

LISR: Learning Linear 3D Implicit Surface Representation Using Compactly Supported Radial Basis Functions

Atharva Pandey, Vishal Yadav, Rajendra Nagar, Santanu Chaudhury

Indian Institute of Technology Jodhpur, India
{pandey.19, yadav.40, rn, santanuc}@iitj.ac.in

Abstract

Implicit 3D surface reconstruction of an object from its partial and noisy 3D point cloud scan is the classical geometry processing and 3D computer vision problem. In the literature, various 3D shape representations have been developed, differing in memory efficiency and shape retrieval effectiveness, such as volumetric, parametric, and implicit surfaces. Radial basis functions provide memory-efficient parameterization of the implicit surface. However, we show that training a neural network using the mean squared error between the ground-truth implicit surface and the linear basis-based implicit surfaces does not converge to the global solution. In this work, we propose locally supported compact radial basis functions for a linear representation of the implicit surface. This representation enables us to generate 3D shapes with arbitrary topologies at any resolution due to their continuous nature. We then propose a neural network architecture for learning the linear implicit shape representation of the 3D surface of an object. We learn linear implicit shapes within a supervised learning framework using ground truth Signed-Distance Field (SDF) data for guidance. The classical strategies face difficulties in finding linear implicit shapes from a given 3D point cloud due to numerical issues (requires solving inverse of a large matrix) in basis and query point selection. The proposed approach achieves better Chamfer distance and comparable F-score than the state-of-the-art approach on the benchmark dataset. We also show the effectiveness of the proposed approach by using it for the 3D shape completion task.

Introduction

3D surface reconstruction of real-world objects is a fundamental research problem in 3D Computer Vision, Geometry Processing, and Computer Vision with multiple applications such as Gaming, Animation, AR/VR, simulation, etc (Berger et al. 2017) (Huang et al. 2022). There exist various variants of this problem like 3D Reconstruction from single view images, multiple view images, single view/multiple view partial 3D scans, etc. (Berger et al. 2017; Huang et al. 2022; Xiu et al. 2022; Tiwari et al. 2021). The 3D Reconstruction problem can further be classified based on the representation of the 3D shape, differing in their memory efficiency and the effectiveness of shape retrieval. The representations can be broadly classified into four categories: point

clouds, mesh-based representations, voxel-based representations, and implicit shape representations (Berger et al. 2017). Implicit shape representation’s primary superiority lies in its ability to generate arbitrary shapes and topologies at any resolution due to their continuous nature (Macêdo, Gois, and Velho 2011). They are modeled using parameterized functions, and learning the shape involves predicting these parameters (Park et al. 2019). Parameterization of the implicit representation defines the efficiency of the implicit surfaces. There exist works in the literature that aim to parameterize implicit surfaces (Park et al. 2019; Macêdo, Gois, and Velho 2011; Yavartanoo et al. 2021). One category of works aims at modeling the implicit representation by a function parameterized by the neural networks known as neural implicit (Park et al. 2019; Saito et al. 2019; Michalkiewicz et al. 2019b; Xu et al. 2022). Yavartanoo *et al.* represents the implicit surface by a finite low-degree algebraic polynomial and learns the co-efficient of the algebraic polynomial (Yavartanoo et al. 2021). This suffers from the fact that the low-degree polynomials fail to model the high curvature region of the surface. The radial basis functions have been used classically to model complex implicit surfaces (Macêdo, Gois, and Velho 2011; Liu et al. 2016; Xu et al. 2022).

In this work, we propose a radial basis function-based representation (Liu et al. 2016; Macêdo, Gois, and Velho 2011) to model implicit representation and learn the parameter of this linear representation using a neural network. The proposed Linear implicit shape model has significantly less number of parameters when compared to neural implicit (Park et al. 2019) and algebraic implicit representation (Yavartanoo et al. 2021) and is a great choice for a task requiring memory efficiency due to less number of parameters. We learn linear implicit surface representation (LISR) from a partial point cloud using a based learning approach which is supervised by the ground truth Signed-Distance Field (SDF) of the underlying shape. In the methodology section, we show that directly using the RBF framework for training a neural network to predict the SDF does not work. For the convergence to the optimal solution, We identify a constraint on the choice of basis and the query points for SDF evaluation to find the loss, which facilitates the learning of linear implicit surfaces. We show that the compactly-supported RBF as basis functions, along with a strategical

choice of query points at the time of training, satisfy the desired constraint, thereby converging to the optimal solution. We run an experiment to show the validity of the constraint and train a neural network with the given constraint on the task of shape prediction and completion from a raw point cloud. We use the proposed approach for the shape completion task to show the application of the proposed approach.

