

HotSpot: Signed Distance Function Optimization with an Asymptotically Sufficient Condition

Zimo Wang* Cheng Wang* Taiki Yoshino Sirui Tao Ziyang Fu Tzu-Mao Li
 UC San Diego




Figure 1. We propose HOTSPOt, a neural signed distance function optimization method that establishes an asymptotic sufficient condition to guarantee convergence to a true distance function, enabling precise surface reconstruction and level set representation for complex shapes. Here we show a reconstruction from a point cloud sampled from the reference bunny (taken from Mehta et al. [1]) on the right. In the inset, we visualize the recovered signed distance function on a horizontal slice, using warm colors for positive values and cool for negative (zoom in for details). Our reconstruction is significantly more accurate than prior works (SAL [2], DiGS [3], and StEik [4]).

Abstract

We propose a method, HOTSPOt, for optimizing neural signed distance functions. Existing losses, such as the eikonal loss, act as necessary but insufficient constraints and cannot guarantee that the recovered implicit function represents a true distance function, even if the output minimizes these losses almost everywhere. Furthermore, the eikonal loss suffers from stability issues in optimization. Finally, in conventional methods, regularization losses that penalize surface area distort the reconstructed signed distance function. We address these challenges by designing a loss function using the solution of a screened Poisson equation. Our loss, when minimized, provides an asymptotically sufficient condition to ensure the output converges to a true distance function. Our loss also leads to stable optimization and naturally penalizes large surface areas. We present theoretical analysis and experiments on both challenging 2D and 3D datasets and show that our method provides better surface reconstruction and a more accurate distance approximation.

1. Introduction

Learning neural signed distance function requires designing loss functions so that the zero crossings of the implicit

function are at the desired locations (e.g., adhere to a sparse input point cloud), and that the implicit function returns the correct signed distance. Designing the loss function proves to be challenging for shapes with complex topology and details (Fig. 1). We propose HOTSPOt, a method based on a relation between the screened Poisson equation and distance that enjoys both theoretical and practical benefits.

A major challenge in signed distance function optimization is ensuring the implicit function follows the true distance. A common regularization loss used is the eikonal equation, which constrains the gradient norm of an implicit function to be 1 almost everywhere. If the implicit function is a signed distance function, then it satisfies the eikonal equation. However, the converse is not true. That is, satisfying the eikonal equation is only a *necessary condition* for the signed distance function, not a *sufficient condition* [5]. Fig. 2 presents the logic and Fig. 3 shows an example.

Another key challenge lies in the optimization. It is known that the eikonal equation regularization is *unstable* [4], and the instability can make optimization converge to suboptimal results or even diverge.

Furthermore, to address optimization ill-posedness and instability, existing methods penalize the surface area of the shape. We prove this degrades the reconstruction quality and distorts the output level sets, thereby leading to a delicate balance between over- and under-smoothing.

*Equal contribution.

$$\min_{\lambda \rightarrow \infty} \{\text{heat loss}\} \Leftrightarrow \begin{cases} \text{Distance Function} \\ \text{a.e.} \\ \Rightarrow \\ \text{Eikonal condition} \\ \text{Zero gradient divergence} \\ \text{Singular Hessian matrix} \\ \dots \end{cases}$$

Figure 2. Existing constraints are only **necessary** but **insufficient** conditions. That is, there exist solutions, like Fig. 3, that satisfy them almost everywhere yet fail to be distance functions. In contrast, we use a heat loss, modeled by a screened Poisson equation with parameter λ . As the parameter λ goes to infinity, the minimizer converges to the distance function and excludes other solutions. Thus, our method is asymptotically both necessary and sufficient.

We propose a simple model that can address all the challenges above. Our model, HOTSPOT, is based on a screened Poisson equation and a classical relation between heat transfer and distance [6] that was popular for approximating both Euclidean and geodesic distance [7–9]. Using this relation, we design a loss function for optimizing neural signed distance functions. As our contributions, we show, both theoretically and empirically, that our model

- serves as a sufficient condition ensuring the difference between the minimizer and the actual distance is bounded and decreases linearly as the parameter λ increases (Fig. 2 and Fig. 3, right).
- is *stable* both spatially and temporally, in the sense that a slight perturbation of the implicit function only introduces changes in a local region, and that the dynamical system formed by the gradient flow will converge in the long run.
- penalizes the surface area naturally, while remaining an undistorted signed distance function.
- works well in both 2D and 3D shapes, and it excels at approaching signed distance functions while optimizing for complex, high-genus shapes.

2. Related Work

Implicit surfaces. Modeling and reconstructing surfaces using implicit functions has a long history in computer vision and graphics [11–15]. Recently, there is a surge of interest to model surfaces using *neural* implicit functions [16–20]. Neural implicit representations are favored for their compactness, ability to represent detailed surfaces, and compatibility to gradient-based optimization [1].

Neural signed distance functions. Among the classes of implicit surfaces, *signed distance functions*, where the implicit function outputs the distance to the closest point on the surface, and the sign represents whether the point is inside or outside the solid object, is one of the most popular variants. Signed distance functions enable an efficient rendering algorithm called sphere tracing [21–24] and helps geometry compression [25]. They can also be used for collision detection [26–28] and numerical simulation [29]. Traditionally, signed distance functions are often represented

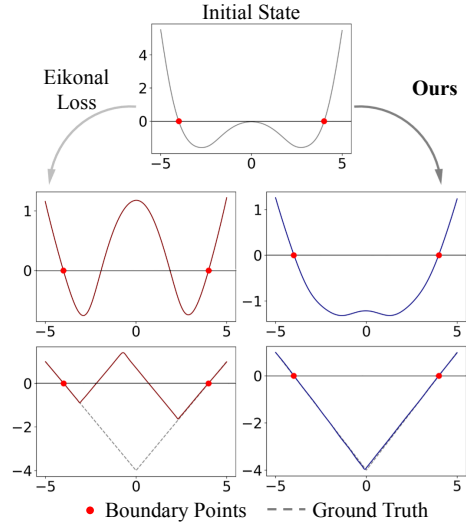


Figure 3. We illustrate a 1D neural signed distance function optimization using the classical eikonal loss [10] and our model. The x axis is the domain and the y axis shows the output of the implicit function. The middle and bottom rows show the intermediate and final states of the optimization respectively. The eikonal loss, as a necessary but insufficient condition, is incapable of converging to the actual signed distance function (dashed line), even when the function satisfies the eikonal equation almost everywhere.

as grids [30–32]. More recently, neural signed distance functions have become an appealing geometric representation [33–35]. (Neural) Signed distance functions are also often used in inverse rendering [36–42], for their flexibility to obtain surfaces with complex topology.

