Surface Reconstruction from Unoriented Points via Green's 3rd Identity

ZHONGHAO WU, University of Science and Technology of China, China and Shanghai Innovation Institute, China
DONG XIAO, University of Science and Technology of China, China
RENJIE CHEN*, University of Science and Technology of China, China

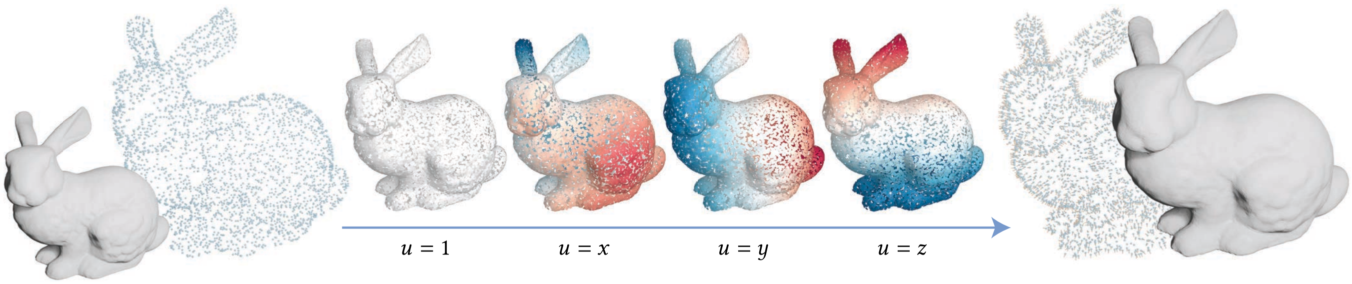

Fig. 1. Overview of G3R. Starting from an input unoriented point cloud (left), we generalize the winding-number approach by evaluating Green's third identity on a basis of harmonic functions (e.g. $u \in \{1, x, y, z\}$, middle). By stacking the integral constraints from multiple harmonic functions, we construct a well-conditioned linear system that directly yields globally consistent normals and high-quality surface reconstruction (right).

Surface reconstruction from unoriented point clouds is a fundamental problem in geometry processing. While recent winding-number-based approaches, such as Parametric Gauss Reconstruction (PGR) and Anisotropic Gauss Reconstruction (AGR), provide elegant implicit formulations, they often suffer from underdetermined or ill-conditioned linear systems. In this work, we analyze the theoretical root of this instability: the winding number formula is oblivious to a broad family of tangential perturbations of the normal field, which introduces a non-trivial nullspace in the resulting integral equations. To address this, we present G3R, a simple yet effective method based on Green's third identity for harmonic functions. This approach eliminates the tangential nullspace with a basis of harmonic functions, resulting in a well-conditioned linear system which can be solved directly without additional regularization. Furthermore, we apply a Gaussian mollification of the fundamental solution, which handles near-field singularities robustly. We adopt a matrix-free implementation and accelerate it with a hierarchical scheme to scale our method to large point clouds. Experiments demonstrate that our approach achieves highly competitive performance in both normal orientation accuracy and surface reconstruction quality. Our code is publicly available at https://github.com/ne0-wu/G3R.

CCS Concepts: • **Computing methodologies** → **Shape modeling**.

ACM Reference Format:
Zhonghao Wu, Dong Xiao, and Renjie Chen. 2026. Surface Reconstruction from Unoriented Points via Green's 3rd Identity. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3799902.3811154

*Corresponding author.

Authors' Contact Information: Zhonghao Wu, School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui, China and Shanghai Innovation Institute, Shanghai, China, wuzhonghao@mail.ustc.edu.cn; Dong Xiao, School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui, China, xiaodong@ustc.edu.cn; Renjie Chen, School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui, China, renjiec@ustc.edu.cn.

SIGGRAPH Conference Papers '26, Los Angeles, CA, USA

ACM ISBN 979-8-4007-2554-8/26/07
https://doi.org/10.1145/3799902.3811154

1 Introduction

Reconstructing high-quality surface meshes from raw, unorganized point clouds is a cornerstone problem in computer graphics and geometry processing. An effective class of approaches in this category is based on the winding number or Gauss formula, which expresses the indicator function of a region (evaluates to 1 inside, 0 outside, and $1/2$ on the boundary) as a boundary integral over the surface. PGR [Lin et al. 2022] and its anisotropic extension AGR [Ma et al. 2025a] utilize the constraints that the winding number formula evaluates to $1/2$ at surface samples, leading to a linear system for the unknown area-weighted normals.

In this paper, we revisit these formulations from the viewpoint of boundary integral equations. The winding number can be seen as a special case of Green's third identity evaluated with the constant harmonic function $u = 1$. This makes the constraints *insensitive* to a broad family of tangential perturbations (in particular, surface-divergence-free fields), thereby introducing a large nullspace in the resulting integral equations, leading to under-determinacy and poor conditioning observed in PGR and AGR. Motivated by this insight, we propose a new formulation based on Green's third identity that evaluates the identity on a basis of harmonic functions, thereby eliminating the tangential nullspace and offering a way to construct a stable linear system.

Specifically, we build equations from Green's third identity using a set of harmonic functions $\{u_k\}_{k=1}^{K}$. Each harmonic function yields a linear constraint on the unknown normals, and stacking these constraints produces a well-conditioned system to be solved in the least-squares sense. In contrast to prior winding-number-based methods, the resulting normal equations exhibit substantially improved numerical conditioning, and allow for direct solution without additional regularization.

To handle the near-field singularities in the integral kernels for stable evaluation, we replace the Dirac delta source in the fundamental solution with a smooth Gaussian mollifier, which removes the singularity and enables robust near-field evaluation of both the single and double layer potentials.

For moderate-sized point clouds, the resulting least-squares system can be assembled explicitly and solved with direct solvers whenever memory permits, providing a simple and reliable implementation path. To scale to large inputs, we further adopt iterative solvers and implement our operators in a matrix-free fashion. We accelerate the time-consuming matrix-vector products using an octree-based scheme, leveraging the far-field decay of the underlying kernels.

Our main contributions are summarized as follows:

- An analysis of the nullspace problem in winding-number-based normal orientation methods.
- A stable, well-conditioned linear system to solve for normals by evaluating Green's third identity on a basis of harmonic functions and eliminating the tangential nullspace.
- A scalable, matrix-free implementation using a hierarchical structure, enabling efficient processing of large point clouds.
- Highly competitive performance on normal orientation accuracy and surface reconstruction quality, demonstrated through extensive experiments.

2 Related Work

2.1 Surface Reconstruction from Oriented Point Clouds

A substantial body of prior work assumes that oriented points—positions together with consistently oriented normals—are available, and focuses on reconstructing an implicit surface from such input. Poisson Surface Reconstruction (PSR) [Kazhdan et al. 2006] and its screened variant SPSR [Kazhdan and Hoppe 2013] are classical methods in this category. Given a set of oriented points, these methods interpret the input normals as samples of the gradient of an indicator function, and recover a scalar field by solving a (screened) Poisson equation on a volumetric grid using a finite element discretization. The resulting scalar field can be viewed as a smooth approximation of the indicator function, from which the surface is extracted as an iso-surface. Sellán and Jacobson [2022] recently proposes a stochastic interpretation of Poisson reconstruction by viewing the indicator function as the posterior mean of a Gaussian process conditioned on the input normal observations, enabling Bayesian uncertainty quantification for the reconstruction.

In parallel, a family of methods derives implicit formulations from the generalized winding number [Barill et al. 2018; Jacobson et al. 2013], also referred to as the Gauss formula [Lu et al. 2018]. These methods express the indicator function of a shape as the total solid angle subtended by the surface at a query point, and approximate the boundary integral by discrete sums over the input points.

