
Phase Transitions, Distance Functions, and Implicit Neural Representations

Supplementary material

Yaron Lipman^{1,2}

1. Γ -convergence

Proof (Theorem 2). We will adapt the proof of Theorem 13.6 in (Rindler, 2018) to the surface reconstruction loss. We want to prove Γ -convergence of $\epsilon^{-1/2}\mathcal{F}_\epsilon$ to \mathcal{F}_0 . Note that $\epsilon^{-1/2}\mathcal{F}_\epsilon$ is equivalent to the functional

$$\mathcal{F}_\epsilon(u) = \lambda\mathcal{L}(u) + \begin{cases} \int_\Omega \epsilon \|\nabla u\|^2 + \frac{1}{\epsilon}W(u) & u \in W^{1,2}(\Omega) \\ +\infty & \text{otherwise} \end{cases} \quad (1)$$

where now $\lambda(\epsilon) = \epsilon^{-1}\tilde{\lambda}(\epsilon^2)$, where $\tilde{\lambda}$ represents the dependence in equation 9 (in the main paper). So we will prove Γ -convergence of this \mathcal{F}_ϵ to \mathcal{F}_0 , where $\lambda \rightarrow \infty$ and $\lambda\sqrt{\epsilon} \rightarrow 0$, as $\epsilon \downarrow 0$. The proof of Γ -convergence requires showing the \liminf and *recovery* properties from Section 2.3 in the main paper. Let us denote:

$$\mathcal{E}_\epsilon(u) = \begin{cases} \int_\Omega \epsilon \|\nabla u\|^2 + \frac{1}{\epsilon}W(u) & u \in W^{1,2}(\Omega) \\ +\infty & \text{otherwise} \end{cases}$$

$$\mathcal{E}_0(u) = \begin{cases} \sigma_0 \text{per}_\Omega(\mathcal{I}) & u \in \text{BV}(\Omega; \{-1, 1\}) \\ +\infty & \text{otherwise} \end{cases}$$

For the reader's convenience we also repeat the definition of the limit functions \mathcal{F}_0 :

$$\mathcal{F}_0(u) = \begin{cases} \sigma_0 \text{per}_\Omega(\mathcal{I}) & u \in \text{BV}(\Omega, \{-1, 1\}), \\ & \text{and } \mathcal{L}(u) = 0 \\ +\infty & \text{otherwise} \end{cases} \quad (2)$$

Lim inf part. In this part we need to consider $u_\epsilon \rightarrow u$ in $L^1(\Omega)$; we abuse notation a bit and let $\epsilon \downarrow 0$ denote some particular sequence $\epsilon_k \rightarrow 0$ as $k \rightarrow \infty$. We need to show that $\liminf_{\epsilon \downarrow 0} \mathcal{F}_\epsilon(u_\epsilon) \geq \mathcal{F}_0(u)$. If $\liminf_{\epsilon \downarrow 0} \mathcal{F}_\epsilon(u_\epsilon) = \infty$ the statement holds trivially, therefore we assume $\liminf_{\epsilon \downarrow 0} \mathcal{F}_\epsilon(u_\epsilon) < \infty$.

In the proof of Theorem 13.6 in Rindler (2018) it is shown that $\int_\Omega W(u(\mathbf{x}))d\mathbf{x} = 0$ and consequently $u(\mathbf{x}) \in$

$\{-1, +1\}$ almost everywhere, i.e., equation 10 (in the main paper) holds for u . Furthermore, it is shown that

$$\liminf_{\epsilon \downarrow 0} \mathcal{E}_\epsilon(u_\epsilon) \geq \mathcal{E}_0(u), \quad (3)$$

where σ_0 is some constant depending on W alone:

$$\sigma_0 = 2 \int_{-1}^1 \sqrt{W(s)}ds. \quad (4)$$

Since $\mathcal{E}_0(u) < \infty$, $u \in \text{BV}(\Omega; \{-1, 1\})$ (this can be seen directly from the definitions in equations 2 and 3 in the main paper).

