
Implicit Geometric Regularization for Learning Shapes

Amos Gropp¹ Lior Yariv¹ Niv Haim¹ Matan Atzmon¹ Yaron Lipman¹

Abstract

Representing shapes as level sets of neural networks has been recently proved to be useful for different shape analysis and reconstruction tasks. So far, such representations were computed using either: (i) pre-computed implicit shape representations; or (ii) loss functions explicitly defined over the neural level sets.

In this paper we offer a new paradigm for computing high fidelity implicit neural representations directly from raw data (i.e., point clouds, with or without normal information). We observe that a rather simple loss function, encouraging the neural network to vanish on the input point cloud and to have a unit norm gradient, possesses an implicit geometric regularization property that favors smooth and natural zero level set surfaces, avoiding bad zero-loss solutions.

We provide a theoretical analysis of this property for the linear case, and show that, in practice, our method leads to state of the art implicit neural representations with higher level-of-details and fidelity compared to previous methods.

1. Introduction

Recently, level sets of neural networks have been used to represent 3D shapes (Park et al., 2019; Atzmon et al., 2019; Chen & Zhang, 2019; Mescheder et al., 2019), i.e.,

$$\mathcal{M} = \{ \mathbf{x} \in \mathbb{R}^3 \mid f(\mathbf{x}; \theta) = 0 \}, \quad (1)$$

where $f : \mathbb{R}^3 \times \mathbb{R}^m \rightarrow \mathbb{R}$ is a multilayer perceptron (MLP); we call such representations *implicit neural representations*. Compared to the more traditional way of representing surfaces via implicit functions defined on volumetric grids (Wu

¹Department of Computer Science & Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel. Correspondence to: Amos Gropp <amos.gropp@weizmann.ac.il>.

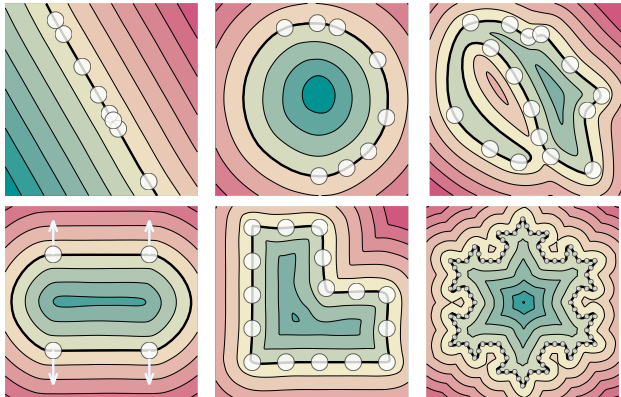


Figure 1. Learning curves from 2D point clouds (white disks) using our method; black lines depict the zero level sets of the trained neural networks, \mathcal{M} . The implicit geometric regularization drives the optimization to reach plausible explanation of the data.

et al., 2016; Choy et al., 2016; Dai et al., 2017; Stutz & Geiger, 2018), neural implicit representations have the benefit of relating the degrees of freedom of the network (i.e., parameters) directly to the shape rather than to the fixed discretization of the ambient 3D space. So far, most previous works using implicit neural representations computed f with 3D supervision; that is, by comparing f to a known (or pre-computed) implicit representation of some shape. (Park et al., 2019) use a regression loss to approximate a pre-computed signed distance functions of shapes; (Chen & Zhang, 2019; Mescheder et al., 2019) use classification loss with pre-computed occupancy function.

In this work we are interested in working directly with raw data: Given an input point cloud $\mathcal{X} = \{ \mathbf{x}_i \}_{i \in I} \subset \mathbb{R}^3$, with or without normal data, $\mathcal{N} = \{ \mathbf{n}_i \}_{i \in I} \subset \mathbb{R}^3$, our goal is to compute θ such that $f(\mathbf{x}; \theta)$ is approximately the signed distance function to a plausible surface \mathcal{M} defined by the point data \mathcal{X} and normals \mathcal{N} .

Some previous works are constructing implicit neural representations from raw data. In (Atzmon et al., 2019) no 3D supervision is required and the loss is formulated directly on the zero level set \mathcal{M} ; iterative sampling of the zero level set is required for formulating the loss. In a more recent work, (Atzmon & Lipman, 2020) use unsigned regression to introduce good local minima that produces useful implicit neural representations, with no 3D supervision and no zero level set

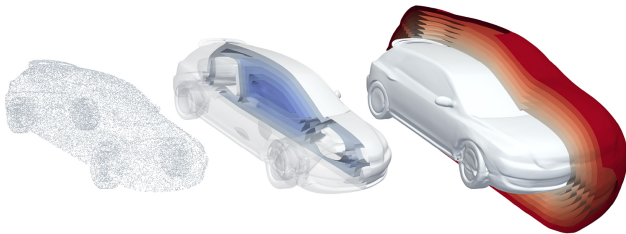


Figure 2. The level sets of an MLP (middle and right) trained with our method on an input point cloud (left); positive level sets are in red; negative are in blue; the zero level set, representing the approximated surface is in white.

sampling. Both of these works, however, explicitly enforce some regularization on the zero level set. Another possible solution is to compute, as a pre-training stage, some implicit representation of the surface using a classical surface reconstruction technique and use one of the previously mentioned methods to construct the neural implicit representation. This approach has two drawbacks: First, finding implicit surface representation from raw data is a notoriously hard (Berger et al., 2017); second, decoupling the reconstruction from the learning stage would hinder *collective* learning and reconstruction of shapes. For example, information from one shape will not be used to improve reconstruction of a different, yet a similar shape; nor consistent reconstructions will be produced.

The goal of this paper is to introduce a novel approach for learning neural shape representations directly from raw data using *implicit geometric regularization*. We show that state of the art implicit neural representations can be achieved without 3D supervision and/or a direct loss on the zero level set \mathcal{M} . As it turns out, stochastic gradient optimization of a simple loss that fits an MLP f to a point cloud data \mathcal{X} , with or without normal data \mathcal{N} , while encouraging unit norm gradients $\nabla_{\mathbf{x}}f$, consistently reaches good local minima, favoring smooth, yet high fidelity, zero level set surfaces \mathcal{M} approximating the input data \mathcal{X} and \mathcal{N} . Figure 1 shows several implicit neural representation of 2D data computed using our method; note that although there is an infinite number of solutions with neural level sets interpolating the input data, the optimization reaches solutions that provide natural and intuitive reconstructions. Figure 2 shows the level sets of an MLP trained with our method from a point cloud shown on the left.

The preferable local minima found by the optimization procedure could be seen as a geometric version of the well known *implicit regularization* phenomenon in neural network optimization (Neyshabur et al., 2014; Neyshabur, 2017; Soudry et al., 2018). Another, yet different, treatment of geometric implicit regularization was discussed in (Williams et al., 2019b) where reducing an entropic regularization in the loss still maintains consistent and smooth neural chart representation.