2.2 Reconstruction from Unoriented Point Clouds

In many practical scenarios, raw point clouds lack consistently oriented normals. Building upon the success of winding-number- and Poisson-based methods for oriented inputs, recent advances leverage the structural properties of implicit fields as geometric priors to tackle the normal orientation problem. PGR [Lin et al. 2022] directly enforces the winding number constraint at the input surface samples. However, this yields an under-determined system, with only N equations for $3N$ unknowns, and thus requires artificial regularization to solve. AGR [Ma et al. 2025a] resolves this problem by deriving a family of anisotropic kernels and stacking the resulting equations to construct an over-determined linear system. Concurrently, Convection Augmented Gauss Reconstruction [Ma et al. 2025b] introduces additional convection terms to address this issue. Globally Consistent Normal Orientation (GCNO) [Xu et al. 2023] takes a different route and enforces the value of winding number away from the surface to be 0 or 1 through a non-linear energy. Metzer et al. [2021] trains a neural network to learn coherent normal field on patches, and attains global consistency via a dipole propagation. Closest to the Poisson framework, Kohlbrenner et al. [2025] formulates the problem on tensor fields rather than vector fields and squares the Poisson energy to handle unoriented inputs. Leveraging the Gaussian Process view of PSR, Huang et al. [2024] incorporates uncertainty quantification into the normal orientation process. Xiao et al. [2023] propose to simultaneously optimize the implicit field and point normals by establishing isovalue constraints on the point cloud.

A particularly successful class of recent methods formulates normal orientation as a fixed-point problem, and solves it via an iterative scheme. Iterative Poisson Surface Reconstruction (iPSR) [Hou et al. 2022] extends SPSR to unoriented inputs by repeatedly solving the screened Poisson equation and updating the normals using the iso-surface's normal vectors. Winding Number Normal Consistency (WNNC) [Lin et al. 2024] states that the negative gradient of the winding-number field at a point should be codirectional with the normal used to construct the field. The accompanying algorithm alternates between minimizing the PGR residual and realigning the normals to match the gradients of the winding number. Liu et al. [2025] also employ this gradient-alignment strategy and applies it to the screened version of winding number.

Besides winding-number- and Poisson-based methods, more geometric approaches have been proposed. Huang et al. [2019] defines an implicit function as the minimizer of a Duchon's energy and excels at sparse data. Xia and Ju [2025] extend this framework to cater to large point clouds by adopting the natural neighbors and rewriting the global formulation as a local one. Gotsman and Hormann [2024] propose a linear method derived from Stokes' theorem, yielding a simple global linear solve for consistent orientation. Fu et al. [2025] transform the binary normal orientation problem into a relaxed sign field optimization on a Delaunay graph and convert it into a sparse linear system. Scrivener et al. [2025] model the input as a conductive Faraday cage and solve a volumetric Poisson equation for the electric potential with regard to multiple linear external fields and derive surface normals from the maximum field strength.

Neural implicit methods have also been explored for reconstruction from unoriented point clouds. Ma et al. [2022] learn local context priors from shape datasets and fit them to the input point cloud by optimizing predictive queries. Instead of relying on dataset-level priors, Ben-Shabat et al. [2022] directly optimize the neural implicit representation using divergence-based regularization, while Wang et al. [2023] enforce the singularity of the Hessian near the surface.

3 Preliminaries

Notations. Let $\Omega \subset \mathbb{R}^3$ be a bounded open set with a sufficiently smooth boundary $\partial\Omega$. We assume that the input point cloud $\{p_i\}_{i=1}^N$ uniformly samples $\partial\Omega$. Let $x, y \in \mathbb{R}^3$ be generic points in space. In the context of boundary integrals, x typically denotes the evaluation point, and y denotes the integration variable. They also represent points in the point cloud when written with subscripts. We use n to represent the outward unit normal field on $\partial\Omega$, while n_i stands for the normal at point p_i. We use η to denote a general vector field on $\partial\Omega$, and η_i for its discrete area-weighted counterpart associated with point p_i. The latter is the unknowns to be solved for in the orientation problem. By a slight abuse of notation, we also use x, y, z to denote the coordinate components of a point $x \in \mathbb{R}^3$ when the context is clear.

3.1 Green's Third Identity

Given a bounded open set $\Omega \subset \mathbb{R}^3$ with a sufficiently smooth boundary $\partial\Omega$, Green's third identity states that for a function u harmonic in Ω, its value at a point x can be represented as a boundary integral. The integral is given by:

$$\hat{u}(x) = \int_{\partial\Omega} \left(\Phi(y-x) \frac{\partial u}{\partial n}(y) - u(y) \frac{\partial \Phi}{\partial n}(y-x) \right) dS(y),$$

where $\Phi(x) = 1/(4\pi|x|)$ is the fundamental solution of the Laplace equation in $\mathbb{R}^3$. The first part of the integrand is known as the single layer potential, while the second part is called the double layer potential. This integral formula remarkably yields the value of $u(x)$ depending on the location of x:

$$\hat{u}(x) = \begin{cases} u(x) & x \in \Omega \\ u(x)/2 & x \in \partial\Omega \\ 0 & x \in \bar{\Omega}^c \end{cases}.$$

In texts [Kress 2014; Steinbach 2008] this formula is often presented for only $x \in \Omega$ or $x \in \bar{\Omega}^c$; for the boundary case, we provide a brief derivation in the supplementary material.

When we choose u to be the constant function $u = 1$, we obtain the classical winding number formula:

$$\hat{u}(x) = \int_{\partial\Omega} \left(\Phi(y-x) \cdot 0 - 1 \cdot \frac{\partial \Phi}{\partial n}(y-x) \right) dS(y) = \begin{cases} 1 & x \in \Omega \\ 1/2 & x \in \partial\Omega \\ 0 & x \in \bar{\Omega}^c \end{cases}.$$

In the orientation problem, the normal vector field $n(y)$ on $\partial\Omega$ is unknown and to be solved for. We consider η as a general vector field on $\partial\Omega$, and rewrite Green's third identity and the winding number formula as linear operators acting on η, mapping from vector fields on $\partial\Omega$ to scalar fields in $\mathbb{R}^3$:

$$(\mathcal{A}[u]\eta)(x) := \int_{\partial\Omega} \left(\Phi(y-x)\nabla u(y) - u(y)\nabla_y \Phi(y-x) \right) \cdot \eta(y)\, dS(y),$$

$$(\mathcal{W}\eta)(x) := (\mathcal{A}[1]\eta)(x) = \int_{\partial\Omega} -\nabla_y \Phi(y-x) \cdot \eta(y)\, dS(y).$$

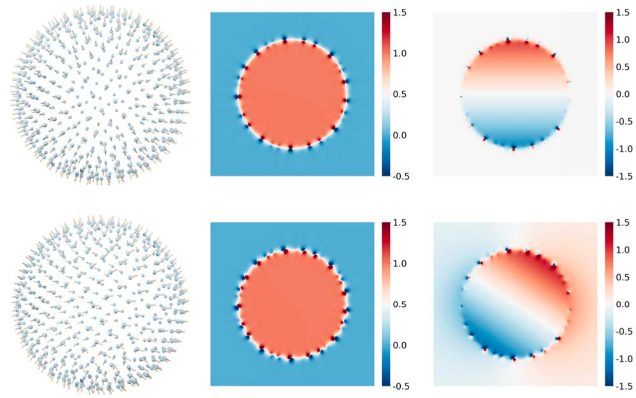

Fig. 2. Visualization of a surface-divergence-free tangential perturbation τ that is *invisible* to the winding number operator $\mathcal{W}$ but *visible* to Green's third-identity operator $\mathcal{A}[u]$. Top: ground-truth normal field $\eta = n$ on the unit sphere. Bottom: perturbed field $\eta = n + \tau$, with τ being tangential and surface-divergence-free. Left to right: (1) η on the point cloud; (2) winding-number scalar field (unchanged by τ); (3) Green's-third-identity scalar field with harmonic $u = y$ (changes under τ), showing $\tau \in \text{null}(\mathcal{W})$ but $\tau \notin \text{null}(\mathcal{A}[u])$. Integrals are approximated by uniformly sampling the unit sphere, and the resulting scalar fields are visualized on the xy-plane; $\tau(x, y, z) = (0, 0, 1) \times (x, y, z) = (-y, x, 0)$. The ripples on the boundary come from near-field behaviors of discretized dipoles and are not of interest.