Now, for our reconstruction loss, since $u_\epsilon \rightarrow u$ in $L^1(\Omega)$ we have that for every $\mathbf{x} \in \mathcal{X}$, $\left| \int_{B_{\mathbf{x}}} u_\epsilon \right| \rightarrow \left| \int_{B_{\mathbf{x}}} u \right|$, as $\epsilon \downarrow 0$. By Fatou's lemma applied for the functions $\mathbf{x} \mapsto \left| \int_{B_{\mathbf{x}}} u_\epsilon \right|$ and the limit function $\mathbf{x} \mapsto \left| \int_{B_{\mathbf{x}}} u \right|$, and the fact that $\liminf_{\epsilon \downarrow 0} \mathcal{F}_\epsilon(u_\epsilon) < \infty$:

$$\mathcal{L}(u) = \mathbb{E}_{\mathbf{x}} \left| \int_{B_{\mathbf{x}}} u \right| \leq \liminf_{\epsilon \downarrow 0} \mathbb{E}_{\mathbf{x}} \left| \int_{B_{\mathbf{x}}} u_\epsilon \right|$$

$$= \liminf_{\epsilon \downarrow 0} \mathcal{L}(u_\epsilon) \leq \liminf_{\epsilon \downarrow 0} \frac{1}{\lambda} \mathcal{F}_\epsilon(u_\epsilon) = 0.$$

where the last equality is due to $\lambda \rightarrow \infty$ as $\epsilon \downarrow 0$. This means that the limit function u satisfies the reconstruction constraints perfectly, or in other words that the reconstructed surface \mathcal{S} passes through all the balls $B_{\mathbf{x}}$, $\mathbf{x} \in \mathcal{X}$, except possibly a subset of \mathcal{X} of measure zero. However, since $\left| \int_{B_{\mathbf{x}}} u \right|$ is continuous as a function of \mathbf{x} this is true for all balls $B_{\mathbf{x}}$, $\mathbf{x} \in \mathcal{X}$.

In particular $\mathcal{E}_0(u) = \mathcal{F}_0(u)$. Now incorporating this with equation 3 we get

$$\liminf_{\epsilon \downarrow 0} \mathcal{F}_\epsilon(u_\epsilon) = \liminf_{\epsilon \downarrow 0} (\lambda\mathcal{L}(u_\epsilon) + \mathcal{E}_\epsilon(u_\epsilon))$$

$$\geq \liminf_{\epsilon \downarrow 0} \mathcal{E}_\epsilon(u_\epsilon)$$

$$\geq \mathcal{E}_0(u) = \mathcal{F}_0(u)$$

as required.

Recovery sequence part. In this part we need to consider an arbitrary $u \in L^1(\Omega)$ and find a sequence $u_\epsilon \in L^1(\Omega)$ so that $u_\epsilon \rightarrow u$ in $L^1(\Omega)$ and $\lim_{\epsilon \downarrow 0} \mathcal{F}_\epsilon(u_\epsilon) = \mathcal{F}_0(u)$.

^{*}Equal contribution ¹Facebook AI Research ²Weizmann Institute of Science. Correspondence to: Yaron Lipman <yilipman@fb.com, yaron.lipman@weizmann.ac.il>.

Let $u \in L^1(\Omega)$ be arbitrary. If $u \notin \text{BV}(\Omega; \{-1, 1\})$ or $\mathcal{L}(u) > 0$ then $\mathcal{F}_0(u) = \infty$ and there is no need to construct a recovery sequence in this case (see, e.g., Theorem 1 where no recovery sequence for such cases is needed). So we assume $\mathcal{F}_0(u) < \infty$, and u of the form $u = -\mathbf{1}_{\mathcal{I}} + \mathbf{1}_{\Omega \setminus \mathcal{I}}$, where \mathcal{I} is defined as in equation 4.

Next, note that if $\int_{\Omega} u \in \{-|\Omega|, |\Omega|\}$ then $\mathcal{L}(u) > 0$ and again $\mathcal{F}_0(u) = \infty$. Therefore we can assume $\int_{\Omega} u \in (-|\Omega|, |\Omega|)$. In this case, Theorem 13.6 in Rindler (2018) shows the existence of a recovery sequence of functions u_{ϵ} so that $u_{\epsilon} \rightarrow u$ in $L^1(\Omega)$ and $\lim_{\epsilon \downarrow 0} \mathcal{E}_{\epsilon}(u_{\epsilon}) = \mathcal{E}_0(u)$. The main observation in this part is that u_{ϵ} is a recovery sequence also for our surface reconstruction functionals \mathcal{F}_{ϵ} and \mathcal{F}_0 .