Regularizing and initializing neural signed distance functions. A commonly used loss for optimizing neural signed distance functions is the eikonal regularization, which encourages the norm of the gradient of the implicit function to be one [10, 39]. However, the eikonal regularization alone is far from sufficient, as 1) minimizing eikonal loss does not necessarily lead to a distance field [5, 30], and 2) its optimization is unstable [4]. Previous works have proposed initialization to bootstrap the optimization with a simple shape [3, 10]. Recent works show that minimizing the divergence and the surface area of the implicit function can lead to better reconstruction [3, 43] and stabilizes the optimization [4]. Second-order information is also often used [43–46] for regularization. All losses above are based on *necessary* properties of real distance but neither establish convergence nor exclude non-SDF solutions.

Specialized architectures have been proposed to limit the Lipschitz constant of the implicit function [35], satisfy p -Poisson equation [47], or represent high-frequency details [25, 34, 48]. Some research [49, 50] has shown a link between occupancy function optimization and an approximate signed distance field regularization. Our experiments further reveal that some implementations [2, 51, 52]

optimizing SDFs with the closest point may struggle when querying near the interpolated surface.

We propose a new loss and method for optimizing signed distance functions. While derived from very different mathematics, our loss turns out to have a close connection to Lipman’s PHASE loss [49]. We show that our theory leads to significantly more numerically stable optimization, and a different conclusion in the parameter setting that leads to drastically different results.

Distance approximation using heat transfer. Our derivation is inspired by a relation between heat transfer and approximated distance [6]. This relation has been used in computer graphics for efficiently computing the geodesic distance on surfaces [7–9, 53, 54]. We apply the relation for optimizing neural signed distance functions, and analyze the properties of the resulting optimization in this context.

3. Background and Motivation

Given an input point cloud $\mathbf{x}_i \in \mathbb{R}^d$ (we focus on 2D and 3D in this work) with N points without normal information that lies on an unknown surface Γ of a solid object S , our goal is to find a signed distance function $f(\mathbf{x})$, such that

$$f(\mathbf{x}) = \begin{cases} -d_\Gamma(\mathbf{x}) & \text{if } \mathbf{x} \in S \\ d_\Gamma(\mathbf{x}) & \text{otherwise,} \end{cases} \quad (1)$$

where $d_\Gamma(\mathbf{x})$ denotes the closest distance from point \mathbf{x} to the surface Γ . In practice, we approximate the signed distance function $f(\mathbf{x})$ using a neural network $u(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ parameterized by its weights and use gradient-based optimization to figure out the weights. Previous work (e.g., [10]) often uses a linear combination of a boundary loss L_{boundary} and a eikonal loss L_{eikonal} :

$$L_{\text{boundary}} = \frac{1}{N} \sum_{i=1}^N |u(\mathbf{x}_i)|^p, \quad (2)$$

$$L_{\text{eikonal}} = \int_{\Omega} \|\|\nabla u(\mathbf{x})\| - 1\|^p \, d\mathbf{x}, \quad p = 1 \text{ or } 2, \quad (3)$$

where $\Omega \subset \mathbb{R}^d$ is our domain of interest. The boundary loss encourages the implicit function u to output 0 at the boundary, and the eikonal loss is based on that observation that if a function f is a distance function, the L^2 norm of its gradient $\|\nabla f(\mathbf{x})\| = 1$ holds almost everywhere. Thus the approximation u is encouraged to behave the same.

However, as noted, simply optimizing for the two terms above does not guarantee that u would be a distance function even when satisfied almost everywhere (Figs. 2 and 3), and the extra regularization previous work added to minimize the area and divergence [3, 4] could further tamper with the distance approximation quality. We derive an additional loss that can be used together with both of the losses above, and analyze our loss theoretically and empirically.

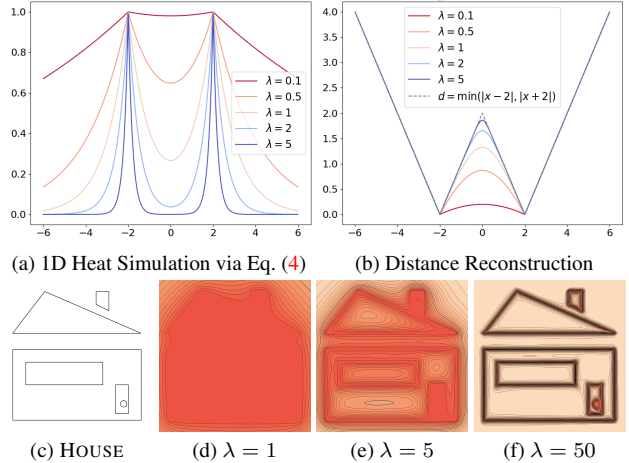


Figure 4. We show an illustration of the relation between screened Poisson equation and a distance field. The top shows a 1D example, where on the left we show a solution to Eq. (4) with different absorption λ , and on the right we show the reconstructed distance $-\frac{1}{\lambda} \ln(h)$ (Eq. (5)). Boundary loss for real boundary points acts as an isothermal heat source, while heat loss diffuses the heat. When λ is increased, the heat decays faster, and the error of the reconstructed function approaches 0. The bottom shows a 2D example.

We will be using a key relation between the screen Poisson equation and the distance. We first define a screened Poisson equation over a *heat field* $h(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$, with an absorption coefficient λ , and set a Dirichlet boundary condition to be 1 on our target surface Γ :

$$\begin{cases} \nabla^2 h(\mathbf{x}) - \lambda^2 h(\mathbf{x}) = 0 & \forall \mathbf{x} \in \mathbb{R}^d \setminus \Gamma \\ \lim_{\|\mathbf{x}\| \rightarrow \infty} h(\mathbf{x}) = 0, \quad h(\mathbf{x}) = 1 & \forall \mathbf{x} \in \Gamma. \end{cases} \quad (4)$$

If a heat field h_λ is a solution to the partial differential equation above, then the following equation holds [6]:

$$\lim_{\lambda \rightarrow \infty} \frac{1}{\lambda} \ln(h_\lambda(\mathbf{x})) = -d_\Gamma(\mathbf{x}). \quad (5)$$

Fig. 4 shows the intuition. The screened Poisson equation simulates heat transfer until an equilibrium is reached, under absorption controlled by the λ coefficient. As we increase the absorption, the reconstructed distance $\frac{1}{\lambda} \ln(h_\lambda)$ would approach the real distance. In practice, we could only obtain a discrete point set Γ_0 . Provided that the neighbor points are close enough, if there are some boundaries connecting them, the boundary condition of Eq. (4) still holds. This can be interpreted as a superset Γ replacing Γ_0 , effectively transforming the boundary from a discrete point cloud into a continuous surface, which is required and precisely the goal of the signed distance function reconstruction. Obtaining a distance function to the point cloud rather than the interpolated surface is not our objective. The reconstructed distance can still faithfully represent the distance to Γ because the approximation in Eq. (5) still holds. Later, we will prove that the measure of Γ is bounded by our loss.