In summary, our contributions are as follows:

- We parameterize the implicit 3D surface of an object using linear and compactly supported radial basis functions and identify a constraint on the choice of basis and query points for enabling efficient learning of linear implicit shapes.
- We show that the proposed learning-based Compactly Supported-RBF (CSRBF) representation along with the proposed query point selection strategy satisfies the constraint on the choice of the basis functions.
- We propose a neural network to learn the coefficient of linearity for CSRBF-based implicit surface representation.
- We experimentally validate that the proposed approach can learn high-quality shape representations over distribution of shapes and show its applicability for the shape completion task.

The code for this paper is available at <https://github.com/Atharvap14/LISR>.

Related Works

Volumetric Representation: Volumetric representation models the surface surface using voxels in a unit cube. The learning-based approaches learn voxel occupancy from the given input point cloud/image of an object (Wu et al. 2015, 2016; Peng et al. 2020; Zhang, Nießner, and Wonka 2022; Wang et al. 2023; Chibane, Alldieck, and Pons-Moll 2020). The accuracy of the reconstructed surfaces depends on the voxel size which is a cubic memory requirements in terms of the levels in the tree-based representation (Dai, Ruizhongtai Qi, and Nießner 2017; Riegler, Osman Ulusoy, and Geiger 2017). However, volumetric representation-based methods still suffer in terms of huge memory requirements to represent high curvature regions of the surface.

Neural Implicit Representation: Recently, many algorithms were proposed to predict the signed distance field and the occupancy from a single point cloud using a neural network (Park et al. 2019; Erler et al. 2020; Michalkiewicz et al. 2019a,b; Mescheder et al. 2019; Niemeyer et al. 2020; Chabra et al. 2020; Liu et al. 2021; Venkatesh et al. 2021; Chen and Zhang 2019). These approaches use the continuous representation of the implicit surface, which they parameterize using the neural networks as compared to the volumetric representation, where the surface is represented by voxels by discretizing the 3D unit cube. There exist approaches that reconstruct the implicit surface from the single image of the object (Liu et al. 2019; Sitzmann, Zollhöfer, and Wetzstein 2019; Yavartanoo et al. 2021; Yu et al. 2022; Saito et al. 2019; Xu et al. 2019). Recently, many methods have decomposed the 3D surfaces into multiple primitives

and learned implicit representation for each of the primitives from the local point cloud patches (Genova et al. 2019, 2020; Deng et al. 2020b; Jiang et al. 2020; Darmon et al. 2022; Wu et al. 2022; Deng et al. 2020a; Nan and Wonka 2017). These approaches do not solve the point completion problem and rely on a clean and complete point cloud of the object, and may not generalize well to unseen categories as they are trained in the global sense.

Parametric Implicit Surfaces : Another categories of algorithms parameterize the implicit surface representation and find optimal parameters either using classical optimization approaches (Huang, Carr, and Ju 2019; Blane et al. 2000; Rouhani and Sappa 2010; Macêdo, Gois, and Velho 2011; Zhao et al. 2021) or using learning-based approaches (Yavartanoo et al. 2021; Xu et al. 2022). Yavartanoo *et al.* represents the implicit surface by a finite low-degree algebraic polynomial and learns the co-efficient of the algebraic polynomial (Yavartanoo et al. 2021). This approach suffers from the fact that the low-degree polynomials fail to model the high curvature region of the surface. The radial basis functions have been used classically to model complex implicit surfaces (Macêdo, Gois, and Velho 2011; Liu et al. 2016; Xu et al. 2022) however have not been utilized for surface reconstruction in a learning-based framework. In this work, we propose an RBF-based representation to model implicit representation and learn the parameters of this linear representation using a neural network. The proposed Linear implicit shape models have significantly less number of parameters when compared to neural implicit (Park et al. 2019) and algebraic implicit representation (Yavartanoo et al. 2021).

