraint of Eq. (4) by modifying the data term E_a for auxiliary edges. The modified energy is

$$E_{disc}(\mathbf{g}', \mathbf{d}, \mathbf{w}) = \sum_{e \in \mathcal{E}_a} w_e \|n'_z(e)(D_e \mathbf{d} - \mathbf{g}'(e))\|^2. \quad (8)$$

As directly minimizing $\|\mathbf{g}'\|_0$ is hard, we instead iteratively build sparse gradients \mathbf{g}' based on depth \mathbf{d} obtained with sparsity-regularized optimization. The variable weight w_e enables sparsity regularization for the residuals in E_{disc} using iteratively reweighted least squares (IRLS) [6], which is often employed to obtain sparse solutions that minimize l_p -norm for $0 < p < 1$. With IRLS, optimization of \mathbf{d} could produce non-zero residuals in E_{disc} at sparse discontinuities, where \mathbf{d} is allowed to be less affected by the values of \mathbf{g}' that may not represent the desired sizes of jumps yet. We then update \mathbf{g}' with explicitly filtered gradients of \mathbf{d} to modulate jump sizes and introduce additional discontinuities. The iterative update is described in the next section.

5. Discontinuity Optimization via Iterative Reweighting and Gradient Filtering

Our iterative optimization involves three variables: the depth map \mathbf{d} , the weights \mathbf{w} , and the gradient editing \mathbf{g}' on the auxiliary edges. The variables are updated in an alternating fashion using Algorithm 1.

Algorithm 1 Update of \mathbf{d} , \mathbf{w} , and \mathbf{g}'

```

n ← 0
g' ← 0
w ← 1
λc = 0.5(λsoft + λhard)
do
  for λ ∈ [λsoft, λc, λhard, λc] do
    d ← argmin Ev(ĝ, d) + λEdisc(g', d, w)
    w ← reweight(d)
    g' ← filter(d)
    n ← n + 1
  end for
while n < Nmax

```




Figure 4. Visualization of \mathbf{g}' and \mathbf{d} across iterative optimization. The gradients \mathbf{g}' on auxiliary edges are visualized as color; bright color means higher magnitude. The depth \mathbf{d} is shown as the height on the surface. The numbers at the bottom show the number of iterations n .

Solving for depth The minimization of $E_v + \lambda E_{disc}$ can be formulated as a weighted least squares (WLS) problem, which forms a linear equation of the form $\mathbf{A}^T \mathbf{W} \mathbf{A} \mathbf{d} = \mathbf{A}^T \mathbf{b}$ such that $\mathbf{A} \mathbf{d} - \mathbf{b}$ represents an array of residuals from all edges. The linear equation is solved using a conjugate gradient solver [9].

Initially, we start with $\mathbf{g}' = \mathbf{0}$ and $\mathbf{w} = 1$ to solve for \mathbf{d} . In this case, we are simply solving an ordinary least squares problem as in the Poisson-equation-based method [10]. Therefore, the initial depth would prefer dense and wide-spread residuals of E_{disc} so that the l_2 norm of $D_e \mathbf{d}$ is minimized, which means most pixels would have a small discontinuity nearby.

Given an intermediate solution \mathbf{d} , the **reweight** and **filter** steps are performed to update the weights \mathbf{w} and the gradients \mathbf{g}' for obtaining a more desirable solution in the next iteration. The updates of \mathbf{g}' and \mathbf{d} through this iterative process are visualized in Fig. 4.

Reweighting Using the depth map \mathbf{d} , we calculate the weights for the next step using directional derivatives D_e :

$$w_e = \min \left[\frac{1}{(D_e \mathbf{d})^2}, 1 \right]. \quad (9)$$

Based on the values $D_e \mathbf{d}$, the reweighting assigns higher weights for small gradients and lower weights for high gradients. In our case, most auxiliary edges have low jumps with small $D_e \mathbf{d}$ and are assigned high weights so that the small jumps can be suppressed by lowering the differences of $D_e \mathbf{d}$ from \mathbf{g}' , which should be zero almost everywhere. On the other hand, edges with large $D_e \mathbf{d}$ are allowed to produce even bigger jumps in the next iteration because of the lowered weights.