4 Method

4.1 Motivation

Existing winding-number-based methods solve over the space of vector fields η on the surface, and apply constraints by fitting $\mathcal{W}\eta$ to the true winding number function. However, a fundamental problem with this formulation is the existence of a non-trivial nullspace in the winding-number operator $\mathcal{W}$. Let τ be a smooth tangential vector field on $\partial\Omega$, then

$$\begin{aligned} (\mathcal{W}\tau)(x) &= \int_{\partial\Omega} -\nabla_y \Phi(y-x) \cdot \tau(y)\, dS(y) \\ &= \int_{\partial\Omega} -\nabla_{\partial\Omega} \Phi(y-x) \cdot \tau(y)\, dS(y) \\ &= \int_{\partial\Omega} \Phi(y-x)\, \text{div}_{\partial\Omega}\, \tau(y)\, dS(y), \end{aligned}$$

where $\nabla_{\partial\Omega}$ and $\text{div}_{\partial\Omega}$ denote the surface gradient and surface divergence, respectively. Since τ is tangential, only the tangential component of $\nabla\Phi$ contributes to the dot product, which gives the second line. The third line then follows from integration by parts on the surface [Lee 2012]. Consequently, any tangent vector field τ that is divergence-free on the surface satisfies

$$\int_{\partial\Omega} -\nabla_y \Phi(y-x) \cdot \tau(y)\, dS(y) = 0, \quad \forall x \in \mathbb{R}^3.$$

This family of fields constitutes a nullspace of the winding-number operator. In practice, the presence of this nullspace leads to ill-conditioning, necessitating artificial regularizations.

In contrast, consider the Green's third identity operator $\mathcal{A}[u]$. We examine

$$\begin{aligned}(\mathcal{A}[u]\tau)(x) &= \int_{\partial\Omega} \big(\Phi(y-x)\nabla u(y) - u(y)\nabla_y \Phi(y-x)\big)\cdot\tau(y)\,dS(y)\\ &= \int_{\partial\Omega} \Phi(y-x)(\nabla_{\partial\Omega} u\cdot\tau + \operatorname{div}_{\partial\Omega}(u\tau))(y)\,dS(y)\\ &= \int_{\partial\Omega} \Phi(y-x)(2\nabla_{\partial\Omega} u\cdot\tau + u\operatorname{div}_{\partial\Omega}\tau)(y)\,dS(y).\end{aligned}$$

The nullspace of $\mathcal{A}[u]$ depends on the choice of harmonic function u, and is no longer described by a simple geometric condition. To intuitively illustrate this difference of nullspaces, we show an example of $\mathcal{W}$ and $\mathcal{A}[u]$ acting on the normal field and that perturbed by a divergence-free tangential vector field in Figure 2. We further discuss the elimination of the tangential nullspace by an appropriate choice of harmonic functions in Section 4.3.

4.2 Construction of the Linear System

With the integral equation arising from Green's third identity, we can choose any harmonic function u, and construct a linear system to solve for the unknown normals. PGR [Lin et al. 2022] and AGR [Ma et al. 2025a] derive their equations from the value of the winding number on the boundary:

$$(\mathcal{W}\eta)(x) = \frac{1}{2}, \quad \forall x\in\partial\Omega.$$

Following this approach, we evaluate Green's third identity for a harmonic function u on the boundary:

$$(\mathcal{A}[u]\eta)(x) = \frac{u(x)}{2}, \quad \forall x\in\partial\Omega.$$

Later in Section 4.3, we will show that this choice helps eliminate the tangential nullspace.

We then discretize the integral equation by approximating the surface integral with a discrete sum over the input point cloud $\{y_i\}_{i=1}^N$, where each point y_i represents a local surface patch. Let η_i be the area-weighted vector field value at y_i, we define the discretized operator $A[u]$ as follows:

$$\begin{aligned}(\mathcal{A}[u]\eta)(x) &= \int_{\partial\Omega} \big(\Phi(y-x)\nabla u(y) - u(y)\nabla_y \Phi(y-x)\big)\cdot\eta(y)\,dS(y)\\ &\approx \sum_{i=1}^N \big(\Phi(y_i-x)\nabla u(y_i) - u(y_i)\nabla_y\Phi(y_i-x)\big)\cdot\eta_i =: (A[u]\eta)(x).\end{aligned}$$

In this discretization, Φ and its gradient become numerically problematic when y_i approaches x. We handle this issue using the Gaussian mollification described in Section 4.4. By evaluating the equation at each input point, we obtain a linear system for the unknowns η_i:

$$A[u]\eta = b[u] := (u(x_1)/2, u(x_2)/2, \ldots, u(x_N)/2)^\top.$$

This linear system is underdetermined with only N equations for $3N$ unknowns. To address this, we generate additional constraints by introducing a set of linearly independent harmonic functions $\{u_k\}_{k=1}^K$. Stacking the resulting equations for each u_k, a $KN\times 3N$ system is formed:

$$A[u_k]\eta = b[u_k], \quad k=1,2,\ldots,K,$$

which can be solved in the least-squares sense as:

$$\eta = \arg\min_\eta \sum_{k=1}^K |A[u_k]\eta - b[u_k]|^2.$$

4.3 Choice of Harmonic Functions

To discuss the choice of harmonic functions $\{u_k\}_{k=1}^K$, we first return to the continuous setting and analyze the nullspace of the integral equation $\mathcal{A}[u]\eta = u/2$ on $\partial\Omega$. Following Section 4.1, substituting η with a tangential vector field τ into $\mathcal{A}[u]\eta = 0$ on $\partial\Omega$ yields

$$\int_{\partial\Omega} \Phi(y-x)(2\nabla_{\partial\Omega} u\cdot\tau + u\operatorname{div}_{\partial\Omega}\tau)(y)\,dS(y) = 0, \quad \forall x\in\partial\Omega.$$

Let $\varphi = 2\nabla_{\partial\Omega} u\cdot\tau + u\operatorname{div}_{\partial\Omega}\tau$ be the induced scalar density on the surface. The above equation can be recognized as a single layer operator applied to φ:

$$(\mathcal{S}\varphi)(x) := \int_{\partial\Omega} \Phi(y-x)\varphi(y)dS(y) = 0, \quad \forall x\in\partial\Omega.$$

Since the single layer operator $\mathcal{S}$ is injective on $C(\partial\Omega)\to C(\partial\Omega)$ [Kress 2014, Theorem 7.39], τ lies in the nullspace if and only if $\varphi\equiv 0$. By choosing a collection of harmonic functions $u\in\{1,x,y,z\}$, i.e. the constant and coordinate functions, τ is forced to be zero on the whole surface, eliminating the tangential nullspace.

To further improve the conditioning of the system, we introduce more harmonic functions to generate additional equations. A simple yet effective choice that extends $\{1,x,y,z\}$ is the set of harmonic polynomials up to a certain degree. In our implementation, we use the real regular solid harmonics u_l^m derived from spherical harmonics Y_l^m:

$$u_l^m(x) = \begin{cases}\sqrt{2}r^l\operatorname{Re}(Y_l^m(\theta,\phi)) & m>0\\ r^l Y_l^0(\theta,\phi) & m=0\\ \sqrt{2}r^l\operatorname{Im}(Y_l^{|m|}(\theta,\phi)) & m<0\end{cases},$$

where $\{u_l^m, -l\le m\le l, 0\le l\le L\}$ forms an orthonormal basis on $L^2(S^2)$ for harmonic polynomials with degree up to L (e.g. $u_0^0=1$, $u_1^{-1}=y$, $u_1^0=z$, $u_1^1=x$, etc. up to normalization). In practice, we find that using harmonic polynomials up to degree $L=2$ provides a good balance between numerical stability and computational cost.

Singular Value Analysis. Figure 3 compares the spectra of the stacked normal matrices $\sum A^\top A$ produced by PGR, AGR and our method on a point cloud of $N=2000$ points sampled from the Stanford Bunny model. PGR is underdetermined ($N\times 3N$) and thus has at most N non-zero singular values. AGR expands this to a $KN\times 3N$ system, but its smallest singular values still nearly vanish regardless of K, and the normal equations remain poorly conditioned. Our experiments show that solving PGR and AGR requires artificial regularization like Tikhonov damping or priors like zero initialization in iterative solvers. We provide detailed discussions in Section 3 of the supplementary material.