4. Method

In this section, we introduce our method, HOTSPOt, based on the screened Poisson equation (Eq. (4)) and derive a theoretical analysis. We use the screened Poisson equation to derive a loss we can directly apply to an implicit function u , and prove that there is a bound to the solution of the screened Poisson equation to the true distance (Section 4.1). We further derive a stability analysis of our loss on both the spatial stability over small perturbations and the temporal convergence (Section 4.2). Next, we show that our loss can penalize surfaces with large area (Section 4.3). Finally, we discuss the relation to prior work, PHASE (Section 4.4).

4.1. Screened Poisson Equation Informed Signed Distance Functions Optimization

Given an implicit function u , we want to derive a loss based on the screened Poisson equation (Eq. (4)) and its relation to the distance (Eq. (5)). We achieve this by applying the following substitution based on Eq. (5):

$$h(\mathbf{x}) = e^{-\lambda|u(\mathbf{x})|}. \quad (6)$$

We want to design a loss such that its minimizer satisfies the screened Poisson equation. This can be achieved by minimizing the energy functional $\frac{1}{2} \int_{\Omega} \|\nabla h(\mathbf{x})\|^2 + \lambda^2 h(\mathbf{x})^2 d\mathbf{x}$. After substitution and removing of the constant factor λ , we obtain our heat loss:

$$L_{\text{heat}} = \frac{1}{2} \int_{\Omega} e^{-2\lambda|u(\mathbf{x})|} (\|\nabla u(\mathbf{x})\|^2 + 1) d\mathbf{x}. \quad (7)$$

If we take the derivative of L_{heat} with respect to the heat field h , we recover the screened Poisson equation $\nabla^2 h - \lambda^2 h = 0$. The boundary condition can be enforced by the standard boundary loss (Eq. (2)), by observing that $u(\mathbf{x}) = 0$ implies $h(\mathbf{x}) = 1$. Together with the eikonal loss, we optimize:

$$L = w_b L_{\text{boundary}} + w_e L_{\text{eikonal}} + w_h L_{\text{heat}}. \quad (8)$$

While prior work [6] has studied the limiting behavior of the absorption coefficient λ , to our knowledge, the bounds of the approximation and the convergence speed remain unknown. In the supplementary material A.3, we prove that, for a boundary Γ formed by a discrete set of points and a small volume around them, the distance approximation obtained from the solution of the screened Poisson equation converges linearly with respect to λ :

$$\frac{1}{\lambda} C_{\text{lower}}(\mathbf{x}) \leq \frac{1}{\lambda} \ln(h_{\lambda}(\mathbf{x})) + d_{\Gamma}(\mathbf{x}) \leq \frac{1}{\lambda} C_{\text{upper}}(\mathbf{x}). \quad (9)$$

Here, both $C_{\text{upper}}(\mathbf{x})$ and $C_{\text{lower}}(\mathbf{x})$ are scalar fields independent of λ . The inequalities are tight, which means that one can find a case where the equality holds.

Compared to the eikonal loss (Eq. (3)), where the distance error can be unbounded even when the equation holds




Figure 5. We visualize the direction of the changes of the implicit function during one iteration of optimization. Left shows the result with only the eikonal loss and right shows the result after adding our loss. Our loss helps to escape the local minimum in Fig. 3. The heat diffusion encourages the reconstructed absolute distance to increase monotonically away from the boundary.

almost everywhere, Eq. (9) shows that even with a finite absorption λ , our loss can lead to much more accurate distance. In fact, any loss that only depends on the norm of the gradients $\|\nabla u\|$ is not a sufficient condition of the distance function (Fig. 2), due to the discontinuous jumps of the gradient. Adding our loss (Eq. (7)) ensures that the absolute value of the implicit function converges to the actual distance. Fig. 5 visualizes the direction of how an implicit function would evolve when using only the eikonal loss and after adding our loss. Our loss ensures that the implicit function evolves towards the actual distance.

4.2. Stability Analysis

Our heat loss enjoys other theoretical benefits apart from the convergence to an actual distance. Here, we show that our loss leads to easier optimization. We separately analyze *spatial* stability and *temporal* stability, where the spatial stability studies the effect of small perturbation error added to the corresponding partial differential equation, and the *temporal* stability studies the convergence of the gradient flow.

4.2.1 Spatial Stability

We show that for the screened Poisson equation, an error introduced to a single point in the solution decays exponentially and does not propagate to infinitely distant positions. On the contrary, for the eikonal equation, the error can propagate to infinitely far. The property of the screened Poisson equation ensures that the perturbation introduced during optimization does not lead to a drastically different target and stabilizes the optimization.

First, we show that for the eikonal equation, an additive error introduced to a point can persist along a line segment or a ray.

Proposition 1. *Given a field $u(\mathbf{x})$ that satisfies the eikonal equation $\|\nabla u\| = 1$, we introduce an additive error u_{0e} at the point \mathbf{x}_0 where $u(\mathbf{x}_0) = u_0$ to obtain a perturbed field u' where $u'(\mathbf{x}_0) = u_0 + u_{0e}$ and $\|\nabla u'\| = 1$. In the*

perturbed field u' , there exists a line segment parameterized by $\mathbf{x}(s)$ starting from \mathbf{x}_0 and $s \in [0, M]$, where $M > 0$ and can be infinite, such that every point along the line has the same error:

$$u'(\mathbf{x}(s)) - u(\mathbf{x}(s)) = u_{0e} \quad (10)$$

The boundary points can contain errors when approximating the true boundaries, and these errors remain constant along the line segment. As a result, the propagated error can become significant and potentially detrimental.

Next, we demonstrate that the error field h_e in Eq. (4) decays exponentially from the perturbed point \mathbf{x}_0 .