Although we do not provide a full theory supporting the implicit geometric regularization phenomenon, we analyze the linear case, which is already non-trivial due to the non-convex unit gradient norm term. We prove that if the point cloud \mathcal{X} is sampled (with small deviations) from a hyperplane \mathcal{H} and the initial parameters of the linear model are randomized, then, with probability one, gradient descent converges to the (approximated) signed distance function of the hyperplane \mathcal{H} , avoiding bad critical solutions. We call this property *plane reproduction* and advocate it as a useful geometric manifestation of implicit regularization in neural networks.

In practice, we perform experiments with our method, building implicit neural representations from point clouds in 3D and learning collections of shapes directly from raw data. Our method produces state of the art surface approximations, showing significantly more detail and higher fidelity compared to alternative techniques. Our code is available at <https://github.com/amosgropp/IGR>.

In summary, our paper’s main contribution is two-fold:

- Suggesting a new paradigm, building on implicit geometric regularization, for computing high fidelity implicit neural representations, directly from raw data.
- Providing a theoretical analysis for the linear case showing gradient descent of the suggested loss function avoids bad critical solutions.

2. Method

Given an input point cloud $\mathcal{X} = \{\mathbf{x}_i\}_{i \in I} \subset \mathbb{R}^3$, with or without normal data, $\mathcal{N} = \{\mathbf{n}_i\}_{i \in I} \subset \mathbb{R}^3$, our goal is to compute parameters θ of an MLP $f(\mathbf{x}; \theta)$, where $f : \mathbb{R}^3 \times \mathbb{R}^m \rightarrow \mathbb{R}$, so that it approximates a signed distance function to a plausible surface \mathcal{M} defined by the point cloud \mathcal{X} and normals \mathcal{N} .

We consider a loss of the form

$$\ell(\theta) = \ell_{\mathcal{X}}(\theta) + \lambda \mathbb{E}_{\mathbf{x}} (\|\nabla_{\mathbf{x}}f(\mathbf{x}; \theta)\| - 1)^2, \quad (2)$$

where $\lambda > 0$ is a parameter, $\|\cdot\| = \|\cdot\|_2$ is the euclidean 2-norm, and

$$\ell_{\mathcal{X}}(\theta) = \frac{1}{|I|} \sum_{i \in I} (|f(\mathbf{x}_i; \theta)| + \tau \|\nabla_{\mathbf{x}}f(\mathbf{x}_i; \theta) - \mathbf{n}_i\|)$$

encourages f to vanish on \mathcal{X} and, if normal data exists (i.e., $\tau = 1$), that $\nabla_{\mathbf{x}}f$ is close to the supplied normals \mathcal{N} .

The second term in equation 2 is called the *Eikonal term* and encourages the gradients $\nabla_{\mathbf{x}}f$ to be of unit 2-norm. The expectation is taken with respect to some probability distribution $\mathbf{x} \sim \mathcal{D}$ in \mathbb{R}^3 .

The motivation for the Eikonal term stems from the Eikonal partial differential equation: A solution f (in the sense of (Crandall & Lions, 1983)) to the Eikonal equation, i.e.,

$$\|\nabla_x f(x)\| = 1, \quad (3)$$

where f vanishes on \mathcal{X} , with gradients \mathcal{N} , will be a signed distance function and a global minimum of the loss in equation 2. Note however, that for point boundary data \mathcal{X}, \mathcal{N} the solution to equation 3 is not unique.

Implicit geometrical regularization. When optimizing the loss in equation 2, two questions immediately emerge: First, why a critical point θ_* that is found by the optimization algorithm leads $f(x; \theta_*)$ to be a signed distance function? Usually, adding a quadratic penalty with a finite weight is not guaranteed to provide *feasible* critical solutions (Nocedal & Wright, 2006), i.e., solutions that satisfy the desired constraint (in our case, unit length gradients). Second, even if the critical solution found is a signed distance function, why would it be a plausible one? There is an infinite number of signed distance functions vanishing on arbitrary discrete sets of points \mathcal{X} with arbitrary normal directions \mathcal{N} .

Remarkably, optimizing equation 2 using stochastic gradient descent (or a variant thereof) results in solutions that are close to a signed distance function with a smooth and surprisingly plausible zero level set. For example, Figure 1 depicts the result of optimizing equation 2 in the planar case ($d = 2$) for different input point clouds \mathcal{X} , with and without normal data \mathcal{N} ; the zero level sets of the optimized MLP are shown in black. Note that the optimized MLP is close to a signed distance function as can be inspected from the equidistant level sets.

The inset shows an alternative signed distance function that would achieve zero loss in equation 2 for the top-left example in Figure 1, avoided by the optimization algorithm that chooses to reconstruct a straight line in this case. This is a consequence of the *plane reproduction* property. In Section 4 we provide a theoretical analysis of the plane reproduction property for the linear model case $f(x; w) = w^T x$ and prove that if \mathcal{X} is sampled with small deviations from a hyperplane \mathcal{H} , then gradient descent provably avoids bad critical solutions and converges with probability one to the approximated signed distance function to \mathcal{H} .



Computing gradients. Incorporating the gradients $\nabla_x f$ in the loss above could be done using numerical estimates of the gradient. A better approach is the following: every layer of the MLP f has the form $y^{\ell+1} = \sigma(Wy^\ell + b)$, where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a non-linear differentiable activation (we use Softplus) applied entrywise, and W, b are the layer's

learnable parameters. Hence, by the chain-rule the gradients satisfy

$$\nabla_x y^{\ell+1} = \text{diag}(\sigma'(Wy^\ell + b)) W \nabla_x y^\ell, \quad (4)$$

where $\text{diag}(z)$ is arranging its input vector $z \in \mathbb{R}^k$ on the diagonal of a square matrix $\mathbb{R}^k \times \mathbb{R}^k$ and σ' is the derivative of σ . Equation 4 shows that $\nabla_x f(x; \theta)$ can

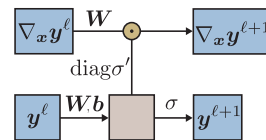


Figure 3. $(f, \nabla_x f)$ network.

be constructed as a neural-network in conjunction with $f(x; \theta)$, see Figure 3 for illustration of a single layer of a network computing both $f(x; \theta)$ and $\nabla_x f(x; \theta)$. In practice, implementing $\nabla_x f(x; \theta)$ using *Automatic Differentiation* packages seems to be a simple alternative.

3. Previous work and discussion

3.1. Deep learning of 3D shapes

There are numerous deep learning based methods applied to 3D shapes. Here we review the main approaches, emphasizing the 3D data representation being used and discuss relations to our approach.