To show that u_{ϵ} is a recovery sequence also in our settings it is enough to show that $\lambda \mathcal{L}(u_{\epsilon}) \rightarrow 0$ as $\epsilon \downarrow 0$. Indeed, if this is the case,

$$\begin{aligned} \lim_{\epsilon \downarrow 0} \mathcal{F}_{\epsilon}(u_{\epsilon}) &= \lim_{\epsilon \downarrow 0} (\mathcal{E}_{\epsilon}(u_{\epsilon}) + \lambda \mathcal{L}(u_{\epsilon})) \\ &= \mathcal{E}_0(u) = \mathcal{F}_0(u) \end{aligned}$$

To show that $\lambda \mathcal{L}(u_{\epsilon}) \rightarrow 0$ we need to use a bit of extra information on u_{ϵ} : u_{ϵ} is constructed to approximate $w = -\mathbf{1}_{G \cap \Omega} + \mathbf{1}_{\Omega \setminus G}$, that is

$$\int_{\Omega} |u_{\epsilon} - w| \rightarrow 0,$$

where $G \subset \mathbb{R}^d$ is open, bounded with smooth boundary, and $|(G \Delta \mathcal{I}) \cap \Omega|$ can be made arbitrary small (see Lemma 13.7 in Rindler (2018), or Lemma 1 in Modica (1987)). This means that for arbitrary $\delta > 0$ we can choose G so that

$$\int_{\Omega} |u - w| \leq |(G \Delta \mathcal{I}) \cap \Omega| \leq \delta$$

Remember that $\mathcal{L}(u) = 0$ and using this last inequality we get that for arbitrary $x \in \mathcal{X}$,

$$\begin{aligned} \left| \int_{B_x} u_{\epsilon} \right| &= \left| \int_{B_x} u_{\epsilon} - \int_{B_x} w \right| + \left| \int_{B_x} w - \int_{B_x} u \right| \\ &\leq \int_{\Omega} |u_{\epsilon} - w| + \int_{\Omega} |w - u| \leq c\sqrt{\epsilon} \end{aligned} \quad (5)$$

where $c > 0$ is some constant, and the last inequality is due to the fact that we can make the choice $\delta = \sqrt{\epsilon}$ and the following bound shown in the proof of Theorem 13.6 in Rindler (2018):

$$\int_{\Omega} |u_{\epsilon} - w| \leq 4\sqrt{\epsilon} \sup_{-2\sqrt{\epsilon} \leq t \leq 2\sqrt{\epsilon}} \mathcal{H}^{d-1}(\mathcal{S}_t \cap \Omega),$$

where $\mathcal{S}_t = \{x \in \mathbb{R}^d \mid d_{\mathcal{S}}(x) = t\}$, $d_{\mathcal{S}}$ is the signed distance function defined in equation 12, \mathcal{S} is defined as in

equation 4 for w , and Lemma 13.9 in Rindler (2018) shows that the Hausdorff measure of $\mathcal{S}_t \cap \Omega$ converges to that of $\partial \mathcal{I} \cap \Omega$ as $\epsilon \downarrow 0$, and therefore is bounded.

Equation 5 implies that

$$\lambda \mathcal{L}(u_{\epsilon}) \leq c\lambda\sqrt{\epsilon}.$$

Lastly, remember that $\lambda\sqrt{\epsilon} \rightarrow 0$ and therefore u_{ϵ} is a recovery sequence as desired. \square

2. Distance functions

Theorem 3. *Let $u_{\epsilon} \in W^{1,2}(\Omega)$ be a (local) minimizer of \mathcal{F}_{ϵ} , and $O \subset \Omega \setminus \cup_{x \in \mathcal{X}} B_x$ a set where $u_{\epsilon} \neq 0$. Then, u_{ϵ} is smooth in the classical sense in O and satisfies*

$$-\epsilon \Delta u_{\epsilon} + u_{\epsilon} - \text{sign}(u_{\epsilon}) = 0. \quad (6)$$

Proof. We start by applying the Euler-Lagrange (EL) conditions (see e.g., Theorem 3.1 in (Rindler, 2018)) for a minimizer u of \mathcal{F}_{ϵ} in O . Let

$$f(x, v, \nu) = \epsilon \|\nu\|^2 + v^2 - 2|v| + 1,$$

be our integrand. That is $\mathcal{F}_{\epsilon}(u) = \int_{\Omega} f(x, u, \nabla u)$. The EL conditions are:

$$-\text{div} [\nabla_{\nu} f(x, u, \nabla u)] + \frac{d}{dv} f(x, u, \nabla u) = 0.$$