Proposition 2. *In 3D space, consider an additive error h_{0e} within a small ball $B(\mathbf{x}_0, \epsilon)$ around the point \mathbf{x}_0 , where the original field $h(\mathbf{x})$ and the perturbed field $h'(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^n \setminus B$ both satisfy the screened Poisson equation in Eq. (4). The resulting error field $h_e = h' - h$ is radially symmetric with respect to \mathbf{x}_0 within some maximum radius R , and is given by*

$$h_e(r) = \frac{\epsilon}{r} h_{0e} e^{\lambda(\epsilon-r)}, \quad \forall r > \epsilon, \quad (11)$$

where $r = \|\mathbf{x} - \mathbf{x}_0\|$ is the radial distance from \mathbf{x}_0 .

The proofs of these two propositions are in the supplementary material A.1 and A.2. By comparing the error fields, we observe that the partial differential equation governed by the screened Poisson equation exhibits greater spatial robustness than the one based on the eikonal equation.

4.2.2 Temporal Stability

We next demonstrate the temporal stability of the heat loss and show that the dynamical system of the gradient update is stable, that is, the solution of the system converges as the time t approaches infinity. This is supported by von Neumann stability analysis, similar to Yang et al.'s work [4].

Yang et al. showed that for the eikonal loss (Eq. (3)), the temporal gradient flow of the optimization is:

$$\begin{aligned} \frac{\partial u}{\partial t} &= -\frac{\delta L_{\text{eikonal}}}{\delta u} = \nabla \cdot (\kappa_e \nabla u) \approx \kappa_e \nabla^2 u, \text{ where} \\ \kappa_e(\mathbf{x}) &= \begin{cases} \frac{1}{|\nabla u(\mathbf{x})|} \operatorname{sgn}(|\nabla u(\mathbf{x})| - 1) & p = 1 \\ 1 - |\nabla u(\mathbf{x})|^{-1} & p = 2 \end{cases} \end{aligned} \quad (12)$$

They analyze the stability by applying Fourier transform:

$$\frac{\partial \hat{u}}{\partial t}(t, \omega) = -\kappa_e |\omega|^2 \hat{u}(t, \omega) \Rightarrow \hat{u}(t, \omega) \propto e^{-\kappa_e |\omega|^2 t}, \quad (13)$$

where $\omega = (\omega_1, \dots, \omega_n)$ denotes the frequency variable and \hat{u} represents the Fourier transform of u . In the frequency domain, if the original function converges, the term $\frac{\partial \hat{u}}{\partial t}$ should also converge. However, when $\kappa_e < 0$, this process corresponds to backward diffusion, which diverges and is inherently unstable regardless of the numerical implementation.

We extend the analysis to our heat loss (Eq. (7)). The gradient flow is

$$\frac{\partial h}{\partial t} = -\frac{\delta L_{\text{heat}}}{\delta h} = \nabla^2 h - \lambda^2 h \quad (14)$$

This is exactly the heat equation, and is known to be stable. We can further verify the stability by applying Fourier transform to both sides:

$$\frac{\partial \hat{h}}{\partial t}(t, \omega) = -(|\omega|^2 + \lambda^2) \hat{h}(t, \omega) \Rightarrow \hat{h}(t, \omega) \propto e^{-(|\omega|^2 + \lambda^2)t} \quad (15)$$

This analysis reveals that the original process $h(\mathbf{x}, t)$ converges over time and remains stable. Due to the constraint $|u| = -\ln(h)/\lambda$, the term $|u(\mathbf{x}, t)|$ also converges during the optimization process. Since u is continuous, it converges as well. Thus, our method exhibits greater spatial and temporal stability compared to the eikonal equation.

4.3. Surface Area Regularization

We next show that our heat loss (Eq. (7)) penalizes large area of the reconstructed surfaces, by showing its relation to the coarea loss used by Pumarola et al. [30]. Pumarola et al. showed that the loss $\int_{\Omega} e^{-\lambda|u|} \|\nabla u\| d\mathbf{x}$ approximates the surface area of the zero level set of the implicit function u . We then form the following inequality:

$$\begin{aligned} \int_{\Omega} e^{-\lambda|u|} \|\nabla u\| d\mathbf{x} &< \int_{\Omega} \sqrt{e^{-2\lambda|u|} (\|\nabla u\|^2 + 1)} d\mathbf{x} \\ &\leq \sqrt{|\Omega|} \cdot \sqrt{\int_{\Omega} e^{-2\lambda|u|} (\|\nabla u\|^2 + 1) d\mathbf{x}} \end{aligned} \quad (16)$$

Therefore, minimizing the right hand side (our heat loss) leads to a reduction in the upper bound of the left hand side.

Furthermore, we can show that the coarea loss is too *aggressive* for our purpose of signed distance field reconstruction. If we directly minimize the coarea loss (the left hand side in Eq. (16)), according to Euler-Lagrange equation, the minimizer would satisfy:

$$1 - \nabla^2 u = 0, \quad (17)$$

which does not relate to the regularization parameter λ . Unfortunately, this partial differential equation system formed by the coarea loss above has a unique solution under Dirichlet boundary condition, and the gradient of its solution does not have a unit norm. This makes the coarea loss unable to accurately shape a signed distance function. By contrast, our heat loss has bounds for its distance approximation and converges to the true distance as we increase the parameter λ (Eq. (9)). Another simple way to verify our convergence and the compatibility with the eikonal equation is to rewrite Eq. (4) in terms of u following Eq. (6):

$$\|\nabla u\|^2 - 1 - \frac{\operatorname{sgn}(u)}{\lambda} \nabla^2 u = 0. \quad (18)$$




Figure 6. We reproduce the PHASE model [49] and show the log-transformed occupancy field, i.e., the optimized signed distance function. PHASE recommends boundary weight range $w_b \in (0.10, 0.32)$ based on their Γ -convergence theory. But in practice, only a stronger boundary weight constrains the boundary to follow the input points. We found that this phenomenon often occurs and show more examples in the supplementary material.

4.4. Relation to PHASE [49]

While derived from different mathematical principles and motivations, our method is closely related to the PHASE model proposed by Lipman. It is motivated by smoothly approximating an *occupancy function* $o(\mathbf{x})$, which is defined as 1 outside, -1 inside, and 0 on the surface.