RGVF: regular grid volumetric function. Maybe the most popular representation for 3D shapes is via a scalar function defined over a regular volumetric grid (RGVF); the shape is then defined as the zero level set of the function. One option is to use an indicator function (Choy et al., 2016; Girdhar et al., 2016; Wu et al., 2016; Yan et al., 2016; Tulsiani et al., 2017; Yang et al., 2017). This is a natural generalization of images to 3D, thus enabling easy adaptation of the successful Convolutional Neural Networks (CNNs) architectures. Tatarchenko et al. (2017) addressed the computation efficiency challenge stemming from the cubic grid size. In (Wu et al., 2016), a variant of generative adversarial network (Goodfellow et al., 2014) is proposed for 3D shapes generation. More generally, researchers have suggested using other RGVFs (Dai et al., 2017; Riegler et al., 2017; Stutz & Geiger, 2018; Liao et al., 2018; Michalkiewicz et al., 2019; Jiang et al., 2019). In (Dai et al., 2017), the RGVF models a signed distance function to a shape, where an encoder-decoder network is trained for the task of shape completion. Similar RGVF representation is used in (Jiang et al., 2019) for the task of multi-view 3D reconstruction. However, they learn the signed distance function representation based on differentiable rendering technique, without requiring pre-training 3D supervision. Another implicit RGVF representation is used in (Michalkiewicz et al., 2019) for the task of image to shape prediction; they introduced a loss function inspired by the *level set method* (Osher et al., 2004) based surface reconstruction techniques (Zhao et al., 2000; 2001), operating directly on level sets of RGVF.

The RGVF has several shortcomings compared to the implicit neural representations in general and our approach in particular: (i) The implicit function is defined only at grid points, requiring an interpolation scheme to extend it to interior of cells; normal approximation would require divided differences; (ii) It requires cubic-size grid and is not data-dependent, i.e., it does not necessarily adhere to the specific geometry of the shapes one wishes to approximate. (iii) A version of the Eikonal regularization term (second term in equation 2) was previously used with RGVF representations in Michalkiewicz et al. (2019); Jiang et al. (2019). These works incorporated Eikonal regularization as a normalization term, in combination with other explicit reconstruction and/or regularization terms. The key point in our work is that the Eikonal term *alone* can be used for (implicitly) regularizing the zero level set. Furthermore, it is not clear whether the implicit regularization property holds, and to what extent, in the fixed volumetric grid scenario. Lastly, in the RGVF setting the gradients are computed using finite differences or via some fixed basis function.

Neural parametric surfaces. Surfaces can also be described explicitly as a collection of charts (in an atlas), where each chart $f : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is a local parametrization. However, finding a consistent atlas covering of the surface could be challenging. (Groueix et al., 2018; Deprelle et al., 2019) suggested modeling charts as MLPs; (Williams et al., 2019b) focused on an individual surface reconstruction, introducing a method to improve atlas consistency by minimizing disagreement between neighboring charts. Some works have considered global surface parametrizations (Sinha et al., 2016; 2017; Maron et al., 2017). Global parametrizations produce consistent coverings, however at the cost of introducing parametrizations with high distortion.

The benefit of parametric representations over implicit neural representations is that the shape can be easily sampled; the main shortcoming is that it is very challenging to produce a set of perfectly overlapping charts, a property that holds by construction for implicit representations.

Hybrid representations. Deng et al. (2019); Chen et al. (2019); Williams et al. (2019a) suggested representations based on the fact that every solid can be decomposed into a union of convex sets. As every convex set can be represented either as an intersection of hyper-planes or a convex-hull of vertices, transforming between a shape explicit and implicit representation can be done relatively easily.

3.2. Solving PDEs with neural networks

Our proposed training objective equation 2 can be interpreted as a quadratic penalty method for solving the *Eikonal Equation* (equation 3). However, the Eikonal equation is a non-linear wave equation and requires boundary conditions

of the form

$$f(x) = 0, \quad x \in \partial\Omega,$$

where $\Omega \subset \mathbb{R}^3$ is a well-behaved open set with boundary, $\partial\Omega$. The viscous solution (Crandall & Lions, 1983; Crandall et al., 1984) to the Eikonal equation is unique in this case. Researchers are trying to utilize neural networks to solve differential equations (Yadav et al., 2015). Perhaps, more related to our work are Sirignano & Spiliopoulos (2018); Raissi et al. (2017a;b) suggesting deep neural networks as a non-linear function space for approximating PDE solutions. Their training objective, similarly to ours, can be seen as a penalty version of the original PDE.

The main difference from our setting is that in our case the boundary conditions of the Eikonal equation do not hold, as we use a *discrete* set of points \mathcal{X} . In particular, any well-behaved domain Ω that contains \mathcal{X} in its boundary, i.e., $\mathcal{X} \subset \partial\Omega$ would form a valid initial condition to the Eikonal equation with $\partial\Omega$ as the zero level set of its solution. Therefore, from PDE theory point of view, the problem equation 2 is trying to solve is ill-posed with infinite number of solutions. The main observation of this paper is that implicit geometry regularization in fact *chooses a favorable solution* out of this solution space.

4. Analysis of the linear model and plane reproduction

In this section we provide some justification for using the loss in equation 2 by analyzing the linear network case. That is, we consider a linear model $f(x; w) = w^T x$ where the loss in equation 2 takes the form

$$\ell(w) = \sum_{i \in I} (w^T x_i)^2 + \lambda (\|w\|^2 - 1)^2, \quad (5)$$

where for simplicity we used squared error and removed the term involving normal data; we present the analysis in \mathbb{R}^d rather than \mathbb{R}^3 .

We are concerned with the *plane reproduction* property, namely, assuming our point cloud \mathcal{X} is sampled approximately from a plane \mathcal{H} , then gradient descent of the loss in equation 5 converges to the approximate signed distance function to \mathcal{H} .

To this end, assume our point cloud data $\mathcal{X} = \{x_i\}_{i \in I}$ satisfies $x_i = y_i + \varepsilon_i$, where y_i , span some $d-1$ -dimension hyperplane $\mathcal{H} \subset \mathbb{R}^d$ that contains the origin, and ε_i are some small deviations satisfying $\|\varepsilon_i\| < \epsilon$. We will show that: (i) For $\lambda > \frac{c\epsilon}{2}$, where c is a constant depending on y_i , the loss in equation 5 has two global minima that correspond to the (approximated) signed distance functions to \mathcal{H} (note there are two signed distance functions to \mathcal{H} differing by a sign); the rest of its critical points are either saddle points

or local maxima. (ii) Using the characterization of critical points and known properties of gradient descent (Ge et al., 2015; Lee et al., 2016) we can prove that applying gradient descent

$$\mathbf{w}^{k+1} = \mathbf{w}^k - \alpha \nabla_{\mathbf{w}} \ell(\mathbf{w}^k), \quad (6)$$

from a random initialization \mathbf{w}^0 and sufficiently small step-size $\alpha > 0$, will converge, with probability one, to one of the global minima, namely to the approximated signed distance function to \mathcal{H} .