Gradient filtering However, applying the update to \mathbf{w} alone does not guarantee strong enough sparsity for the gradient of \mathbf{d} to properly express surface discontinuities (Sec. 6.3). We additionally perform filtering of the gradients $D_e \mathbf{d}$ in the update of \mathbf{g}' to build sparse gradients. This direct manipulation of the gradients of auxiliary edges enables more aggressive control of sparsity on the discontinuities in the depth map.

We build a new sparse gradient map \mathbf{g}' by filtering $D_e \mathbf{d}$ with a filter function f defined on auxiliary edges:

$$\mathbf{g}'(e) = f(e) \cdot D_e \mathbf{d}, \quad (10)$$

where $e \in \mathcal{E}_a$. We design $f(e)$ to produce a more sparse version of $D_e \mathbf{d}$, with non-zero values located at discontinuities. First, we expect the direction crossing the discontinuity to introduce abrupt change in the tangents to the view direction. We measure this quality using a function $s(e)$ that calculates the degree of local change in tangents:

$$s(e) = [n_z(e'_v) - n_z(e''_v)]^2 + \tau, \quad (11)$$

for a small value τ and non-auxiliary neighbor edges e'_v , e''_v of e along the direction e . n_z is used to capture the tangents with the view direction; a small value means the surface is more tangent to the view direction. Using this scaling function, we calculate scaled gradients $G(e)$:

$$G(e) = s(e) \cdot D_e \mathbf{d}. \quad (12)$$

Second, in addition to considering the tangents to the view direction, we also suppress gradients with non-maximal magnitudes; the jump magnitudes for the discontinuities are expected to be locally maximum. Therefore, we calculate the local-maximumness of the squared magnitude $[G(e)]^2$ using the Laplacian $L(e)$:

$$L(e) = 2[G(e)]^2 - [G(e'_a)]^2 - [G(e''_a)]^2 \quad (13)$$

for two neighboring auxiliary edges e'_a and e''_a along the direction of e .

Finally, the filter function $f(e)$ is designed as the thresholding on the local-maximumness to enforce sparsity:

$$f(e) = \begin{cases} 1/(1 + \exp(-kL(e))) & L(e) > 0 \\ 0 & L(e) \leq 0 \end{cases} \quad (14)$$

for a large value k . This filtering suppresses most gradients from $D_e \mathbf{d}$ except sparse jumps across discontinuities.




Figure 5. Qualitative comparison with previous normal integration methods using the DiLiGeNT dataset. Our method can reconstruct challenging and small discontinuities as well as large discontinuities that are handled by BiNI [5].

Two-pass update with λ_{soft} and λ_{hard} Aggressive control of sparsity in \mathbf{g}' via direct gradient manipulation may result in incorrect discontinuities: a high value of \mathbf{g}' may occur inside a continuous surface patch. Still, such erroneous large jumps would be healed automatically via the next depth solve step that produces a depth \mathbf{d} globally consistent with E_v , penalizing the discrepancy from the gradients $\hat{\mathbf{g}}$. On the other hand, the depth solve step may also remove valid non-zero values in \mathbf{g}' if the non-zero jump values are not spread wide enough to form a cut that allows a coherent large jump in depth values across the occlusion boundary.

To heal the errors while keeping the valid newborn discontinuities from vanishing, we design a two-pass update scheme, which is inspired by cosine annealing [12] that is known to facilitate the escape from local optima. We alternate between small weight λ_{soft} and large weight λ_{hard} for E_{disc} . In the first pass, we run the update process with λ_{soft} to heal any defective gradients in \mathbf{g}' . Then, the second pass reinforces the discontinuities obtained from the first step by updating with λ_{hard} : optimizing with a strong weight λ_{hard} widens the discontinuities in \mathbf{g}' , as strong influences from non-zero gradients in \mathbf{g}' promote nearby auxiliary edges to form discontinuities accordingly. Each pass is followed by another pass using the mean of λ_{soft} and λ_{hard} similarly to cosine annealing. This reinforcement step helps the next iteration with λ_{soft} to retain valid discontinuities.