Plugging our f and noting that $W'(s) = 2s - 2\text{sign}(s)$, which is differentiable in O , we get

$$-2\epsilon \Delta u + 2u - 2\text{sign}(u) = 0.$$

In particular, equation 6 will be satisfied in the weak sense with any test function $\psi \in C_c^{\infty}(O)$. Second, regularity results for elliptic operators (e.g., Corollary 8.11 in (Gilbarg & Trudinger, 2015)) show that u is smooth in O and satisfies equation 6 in the classical sense. \square

Theorem 4. *Let $O \subset \Omega$ be a domain as defined in Theorem 3. Then, over O , w_{ϵ} satisfies*

$$-\sqrt{\epsilon} \Delta w_{\epsilon} + \text{sign}(u_{\epsilon})(\|\nabla w_{\epsilon}\|^2 - 1) = 0 \quad (7)$$

Proof. For brevity, we denote $u = u_{\epsilon}$, and assume $u > 0$ in O . Then, equation 14 (in the main paper) in this case is

$$w = -\sqrt{\epsilon} \log(1 - u).$$

Let $\mathbf{x} = (x_1, x_2, \dots, x_d)$. Now,

$$\begin{aligned} \frac{\partial w}{\partial x_i} &= \sqrt{\epsilon} \frac{1}{1-u} \frac{\partial u}{\partial x_i} \\ \frac{\partial^2 w}{\partial x_i^2} &= \sqrt{\epsilon} \frac{1}{(1-u)^2} \left[\frac{\partial u}{\partial x_i} \right]^2 + \sqrt{\epsilon} \frac{1}{1-u} \frac{\partial^2 u}{\partial x_i^2} \end{aligned}$$

Plugging in the l.h.s. of equation 7 we get

$$\frac{-\epsilon}{(1-u)^2} \|\nabla u\|^2 + \frac{-\epsilon}{1-u} \Delta u + \frac{\epsilon}{(1-u)^2} \|\nabla u\|^2 - 1$$

and this term vanishes in view of equation 6.

If $u < 0$ in O , then equation 14 (in the main paper) in this case is

$$w = \sqrt{\epsilon} \log(1+u).$$

Similar to before:

$$\begin{aligned} \frac{\partial w}{\partial x_i} &= \sqrt{\epsilon} \frac{1}{1+u} \frac{\partial u}{\partial x_i} \\ \frac{\partial^2 w}{\partial x_i^2} &= -\sqrt{\epsilon} \frac{1}{(1+u)^2} \left[\frac{\partial u}{\partial x_i} \right]^2 + \sqrt{\epsilon} \frac{1}{1+u} \frac{\partial^2 u}{\partial x_i^2} \end{aligned}$$

Again, plugging in the l.h.s. of equation 7 we get

$$\frac{\epsilon}{(1+u)^2} \|\nabla u\|^2 + \frac{-\epsilon}{1+u} \Delta u + 1 - \frac{\epsilon}{(1+u)^2} \|\nabla u\|^2$$

that again vanishes in view of equation 6. \square

Theorem 5. *Let O be an open set, and u_ϵ a solution to equation 6 in O , $u_\epsilon = 0$ on ∂O and $u_\epsilon \neq 0$ in O . Then $w_\epsilon \rightarrow \text{sign}(u_\epsilon) d_{\partial O}$ pointwise uniformly in any compact subset $O \cup \partial O$.*

Note that $\text{sign}(u_\epsilon)$ is well defined in O since we assume that O does not vanish in O .

Proof. First, assume $\text{sign}(u_\epsilon) > 0$ in O . Then, u_ϵ satisfies:

$$\begin{aligned} -\epsilon \Delta u_\epsilon + u_\epsilon - 1 &= 0 && \text{in } O \\ u_\epsilon &= 0 && \text{in } \partial O \end{aligned}$$

Now the change of variables $v_\epsilon = 1 - u_\epsilon$ leads to

$$\begin{aligned} \frac{1}{2} \Delta v_\epsilon &= \frac{1}{2\epsilon} v_\epsilon && \text{in } \Omega \\ v_\epsilon &= 1 && \text{in } \partial \Omega \end{aligned} \quad (8)$$

Therefore, Theorem 2.3 in (Varadhan, 1967) with $\lambda = \frac{1}{2\epsilon}$ now implies

$$-\sqrt{\epsilon} \log(v_\epsilon) = -\sqrt{\epsilon} \log(1 - u_\epsilon) \xrightarrow{\epsilon \rightarrow 0} d_{\partial O}$$

uniformly in compact subsets of $O \cup \partial O$.