Change of coordinates. We perform a change of coordinates: Let $\sum_{i \in I} \mathbf{x}_i \mathbf{x}_i^T = \mathbf{U} \mathbf{D} \mathbf{U}^T$, $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_d)$, $\mathbf{D} = \text{diag}(\lambda_1, \dots, \lambda_d)$ be a spectral decomposition, and $0 \leq \lambda_1 < \lambda_2 \leq \dots \leq \lambda_d$. Using perturbation theory for eigenvalues and eigenvectors of hermitian matrices one proves:

Lemma 1. *There exists constants $c, c' > 0$ depending on $\{\mathbf{y}_i\}_{i \in I}$ so that $\lambda_1 \leq c\epsilon$ and $\|\mathbf{u}_1 - \mathbf{n}\| \leq c'\epsilon$, where \mathbf{n} is a normal direction to \mathcal{H} .*

Proof (Lemma 1). Let $\sum_{i \in I} \mathbf{x}_i \mathbf{x}_i^T = \sum_{i \in I} \mathbf{y}_i \mathbf{y}_i^T + \mathbf{y}_i \boldsymbol{\epsilon}_i^T + \boldsymbol{\epsilon}_i \mathbf{y}_i^T + \boldsymbol{\epsilon}_i \boldsymbol{\epsilon}_i^T = \sum_{i \in I} \mathbf{y}_i \mathbf{y}_i^T + \mathbf{E}$. Now use hermitian matrix eigenvalue perturbation theory, e.g., (Stewart, 1990), Section IV:4, and perturbation theory for simple eigenvectors, see (Stewart, 1990) Section V:2.2-2.3, to conclude. \square

Then, performing the change of coordinates, $\mathbf{q} = \mathbf{U}^T \mathbf{w}$, in equation 5 leads to the diagonalized form

$$\ell(\mathbf{q}) = \mathbf{q}^T \mathbf{D} \mathbf{q} + \lambda \left(\|\mathbf{q}\|^2 - 1 \right)^2, \quad (7)$$

where $\|\mathbf{w}\| = \|\mathbf{q}\|$ due to the invariance of the euclidean norm to orthogonal transformations. The plane \mathcal{H} in the transformed coordinates is \mathbf{e}_1^\perp , where $\mathbf{e}_1 \in \mathbb{R}^d$ is the first standard basis vector.

Classification of critical points. Next we classify the critical points of our loss. The gradient of the loss in equation 7 is

$$\nabla_{\mathbf{q}} \ell(\mathbf{q})^T = 2 \left(\mathbf{D} + 2\lambda(\|\mathbf{q}\|^2 - 1) \mathbf{I} \right) \mathbf{q}, \quad (8)$$

where \mathbf{I} is the identity matrix. The Hessian takes the form

$$\nabla_{\mathbf{q}}^2 \ell(\mathbf{q}) = 2\mathbf{D} + 4\lambda \left(\|\mathbf{q}\|^2 - 1 \right) \mathbf{I} + 8\lambda \mathbf{q} \mathbf{q}^T. \quad (9)$$

We prove:

Theorem 1. *If $\lambda > \frac{\lambda_1}{2}$, then the loss in equation 7 (equivalently, equation 5) has at-least 3 and at-most $2d + 1$ critical points. Out of which, the following two correspond to the approximated signed distance functions to the plane \mathbf{e}_1^\perp , and are global minima:*

$$\pm \mathbf{q} = \pm \sqrt{1 - \frac{\lambda_1}{2\lambda}} \mathbf{e}_1.$$

The rest of the critical points are saddle points or local maxima.

Before proving this theorem we draw some conclusions. The global minima in the original coordinate system are $\pm \mathbf{w} = \pm \sqrt{1 - \frac{\lambda_1}{2\lambda}} \mathbf{u}_1$ that correspond to the approximate signed distance function to \mathcal{H} . Indeed, Lemma 1 implies that $\lambda_1/2\lambda \leq c\epsilon/2\lambda$ and $\|\mathbf{u}_1 - \mathbf{n}\| \leq c'\epsilon$. Therefore

$$\|\mathbf{w} - \mathbf{n}\| \leq \left(\frac{c}{2\lambda} + c' \right) \epsilon,$$

where we used the triangle inequality and $\sqrt{1-s} \geq 1-s$ for $s \in [0, 1]$. This shows that $\lambda > 0$ should be chosen sufficiently large compared to ϵ , the deviation of the data from planarity. In the general MLP case, one could consider this analysis locally noting that locally an MLP is approximately linear and the deviation from planarity is quantified locally by the *curvature* of the surface, e.g., the two principle surface curvatures σ_1, σ_2 that represent the reciprocal radii of two osculating circles (Do Carmo, 2016).

Proof (Theorem 1). First, let us find all critical points. Clearly, $\mathbf{q} = \mathbf{0}$ satisfies $\nabla_{\mathbf{q}} \ell(\mathbf{0}) = 0$. Now if $\mathbf{q} \neq \mathbf{0}$ then the only way equation 8 can vanish is if $2\lambda(\|\mathbf{q}\|^2 - 1) = -\lambda_j$, for some $j \in [d]$. That is, $\|\mathbf{q}_j\|^2 = 1 - \frac{\lambda_j}{2\lambda}$, and if the r.h.s. is strictly greater than zero then $\mathbf{q}_j = \sqrt{1 - \frac{\lambda_j}{2\lambda}} \mathbf{e}_j$ is a critical point, where \mathbf{e}_j is the j -th standard basis vector in \mathbb{R}^d . Note that also $-\mathbf{q}_j, j \in [d]$ are critical points. So in total we have found at-least 3 and up-to $2d + 1$ critical points: $\mathbf{0}, \pm \mathbf{q}_j, j \in [d]$.

Plugging these critical solutions into the Hessian formula, equation 9 we get

$$\nabla_{\mathbf{q}}^2 \ell(\pm \mathbf{q}_j) = 2\text{diag}(\lambda_1 - \lambda_j, \dots, \lambda_d - \lambda_j) + 8\left(\lambda - \frac{\lambda_j}{2}\right) \mathbf{e}_j \mathbf{e}_j^T.$$

From this equation we see that if $\lambda > \frac{\lambda_1}{2}$ then $\pm \mathbf{q}_1$ are local minima; and for all $\lambda > 0$, $\pm \mathbf{q}_j, j \geq 2$ are saddle points or local maxima (in particular, \mathbf{q}_d for small λ); i.e., the Hessian $\nabla_{\mathbf{q}}^2 \ell(\mathbf{q}_j), j \geq 2$, has at-least one strictly negative eigenvalue. Since $\ell(\mathbf{q}) \rightarrow \infty$ as $\|\mathbf{q}\| \rightarrow \infty$ we see that $\pm \mathbf{q}_1$ are also global minima. \square

Convergence to global minima. Given the classification of the critical points in Theorem 1 we can prove that

Theorem 2. *The gradient descent in equation 6, with sufficiently small step-size $\alpha > 0$ and a random initialization \mathbf{w}^0 will avoid the bad critical points of the loss function ℓ , with probability one.*

Since $\ell(\mathbf{w}) \rightarrow \infty$ as $\|\mathbf{w}\| \rightarrow \infty$ and $\ell(\mathbf{w}) \geq 0$ everywhere, an immediate consequence of this theorem is that equation 6 converges (up-to the constant step-size) with probability one

to one of the two global minima that represent the signed distance function to \mathcal{H} .