In contrast, our method with $u\in\{x,y,z\}$ yields a spectrum whose minimum singular value is several orders of magnitude larger than that of AGR, leading to a dramatically smaller condition number. As an additional evidence to our analysis, if we keep $K=3$ but replace one degree-1 polynomial by $u=1$, the tangential nullspace

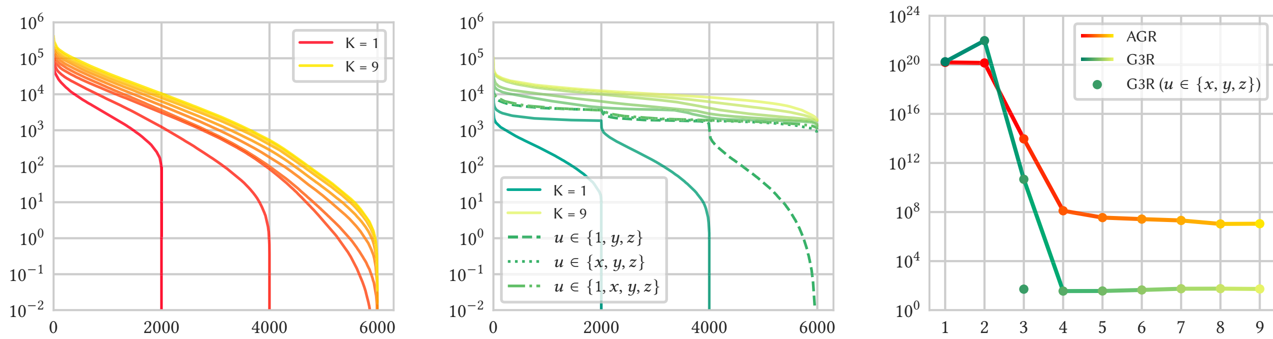

(a) Singular value distributions of AGR (b) Singular value distributions of G3R (c) Condition number vs. number of equations

Fig. 3. **Spectral properties of the normal matrices $\sum A^\top A$ for AGR and G3R.** The $K = 1$ case of both methods corresponds to PGR. (a) AGR remains ill-conditioned as K increases, with many near-zero singular values. (b) G3R lifts the spectrum tail once the tangential nullspace is removed. (c) Accordingly, G3R quickly attains a small condition number, while AGR stays highly ill-conditioned. For G3R, K counts the first K real solid harmonics ordered, e.g. $\{1\}$, $\{1, y\}$, $\{1, y, z\}$, $\{1, y, z, x\}$, . . ., and we additionally include the $u \in \{x, y, z\}$ case which is not part of this sequence. In our analysis we used $\{1, x, y, z\}$ to eliminate the tangential nullspace, while experiments show that using $\{x, y, z\}$ alone suffices. We include discussion on this in supplementary Section 2.

reappears and the spectrum tail collapses toward zero. As a result of eliminating the tangential nullspace, G3R produces well-conditioned normal equations, which can be solved robustly using direct solvers without regularization.

4.4 Mollifying the Fundamental Solution

As noted in Section 4.2, the fundamental solution $\Phi(y - x) = 1/(4\pi|y-x|)$ and its gradient $\nabla_y \Phi(y-x) = (x-y)/(4\pi|y-x|^3)$ are both singular when y approaches x. Most winding-number-based methods [Barill et al. 2018; Jacobson et al. 2013; Lin et al. 2022; Lu et al. 2018; Ma et al. 2025a] therefore rely on ad-hoc near-field treatments of the double layer kernel, typically by clipping or truncating contributions from a small neighborhood of $y = x$. This heuristic is often effective because the double layer kernel vanishes on the local tangent plane at $y = x$. In our formulation based on Green's third identity, however, the single layer term does not enjoy the same cancellation and generally contributes a non-negligible near-field component. Naively truncating it can therefore bias the integral and degrade numerical stability, which motivates a mollification of the fundamental solution. Figure 4 visualizes this qualitative difference.

Recently, Koneputugodage et al. [2025] proposed a bounded approximation of the point-cloud winding number to alleviate the near-field pathology of the standard formulation. However, this treatment makes the resulting equations nonlinear in the unknown area-weighted normals η. Thus rather than relying on near-field truncation or a nonlinear reformulation, we follow the approach proposed by Chen et al. [2024] and replace the standard fundamental solution with a mollified version:

$$\Phi_\epsilon(x) = \frac{\operatorname{erf}\left(\frac{|x|}{\sqrt{2}\epsilon}\right)}{4\pi|x|}.$$

The negative Laplacian of the standard fundamental solution is the Dirac delta function, while that of the mollified version is the Gaussian kernel with standard deviation ϵ and integral 1.

$$-\Delta\Phi(x) = \delta(x), \quad -\Delta\Phi_\epsilon(x) = \delta_\epsilon(x) = \frac{1}{(\sqrt{2\pi}\epsilon)^3} \exp\left(-\frac{|x|^2}{2\epsilon^2}\right).$$

This mollification smooths out the singularity, enabling stable numerical evaluation. For the corresponding mollified version of Green's third identity, we apply Green's second identity:

$$\begin{aligned}\hat{u}_\epsilon(x) &= \int_{\partial\Omega} \left(\Phi_\epsilon(y-x)\frac{\partial u}{\partial n}(y) - u(y)\frac{\partial \Phi_\epsilon}{\partial n}(y-x)\right) dS(y) \\ &= \int_\Omega u(y)\left(-\Delta_y \Phi_\epsilon(y-x)\right) dV(y) = \int_\Omega u(y)\delta_\epsilon(y-x) dV(y).\end{aligned}$$

Since the solid harmonics u_l^m we chose are harmonic in $\mathbb{R}^3$, the mean value property implies that it is invariant under Gaussian mollification. Moreover, when x lies on a smooth boundary $\partial\Omega$ and

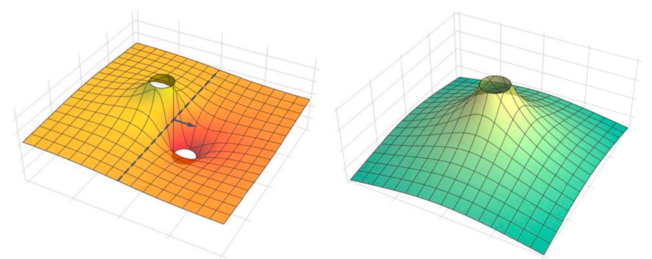

(a) Double layer kernel (b) Single layer kernel

Fig. 4. **Double vs. single layer kernels near the singularity.** For visual clarity we plot height fields with the 2D analogue of the fundamental solution Φ over the offset $(y - x)$. (a) The double layer kernel $-\partial\Phi/\partial n$ vanishes on the local tangent space (dashed; arrow shows the inward normal), making near-field truncation comparatively benign. (b) The single layer kernel Φ is positive and sharply peaked at $y = x$.

ϵ is sufficiently small, the radially symmetric kernel $\delta_\epsilon(y-x)$ has approximately half of its mass inside Ω and half outside. Thus, we can conclude that for x on $\partial\Omega$,

$$\hat{u}_\epsilon(x) \approx \frac{1}{2}\int_{\mathbb{R}^3} u(y)\delta_\epsilon(y-x)dV(y) = \frac{u(x)}{2}.$$

Therefore, our integral equation remains unchanged for the mollified version.

4.5 Octree-based Iterative Solver

While for moderate-sized point clouds, we can directly assemble and solve the normal equations, scaling to large point clouds with tens or hundreds of thousands of points poses significant computational and memory challenges. The normal equation results in a dense $3N \times 3N$ matrix. For point clouds with tens of thousands of points, storing this matrix alone exceeds memory capacity. Moreover, assembling and directly solving such matrices leads to cubic time complexity.

A common strategy to address this challenge is to use iterative solvers like steepest descent or conjugate gradient, which only require matrix-vector products. Crucially, the kernel in our formulation behaves similarly to the classical Gauss kernel: the interaction between distant points x and y decays rapidly. We therefore leverage well-established octree acceleration techniques for fast kernel summation [Barnes and Hut 1986], reducing the time complexity of each matrix-vector product from $O(N^2)$ to $O(N\log N)$.

5 Experiments

In this section, we first describe the experimental setup, including implementation details and evaluation metrics. Subsequently, we compare our method with state-of-the-art normal orientation approaches across various datasets. Finally, demonstrate the robustness of our method in handling challenging scenarios.