In the supplement, we show that there is a close relation between our heat field h and the smoothed occupancy function o . We also show that our theory and derivation provide additional insights and significant differences in implementation and results. In particular, PHASE’s theory introduces a condition stating that the boundary loss weight w_b should be small, whereas our theory emphasizes the need for a firm boundary condition (Eq. (4) and Fig. 6).

We further found that optimizing occupancy field in the PHASE model, when used together with the eikonal loss, can lead to unbounded network weight derivatives, impeding the optimization. Our method avoids this issue.

5. Experiments

We experiment on both 2D and 3D datasets. We construct a 2D dataset containing 14 shapes, ranging from simple geometric shapes to complex nested structures. We also evaluate on two 3D benchmarks: a subset of ShapeNet [55] and the Surface Reconstruction Benchmark (SRB) [56]. Additionally, we generate a dataset with high-genus geometries from Mehta et al. [1]. In all of our experiments, we use only the position information of the point cloud.

We report the Intersection over Union (IoU), Chamfer distance, and Hausdorff distance between reconstructed shapes and ground truth shapes as metrics for surface reconstruction. In addition, we report Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Symmetric Mean Absolute Percentage Error (SMAPE) as metrics for distance field evaluation. We further evaluate non-zero level set quality near the surface, and render the zero level

set of the signed distance functions via sphere tracing.

We use a network with 5 hidden layers and 128 hidden channels for our method in all experiments. The hyperparameters are kept consistent within each dataset. We additionally apply importance sampling to our heat loss integral. Full training details can be found in the supplement.

5.1. 2D Dataset

An overview of our 2D dataset is attached in the supplementary material. For 2D experiments, we simply use a linear network and geometric initialization, as in DiGS [3]. We compare our method with DiGS [3] and StEik [4] with the hyperparameters and training configuration provided in their released code, and report the surface reconstruction metrics and distance field metrics in Table 1. Our method outperforms theirs across all metrics with only 8,192 sample points, whereas both DiGS [3] and StEik [4] use 15,000 sample points based on their default settings.

5.2. Ablation Study on 2D Dataset

To more comprehensively compare over different losses, we conduct an ablation study on our 2D dataset using the same network architecture, parameter initialization, and sampling strategy. The only difference lies in the loss functions used. Quantitative results are shown in Table 2.

The last row corresponds to our HOTSPOT model, combining boundary, eikonal, and heat losses. Our method reconstructs the correct topology for all 14 shapes, while others struggle with the harder ones. It outperforms across nearly all metrics, demonstrating superior reconstruction quality and distance query accuracy. Although HOTSPOT slightly trails SAL [2] in RMSE and MAE, it captures the complex topology where SAL fails. A visualization is shown in Fig. 7. Full qualitative results and detailed experiment parameters are provided in the supplementary material. Models relying heavily on area losses [3, 4] often face a trade-off between preserving original details and eliminating redundant boundaries, which distorts the distance field (Section 4.3). Moreover, methods without our heat loss fall in local optimum easily.

5.3. 3D Datasets

We evaluate our method on a processed subset [16, 19, 57] of ShapeNet [55], which provides surface points sampled on preprocessed watertight meshes for 260 shapes across 13 categories. To ensure a reasonable spatial density for passing the heat, we scale the point cloud so that 70% of the points lie within a sphere centered at the origin with a radius of 0.45. For a fair comparison, we transform our outcomes back into the coordinate system used in StEik [4] and DiGS [3]. Additionally, we implement a scheduler to gradually increase the absorption parameter λ , which enhances the representation of level set details.

	IoU \uparrow	d_C \downarrow	d_H \downarrow
DiGS [3]	0.7882	0.0055	0.1267
StEik [4]	0.6620	0.0073	0.1425
Ours	0.9870	0.0014	0.0153

	RMSE \downarrow	MAE \downarrow	SMAPE \downarrow
DiGS [3]	0.0597	0.0315	0.3363
StEik [4]	0.0725	0.0419	0.4222
Ours	0.0199	0.0101	0.0699

Table 1. Comparison of 2D reconstruction and distance query. d_C and d_H stand for the Chamfer and Hausdorff distances respectively. Our method outperforms DiGS [3] and StEik [4] on all metrics.

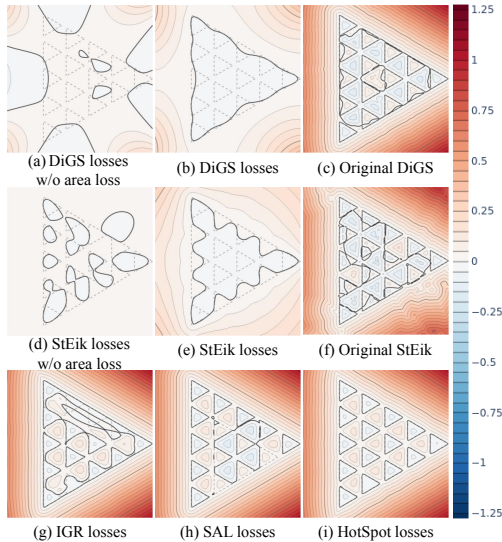


Figure 7. Our heat loss achieves the best reconstruction of triangle fragments. Dashed lines show real boundaries, while the boldest line represents reconstructed ones. (c) and (f) are from the original models by DiGS [3] and StEik [4]; the rest are from our ablation study with different loss combinations. Other combinations fail to remove extra boundaries or accurately capture the topology. More fundamentally, (c), (f), (g), and (h) fall into local optima, leading to incorrect depressions and elevations.

We present the quantitative results in Table 3 and visual results in Fig. 9. More baselines [46, 51] are compared in the supplement. Our HOTSPOT model outperforms current state-of-the-art models across all surface reconstruction metrics, and achieves near-top accuracy in distance queries. While SAL [2] uses the closest point distance in their loss function, leading to slightly lower total absolute error, their performance drops significantly in relative error and near-surface metrics. This suggests that the closest point distance alone is unreliable near the surface. In contrast, as mentioned at the end of Section 3, after the interpolation driven by spectral bias [58], our loss remains effective in faithfully capturing the distance to the interpolated surface.