To prove Theorem 2 note that Theorem 1 shows that the Hessian (equation 9) evaluated at all bad critical points (i.e., excluding the two that correspond to the signed distance function) have at-least one strictly negative eigenvalue. Such saddle points are called *strict saddle points*. Theorem 4.1 in (Lee et al., 2016) implies that gradient descent will avoid all these strict saddle points. Since the loss is non-negative and blows-up at infinity the proof is concluded. \square

5. Implementation and evaluation details

Architecture. For representing shapes we used level sets of MLP $f(\mathbf{x}; \theta)$; $f: \mathbb{R}^3 \times \mathbb{R}^m \rightarrow \mathbb{R}$, with 8 layers, each contains 512 hidden units, and a single skip connection from the input to the middle layer as in (Park et al., 2019). The weights $\theta \in \mathbb{R}^m$ are initialized using the geometric initialization from (Atzmon & Lipman, 2020). We set our loss parameters (see equation 2) to $\lambda = 0.1, \tau = 1$.

Distribution \mathcal{D} . We defined the distribution \mathcal{D} for the expectation in equation 2 as the average of a uniform distribution and a sum of Gaussians centered at \mathcal{X} with standard deviation equal to the distance to the k -th nearest neighbor (we used $k = 50$). This choice of \mathcal{D} was used throughout all experiments.

Level set extraction. We extract the zero (or any other) level set of a trained MLP $f(\mathbf{x}; \theta)$ using the *Marching Cubes* meshing algorithm (Lorenson & Cline, 1987) on a uniform sampled grids of size ℓ^3 , where $\ell \in \{256, 512\}$.

Evaluation metrics. Our quantitative evaluation is based on the following collection of standard metrics of two point sets $\mathcal{X}_1, \mathcal{X}_2 \subset \mathbb{R}^3$: the Chamfer and Hausdorff distances,

$$d_C(\mathcal{X}_1, \mathcal{X}_2) = \frac{1}{2} (d_C^\rightarrow(\mathcal{X}_1, \mathcal{X}_2) + d_C^\rightarrow(\mathcal{X}_2, \mathcal{X}_1)) \quad (10)$$

$$d_H(\mathcal{X}_1, \mathcal{X}_2) = \max\{d_H^\rightarrow(\mathcal{X}_1, \mathcal{X}_2), d_H^\rightarrow(\mathcal{X}_2, \mathcal{X}_1)\} \quad (11)$$

where

$$d_C^\rightarrow(\mathcal{X}_1, \mathcal{X}_2) = \frac{1}{|\mathcal{X}_1|} \sum_{\mathbf{x}_1 \in \mathcal{X}_1} \min_{\mathbf{x}_2 \in \mathcal{X}_2} \|\mathbf{x}_1 - \mathbf{x}_2\|, \quad (12)$$

$$d_H^\rightarrow(\mathcal{X}_1, \mathcal{X}_2) = \max_{\mathbf{x}_1 \in \mathcal{X}_1} \min_{\mathbf{x}_2 \in \mathcal{X}_2} \|\mathbf{x}_1 - \mathbf{x}_2\| \quad (13)$$

are the one-sided Chamfer and Hausdorff distances (resp.).

6. Model evaluation

Signed distance function approximation. We start our evaluation by testing the ability of our trained model f to reproduce a signed distance function (SDF) to known

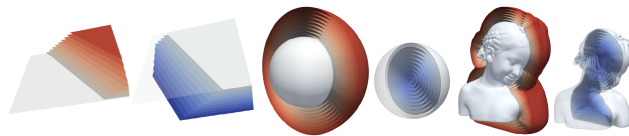


Figure 4. Level sets of MLPs trained with our method.

manifold surfaces. We tested: a plane, a sphere, and the Bimba model. In this experiment we used no normals and took the sample point cloud \mathcal{X} to be of infinite size (i.e., draw fresh point samples every iteration).

For each surface we separately train an MLP $f(\mathbf{x}; \theta)$ with sampling distribution \mathcal{D} . Table 1 logs the results, where we report mean \pm std of the relative error mea-

	Relative Error
Plane	0.003 ± 0.04
Sphere	0.004 ± 0.08
Bimba	0.008 ± 0.11

Table 1. SDF approximation.

sured at 100k random points. The relative error is defined by $\frac{|f(\mathbf{x}; \theta) - s(\mathbf{x})|}{|s(\mathbf{x})|}$, where $s: \mathbb{R}^3 \rightarrow \mathbb{R}$ is the ground truth signed distance function. Figure 4 provides a visual validation of the quality of our predictions, where equispaced positive (red) and negative (blue) level sets of the trained f are shown; the zero level sets are in white.

Fidelity and level of details. As mentioned above, previous works have suggested learning shapes as level sets of implicit neural networks (see equation 1) via regression or classification (Park et al., 2019; Mescheder et al., 2019; Chen & Zhang, 2019).

To test the faithfulness or fidelity of our learning method in comparison to regression we considered two raw scans (i.e., triangle soups) of a human, \mathcal{S} , from the D-Faust (Bogo et al., 2017) dataset. For each model we took a point sample $\mathcal{X} \subset \mathcal{S}$ of 250k points, and a corresponding normal sample \mathcal{N} (from the triangles). We used this data \mathcal{X}, \mathcal{N} to train an MLP with our method.

For regression, we trained an MLP with the same architecture using an approximated SDF data pre-computed using a standard local SDF approximation. Namely, $s(\mathbf{x}) = \mathbf{n}_*^T (\mathbf{x} - \mathbf{y}_*)$, where $\mathbf{y}_* = \arg \min_{\mathbf{y} \in \mathcal{S}} \|\mathbf{y} - \mathbf{x}\|_2$ and \mathbf{n}_* is the unit normal of the triangle containing \mathbf{y}_* . We trained the MLP with an L_1 regression loss $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}'} |f(\mathbf{x}; \theta) - s(\mathbf{x})|$, where \mathcal{D}' is a discrete distribution of 500k points (2 points for every point in \mathcal{X}) defined as in (Park et al., 2019). Figure 5 shows the zero level sets of the trained networks. Note that our method produced considerably more details than the regression approach. This improvement can be potentially attributed to two properties: First, our loss incorporates only points contained in the surface, while regression approximates the SDF nearby the actual surface. Second, we believe the implicit regularization property improves the fidelity of the learned level sets.

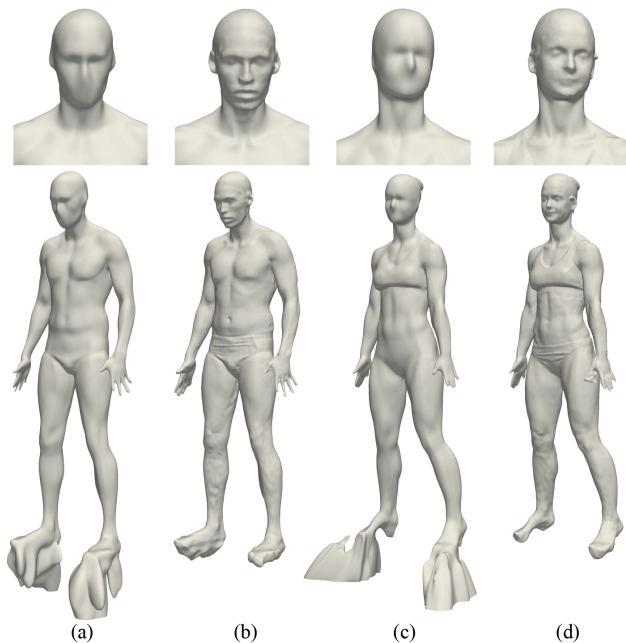


Figure 5. Level of details comparison. The zero level sets of an MLP trained with our method in (b) and (d); and using regression loss in (a) and (c), respectively.