5.1 Experiment Setup

Implementation. We build the iterative version of our method upon the octree-based implementation of PGR by Lin et al. [2024], and implement both CPU and GPU versions of G3R. We provide detailed explainations for octree-based acceleration and pseudocode for the complete pipeline in the supplementary material. Our method can run on a moderate machine and handle point clouds with millions of points.

For noise-free inputs, we set the mollification parameter ϵ to be a value proportional to sampling density. Specifically, we set it to be 0.8 times the side length of the leaf node in the octree throughout the experiments. For noisy inputs, this heuristic is less effective; therefore, we determine ϵ empirically using fixed absolute values tailored to each noise level. We use $\{u_l^m \mid -l \le m \le l,\ 0 \le l \le 2\}$ as the set of harmonic functions, i.e., all nine harmonic polynomials of degree up to 2. For the iterative solver, we use 20 steps of steepest descent and 60 steps of conjugate gradient.

Metrics. We evaluate both normal orientation and mesh reconstruction quality. With the ground-truth normals $\{n_i^{\mathrm{gt}}\}$ and estimated normals $\{n_i\}$, we report the average point-cloud angular error and the percentage of correctly oriented normals

$$\mathrm{AE}_p = \frac{1}{N}\sum_{i=1}^{N}(1-n_i^{\mathrm{gt}}\cdot n_i)/2\,, \quad \mathrm{PCO} = \#\{\ i : n_i^{\mathrm{gt}}\cdot n_i > 0\ \}/N.$$

To assess reconstruction quality, we generate meshes from the oriented point clouds using SPSR [Kazhdan and Hoppe 2013] with a maximum depth of 10 for all methods. We measure the Chamfer distance and the symmetric Hausdorff distance between the ground-truth mesh $\mathcal{M}_{\mathrm{gt}}$ and the reconstructed mesh $\mathcal{M}$:

$$\mathrm{CD}(\mathcal{M}_{\mathrm{gt}},\mathcal{M}) = \frac{1}{2N}\left(\sum_{i=1}^{N}|x_i-\hat{x}_i| + \sum_{i=1}^{N}|y_i-\hat{y}_i|\right),$$

$$\mathrm{HD}(\mathcal{M}_{\mathrm{gt}},\mathcal{M}) = \max\left\{\sup_{x\in\mathcal{M}_{\mathrm{gt}}} d(x,\mathcal{M}),\ \sup_{y\in\mathcal{M}} d(y,\mathcal{M}_{\mathrm{gt}})\right\},$$

where $\{x_i\}$ is a uniform sampling of $\mathcal{M}_{\mathrm{gt}}$ (resp. $\{y_i\}$ for $\mathcal{M}$), and $\hat{x}_i, \hat{y}_i$ denote the closest point on the other mesh. We take 200K samples for N in the Chamfer distance computation, and use the code by Sacht and Jacobson [2024] to compute the Hausdorff distance. Finally, we report the mesh normal error AE_m on these closest-point pairs and averaged in both directions.

Table 1. Quantitative comparisons of our method with other state-of-the-art methods on normal orientation and mesh reconstruction from uniformly sampled point clouds with 50 K points. († indicates downsampling to 5 K due to excessive time requirements.) Best results are highlighted in bold. Our method achieves competitive performance on normal orientation (AE_p, PCO) and best or near-best performance on mesh reconstruction (CD, AE_m, HD). We additionally report reconstruction results from ground-truth normals as a reference (excluded from the boldface comparison).

Metric	Dataset	GCNO †	G3R †	PGR	iPSR	WNNC	AGR	G3R	GT
$\mathrm{AE}_p(10^{-2})\downarrow$	Famous	11.04	**9.98**	9.05	2.41	1.31	**1.30**	4.94	–
	ABC	8.81	**8.21**	6.85	3.93	**2.51**	2.55	4.71	–
	Thingi10K	6.33	**6.05**	5.01	2.31	**1.24**	1.38	3.00	–
$\mathrm{PCO}(\%)\uparrow$	Famous	92.53	**95.25**	95.92	99.36	99.63	**99.74**	98.54	–
	ABC	94.01	**96.45**	96.54	97.46	98.18	**98.24**	97.95	–
	Thingi10K	96.41	**97.75**	98.00	98.69	**99.25**	99.22	98.98	–
$\mathrm{CD}(10^{-4})\downarrow$	Famous	39.53	**19.20**	6.75	4.74	**3.27**	3.32	4.10	2.74
	ABC	36.19	**18.59**	5.95	4.93	3.66	**3.44**	3.57	3.91
	Thingi10K	38.70	**16.81**	5.51	10.74	**2.74**	2.87	2.96	3.22
$\mathrm{AE}_m(10^{-2})\downarrow$	Famous	14.42	**10.24**	7.90	6.57	6.09	**6.00**	6.40	5.70
	ABC	8.70	**6.25**	3.79	3.57	2.87	2.80	**2.80**	1.83
	Thingi10K	5.33	**3.93**	2.27	2.04	**1.50**	1.56	1.51	0.81
$\mathrm{HD}(10^{-4})\downarrow$	Famous	535.94	**355.87**	240.38	156.17	**138.14**	143.14	142.47	127.81
	ABC	581.78	**322.46**	220.44	203.99	148.59	143.35	**132.64**	132.04
	Thingi10K	605.60	**432.29**	243.56	287.58	**191.57**	198.82	192.46	204.95

5.2 Comparison with State-of-the-art Methods

We compare our method (G3R) with several state-of-the-art normal orientation methods including GCNO [Xu et al. 2023], PGR [Lin et al. 2022], iPSR [Hou et al. 2022], WNNC [Lin et al. 2024], and AGR [Ma et al. 2025a]. We use the official public implementations provided by the respective authors of each method, and use their recommended parameters for all experiments. Most methods are evaluated on uniformly sampled point clouds containing 50 K points. However, due to the high computational and memory demands of GCNO, we downsample the inputs to 5 K points for comparisons involving this method.