Fig. 1 and Fig. 8 show high-genus surface reconstruction with shapes from Mehta et al. [1]. Existing methods strug-

\mathcal{L}_B	\mathcal{L}_E	\mathcal{L}_A	\mathcal{L}_D	\mathcal{L}_{DD}	\mathcal{L}_S	\mathcal{L}_H	IoU \uparrow	d_C \downarrow	d_H \downarrow	RMSE \downarrow	MAE \downarrow	SMAPE \downarrow
✓	✓						0.8192	0.0068	0.0712	0.0377	0.0158	0.1315
✓					✓		0.8936	0.0029	0.0579	0.0165	0.0081	0.1205
✓	✓		✓				0.3907	0.0843	0.5598	0.2966	0.2388	1.7283
✓	✓	✓	✓				0.6338	0.0566	0.4221	0.2058	0.1616	1.0919
✓	✓			✓			0.2938	0.0377	0.3444	0.2990	0.2355	1.6958
✓	✓	✓		✓			0.6037	0.0414	0.3656	0.1304	0.0937	0.7792
✓	✓					✓	0.9851	0.0016	0.0160	0.0199	0.0105	0.0754

Table 2. Ablation study of different loss combinations and metrics. From top to bottom, the loss combinations are from: IGR [10], SAL [2], DiGS w/o area loss, DiGS [3], StEik w/o loss, StEik [4], and our HOTSPOT. From left to right, the losses are boundary loss, eikonal loss, area loss, divergence loss, directional divergence loss, SAL loss, and heat loss. Bold and underlined data: optimal; bold only: suboptimal. Same below.




Figure 8. Visual examples on high genus dataset. Our method uses significantly fewer iterations but reconstructs the topology correctly for all shapes, whereas others generate extra boundaries.

gle with the complex topology, while our method accurately reconstructs both the surface and the level set. In the supplement, we present evaluations on other high-genus shapes from Mehta et al. and the SRB dataset.

5.4. Sphere Tracing

Fig. 10 visualizes the number of steps a sphere tracer [21] takes to render the shapes obtained by SAL, DiGS, StEik, and our method. Our method has the fewest iterations on average, since it provides the most accurate level sets. More results can be found in the supplement.

6. Discussion

Efficiency. Our method uses a loss function of only the neural network output and its first-order derivatives, and thus can be more efficient than previous methods which involve second-order derivatives. We report the runtime of our method and StEik [4] for training a model from ShapeNet [55] over 100 iterations using a single NVIDIA A10 GPU with varying network sizes, while keeping all other settings identical, as shown in Tab. 4.

Determining absorption coefficient λ . The absorption coefficient λ has the units of $[L]^{-1}$, where L denotes a length unit. Rescaling the space is equivalent to rescaling λ . While a larger λ can improve accuracy, it introduces two challenges: (1) our heat loss includes the factor $e^{-2\lambda|u|}$, which may become smaller than the precision limit of floating-point numbers, leading to round-off errors; (2) even if round-off errors are ignored, a large $\lambda|u|$ may reduce the optimizer’s ability to effectively shape the field. Never-

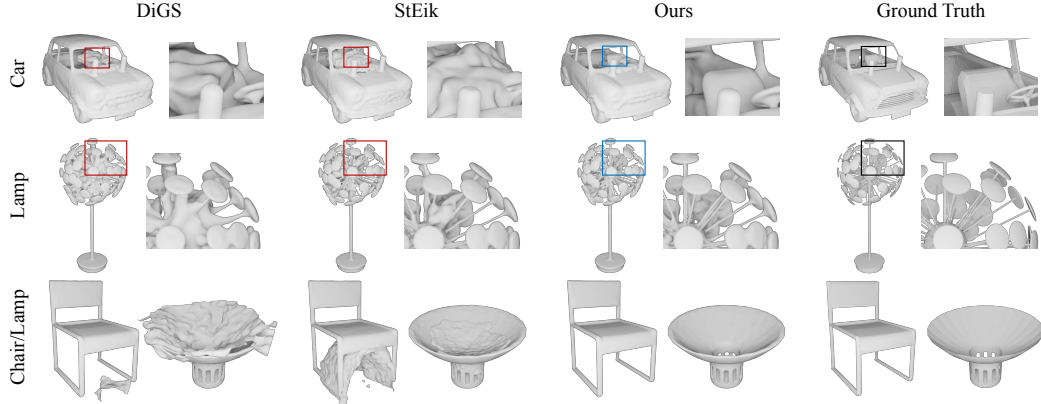


Figure 9. Visual results of DiGS [3], StEik [4], and HOTSPOT on ShapeNet [55]. Previous methods generate extra boundaries and fail to escape local optima, while our method successfully reconstructs the correct topology.

	IoU \uparrow	d_C \downarrow	d_H \downarrow	RMSE \downarrow	MAE \downarrow	SMAPE \downarrow	RMSE _{0.1}	MAE _{0.1}	SMAPE _{0.1}
SAL [2]	0.7400	0.0074	0.0851	0.0251	0.0142	0.1344	0.0245	0.0182	0.6848
SIREN wo/ n [34]	0.4874	0.0051	0.0558	0.5009	0.4261	1.2694	0.0513	0.0382	0.8858
DiGS [3]	0.9636	0.0031	0.0435	0.1194	0.0725	0.2140	0.0152	0.0081	0.1760
StEik [4]	0.9641	0.0032	0.0368	0.0387	0.0248	0.0931	0.0147	0.0081	0.1770
Ours	0.9796	0.0029	0.0250	0.0281	0.0176	0.0540	0.0094	0.0047	0.1206

Table 3. Surface reconstruction metrics on ShapeNet [55], where d_C and d_H stand for the Chamfer and Hausdorff distances, RMSE, MAE, and SMAPE are distance field metrics, and RMSE_{0.1}, MAE_{0.1}, and SMAPE_{0.1} are the same distance field metrics but only for points within distance 0.1 of the surface. See the supplementary material for detailed results with more baselines.

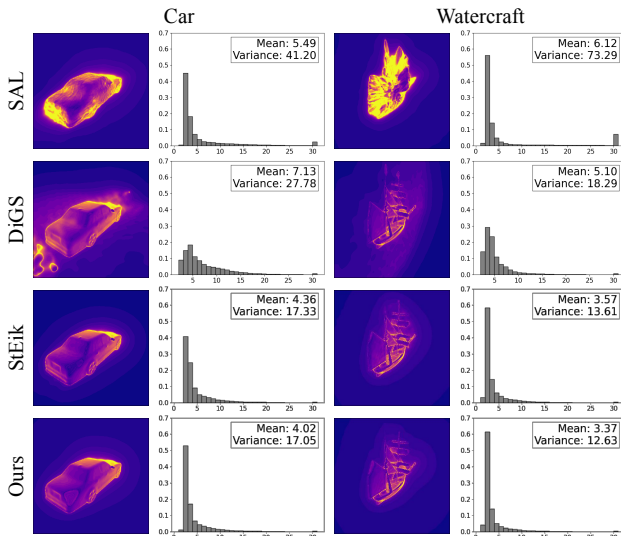


Figure 10. Visualizations and statistics of sphere tracing [21] iteration counts required to locate the surface. In the visual results brighter pixels indicate higher iteration counts. The histograms show the distribution of iterations for all pixels across 10 different camera poses. HOTSPOT finds the surface in fewer steps.

theless, our loss still pushes u away from 0, preventing unwanted surfaces where $\lambda|u|$ should be large. We compensate by combining eikonal loss with heat losses. When $\lambda|u|$ is large, the eikonal loss becomes dominant. Our λ scheduler also effectively shapes the distant regions. Spatially-adaptive parameters may also improve the results.