In the case $\text{sign}(u_\epsilon) < 0$ in O , u_ϵ satisfies:

$$\begin{aligned} -\epsilon \Delta u_\epsilon + u_\epsilon + 1 &= 0 && \text{in } O \\ u_\epsilon &= 0 && \text{in } \partial O \end{aligned}$$

The change of coordinates $v_\epsilon = u_\epsilon + 1$ now leads again to v_ϵ satisfying equation 8 and invoking (Varadhan, 1967) again implies that

$$\sqrt{\epsilon} \log(v_\epsilon) = \sqrt{\epsilon} \log(1 + u_\epsilon) \xrightarrow{\epsilon \rightarrow 0} -d_{\partial O}$$

uniformly in compact subsets of $O \cup \partial O$. Putting the two cases together and comparing to equation 14 (in the main paper) proves the theorem. \square

3. Networks in Sobolev spaces

Let $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$ be a multilayer perceptron (MLP) with ReLU activation. That is, $f(\mathbf{x}; \theta)$ is composed of layers of the form $\mathbf{z} = \sigma(\mathbf{W}\mathbf{y} + \mathbf{b})$, where \mathbf{W} , \mathbf{b} are the parameters, collectively defining $\theta \in \mathbb{R}^p$, and $\sigma(s) = \max\{0, s\}$ is the ReLU applied entry-wise. We consider the functions $f(\cdot; \theta)$, for a fixed θ : these are piecewise linear and continuous. In fact, each linear piece $L_k = f|_{\Omega_k}$ is defined over a polytope $\Omega_k \subset \mathbb{R}^d$. For the analysis of the loss functions we discuss in this paper we require a complete function space that contains these neural functions and allow discussing their "derivatives" and convergence. The most natural such space is $W^{1,p}(\Omega)$, $p \in [1, \infty)$, the Sobolev space of all $L^p(\Omega)$ functions with first weak derivatives also in $L^p(\Omega)$. We note that MLP with Softplus activation is smooth in the classical sense hence in particular belongs to $W^{1,p}(\Omega)$.

Lemma 1. *The functions $f(\cdot; \theta)$, where f is an MLP with ReLU activation belongs to $W^{1,p}(\Omega)$, for all $p \in [1, \infty]$, and Ω Lipschitz domain.*

Proof. Denote $f(\mathbf{x}) := f(\mathbf{x}; \theta)$. To show that $f \in W^{1,p}(\Omega)$ we need to show two things: $f \in L^p(\Omega)$ (which is clear since f is measurable and bounded as well as Ω is bounded), and that it has *weak derivatives* in $L^p(\Omega)$, that is for every $i \in [d]$ there exists a function $g \in L^p(\Omega)$ so that

$$\int_{\Omega} f \frac{\partial \psi}{\partial x_i} = - \int_{\Omega} g \psi, \quad (9)$$

for all compactly supported smooth functions $\psi \in C_c^\infty(\Omega)$. Let $\Omega = \cup_k \Omega_k$ be the decomposition of Ω to the subdomains where $L_k = f|_{\Omega_k}$ is linear. The multivariate integration by parts formula provides:

$$\int_{\Omega_k} f \frac{\partial \psi}{\partial x_i} = \int_{\partial \Omega_k} f \psi \langle \mathbf{e}_i, \mathbf{n} \rangle - \int_{\Omega_k} \frac{\partial f}{\partial x_i} \psi. \quad (10)$$

Therefore, a natural candidate for the weak derivative $\frac{\partial f}{\partial x_i}$ is the piecewise constant g that equals $\frac{\partial L_k}{\partial x_i}$ in the interior of each Ω_k . Indeed, noticing that $\langle \mathbf{e}_i, \mathbf{n} \rangle$ flips sign when the normal \mathbf{n} flips sign, we see that all interior contributions of the boundary integrals $\int_{\partial \Omega_k} f \psi \langle \mathbf{e}_i, \mathbf{n} \rangle$ cancel, and since ψ vanishes on the boundary $\partial \Omega$, summing equation 10 over k we get equation 9. \square




Figure 1. Ablation of ϵ . From left to right $\epsilon = 1.0, 0.1, 0.05, 0.01, 0.005$; $\lambda = c\epsilon^{0.3}$ according to Theorem 2, and $c \in \{1, 10\}$.