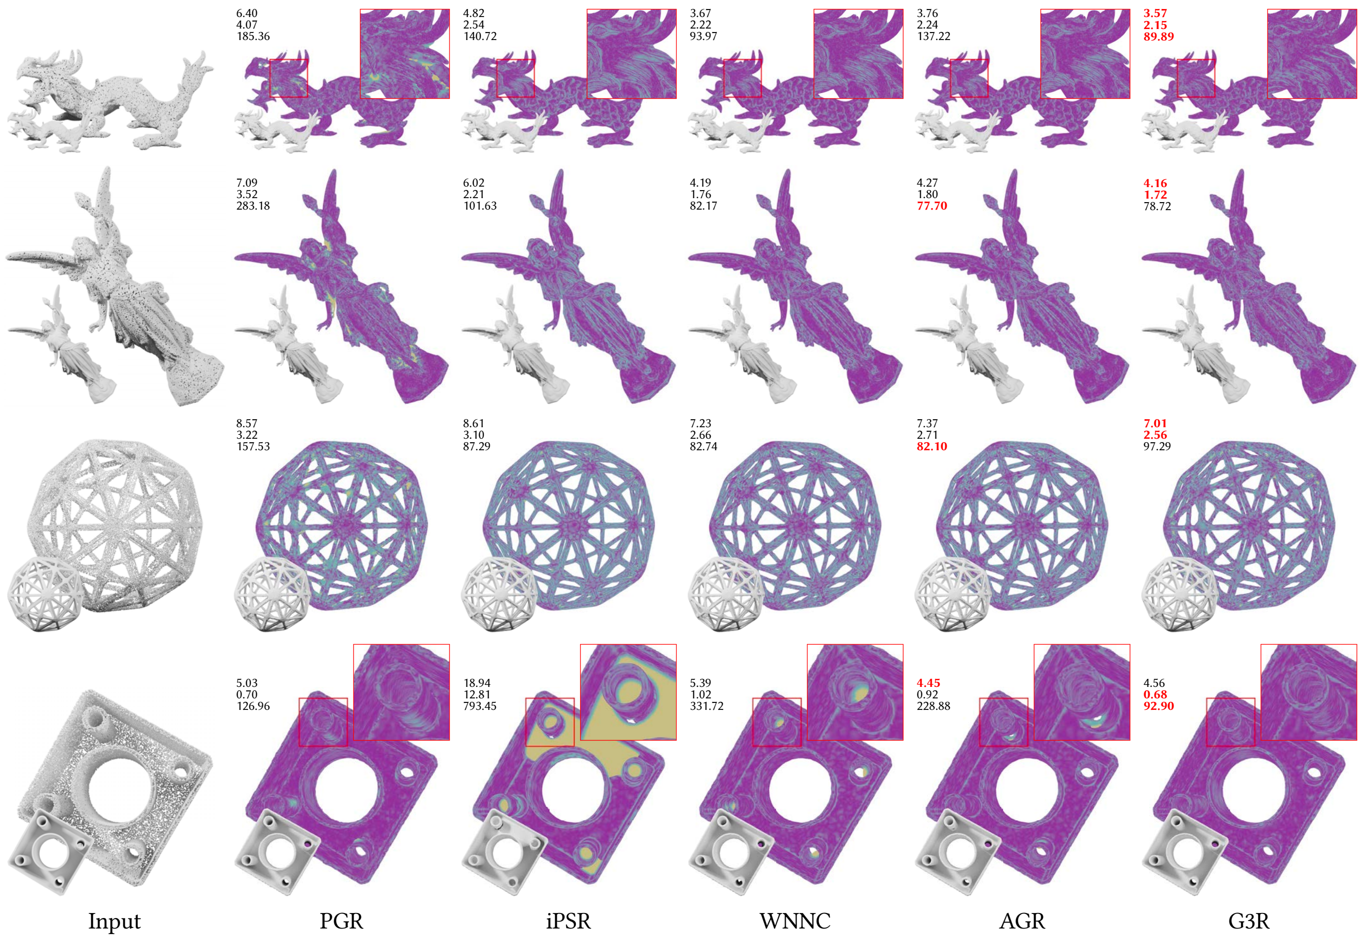

Fig. 5. Qualitative comparisons of our method (G3R) with other state-of-the-art methods on mesh reconstruction. Our method consistently produces high-quality reconstructions across diverse shapes. We report Chamfer distance ($\times 10^{-4}$), average angular error ($\times 10^{-2}$) and Hausdorff distance ($\times 10^{-4}$) between the reconstructed mesh and the ground-truth mesh. Best results are highlighted in bold and red.

For quantitative comparisons, we evaluate our method on several datasets, including 22 models from Famous [Erler et al. 2020], 100 models from ABC [Koch et al. 2019] and 100 models from Thingi10K [Zhou and Jacobson 2016], which are previously used by Ma et al. [2025a], and report average performance over each dataset. Additionally, we reconstruct surfaces using the ground-truth normals. These results serve as a reference for the reconstruction quality achievable under perfectly aligned normals. Table 1 summarizes the average metrics on normal orientation and mesh reconstruction. For normal orientation, our method consistently outperforms integral-equation-based methods, and achieves comparable performance with the state-of-the-art fixed-point-based methods WNNC and AGR. For mesh reconstruction quality, our method achieves state-of-the-art performance, either best or near-best in all metrics across all datasets.

We further provide visual comparisons that showcase the effectiveness of our method. Figure 5 shows visual comparisons of our method with other state-of-the-art methods. Our method consistently produces high-quality reconstructions across diverse shapes.

In particular, we find that our method excels in recovering thin structures. See Figure 6 for examples. Fixed-point iterative methods, such as WNNC and iPSR, often struggle in these scenarios, as they are prone to becoming trapped in local minima with misaligned normals. In contrast, directly solving the integral equation provides a more global perspective that is less sensitive to local ambiguities.

Noisy Inputs. We further compare our method with other state-of-the-art methods on synthetic noisy inputs. We add Gaussian noise with standard deviation 0.5% of the bounding box diagonal to each point in the uniformly sampled point clouds with 50 K points. We set the mollification parameter ϵ to 0.04 for all models. Figure 7 shows that our method maintains the ability to recover the thin structures under noise, whereas other methods struggle. However, regarding average metrics across the entire dataset, our method performs slightly below fixed-point-based methods. Detailed quantitative results for noisy inputs are provided in the supplementary material.

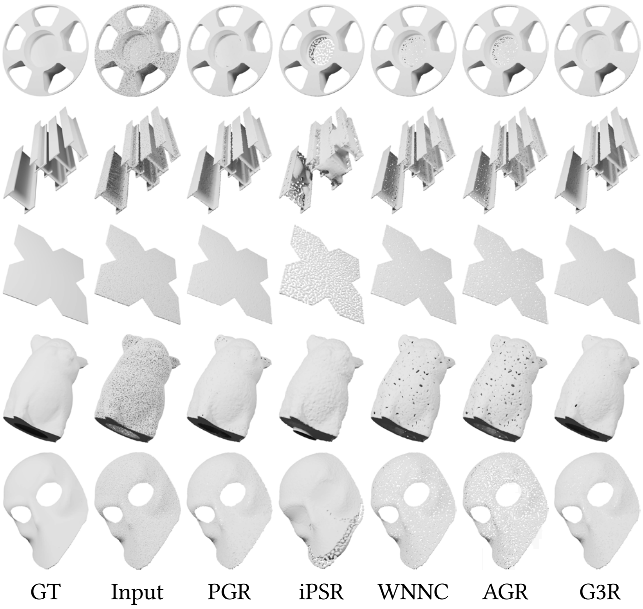

Fig. 6. Qualitative comparison on models with thin structures. Result shows that methods solving a single integral equation system (PGR and our G3R) can successfully recover these features. In contrast, methods employing a fixed-point iteration scheme (iPSR, WNNC, AGR) struggle to do so.

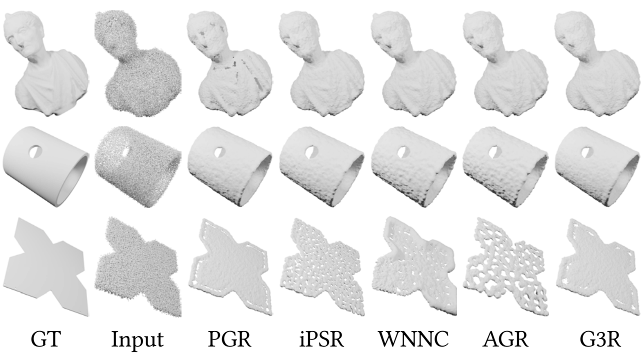

Fig. 7. Qualitative comparison on point clouds with Gaussian noise (std = 0.5% of bounding box diagonal). 50 K points are sampled from each model.

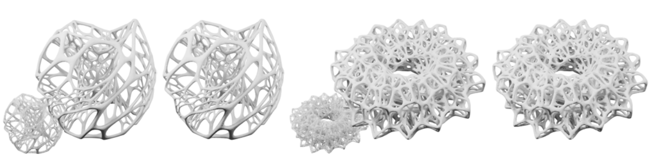

Fig. 8. Reconstruction results on point clouds with highly complex topologies. Left to right: input points, ground-truth mesh, reconstructed mesh.

5.3 Challenging Cases

Highly Complex Topology. We experiment on point clouds with highly complex topologies in Figure 8. We sampled 80 K and 100 K points respectively. Our method successfully orients the normals and captures the complex structures for both models.

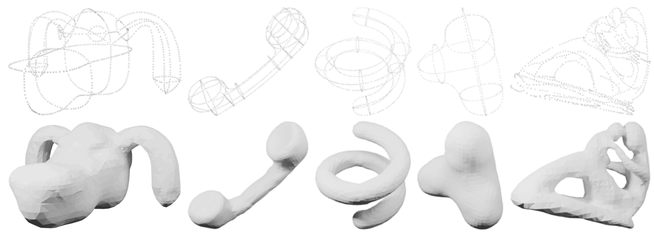

Fig. 9. Reconstruction results on wireframe point clouds.

Fig. 10. Reconstruction results on the Stanford raw range scans. Disconnected redundant components in reconstructed mesh caused by outliers are removed manually for better visualization.

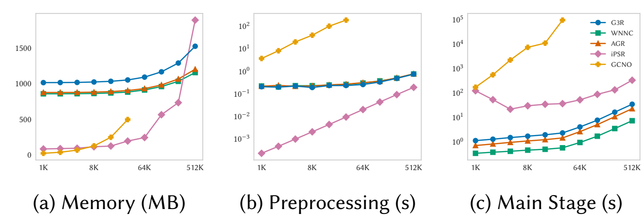

Fig. 11. Memory usage and runtime comparisons across different methods.