	Method	Ground Truth		Scans	
		d_C	d_H	d_C^{\rightarrow}	d_H^{\rightarrow}
Anchor	DGP	0.33	8.82	0.08	2.79
	Ours	0.22	4.71	0.12	1.32
Daratech	DGP	0.2	3.14	0.04	1.89
	Ours	0.25	4.01	0.08	1.59
Dc	DGP	0.18	4.31	0.04	2.53
	Ours	0.17	2.22	0.09	2.61
Gargoyle	DGP	0.21	5.98	0.062	3.41
	Ours	0.16	3.52	0.064	0.81
Lord Quas	DGP	0.14	3.67	0.04	2.03
	Ours	0.12	1.17	0.07	0.98

Table 2. Evaluation on the surface reconstruction benchmark versus DGP (Williams et al., 2019b).

7. Experiments

7.1. Surface reconstruction

We tested our method on the task of surface reconstruction. That is, given a *single* input point cloud $\mathcal{X} \subset \mathbb{R}^3$ with or without a set of corresponding normal vectors $\mathcal{N} \subset \mathbb{R}^3$, the goal is to approximate the surface that \mathcal{X} was sampled from. The sample \mathcal{X} is usually acquired using a 3D scanner, potentially introducing variety of defects and artifacts. We evaluated our method on the surface reconstruction benchmark (Berger et al., 2013), using data (input point clouds \mathcal{X} , normal data \mathcal{N} , and ground truth meshes for evaluation) from (Williams et al., 2019b). The benchmark consists of five shapes with challenging properties such as non trivial topology or details of various feature sizes including sharp features. We compared our performance to the method from

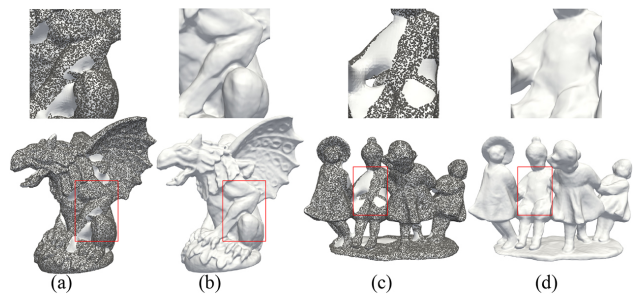


Figure 6. Reconstructions with our method in (b) and (d) versus (Williams et al., 2019b) (DGP) in (a) and (c). Note the charts of DGP do not cover the entire surface area.

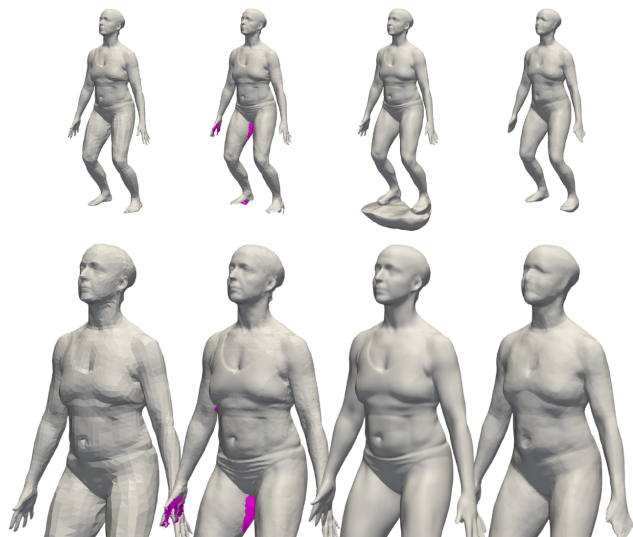


Figure 7. A train result on D-Faust. Left to right: registrations, scans, our results, SAL.

(Williams et al., 2019b) (DGP), which is a recent deep learning chart-based surface reconstruction technique; (Williams et al., 2019b) also provide plethora of comparisons to other surface reconstruction techniques and establish itself as state of the art method. Table 2 logs the performance of both methods using the following metrics: we measure distance of reconstructions to ground truth meshes using the (two-sided) Chamfer distance d_C and the (two-sided) Hausdorff distance d_H ; and distance from input point clouds to reconstructions using the (one-sided) Chamfer distance d_C^{\rightarrow} and the (one-sided) Hausdorff distance d_H^{\rightarrow} . Our method improves upon DGP in 4 out of 5 of the models in the dataset when tested on the ground truth meshes. DGP provides a better fit in average to the input data \mathcal{X} (our method performs better in Hausdorff); this might be explained by the tendency of DGP to leave some uncovered areas on the surface, see e.g., Figure 6 highlighting uncovered parts by DGP.

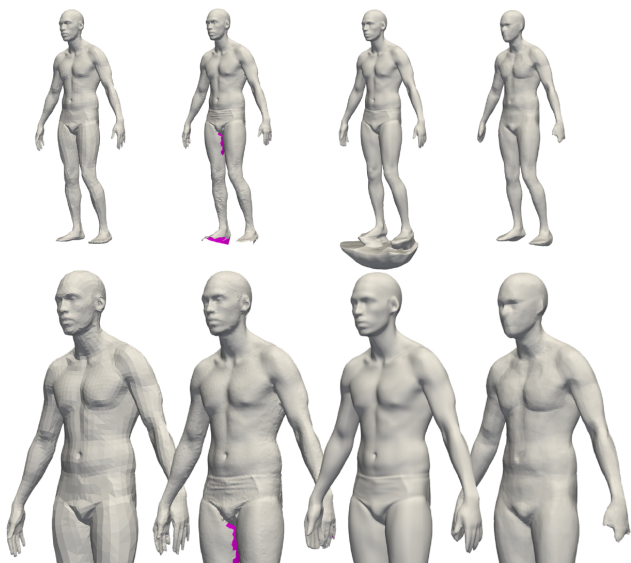


Figure 8. A test result on D-Faust. Left to right: registrations, scans, our results, SAL.

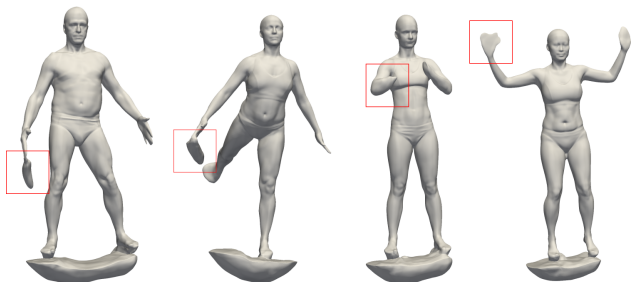


Figure 9. Failures of our method on D-Faust.