Necessity of firm boundary condition. The heat diffu-

Structure	5×128	5×256	8×128	8×256
StEik [4]	5.65s	10.62s	8.78s	17.00s
Ours	4.49s	8.43s	6.79s	12.86s

Table 4. Runtime for 100 iterations with different network sizes.

sion of the screened Poisson equation is based on a well-established boundary condition. We enforce the boundary condition through the boundary loss over discrete points, with neural networks interpolating through spectral bias [58]. When the input points are sparse and the absorption coefficient λ is high, an overly strong heat loss (w_b being too small) can *tear* the boundary, causing the signed distance to collapse into an unsigned one like Fig. 6. Thus, our theory highlights the importance of a high boundary weight w_b , while absorption λ is adjusted to match the input point density or by scaling the point cloud itself. Our experiments show that proper parameter setting and rescaling prevent the collapse. Future research is required for very sparse boundaries, and applications to inverse rendering.

7. Conclusion

We propose a new model for neural signed distance function optimization based on the screened Poisson equation. We analyze our loss theoretically and show that it is an asymptotically sufficient condition to the true distance, is stable both to a small perturbation and in the temporal dynamics, and penalizes large surface area. Our experiments show that we reconstruct both better surfaces and better distance approximations compared to many existing methods, especially on complex and high-genus shapes.

Acknowledgements. This work was supported in part by NSF grants 2127544, 2238839, 2100237, 2120019, and gifts from Adobe and Google. Additionally, we would like to thank Bing Xu, Zilu Li, Yash Belhe, Ishit Mehta, and Jianan Xiao for their suggestions, discussions, proofreading, and providing reference materials.

References

- [1] Ishit Mehta, Manmohan Chandraker, and Ravi Ramamoorthi. A level set theory for neural implicit evolution under explicit flows. In *ECCV*, pages 711–729, 2022. [1](#), [2](#), [6](#), [7](#), [8](#), [10](#)
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2565–2574, 2020. [1](#), [2](#), [6](#), [7](#), [8](#), [3](#), [5](#)
- [3] Yizhak Ben-Shabat, Chamin Hewa Koneputugodage, and Stephen Gould. DiGS: Divergence guided shape implicit neural representation for unoriented point clouds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 19323–19332, 2022. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [5](#)
- [4] Huizong Yang, Yuxin Sun, Ganesh Sundaramoorthi, and Anthony Yezzi. StEik: stabilizing the optimization of neural signed distance functions and finer shape representation. *arXiv preprint arXiv:2305.18414*, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [5] Zoë Marschner, Silvia Sellán, Hsueh-Ti Derek Liu, and Alec Jacobson. Constructive solid geometry on neural signed distance fields. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–12, 2023. [1](#), [2](#), [3](#)
- [6] Sathamangalam R Srinivasa Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics*, 20(2):431–455, 1967. [2](#), [3](#), [4](#)
- [7] Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM TOG*, 32(5):1–11, 2013. [2](#), [3](#)
- [8] Alexander G Belyaev and Pierre-Alain Fayolle. On variational and pde-based distance function approximations. In *Computer Graphics Forum*, volume 34, pages 104–118. Wiley Online Library, 2015.
- [9] Alexander Belyaev and Pierre-Alain Fayolle. An ADMM-based scheme for distance function approximation. *Numerical Algorithms*, 84:983–996, 2020. [2](#), [3](#)
- [10] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. [2](#), [3](#), [7](#)
- [11] James F Blinn. A generalization of algebraic surface drawing. *ACM TOG*, 1(3):235–256, 1982. [2](#)
- [12] Jules Bloomenthal and Brian Wyvill. Interactive techniques for implicit modeling. *Comput. Graph. (Proc. SIGGRAPH)*, 24(2):109–116, 1990.
- [13] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. *Comput. Graph. (Proc. SIGGRAPH)*, 26(2):71–78, 1992.
- [14] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH*, page 303–312, 1996.
- [15] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, and Andrew Davison. KinectFusion: real-time 3D reconstruction and interaction using a moving depth camera. In *ACM Symposium on User Interface Software and Technology*, pages 559–568, 2011. [2](#)
- [16] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4460–4470, 2019. [2](#), [6](#), [7](#)
- [17] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5939–5948, 2019.
- [18] Thomas Davies, Derek Nowrouzezahrai, and Alec Jacobson. On the effectiveness of weight-encoded neural implicit 3d shapes. *arXiv preprint arXiv:2009.09808*, 2020.
- [19] Francis Williams, Matthew Trager, Joan Bruna, and Denis Zorin. Neural Splines: Fitting 3D surfaces with infinitely-wide neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9949–9958, 2021. [6](#)
- [20] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. 41(2):641–676, 2022. [2](#)
- [21] John C Hart. Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *The Vis. Comput.*, 12(10):527–545, 1996. [2](#), [7](#), [8](#)
- [22] Dario Seyb, Alec Jacobson, Derek Nowrouzezahrai, and Wojciech Jarosz. Non-linear sphere tracing for rendering deformed signed distance fields. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 38(6), 2019.
- [23] Eric Galin, Eric Guérin, Axel Paris, and Adrien Peytavie. Segment tracing using local Lipschitz bounds, 2020.
- [24] Nicholas Sharp and Alec Jacobson. Spelunking the deep: guaranteed queries on general neural implicit surfaces via range analysis. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 41(4), 2022. [2](#)
- [25] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11358–11367, 2021. [2](#)
- [26] Arnulph Fuhrmann, Gerrit Sobotka, and Clemens Groß. Distance fields for rapid collision detection in physically based modeling. In *Proceedings of GraphiCon*, volume 2003, pages 58–65, 2003. [2](#)
- [27] Eran Guendelman, Robert Bridson, and Ronald Fedkiw. Nonconvex rigid bodies with stacking. *ACM TOG*, 22(3):871–878, 2003.
- [28] Miles Macklin, Kenny Erleben, Matthias Müller, Nuttapon Chentanez, Stefan Jeschke, and Zach Corse. Local optimization for robust signed distance field collision. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 3(1):1–17, 2020. [2](#)
- [29] Rohan Sawhney and Keenan Crane. Monte Carlo geometry processing: a grid-free approach to PDE-based methods on