4. Experiments and implementation details

In all the experiments we used $\epsilon = 0.01$, and did a parameter search over $\lambda \in \{0.2, 0.3, 0.5, 1, 10, 20\}$ (note that $\epsilon^{1/3} \approx 0.2$, and $1/3$ is in the range suggested by Theorem 2), and similarly for μ . The gradients $\nabla u(\mathbf{x})$ are computed with automatic differentiation. A single training iteration with batch size $16k$ takes 0.16sec, on an NVIDIA Quadro GP100.

4.1. Fourier features

In some experiments we used Fourier features (Tancik et al., 2020) as the first constant layer in the network, $\delta_k : \mathbb{R}^d \rightarrow \mathbb{R}^{2kd}$, where k is a parameter representing the number of frequencies used. For $\mathbf{x} \in \mathbb{R}^d$ the Fourier feature layer is defined to be a vector $\delta_k(\mathbf{x}) \in \mathbb{R}^{2kd}$ with the real and imaginary parts of $\exp(i2^\omega \pi x_j)$, $\omega \in [k]$, $j \in [d]$ as entries.

4.2. Metrics

The distance between two point clouds $\mathcal{X}_1, \mathcal{X}_2$ is computed as in (Williams et al., 2019) using the standard one-sided and double-sided Chamfer and Hausdorff distances. We denote by $d_C^\rightarrow(\mathcal{X}_1, \mathcal{X}_2)$, $d_C(\mathcal{X}_1, \mathcal{X}_2)$ the one-sided and double-sided ℓ_1 Chamfer distance; and by $d_H^\rightarrow(\mathcal{X}_1, \mathcal{X}_2)$, $d_H(\mathcal{X}_1, \mathcal{X}_2)$ the one-sided and double-sided Hausdorff distance. These are defined as follows:

$$\begin{aligned} d_C^\rightarrow(\mathcal{X}, \mathcal{Y}) &= \frac{1}{|\mathcal{X}|} \sum_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{x} - \mathbf{y}\|_2 \\ d_C(\mathcal{X}, \mathcal{Y}) &= \frac{1}{2} (d_C^\rightarrow(\mathcal{X}, \mathcal{Y}) + d_C^\rightarrow(\mathcal{Y}, \mathcal{X})) \\ d_H^\rightarrow(\mathcal{X}, \mathcal{Y}) &= \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{x} - \mathbf{y}\|_2 \\ d_H(\mathcal{X}, \mathcal{Y}) &= \max \{d_H^\rightarrow(\mathcal{X}, \mathcal{Y}) + d_H^\rightarrow(\mathcal{Y}, \mathcal{X})\}, \end{aligned}$$

To measure the distance between a surface and a point cloud, or between two surfaces, we first sample each surface \mathcal{S} densely (i.e., with $10m$ uniformly random points) $\mathcal{Y} \subset \mathcal{S}$, and then measure the distance between the corresponding points clouds as described above.

4.3. 2D evaluation

Figure 1 demonstrates ablation of ϵ ; we show w_ϵ for ϵ in a range of values $\epsilon \in \{1.0, 0.1, 0.05, 0.01, 0.005\}$, where λ is chosen according to Theorem 2, i.e., $\lambda = c\epsilon^{0.3}$, and we take $c \in \{1, 10\}$. Note that as ϵ decreases w_ϵ is closer to a distance function.

4.4. Surface reconstruction benchmark

This benchmark consists of 5 models, the input (train data) to each model is a point cloud \mathcal{X} of size $170k - 290k$, and corresponding normal data $\mathbf{n} : \mathcal{X} \rightarrow \mathcal{S}(\mathbb{R}^3)$, as well as ground-truth test point cloud \mathcal{Y} . In this experiment we used the PHASE loss in equation 20 (in the main paper) with the normal loss equation 18 (in the main paper), and $\mu = \lambda = 10$. The batch size was taken to be $16k$ and we performed $100k$ iterations, which roughly correspond to $5k - 10k$ epochs. Figure 2 shows qualitative results of PHASE.