7.2. Learning shape space

In this experiment we tested our method on the task of learning shape space from raw scans. To this end, we use the D-Faust (Bogo et al., 2017) dataset, containing high resolution raw scans (triangle soups) of 10 humans in multiple poses. For training, we sampled each raw scan, $\mathcal{S}_j, j \in J$, uniformly to extract point and normal data, $\{(\mathcal{X}_j, \mathcal{N}_j)\}_{j \in J}$. We tested our method in two different settings: (i) random 75%-25% train-test split; (ii) generalization to unseen humans - 8 out of 10 humans are used for training and the remaining 2 for testing. In both cases we used the same splits as in (Atzmon & Lipman, 2020). The second column from the left in Figures 7, 8 show examples of input *scans* from the D-Faust dataset. The left column in these figures shows the ground truth *registrations*, $\mathcal{R}_j, j \in J$, achieved using extra data (e.g., color and texture) and human body model fitting (Bogo et al., 2017). Note that we *do not* use the registrations data in our training, we only use the raw scans.

Multi-shape architecture. In order to extend the network architecture described in section 5 for learning multiple

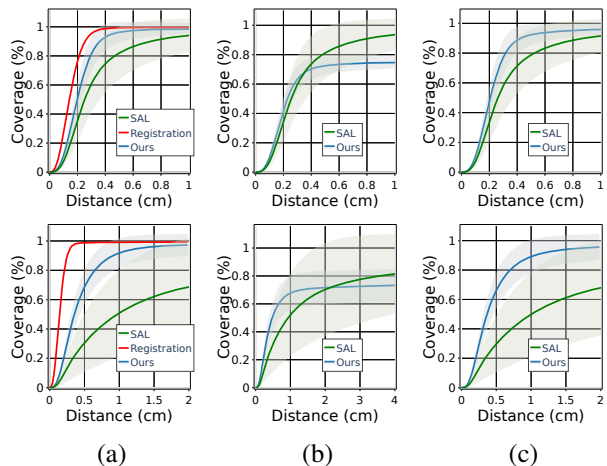


Figure 10. Error versus coverage for D-Faust test shapes with random split (first row) and unseen humans split (second row): (a) one-sided Chamfer distance from scan-to-reconstruction; (b) one-sided Chamfer distance reconstruction-to-registration; (c) one-sided Chamfer distance registration-to-reconstruction.

shapes, we use the auto-decoder setup as in (Park et al., 2019). That is, an MLP $f(x; \theta, z_j)$, where $z_j \in \mathbb{R}^{256}$ is a latent vector corresponding to each training example $j \in J$. The latent vectors $z_j \in \mathbb{R}^{256}$ are initialized to $\mathbf{0} \in \mathbb{R}^{256}$. We optimize a loss of the form $\frac{1}{|B|} \sum_{j \in B} \ell(\theta, z_j) + \alpha \|z_j\|$, where $B \subset J$ is a batch, $\alpha = 0.01$, ℓ defined in equation 2; τ, λ as above.

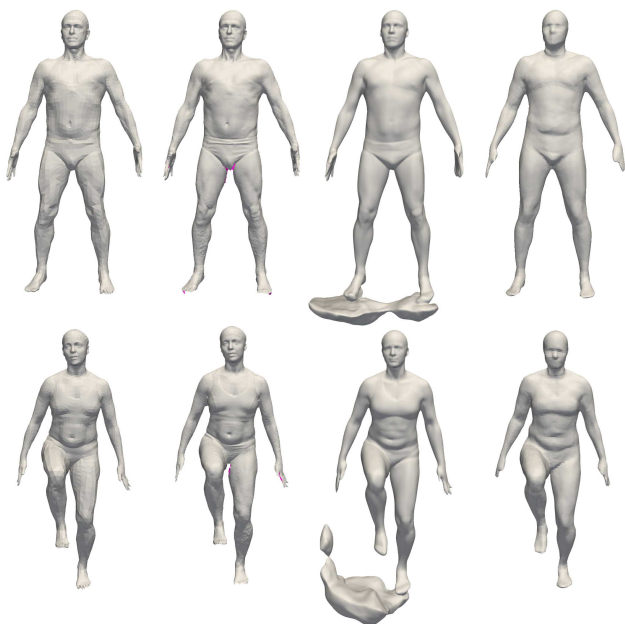


Figure 11. A test result on D-Faust with unseen humans split. Left to right: registrations, scans, our results, SAL.



Figure 12. Averaged shapes: Zero level sets (in blue) using averages of latent vectors of train examples (in gray).

Evaluation. For evaluation, we used our trained model for predicting shapes on the held out test set. As our architecture is an auto-decoder, the prediction for each test shape is obtained by performing 800 iterations of latent vector z optimization of the loss ℓ . We compared our results versus those obtained using SAL (Atzmon & Lipman, 2020), considered as state-of-the-art learning method on this dataset.

Figures 7 and 8 show examples from the train and test random split, respectively. Results for the unseen humans experiment are shown in 11. More results can be found in the appendix. Magenta triangles are back-faces hence indicating holes in the scan. Note that our method produces high level reconstructions with more details than SAL. In the unseen human tests the method provides plausible approximation despite training on only 8 human shapes. We note that our method produces a sort of a common "base" to the models, probably due to parts of the floor in some of the train data.

Figure 10 quantifies coverage as a function of error: Given a point cloud \mathcal{X} , we measure, for each distance value ϵ (X-axis), the fraction of points $x \in \mathcal{X}$ that satisfy $d(x, \mathcal{Y}) < \epsilon$. In (a) $\mathcal{X} \subset \mathcal{S}_j$, and \mathcal{Y} are the reconstructions of the registration, SAL and our method; note that our reconstructions are close in performance to the ground truth registrations. In (b), \mathcal{X} is a sampling of the reconstruction of SAL and our method, $\mathcal{Y} = \mathcal{R}_j$. In (c), \mathcal{X} is a sampling of \mathcal{R}_j and \mathcal{Y} is the reconstructions of SAL and our method. The lines represent mean over $j \in J$ and shades represent standard deviations. Note that we improve SAL except for larger errors in (b), a fact which can be attributed to the "base" reconstructed by our method. Some failure examples are shown in 9, mainly caused by noisy normal data. In addition, note that for the unseen humans split, we get relatively higher reconstruction error rate than the random split. We attribute this to the fact that there are only 8 humans in the dataset.

Shape space exploration. For qualitative evaluation of the learned shape space, we provide reconstructions obtained by interpolating latent vectors. Figure 12 shows humans shapes (in blue) corresponding to average interpolation of latent vectors z_j of training examples (in gray). Notice that the averaged shapes nicely combine and blend body shape and pose.

8. Conclusions

We introduced a method for learning high fidelity implicit neural representations of shapes directly from raw data. The method builds upon a simple loss function. Although this loss function possesses an infinite number of signed distance functions as minima, optimizing it with gradient descent tends to choose a favorable one. We analyze the linear case proving convergence to the approximate signed distance function, avoiding bad critical points. Analyzing non-linear models is a very interesting future work, e.g., explaining the reproduction of lines and circles in 1.

The main limitation of the method seems to be sensitivity to noisy normals as discussed in section 7.2. We believe the loss can be further designed to be more robust to this kind of noise.

Practically, the method produces neural level sets with significant more details than previous work. An interesting research venue would be to incorporate this loss in other deep 3D geometry systems (e.g., differentiable rendering and generative models).