Implicit Surfaces and Shape Completion: Recently, learning-based approaches have been proposed to predict the full 3D object from its partial 3D point cloud scan. The 3D surface completion algorithms differ based on the representation of the 3D objects, e.g. triangle meshes, volumetric, 3D point clouds, and 3D implicit surface functions. The approaches (Choy et al. 2016; Dai et al. 2017; Dai, Ruizhongtai Qi, and Nießner 2017; Häne, Tulsiani, and Malik 2017; Sun et al. 2022) predict the completed 3D object in terms of its volumetric representation. However, These algorithms are computationally expensive for high-resolution volumetric point cloud completion. The triangle mesh-based 3D shape completion methods fail in capturing complex geometric structures (Litany et al. 2018). Implicit surfaces are well known for their interpolation property even in case of incomplete point clouds (Keren, Cooper, and Subrahmonia 1994; Turk and O’Brien 2002). We use this property to solve the shape completion problem using implicit surface reconstruction. There exist various works (Yuan et al. 2018; Yan et al. 2022b; Yu et al. 2021; Xiang et al. 2022; Mittal et al. 2022; Yew and Lee 2022; Yan et al. 2022a) for shape completion task, but there is no work to solve these two problems simultaneously. The algorithm proposed in (Chibane, Alldieck, and Pons-Moll 2020) addresses this problem but mostly for human 3D shapes and also uses an occupancy network to first learn the implicit representation and subsequently solve the completion problem.

Methodology

In this work, we aim to solve the problem of implicit 3D surface reconstruction of an object \mathcal{S} from a scene from its partial and noisy 3D point cloud scan. Let $\mathcal{M} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathcal{S}$ be a 3D point cloud scan of an object. Given \mathcal{M} , our goal is to find an implicit representation of the 3D surface of the underlying object represented by \mathcal{M} . We represent the implicit surface of an object by the signed distance field $f: \mathbb{R}^3 \rightarrow \mathbb{R}$. The value $f(\mathbf{x})$ represents the distance of the point $\mathbf{x} \in \mathbb{R}^3$ from the surface of the object \mathcal{M} . The zero level set of the SDF f represents the surface of the underlying object \mathcal{M} , i.e., $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 \mid f(\mathbf{x}) = 0\}$. In this work, we propose a Linear Implicit Surface Representation (LISR) of f which we discuss in the following Sections.

Linear Implicit Surface Representation

In the proposed Linear Implicit Shape Representation (LISR) framework, we represent the surface of a 3D object as a zero-level set of a scalar field modeled using a linear combination of basis functions. The LISR representation is defined by a set of coefficients and a set of basis functions which is defined in Equation (1).

$$f(\mathbf{x}) = \sum_{i=1}^n \alpha_i \phi_i(\mathbf{x}) \quad (1)$$

Here, $f(\mathbf{x})$ represents the implicit function value at point \mathbf{x} , α_i 's are the coefficients of linear combinations, and the functions ϕ_i are the basis functions evaluated at \mathbf{x} . Given a point cloud \mathcal{M} with its ground truth signed distance field $s_{\text{gt}}(\mathbf{x})$, we formulate an optimization problem where we minimize the L2 loss between the predicted SDF and the ground truth SDF. Formally, we define the optimization problem in Equation (2).

$$\arg \min_{\alpha} \sum_{\mathbf{x} \in \mathcal{Q}} \|f(\mathbf{x}) - s_{\text{gt}}(\mathbf{x})\|_2^2 \quad (2)$$

Here, $\alpha = [\alpha_1 \ \alpha_2 \ \dots \ \alpha_n]^\top \in \mathbb{R}^n$ contains the coefficients to be optimized, \mathcal{Q} represents the domain of the optimization problem, and $\|\cdot\|_2$ denotes the L2 norm. It is easy to observe that the optimization problem defined in Equation (2) can be rephrased as a solution to a linear system $\mathbf{V}^\top \alpha = \mathbf{s}_{\text{gt}}$ using the condition $f(\mathbf{x}) = s_{\text{gt}}(\mathbf{x})$, $\forall \mathbf{x} \in \mathcal{Q}$. Here, \mathbf{s}_{gt} is the vector containing ground truth SDF values for query points $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$ in the domain \mathcal{Q} , $\mathbf{V}^\top \in \mathbb{R}^{m \times n}$ is the matrix of basis functions evaluated at query points (n is the number of basis functions, and m is the number of query

points) and defined as $\mathbf{V}^\top = \begin{bmatrix} \phi_1(\mathbf{x}_1) & \dots & \phi_n(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{x}_m) & \dots & \phi_n(\mathbf{x}_m) \end{bmatrix}$.

We first determine the condition where we can determine the optimal solution to this linear system which we state in the below result. Let α^* represent the solution obtained by solving the optimization problem defined in Equation (1) (We use the gradient descent algorithm).