PHASE and Fourier Features. Training PHASE with Fourier features and normal data leads to comparable results to PHASE alone, but requires order of magnitude less iterations. In Table 1 we show the results of PHASE trained with Fourier Features ($k = 6$), on the surface reconstruction benchmark of (Williams et al., 2019) (i.e., point clouds with normals), with the same parameters $\mu = \lambda = 10$, as before. In this case we trained the model for $10k$ iterations with batch-size $16k$. See Table 1 in the main paper for comparison with other methods. Note that PHASE and PHASE+FF are roughly equivalent according to the distance metric scores, maybe with slight advantage to PHASE without FF. However, PHASE+FF was trained for $10k$ iterations versus $100k$ iterations in the case of PHASE without FF, and it also exhibits slightly more high frequency details not captured by the metric scores, see Figure 2.




Figure 2. Surface reconstruction benchmark (Williams et al., 2019) comparing PHASE and IGR with/without normal data and Fourier features. Note that IGR has a slight tendency to add redundant surface area that is amplified in the more challenging points-only case with Fourier features, where PHASE, due to its minimal surface perimeter property is less susceptible to high frequency artifacts.

Model	Method	Ground Truth		Scans	
		d_C	d_H	d_C^{\rightarrow}	d_H^{\rightarrow}
Anchor	PHASE+FF	0.24	5.30	0.09	1.17
	PHASE	0.21	4.29	0.09	1.23
Daratech	PHASE+FF	0.18	2.53	0.08	1.79
	PHASE	0.18	2.92	0.08	1.80
DC	PHASE+FF	0.15	2.32	0.05	2.77
	PHASE	0.15	2.52	0.05	2.78
Gargoyle	PHASE+FF	0.16	3.68	0.06	0.87
	PHASE	0.16	3.14	0.07	1.09
Lord Quas	PHASE+FF	0.12	0.84	0.04	0.94
	PHASE	0.11	0.96	0.04	0.96

Table 1. PHASE+FF results on the surface reconstruction benchmark of (Williams et al., 2019). See also Table 1 in main paper for more methods.

4.5. Learning from point clouds and Fourier features

In this example we again worked with the surface reconstruction benchmark but explored the more challenging case of removing the normal data and using only the point clouds \mathcal{X} . We wanted to see the bias caused by the high frequency methods and compared PHASE and IGR in the case of adding Fourier features $k = 6$ to both. For PHASE we used $\lambda = 10$ as before, but smaller μ , i.e., $\mu = 0.5$, as we found larger μ is more unstable. For IGR we did a parameter search and chose the one that provides lowest error metrics on the Gargoyle model. Figure 2 shows the results of both. Note that IGR is much more sensitive to the high frequency bias and adds extraneous surface parts. In contrast, PHASE, due to its minimal perimeter property is much less susceptible to high frequency bias, although is not completely immune to this bias for point cloud data alone, as can be inspected in the image. For point cloud and normal data, we found PHASE to be even more robust to the high frequency bias, see PHASE+FF (points and normals) experiment above.

4.6. Parameters scan

Table 2, logs the *entire* range of reconstruction metric errors for *all* hyperparams options $\lambda, \mu \in \{0.2, 1, 10\}$, trained with a smaller network: an MLP with 5 layers of 256 neurons each. Note that the range intervals are rather small (except Anchor, where the worst and best results are shown in Figure 3), and in fact the best results further improve the state of art in some of the cases.

Metric	Anchor	Daratech	DC	Gargoyle	Lord Quas
d_C	[0.229, 0.457]	[0.174, 0.199]	[0.143, 0.161]	[0.158, 0.172]	[0.112, 0.127]
d_H	[5.02, 14.8]	[2.43, 3.68]	[1.61, 2.38]	[3.23, 4.69]	[0.77, 2.32]

Table 2. Reconstruction ranges for *all combinations* of params.




Figure 3. The *worst* and *best* examples in Table 1, Anchor.

References

- Gilbarg, D. and Trudinger, N. S. *Elliptic partial differential equations of second order*. springer, 2015.
- Modica, L. The gradient theory of phase transitions and the minimal interface criterion. *Archive for Rational Mechanics and Analysis*, 98(2):123–142, 1987.
- Rindler, F. *Calculus of Variations*. Springer, 2018.
- Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- Varadhan, S. R. S. 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.
- Williams, F., Schneider, T., Silva, C., Zorin, D., Bruna, J., and Panozzo, D. Deep geometric prior for surface reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10130–10139, 2019.