Acknowledgments

The research was supported by the European Research Council (ERC Consolidator Grant, "LiftMatch" 771136), the Israel Science Foundation (Grant No. 1830/17) and by a research grant from the Carolito Stiftung (WAIC).

References

- Atzmon, M. and Lipman, Y. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2565–2574, 2020.
- Atzmon, M., Haim, N., Yariv, L., Israelov, O., Maron, H., and Lipman, Y. Controlling neural level sets. In *Advances in Neural Information Processing Systems*, pp. 2034–2043, 2019.
- Berger, M., Levine, J. A., Nonato, L. G., Taubin, G., and Silva, C. T. A benchmark for surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(2):1–17, 2013.
- Berger, M., Tagliasacchi, A., Seversky, L. M., Alliez, P., Guennebaud, G., Levine, J. A., Sharf, A., and Silva, C. T. A survey of surface reconstruction from point clouds. In *Computer Graphics Forum*, volume 36, pp. 301–329. Wiley Online Library, 2017.
- Bogo, F., Romero, J., Pons-Moll, G., and Black, M. J. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Chen, Z. and Zhang, H. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5939–5948, 2019.
- Chen, Z., Tagliasacchi, A., and Zhang, H. Bsp-net: Generating compact meshes via binary space partitioning. *arXiv preprint arXiv:1911.06971*, 2019.
- Choy, C. B., Xu, D., Gwak, J., Chen, K., and Savarese, S. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pp. 628–644. Springer, 2016.
- Crandall, M. G. and Lions, P.-L. Viscosity solutions of hamilton-jacobi equations. *Transactions of the American mathematical society*, 277(1):1–42, 1983.
- Crandall, M. G., Evans, L. C., and Lions, P.-L. Some properties of viscosity solutions of hamilton-jacobi equations. *Transactions of the American Mathematical Society*, 282(2):487–502, 1984.
- Dai, A., Ruizhongtai Qi, C., and Nießner, M. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5868–5877, 2017.
- Deng, B., Genova, K., Yazdani, S., Bouaziz, S., Hinton, G., and Tagliasacchi, A. Cvxnets: Learnable convex decomposition. *arXiv preprint arXiv:1909.05736*, 2019.
- Deprelle, T., Groueix, T., Fisher, M., Kim, V. G., Russell, B. C., and Aubry, M. Learning elementary structures for 3d shape generation and matching. *arXiv preprint arXiv:1908.04725*, 2019.
- Do Carmo, M. P. *Differential geometry of curves and surfaces: revised and updated second edition*. Courier Dover Publications, 2016.
- Ge, R., Huang, F., Jin, C., and Yuan, Y. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pp. 797–842, 2015.
- Girdhar, R., Fouhey, D. F., Rodriguez, M., and Gupta, A. Learning a predictable and generative vector representation for objects. In *European Conference on Computer Vision*, pp. 484–499. Springer, 2016.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Groueix, T., Fisher, M., Kim, V. G., Russell, B., and Aubry, M. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- Jiang, Y., Ji, D., Han, Z., and Zwicker, M. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. *arXiv preprint arXiv:1912.07109*, 2019.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. Gradient descent converges to minimizers. *arXiv preprint arXiv:1602.04915*, 2016.
- Liao, Y., Donne, S., and Geiger, A. Deep marching cubes: Learning explicit surface representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2916–2925, 2018.
- Lorensen, W. E. and Cline, H. E. Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH computer graphics*, 21(4):163–169, 1987.
- Maron, H., Galun, M., Aigerman, N., Trope, M., Dym, N., Yumer, E., Kim, V. G., and Lipman, Y. Convolutional neural networks on surfaces via seamless toric covers. *ACM Trans. Graph.*, 36(4):71–1, 2017.
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., and Geiger, A. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4460–4470, 2019.

- Michalkiewicz, M., Pontes, J. K., Jack, D., Baktashmotlagh, M., and Eriksson, A. Deep level sets: Implicit surface representations for 3d shape inference. *arXiv preprint arXiv:1901.06802*, 2019.
- Neyshabur, B. Implicit regularization in deep learning. *arXiv preprint arXiv:1709.01953*, 2017.
- Neyshabur, B., Tomioka, R., and Srebro, N. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.
- Nocedal, J. and Wright, S. *Numerical optimization*. Springer Science & Business Media, 2006.
- Osher, S., Fedkiw, R., and Piechor, K. Level set methods and dynamic implicit surfaces. *Appl. Mech. Rev.*, 57(3): B15–B15, 2004.
- Park, J. J., Florence, P., Straub, J., Newcombe, R., and Lovegrove, S. Deepsdf: Learning continuous signed distance functions for shape representation. *arXiv preprint arXiv:1901.05103*, 2019.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017a.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics informed deep learning (part ii): Data-driven discovery of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10566*, 2017b.
- Riegler, G., Ulusoy, A. O., Bischof, H., and Geiger, A. Octnetfusion: Learning depth fusion from data. In *2017 International Conference on 3D Vision (3DV)*, pp. 57–66. IEEE, 2017.
- Sinha, A., Bai, J., and Ramani, K. Deep learning 3d shape surfaces using geometry images. In *European Conference on Computer Vision*, pp. 223–240. Springer, 2016.
- Sinha, A., Unmesh, A., Huang, Q., and Ramani, K. Surfnet: Generating 3d shape surfaces using deep residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6040–6049, 2017.
- Sirignano, J. and Spiliopoulos, K. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Stewart, G. W. Matrix perturbation theory. 1990.
- Stutz, D. and Geiger, A. Learning 3d shape completion from laser scan data with weak supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1955–1964, 2018.
- Tatarchenko, M., Dosovitskiy, A., and Brox, T. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2088–2096, 2017.
- Tulsiani, S., Zhou, T., Efros, A. A., and Malik, J. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2626–2634, 2017.
- Williams, F., Panozzo, D., Yi, K. M., and Tagliasacchi, A. Voronoinet: General functional approximators with local support. *arXiv preprint arXiv:1912.03629*, 2019a.
- 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, 2019b.
- Wu, J., Zhang, C., Xue, T., Freeman, B., and Tenenbaum, J. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in neural information processing systems*, pp. 82–90, 2016.
- Yadav, N., Yadav, A., Kumar, M., et al. *An introduction to neural network methods for differential equations*. Springer, 2015.
- Yan, X., Yang, J., Yumer, E., Guo, Y., and Lee, H. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *Advances in neural information processing systems*, pp. 1696–1704, 2016.
- Yang, B., Wen, H., Wang, S., Clark, R., Markham, A., and Trigoni, N. 3d object reconstruction from a single depth view with adversarial learning. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 679–688, 2017.
- Zhao, H.-K., Osher, S., Merriman, B., and Kang, M. Implicit and nonparametric shape reconstruction from unorganized data using a variational level set method. 2000.
- Zhao, H.-K., Osher, S., and Fedkiw, R. Fast surface reconstruction using the level set method. In *Proceedings IEEE Workshop on Variational and Level Set Methods in Computer Vision*, pp. 194–201. IEEE, 2001.