6. Experiments

6.1. Implementation details

In the following experiments, we use the hyperparameters $\lambda_{soft} = 0.2$, $\lambda_{hard} = 1.2$, $N_{max} = 5000$, $k = 1000$, and $\tau = 10^{-2}$. Our method is implemented using PyTorch with CUDA enabled. Our unoptimized prototype implementation takes approximately 120 seconds to process a normal map of size 512×512 using a Linux Desktop with Intel i7-10700 CPU and NVIDIA Titan RTX GPU.

6.2. Comparison with previous methods

We compare our methods with the state-of-the-art discontinuity-preserving normal integration method, BiNI [5], which demonstrated significant improvement over previous discontinuity-preserving normal integration methods [14] in terms of depth accuracy. We evaluate the accuracy of depth reconstruction using a dataset with known ground-truth 3D meshes for the objects to be reconstructed. We use DiLiGeNT dataset [16], which contains normal maps and ground-truth depths for 10 objects, including sculptures of animals and teapots. We also construct a 3D human body dataset by collecting 27 human body models from RenderPeople [1]. For the human body dataset, we render the input normal maps and ground-truth depth maps using the ground-truth 3D meshes to build evaluation data. We use a perspective camera with field of view 33.4° for rendering.




Figure 6. Qualitative comparison to previous normal integration methods using the human body dataset.

In Figs. 5 and 6, we qualitatively compare our method with two methods: Poisson solver [10] that does not consider discontinuity in the integration, and the state-of-the-art method BiNI [5], using the hyperparameter $k = 2$ used in the paper. The Poisson solver introduced global distortions to the reconstruction results, leading to large differences from the ground-truth depth map. BiNI successfully handled large and obvious discontinuities. However, small and challenging discontinuities with small jumps are not handled properly. Our method correctly recovered large discontinuities found by BiNI. In addition, our method successfully recovered small discontinuities that are missed by BiNI. Our method can detect small discontinuities by solving for a more sparse solution based on iterative filtering of explicit gradients on auxiliary edges.

Table 1 shows quantitative evaluation results for the DiLiGenT dataset. Compared to the Poisson solver, our method and BiNI reconstruct the surface accurately; compared to BiNI, our method produces a lower error in most of the cases because jumps for small discontinuities are accurately handled by our sparsity regularization.

Table 1. Quantitative evaluation results for the DiLiGenT dataset. Mean absolute depth errors from the ground-truth are presented. We compare mainly with BiNI [5] because this state-of-the-art method demonstrated large improvements over previous discontinuity-preserving methods. Refer to BiNI [5] for the quantitative evaluation results of other methods.

	Bear	Buddha	Cat	Cow	Harvest	Pot1	Pot2	Reading	Goblet
Poisson [10]	1.20	3.71	1.60	0.89	10.09	1.51	0.75	6.62	11.64
BiNI [5]	0.49	0.76	0.11	0.07	3.30	0.63	0.24	0.32	8.65
Ours	0.45	0.67	0.24	0.06	2.44	0.57	0.19	0.15	9.02

6.3. Ablation study

Fig. 7 shows ablation study results for our iterative optimization. Without the reweighting, the optimization of the discontinuity takes a longer time because the updates for the jump magnitudes become more conservative due to remaining high weights for residuals on auxiliary edges at the discontinuity. Without gradient filtering, the iterative reweighting alone would not effectively create sparse discontinuities. Without two-pass update, the optimization fails to create discontinuities successfully for complex shapes.