Theorem 1. *The solution α^* will be equal to the optimal solution of the linear system $\mathbf{V}^\top \alpha = \mathbf{s}_{\text{gt}}$ if and only if $\mathbf{V}\mathbf{V}^\top$ is a full rank matrix.*

Proof. To prove this result, we first show that the gradient of the SDF loss function concerning α is equal $\mathbf{V}\mathbf{V}^\top(\alpha - (\mathbf{V}\mathbf{V}^\top)^{-1}\mathbf{V}\mathbf{s}_{\text{gt}})$. Consider the the SDF Loss function $\ell(\alpha) = \sum_{\mathbf{x} \in \mathcal{Q}} \|f(\mathbf{x}) - s_{\text{gt}}(\mathbf{x})\|_2^2$ as defined in Equation (1). Here, $f(\mathbf{x}) = \mathbf{V}^\top(\mathbf{x})\alpha$ and $\mathbf{V}(\mathbf{x}) = [\phi_1(\mathbf{x}) \ \phi_2(\mathbf{x}) \ \dots \ \phi_n(\mathbf{x})]^\top$ is the vector of values of basis functions evaluated at query point $\mathbf{x} \in \mathcal{Q}$. Now, using the chain rule, we can easily show that the gradient of the SDF loss with respect to α can be defined as:

$$\begin{aligned} \nabla_{\alpha} \ell &= 2 \sum_{\mathbf{x} \in \mathcal{Q}} (\mathbf{V}^\top(\mathbf{x})\alpha - s_{\text{gt}}(\mathbf{x})) \nabla_{\alpha} (\mathbf{V}^\top(\mathbf{x})\alpha) \quad (3) \\ &= 2\mathbf{V}\mathbf{V}^\top \left(\alpha - (\mathbf{V}\mathbf{V}^\top)^{-1}\mathbf{V}\mathbf{s}_{\text{gt}} \right) \quad (4) \end{aligned}$$

We observe that $\alpha^* = (\mathbf{V}\mathbf{V}^\top)^{-1}\mathbf{V}\mathbf{s}_{\text{gt}}$ represents the optimal solution for the linear system $\mathbf{V}^\top \alpha = \mathbf{s}_{\text{gt}}$. Hence, the optimal solution of the gradient descent algorithm will be $\alpha_0 + \mathbf{V}\mathbf{V}^\top \gamma^*$. Here, $\gamma^* = \arg \min_{\gamma} \|\alpha_0 + \mathbf{V}\mathbf{V}^\top \gamma - \alpha^*\|_2^2$.

Here, α_0 is the initialized values and $\alpha_0 + \mathbf{V}\mathbf{V}^\top \gamma^*$ represents that optimal α lies in the subspace of $\mathbf{V}\mathbf{V}^\top$ centered at α_0 . We can easily deduce (using least square solution) that the optimal γ^* is the solution to a linear system $\mathbf{V}\mathbf{V}^\top \gamma = (\alpha^* - \alpha_0)$. Unique value of γ exists if and only if $\mathbf{V}\mathbf{V}^\top$ is a full-rank matrix and hence the gradient descent algorithm will be able to converge if and only if $\mathbf{V}\mathbf{V}^\top$ is a full-rank matrix. \square

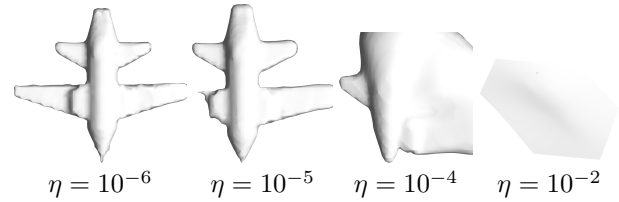


Figure 1: Perturbation of HRBF coefficients: Noise was added to each β value and the noise was uniformly sampled from the range $[-\eta, \eta]$. We observe that the shape is too sensitive to β as even for $\eta = 10^{-5}$ one of the wings of the plane has been clipped.

Compactly Supported Radial Basis Functions

In the LISR, we experimentally observe that the basis functions ϕ_i 's with global support, such as polynomials (Blane et al. 2000; Rouhani and Sappa 2010) and Hermite-RBFs (Macêdo, Gois, and Velho 2011), do not meet the rank criteria defined in Theorem 1 for all the query points $\mathbf{x} \in \mathcal{Q}$. The shapes generated using these basis functions are highly sensitive to coefficients, as shown in Figure 1. This makes it harder to learn the distribution of coefficients over a class of shapes. In contrast, the use of compactly supported basis functions reduced the sensitivity of the entire shape to coefficient perturbations and enhanced the correlation between coefficients and their respective basis functions. Intuitively, this implies that only a specific subset of basis functions

and their corresponding coefficients control the shape representation in localized regions of the 3D space. Thereby reducing the inter-dependence of coefficients and increasing the matrix rank. Therefore, we propose to use the compactly supported radial basis functions as our basis functions ϕ_i 's. We define the local support for each basis function using the Voronoi space partition of kernel points. We modify the LISR defined in Equation (1) and define the modified linear implicit surface representation in Equation (5).