- volumetric domains. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 39(4), 2020. [2](#)
- [30] Albert Pumarola, Arsiom Sanakoyeu, Lior Yariv, Ali Thabet, and Yaron Lipman. Visco grids: Surface reconstruction with viscosity and coarea grids. *Advances in Neural Information Processing Systems*, 35:18060–18071, 2022. [2](#), [5](#)
- [31] Songyou Peng, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape as points: A differentiable poisson solver. *Advances in Neural Information Processing Systems*, 34:13032–13044, 2021.
- [32] Chao Chen, Yu-Shen Liu, and Zhizhong Han. Gridpull: Towards scalability in learning implicit representations from 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18322–18334, 2023. [2](#)
- [33] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 165–174, 2019. [2](#)
- [34] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020. [2](#), [8](#), [5](#), [7](#)
- [35] Guillaume Coiffier and Louis Béthune. 1-Lipschitz neural distance fields. *Comput. Graph. Forum (Eurographics STARs)*, 43(5):e15128, 2024. [2](#)
- [36] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. SDFDiff: Differentiable rendering of signed distance fields for 3D shape optimization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1251–1261, 2020. [2](#)
- [37] Petr Kellnhofer, Lars C Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4287–4297, 2021.
- [38] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3504–3515, 2020.
- [39] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *Advances in Neural Information Processing Systems*, volume 33, pages 2492–2502, 2020. [2](#), [8](#)
- [40] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhysSG: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5453–5462, 2021.
- [41] Sai Bangaru, Michael Gharbi, Tzu-Mao Li, Fujun Luan, Kalyan Sunkavalli, Milos Hasan, Sai Bi, Zexiang Xu, Gilbert Bernstein, and Fredo Durand. Differentiable rendering of neural sdfs through reparameterization. In *SIGGRAPH Asia Conference Proceedings*, 2022.
- [42] Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Differentiable signed distance function rendering. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 41(4):1–18, 2022. [2](#)
- [43] Jingyang Zhang, Yao Yao, Shiwei Li, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Critical regularizations for neural surface reconstruction in the wild. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6270–6279, 2022. [2](#)
- [44] Ruian Wang, Zixiong Wang, Yunxiao Zhang, Shuangmin Chen, Shiqing Xin, Changhe Tu, and Wenping Wang. Aligning gradient and hessian for neural signed distance function. In *Advances in Neural Information Processing Systems*, volume 36, 2024.
- [45] Zixiong Wang, Yunxiao Zhang, Rui Xu, Fan Zhang, Peng-Shuai Wang, Shuangmin Chen, Shiqing Xin, Wenping Wang, and Changhe Tu. Neural-singular-hessian: Implicit neural representation of unoriented point clouds by enforcing singular hessian. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023. [5](#), [7](#), [8](#)
- [46] Qiuji Dong, Rui Xu, Pengfei Wang, Shuangmin Chen, Shiqing Xin, Xiaohong Jia, Wenping Wang, and Changhe Tu. Neurcadrecon: Neural representation for reconstructing cad surfaces by enforcing zero gaussian curvature. *arXiv preprint arXiv:2404.13420*, 2024. [2](#), [7](#)
- [47] Yesom Park, Taekyung Lee, Jooyoung Hahn, and Myungjoo Kang. p -poisson surface reconstruction in curl-free flow from point clouds. In *Advances in Neural Information Processing Systems*, volume 36, 2024. [2](#)
- [48] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8456–8465, 2023. [2](#)
- [49] Yaron Lipman. Phase transitions, distance functions, and implicit neural representations. In *International Conference on Machine Learning*, volume 139, pages 6702–6712, 2021. [2](#), [3](#), [6](#), [8](#), [9](#)
- [50] Angela Dai and Matthias Nießner. Neural poisson: Indicator functions for neural fields. *arXiv preprint arXiv:2211.14249*, 2022. [2](#)
- [51] Baorui Ma, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Neural-pull: Learning signed distance functions from point clouds by learning to pull space onto surfaces. *arXiv preprint arXiv:2011.13495*, 2020. [2](#), [7](#), [3](#), [5](#), [8](#)
- [52] Matan Atzmon and Yaron Lipman. SALD: Sign agnostic learning with derivatives. In *ICLR*, 2021. [2](#), [3](#)
- [53] Nicholas Sharp, Yousuf Soliman, and Keenan Crane. The vector heat method. *ACM TOG*, 38(3):1–19, 2019. [3](#)
- [54] Nicole Feng and Keenan Crane. A heat method for generalized signed distance. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 43(4), 2024. [3](#)
- [55] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [6](#), [7](#), [8](#), [5](#), [9](#)
- [56] Matthew Berger, Joshua A Levine, Luis Gustavo Nonato, Gabriel Taubin, and Claudio T Silva. A benchmark for surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(2):1–17, 2013. [6](#), [8](#)

- [57] David Stutz and Andreas Geiger. Learning 3d shape completion under weak supervision. *International Journal of Computer Vision*, 128:1162–1181, 2020. [6](#), [7](#)
- [58] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310, 2019. [7](#), [8](#), [3](#), [5](#), [11](#)
- [59] Dmitry S Kulyabov, Anna V Korolkova, Tatiana R Velieva, and Migran N Gevorkyan. Numerical analysis of eikonal equation. In *Saratov Fall Meeting 2018: Laser Physics, Photonic Technologies, and Molecular Modeling*, volume 11066, pages 188–195. SPIE, 2019. [1](#)
- [60] Zhen-Hang Yang and Yu-Ming Chu. On approximating the modified bessel function of the second kind. *Journal of inequalities and applications*, 2017:1–8, 2017. [2](#)
- [61] Luciano Modica. The gradient theory of phase transitions and the minimal interface criterion. *Archive for Rational Mechanics and Analysis*, 98:123–142, 1987. [10](#)
- [62] Peter Sternberg. The effect of a singular perturbation on non-convex variational problems. *Archive for Rational Mechanics and Analysis*, 101:209–260, 1988. [10](#)