Wireframe Sampling. 3D wireframe point clouds, typically generated by VR/AR sketching tools, pose significant challenge due to their low sampling rates. We test our method on wireframe point clouds from Huang et al. [2019], each consisting of around 1 K points. In Figure 9 we show that our method can robustly orient normals and reconstruct surfaces.

Real-world Scans. We present two results from raw range scans from the Stanford 3D Scanning Repository [Krishnamurthy and Levoy 1996; Turk and Levoy 1994]. The dataset comes from aligning multiple range scans, and thus contain noise, imperfect alignment and outliers. We directly apply our method on the raw point clouds without preprocessing. The Armadillo scan contains 2,246,582 points, while the Bunny scan contains 362,272 points. We set the mollification parameter ϵ to 0.01 for both models. As shown in Figure 10, our method successfully orients normals and reconstructs high-quality meshes from large-scale inputs with real-world imperfections.

5.4 Memory and Time Usage Comparisons

In Figure 11, we compare the peak memory usage (CPU + GPU), preprocessing time, and main-stage runtime of our method with

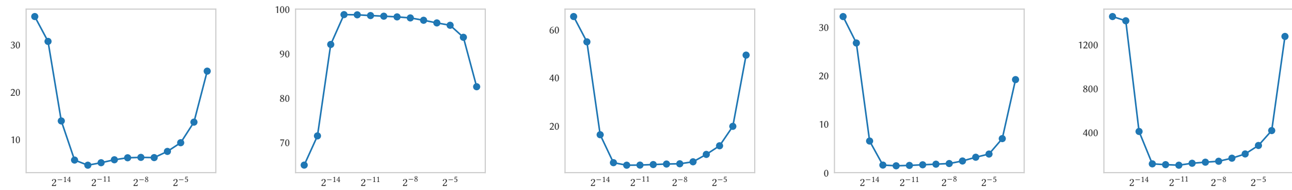

(a) $\mathrm{AE}_p\,(10^{-2}) \downarrow$ (b) $\mathrm{PCO}(\%) \uparrow$ (c) $\mathrm{CD}\,(10^{-4}) \downarrow$ (d) $\mathrm{AE}_m\,(10^{-2}) \downarrow$ (e) $\mathrm{HD}\,(10^{-4}) \downarrow$

Fig. 12. Ablation study on the mollifier parameter ϵ on 50 K points sampled from the 3DBenchy model.

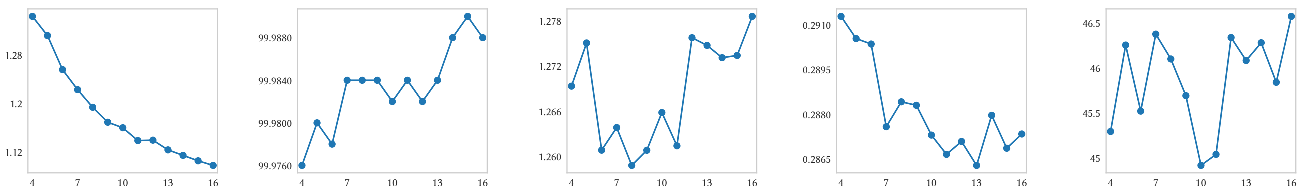

(a) $\mathrm{AE}_p\,(10^{-2}) \downarrow$ (b) $\mathrm{PCO}(\%) \uparrow$ (c) $\mathrm{CD}\,(10^{-4}) \downarrow$ (d) $\mathrm{AE}_m\,(10^{-2}) \downarrow$ (e) $\mathrm{HD}\,(10^{-4}) \downarrow$

Fig. 13. Ablation study on the number of harmonic functions on 50 K points sampled from the Stanford Bunny model.

other state-of-the-art methods. For our method, WNNC, AGR, and iPSR, preprocessing corresponds to octree construction, while the main stage solves for normals. For GCNO, preprocessing corresponds to Delaunay triangulation, and the main stage corresponds to the L-BFGS optimization. For all methods, file I/O and the final reconstruction by SPSR are excluded.

The experiment is conducted on point clouds uniformly sampled from the Armadillo model with varying sizes. G3R, WNNC, and AGR are GPU-accelerated, whereas GCNO and iPSR are CPU-only. All experiments are performed on a machine with an Intel Xeon Gold 6530 CPU and an NVIDIA GeForce RTX 4090 GPU.

Overall, our method is slightly more expensive than WNNC and AGR, due to the additional constraints it introduces and the cost of evaluating the Green's third-identity operator. Nevertheless, it remains substantially more efficient than GCNO and iPSR, especially for larger inputs.

5.5 Ablation Study

Figure 12 shows the effect of the mollification parameter ϵ on normal orientation accuracy and mesh reconstruction quality, with ϵ ranging from 2^{-16} to 2^{-3}. The results show that $\epsilon = 2^{-12}$ provides the best overall performance across the evaluated metrics, while values from 2^{-14} to 2^{-8} still yield acceptable results.

Figure 13 shows the effect of increasing the number of harmonic functions. We start from the first four regular solid harmonics (i.e. $u \in \{1, x, y, z\}$) and then add degree-2 and degree-3 harmonics one by one. The results show that the normal metrics (AE_p, PCO, AE_m) improve gradually as more functions are included. Meanwhile, Chamfer and Hausdorff distances vary only marginally, as they are already very low. Harmonic polynomials up to degree 2 (i.e. 9 functions in total) provides the best overall trade-off. We therefore use this configuration as the default.

6 Conclusion

In this work, we reviewed the winding-number-based normal orientation methods from the perspective of boundary integral equations, and identified the tangential nullspace problem in prior formulations. To address this problem, we proposed a novel formulation that leverages multiple harmonic functions to eliminate the tangential nullspace, resulting in well-conditioned linear systems that can be solved directly. Experiments demonstrated that our method achieves state-of-the-art performance on surface reconstruction tasks.

Limitations and Future Work. We have shown that the tangential nullspace problem can be resolved by introducing multiple harmonic functions and posing constraints on the boundary only, it remains to be seen whether the nullspace within the normal direction is zero. Numerical experiments suggest that this is indeed the case, but a rigorous proof is still lacking.

Our formulation still follows that of PGR, where constraints are only enforced at the input points (i.e. on the boundary), which makes it vulnerable to noise and outliers. In future work, combining G3R with fixed-point iterative schemes could further improve robustness to noise. Another direction is to model noise and outliers explicitly with stochastic differential equations to enhance the resilience of normal orientation. It would also be interesting to explore the use of volumetric constraints to further enhance robustness.

Acknowledgments

We thank the anonymous reviewers for their valuable comments. This research was partially funded by the National Natural Science Foundation of China under Grant No.: 12494552 and the Natural Science Foundation of Anhui Province under Grant No.: 2508085MA001.

References

Gavin Barill, Neil G. Dickson, Ryan Schmidt, David I. W. Levin, and Alec Jacobson. 2018. Fast winding numbers for soups and clouds. *ACM Transactions on Graphics* 37, 4, Article 43 (July 2018), 12 pages. doi:10.1145/3197517.3201337

Josh Barnes and Piet Hut. 1986. A hierarchical O(N log N) force-calculation algorithm. *Nature* 324, 6096 (01 Dec 1986), 446–449. doi:10.1038/324446a0

Yizhak Ben-Shabat, Chamin Hewa Koneputugodage, and Stephen Gould. 2022. DiGS: Divergence guided shape implicit neural representation for unoriented point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, Los Alamitos, CA, USA, 19301–19310. doi:10.1109/CVPR52688.2022.01872

Hanyu Chen, Bailey Miller, and Ioannis Gkioulekas. 2024. 3D reconstruction with fast dipole sums. *ACM Transactions on Graphics* 43, 6, Article 192 (Nov. 2024), 19 pages. doi:10.1145/3687914

Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J. Mitra, and Michael Wimmer. 2020. Points2Surf: learning implicit surfaces from point clouds. In *Computer Vision – ECCV 2020*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer International Publishing, Cham, 108–124. doi:10.1007/978-3-030-58558-7_7