$$f(\mathbf{x}) = \sum_{i=1}^q r_i(\mathbf{x}) \beta_i^\top \nabla (\|\mathbf{x} - \mathbf{p}_i\|_2^3). \quad (5)$$

Here, $\beta_i = [\beta_{i,x} \ \beta_{i,y} \ \beta_{i,z}]^\top \in \mathbb{R}^3$ is a vector formed by three unknown parameters, \mathbf{p}_i is the kernel point and $r_i(\mathbf{x})$ is the function defining the local support for basis with kernel point \mathbf{p}_i , i.e, if the point \mathbf{x} lies in the local support of \mathbf{p}_i then $r_i(\mathbf{x})$ evaluates to one, zero otherwise.

Query Point Selection

Along with a proper choice of basis functions, the selection of query points is a critical step to satisfy the full rank matrix criteria, as all the query points do not satisfy the rank criterion (Theorem 1). Our strategy for query point selection is to select at least three linearly independent query points from each local support. This strategy enables us to make the submatrix of \mathbf{V}^\top corresponding to the particular local support a full rank matrix.

Theorem 2. *The matrix $\mathbf{V}\mathbf{V}^\top$ is a full rank matrix for the linear implicit surface representation using the proposed compactly supported radial basis functions and the query point selection strategy.*

Proof. Consider the CSRBF formulation of linear implicit surface as defined in Equation (5). We rewrite the SDF $f(\mathbf{x})$ by substituting the expanded gradients and coefficients as:

$$f(\mathbf{x}) = \sum_{i=1}^q r_i(\mathbf{x}) \left(\beta_{i,x} \frac{\partial}{\partial x} (\|\mathbf{x} - \mathbf{p}_i\|_2^3) + \beta_{i,y} \frac{\partial}{\partial y} (\|\mathbf{x} - \mathbf{p}_i\|_2^3) + \beta_{i,z} \frac{\partial}{\partial z} (\|\mathbf{x} - \mathbf{p}_i\|_2^3) \right) \quad (6)$$

Now, the proposed CSRBF-based linear implicit surface representation can be written as

$$f(\mathbf{x}) = \sum_{i=1}^{3q} \alpha_i \phi_i(\mathbf{x}). \quad (7)$$

Where, ϕ_i 's and α_i 's are defined as in Equations (8) and (9), respectively.

$$\phi_i(\mathbf{x}) = \begin{cases} r_{\frac{i+2}{3}}(\mathbf{x}) \frac{\partial}{\partial x} (\|\mathbf{x} - \mathbf{p}_k\|_2^3), & \text{if } i = 3k + 1 \\ r_{\frac{i+1}{3}}(\mathbf{x}) \frac{\partial}{\partial y} (\|\mathbf{x} - \mathbf{p}_k\|_2^3), & \text{if } i = 3k + 2 \\ r_{\frac{i}{3}}(\mathbf{x}) \frac{\partial}{\partial z} (\|\mathbf{x} - \mathbf{p}_k\|_2^3), & \text{if } i = 3k \end{cases} \quad (8)$$

$$\alpha_i(\mathbf{x}) = \begin{cases} \beta_{\frac{i+2}{3},x}, & \text{if } i = 3k + 1 \\ \beta_{\frac{i+1}{3},y}, & \text{if } i = 3k + 2 \\ \beta_{\frac{i}{3},z}, & \text{if } i = 3k. \end{cases} \quad (9)$$

Algorithm 1: Faster-Convergence Query Point Selection

Input: Basis Kernel Points $\{\mathbf{p}\}_{i=1}^q$ and their 3D Voronoi Cells $\{\mathcal{S}_i\}_{i=1}^q$

Output: Query Points set \mathcal{Q}

- 1: Let $\mathcal{Q} = \emptyset$.
 - 2: **for** Each basis kernel point \mathbf{p}_i **do**
 - 3: Determine ϵ such that $\epsilon < \text{radius}(\mathcal{S}_i)$.
 - 4: $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{\mathbf{p}_i + \epsilon \hat{i}, \mathbf{p}_i + \epsilon \hat{j}, \mathbf{p}_i + \epsilon \hat{k}\}$
 - 5: **end for**
 - 6: **return** \mathcal{Q}
-