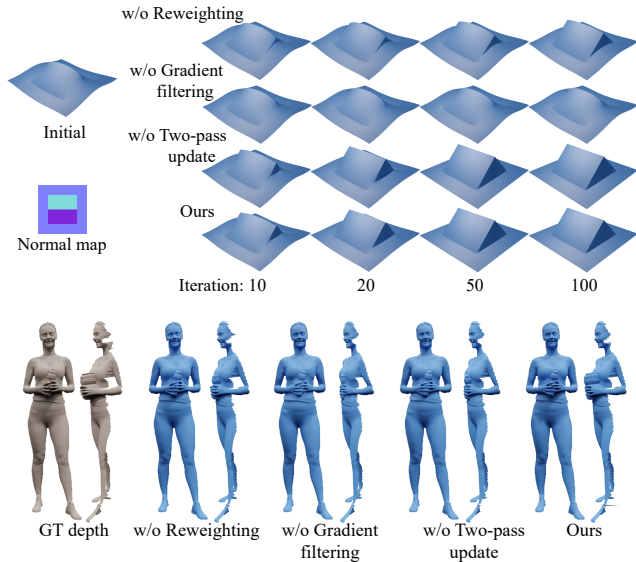


Figure 7. Ablation study of our iterative optimization algorithm. All of the reweighting, the gradient filtering, and the two-pass updates are required to obtain accurate discontinuities.

6.4. Building a mesh for the reconstructed surface

After normal integration, we build the final depth map by averaging four depth vertices in each quadrilateral. As a result, the final reconstructed surface is smoothed. We visually find that this smoothing is also observed when solving the normal integration on the original image grid, e.g., as in BiNI [5]. In contrast, by building a 3D mesh directly using the quadrilaterals in our graph with auxiliary edges, the reconstructed surface can exhibit almost perfect matching to the ground truth surface in terms of appearance. If the output dimension for the normal integration is not restricted to be the input grid, a mesh directly built on our graph would recover details better, as visualized in Fig. 8.

7. Conclusion

In this paper, we introduced a discontinuity-preserving normal integration method based on a novel discretization scheme of the integration domain. To handle surface discontinuities in normal integration, previous approaches implicitly weaken the effects of large gradients across occlusion boundaries. In contrast, we explicitly model the magnitudes of sparse gradients for occlusion boundaries using auxiliary edges. Our approach performs iterative optimization of gradients at auxiliary edges, achieving strong sparsity regularization on the discontinuities.

Other important properties than discontinuity preservation, such as robustness to noise or faster running time discussed in the survey of Quéau *et al.* [13], are left as future work. In addition, a normal map can contain multiple regions completely separated by occlusion boundaries. In this



Figure 8. Compared to the usual approach of normal integration using the input image grid, our integration using quadrilaterals on pixels reconstructs a 3D surface with more accurate normals. From left to right: GT mesh that the GT depth map is sampled from, GT depth map, BiNI [5] result, our result with depth averaging per quadrilateral, and our result without the averaging. The last column shows near-identical shading to the GT mesh and depth.

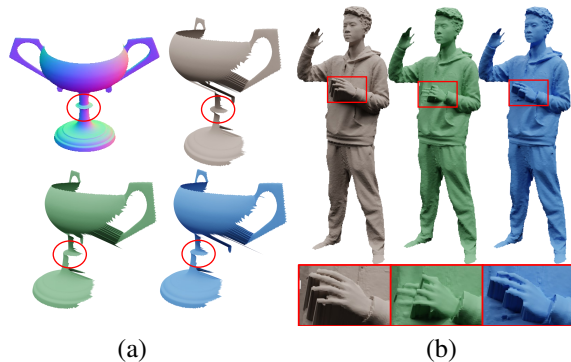


Figure 9. Failure cases. Brown: ground-truth depth map. Green: BiNI [5] results. Blue: ours. (a) Due to the global offset or scale ambiguity of the normal integration framework in general, the depth can be erroneous if the discontinuities completely cut the integration domain. (b) Depth reconstruction of noisy human hand scan fails in both BiNI and ours.

case, due to the ambiguities of normal integration in global translation or scale, large depth errors can occur in such surface islands (Fig. 9a). Handling the depth ambiguity for surface islands can be an interesting research direction.