Rao Fu, Jianmin Zheng, and Liang Yu. 2025. Consistent normal orientation for 3D point clouds via least squares on Delaunay graph. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16932–16942. doi:10.1109/CVPR52734.2025.01578

Craig Gotsman and Kai Hormann. 2024. A linear method to consistently orient normals of a 3D point cloud. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) *(SIGGRAPH '24)*. Association for Computing Machinery, New York, NY, USA, Article 55, 10 pages. doi:10.1145/3641519.3657429

Fei Hou, Chiyu Wang, Wencheng Wang, Hong Qin, Chen Qian, and Ying He. 2022. Iterative Poisson surface reconstruction (iPSR) for unoriented points. *ACM Transactions on Graphics* 41, 4, Article 128 (July 2022), 13 pages. doi:10.1145/3528223.3530096

Guojin Huang, Qing Fang, Zheng Zhang, Ligang Liu, and Xiao-Ming Fu. 2024. Stochastic normal orientation for point clouds. *ACM Transactions on Graphics* 43, 6, Article 188 (Nov. 2024), 12 pages. doi:10.1145/3687944

Zhiyang Huang, Nathan Carr, and Tao Ju. 2019. Variational implicit point set surfaces. *ACM Transactions on Graphics* 38, 4, Article 124 (July 2019), 13 pages. doi:10.1145/3306346.3322994

Alec Jacobson, Ladislav Kavan, and Olga Sorkine-Hornung. 2013. Robust inside-outside segmentation using generalized winding numbers. *ACM Transactions on Graphics* 32, 4, Article 33 (July 2013), 12 pages. doi:10.1145/2461912.2461916

Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Symposium on Geometry Processing*, Alla Sheffer and Konrad Polthier (Eds.). The Eurographics Association, 61–70. doi:10.2312/SGP/SGP06/061-070

Michael Kazhdan and Hugues Hoppe. 2013. Screened Poisson surface reconstruction. *ACM Transactions on Graphics* 32, 3, Article 29 (July 2013), 13 pages. doi:10.1145/2487228.2487237

Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. 2019. ABC: a big CAD model dataset for geometric deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9601–9611.

Maximilian Kohlbrenner, Hongyi Liu, Marc Alexa, and Michael Kazhdan. 2025. Symmetrized Poisson reconstruction. *Computer Graphics Forum* 44, 5 (2025), e70210. doi:10.1111/cgf.70210

Chamin Hewa Koneputugodage, Dylan Campbell, and Stephen Gould. 2025. Leaps and bounds: an improved point cloud winding number formulation for fast normal estimation and surface reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 26116–26125.

Rainer Kress. 2014. *Linear integral equations*. Springer New York, New York, NY. doi:10.1007/978-1-4614-9593-2

Venkat Krishnamurthy and Marc Levoy. 1996. Fitting smooth surfaces to dense polygon meshes. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. Association for Computing Machinery, New York, NY, USA, 313–324. doi:10.1145/237170.237270

John M. Lee. 2012. Integration on Manifolds. In *Introduction to Smooth Manifolds*. Springer New York, New York, NY, 400–439. doi:10.1007/978-1-4419-9982-5_16

Siyou Lin, Zuoqiang Shi, and Yebin Liu. 2024. Fast and globally consistent normal orientation based on the winding number normal consistency. *ACM Transactions on Graphics* 43, 6, Article 189 (Nov. 2024), 19 pages. doi:10.1145/3687895

Siyou Lin, Dong Xiao, Zuoqiang Shi, and Bin Wang. 2022. Surface reconstruction from point clouds without normals by parametrizing the Gauss formula. *ACM Transactions on Graphics* 42, 2, Article 14 (Oct. 2022), 19 pages. doi:10.1145/3554730

Weizhou Liu, Jiaze Li, Xuhui Chen, Fei Hou, Shiqing Xin, Xingce Wang, Zhongke Wu, Chen Qian, and Ying He. 2025. Diffusing winding gradients (DWG): a parallel and scalable method for 3D reconstruction from unoriented point clouds. *ACM Transactions on Graphics* 44, 2, Article 19 (April 2025), 18 pages. doi:10.1145/3727873

Wenjia Lu, Zuoqiang Shi, Jian Sun, and Bin Wang. 2018. Surface reconstruction based on the modified Gauss formula. *ACM Transactions on Graphics* 38, 1, Article 2 (Dec. 2018), 18 pages. doi:10.1145/3233984

Baorui Ma, Yu-Shen Liu, Matthias Zwicker, and Zhizhong Han. 2022. Surface reconstruction from point clouds by learning predictive context priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6316–6327. doi:10.1109/CVPR52688.2022.00622

Yueji Ma, Jialu Shen, Yanzun Meng, Dong Xiao, Zuoqiang Shi, and Bin Wang. 2025a. Anisotropic Gauss reconstruction and global orientation with octree-based acceleration. *Computer Graphics Forum* 44, 5 (2025), e70199. doi:10.1111/cgf.70199

Yueji Ma, Dong Xiao, Zuoqiang Shi, and Bin Wang. 2025b. Convection augmented Gauss reconstruction for unoriented point clouds. *ACM Transactions on Graphics* 44, 5, Article 173 (Sept. 2025), 18 pages. doi:10.1145/3750723

Gal Metzer, Rana Hanocka, Denis Zorin, Raja Giryes, Daniele Panozzo, and Daniel Cohen-Or. 2021. Orienting point clouds with dipole propagation. *ACM Transactions on Graphics* 40, 4, Article 165 (July 2021), 14 pages. doi:10.1145/3450626.3459835

Leonardo Sacht and Alec Jacobson. 2024. Cascading upper bounds for triangle soup Pompeiu-Hausdorff distance. *Computer Graphics Forum* 43, 5 (2024), e15129. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.15129 doi:10.1111/cgf.15129

Daniel Scrivener, Daniel Cui, Ellis Coldren, S. Mazdak Abulnaga, Mikhail Bessmeltsev, and Edward Chien. 2025. Faraday cage estimation of normals for point clouds and ribbon sketches. *ACM Transactions on Graphics* 44, 4, Article 49 (July 2025), 13 pages. doi:10.1145/3731212

Silvia Sellán and Alec Jacobson. 2022. Stochastic Poisson surface reconstruction. *ACM Transactions on Graphics* 41, 6, Article 227 (Nov. 2022), 12 pages. doi:10.1145/3550454.3555441

Olaf Steinbach. 2008. Fundamental Solutions. In *Numerical Approximation Methods for Elliptic Boundary Value Problems: Finite and Boundary Elements*. Springer New York, New York, NY, 89–109. doi:10.1007/978-0-387-68805-3_5

Greg Turk and Marc Levoy. 1994. Zippered polygon meshes from range images. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. Association for Computing Machinery, New York, NY, USA, 311–318. doi:10.1145/192161.192241

Zixiong Wang, Yunxiao Zhang, Rui Xu, Fan Zhang, Peng-Shuai Wang, Shuangmin Chen, Shiqing Xin, Wenping Wang, and Changhe Tu. 2023. Neural-Singular-Hessian: implicit neural representation of unoriented point clouds by enforcing singular Hessian. *ACM Transactions on Graphics* 42, 6, Article 274 (Dec. 2023), 14 pages. doi:10.1145/3618311

Jianjun Xia and Tao Ju. 2025. Variational surface reconstruction using natural neighbors. *ACM Transactions on Graphics* 44, 4, Article 85 (July 2025), 19 pages. doi:10.1145/3731191

Dong Xiao, Zuoqiang Shi, Siyu Li, Bailin Deng, and Bin Wang. 2023. Point normal orientation and surface reconstruction by incorporating isovalue constraints to Poisson equation. *Computer Aided Geometric Design* 103 (2023), 102195. doi:10.1016/j.cagd.2023.102195

Rui Xu, Zhiyang Dou, Ningna Wang, Shiqing Xin, Shuangmin Chen, Mingyan Jiang, Xiaohu Guo, Wenping Wang, and Changhe Tu. 2023. Globally consistent normal orientation for point clouds by regularizing the winding-number field. *ACM Transactions on Graphics* 42, 4, Article 111 (July 2023), 15 pages. doi:10.1145/3592129

Qingnan Zhou and Alec Jacobson. 2016. Thingi10K: a dataset of 10,000 3D-printing models. arXiv:1605.04797