Now, with this modified definition of LISR, we can rewrite the i -th row of the matrix \mathbf{V}^\top as: $[\mathbf{0}_{m_i \times 3i-3} \ \mathbf{M}_i \ \mathbf{0}_{m_i \times 3q-3i}]$ where m_i represents the number of query points present in local support corresponding to \mathbf{p}_i and

$$\mathbf{M}_i = \begin{bmatrix} \phi_{3i-2}(x_{i,1}) & \phi_{3i-1}(x_{i,1}) & \phi_{3i}(x_{i,1}) \\ \vdots & \vdots & \vdots \\ \phi_{3i-2}(x_{i,m_i}) & \phi_{3i-1}(x_{i,m_i}) & \phi_{3i}(x_{i,m_i}) \end{bmatrix}.$$

For \mathbf{V}^\top to be a full-rank matrix, each of the \mathbf{M}_i has to be full-rank. It is trivial to observe that \mathbf{M}_i is a full-rank as the query points are selected along the m_i linearly independent directions (Algorithm 1). Since \mathbf{M}_i is a full-rank linear transformation on query points, the query point matrix must be full-rank (which is 3). Hence there must exist at least three linearly independent points in the set of query points belonging to each local support such that $\mathbf{V}\mathbf{V}^\top$ is a full-rank matrix. Hence, the proposed query point sampling strategy is sufficient for learning the linear implicit shape with a CSRBF basis. \square

To improve the rate of convergence, we sample the points $\{\mathbf{p}_i + \epsilon \hat{i}, \mathbf{p}_i + \epsilon \hat{j}, \mathbf{p}_i + \epsilon \hat{k}\}$ for local support of \mathbf{p}_i instead of sampling any three linearly independent points in local support. The value of ϵ is common across local supports and chosen such that all sampled query points lie in their respective local supports. The gradient of the SDF loss can be written as follows:

$$\nabla_{\alpha} \left(\sum_{\mathbf{x} \in \mathcal{Q}} \|f(\mathbf{x}) - S_{\text{gt}}(\mathbf{x})\|_2^2 \right) = 2\mathbf{V}\mathbf{V}^\top (\alpha - \alpha^*)$$

where $\alpha^* = (\mathbf{V}\mathbf{V}^\top)^{-1} \mathbf{V} \mathbf{s}_{\text{gt}}$ is the solution to the optimization problem. The sampling strategy in Algorithm 1 ensures that $\mathbf{V}\mathbf{V}^\top = c\mathbf{I}$ where c is an arbitrary constant and \mathbf{I} is the identity matrix. This shows that gradient descent always occurs in the direction of global optima. This ensures that all dimensions are similarly scaled allowing the model to train faster.

Learning LISR

Now, we explain our approach for finding the optimal coefficients $\alpha = [\alpha_1 \ \dots \ \alpha_{3q}]^\top$. As shown in Figure 1, the final solution is too sensitive to the coefficients, and the final reconstructed surfaces using the classical algorithms are