Limitations Our iterative optimization recovers most surface discontinuities correctly, but failure cases exist. For a noisy human scan, reconstruction of fingers is challenging, and both BiNI and our method fail to reconstruct the correct discontinuity (Fig. 9b). Incorporating a depth bias from a data-driven model could be a promising approach for resolving both the aforementioned depth ambiguities and challenging failure cases.

Acknowledgements This work was supported by the NRF grant (RS-2023-00280400) and IITP grants (AI Graduate School Program, 2019-0-01906; AI Innovation Hub, 2021-0-02068) funded by Korea government (MSIT) and Samsung Electronics.

References

- [1] Renderpeople-scanned 3d people models provider. <http://renderpeople.com>. 6
- [2] Jens Ackermann, Michael Goesele, et al. A survey of photometric stereo techniques. *Foundations and Trends in Computer Graphics and Vision*, 9(3-4):149–254, 2015. 1
- [3] A. Agrawal, R. Chellappa, and R. Raskar. An algebraic approach to surface reconstruction from gradient fields. In *Proc. ICCV*, Beijing, China, 2005. IEEE. 3
- [4] Amit Agrawal, Ramesh Raskar, and Rama Chellappa. What Is the Range of Surface Reconstructions from a Gradient Field? In *Proc. ECCV*, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. 2, 3
- [5] Xu Cao, Hiroaki Santo, Boxin Shi, Fumio Okura, and Yasuyuki Matsushita. Bilateral normal integration. In *Proc. ECCV*. Springer, 2022. 1, 2, 3, 6, 7, 8
- [6] Rick Chartrand and Wotao Yin. Iteratively reweighted algorithms for compressive sensing. In *Proc. ICASSP*. IEEE, 2008. 2, 4
- [7] R. Fraile and E.R. Hancock. Combinatorial Surface Integration. In *Proc. ICPR*, Hong Kong, China, 2006. IEEE. 3
- [8] Lei He and Scott Schaefer. Mesh denoising via l 0 minimization. *ACM Transactions on Graphics (TOG)*, 32(4):1–8, 2013. 3
- [9] Magnus R Hestenes, Eduard Stiefel, et al. Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards*, 49(6):409–436, 1952. 5
- [10] Berthold KP Horn and Michael J Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33(2):174–208, 1986. 1, 3, 5, 6, 7
- [11] Bilge Karaçalı and Wesley Snyder. Reconstructing discontinuous surfaces from a given gradient field using partial integrability. *Computer Vision and Image Understanding*, 92(1):78–111, 2003. 3
- [12] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *Proc. ICLR*, 2016. 6
- [13] Yvain Quéau, Jean-Denis Durou, and Jean-François Aujol. Normal Integration: A Survey. *Journal of Mathematical Imaging and Vision*, 60(4):576–593, 2018. 2, 8
- [14] Yvain Quéau, Jean-Denis Durou, and Jean-François Aujol. Variational Methods for Normal Integration. *Journal of Mathematical Imaging and Vision*, 60(4):609–632, 2018. 2, 3, 6
- [15] Rafael FV Saracchini, Jorge Stolfi, Helena CG Leitão, Gary A Atkinson, and Melvyn L Smith. A robust multi-scale integration method to obtain the depth from gradient maps. *Computer Vision and Image Understanding*, 116(8):882–895, 2012. 2
- [16] Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proc. CVPR*, pages 3707–3716, 2016. 6
- [17] Tai-Pang Wu and Chi-Keung Tang. Visible Surface Reconstruction from Normals with Discontinuity Consideration. In *Proc. CVPR*, New York, NY, USA, 2006. IEEE. 2, 3
- [18] Wuyuan Xie, Yunbo Zhang, Charlie C. L. Wang, and Ronald C.-K. Chung. Surface-from-Gradients: An Approach Based on Discrete Geometry Processing. In *Proc. CVPR*, 2014. 3
- [19] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans optimized via normal integration. In *Proc. CVPR*, 2023. 1
- [20] Li Xu, Cewu Lu, Yi Xu, and Jiaya Jia. Image smoothing via l 0 gradient minimization. In *Proc. SIGGRAPH Asia*, 2011. 3