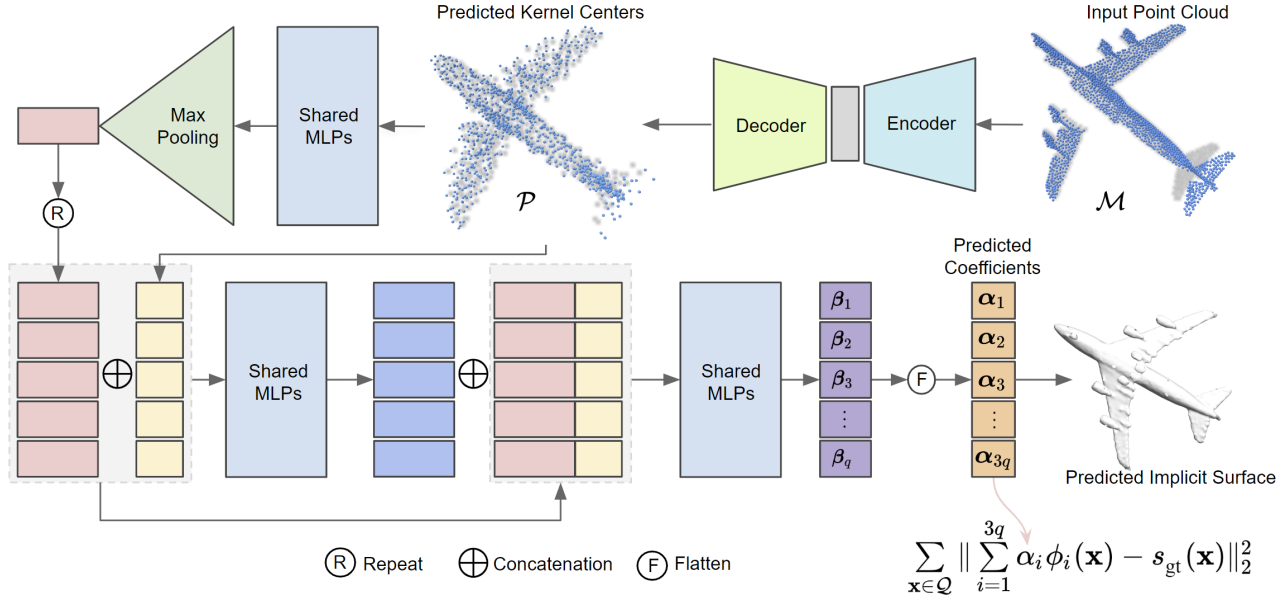


Figure 2: The proposed framework consists of two interconnected networks: the Basis Prediction Network (top) and the Coefficient Prediction Network (bottom). The Basis Prediction Network predicts the kernel points of the Compactly Supported Radial Basis Function (CSRBF) for partial shapes, while for complete shapes, the input point cloud can directly serve as the kernel points for basis. The learned basis kernel points are then passed into the Coefficient Prediction Network, which estimates the coefficients corresponding to each basis. Together, these predictions form a linear implicit shape representation for a given input point cloud.

not robust to missing parts in the input point cloud. Hence, we follow a learning-based approach to find optimal coefficients.

Predicting Kernel Centers: Our neural network architecture consists of two sub-networks. The first network is the Basis Prediction Network, which predicts the kernel points of the Compactly Supported Radial Basis Function (CSRBF) for partial shapes. For partial shapes, the network learns to predict the optimal locations for these kernel points. For complete shapes, however, an interesting optimization comes into play: the input point cloud itself can directly serve as the kernel points for the basis, as they provide the necessary information to encode the shape’s structure. The structure of the proposed basis prediction network is the simple 3D auto-encoder which is inspired by the network proposed by (Yuan et al. 2018). The input to this network is the partial input point cloud $\mathcal{M} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. The output of this network is the set of kernel 3D points $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_q\}$. We use the Chamfer-L1 distance between the predicted kernel centers point cloud \mathcal{P} and the input point cloud \mathcal{M} as the loss function to train the kernel center prediction network defined as the below equation:

$$\ell_1 = \frac{1}{|\mathcal{M}|} \sum_{\mathbf{x} \in \mathcal{M}} \min_{\mathbf{p} \in \mathcal{P}} \|\mathbf{x} - \mathbf{p}\|_1 + \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} \min_{\mathbf{x} \in \mathcal{M}} \|\mathbf{x} - \mathbf{p}\|_1$$

Coefficient Prediction: Once we predict the kernel points using the basis prediction network or obtained from the input point cloud for complete shapes, we pass them into the coefficient prediction network. This network is responsible

for estimating the coefficients α_i corresponding to each basis ϕ_i where $i \in \{1, 2, \dots, 3q\}$.

By combining the kernel points and their corresponding coefficients, the linear implicit surface representation is constructed as $f(\mathbf{x}) = \sum_{i=1}^{3q} \alpha_i \phi_i(\mathbf{x})$. To train the proposed coefficient prediction neural network, we use the following loss between the predicted SDF $f(\mathbf{x}) = \sum_{i=1}^{3q} \alpha_i \phi_i(\mathbf{x})$ and the ground-truth SDF $s_{\text{gt}}(\mathbf{x})$:

$$\ell_2 = \sum_{\mathbf{x} \in \mathcal{Q}} \left\| \sum_{i=1}^{3q} \alpha_i \phi_i(\mathbf{x}) - s_{\text{gt}}(\mathbf{x}) \right\|_2^2 \quad (10)$$

In Figure 2, we present the structure of both the kernel center prediction network and the coefficient prediction network and the overall flow of the proposed approach.

Experiments and Results

In this section, we show the improvement in learning ability with our choice of basis and query points over other basis functions such as Tri-harmonic, Mono-harmonic, and Hermite-RBFs. We evaluate the ability of our model to generalize over a distribution of shapes by training it on various classes of shapes in the ShapeNet dataset (Chang et al. 2015). We also evaluate the reconstruction of the shapes from partial point clouds by coupling the model with the Basis Prediction Network.

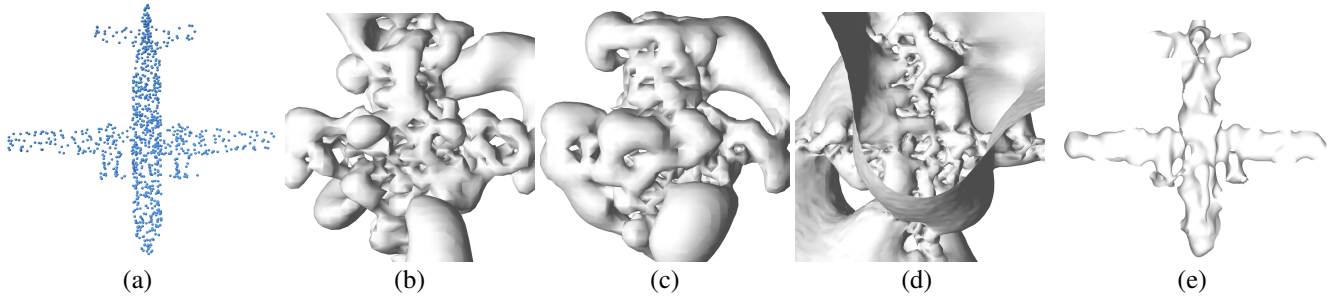


Figure 3: Single Shape optimization solution on the Point Cloud (shown (a)) for (b) Tri-harmonic radial basis, (c) Mono-harmonic radial basis, (d) Hermite-Radial basis with uniform query point selection strategy in the learning-based framework and (e) LISR with query point selection strategy defined in Algorithm 1 over $[-1, 1]^3$ optimization domain and single shape overfitting with only 1000 kernel points. LISR can capture the global shape properties.

Basis Functions	Rank of $\mathbf{V}\mathbf{V}^\top$	Maximum rank
Tri-Harmonic	954	1000
Mono-Harmonic	997	1000
HRBF	13	3000
CSRBF (Proposed)	3000	3000

Table 1: Rank of the matrix $\mathbf{V}\mathbf{V}^\top$ for different choices of the basis functions. For the convergence of the optimization algorithm to train a coefficient prediction network, the $\mathbf{V}\mathbf{V}^\top$ has to be a full rank matrix as shown in Theorem 1. We observe that the proposed radial basis functions with compact support achieve the full rank constraint.

Single Shape Learning

We demonstrate the improved shape representation and learning of our approach by comparing our choice of basis and query points against various choices of basis functions with global support such as Tri-harmonic RBFs, mono-harmonic, and HRBF (Liu et al. 2016) with uniformly sampled query points from the optimization domain $[-1, 1]^3$. We formulate the shape representation as given in Equation (7) and use ADAM optimizer to compute the optimal value for α for the loss function given in Equation (10) for overfitting on a single shape. We compute the rank of the matrix $\mathbf{V}\mathbf{V}^\top$ matrix for each choice of basis and their corresponding query points as given in Table 1. We also perform qualitative analysis by comparing the shape extracted after zero-level iso-surface extraction given in Figure 3. For the convergence of the optimization algorithm to train a coefficient prediction network, the $\mathbf{V}\mathbf{V}^\top$ has to be a full rank matrix as shown in Theorem 1. We observe that the proposed radial basis functions with compact support achieve the full rank constraint. We have used 1000 kernel centers for this experiment. We observe that, in the learning-based framework, the proposed LISR can learn the approximate shape of the actual object even with 1000 kernel centers on overfitting a large network with a single step. Whereas, with other basis functions, it is not able to capture any aspect of the object surface.

3D Surface Reconstruction Evaluation

To check the performance of the proposed approach for the task of 3D implicit surface fitting, we compare the performance of the proposed approach (LISR) with the performance of the state-of-the-art approaches OccNet (Mescheder et al. 2019), ConvNet (Peng et al. 2020), IF-Net (Chibane, Alldieck, and Pons-Moll 2020), and 3DILG (Zhang, Nießner, and Wonka 2022). We use the standard evaluation metrics for this task such as Chamfer-L1 distance (CD) and the F-Score (Zhang, Nießner, and Wonka 2022; Tatarchenko et al. 2019). We use five classes (Chair, Aeroplane, Lamp, Sofa, and Table) from the ShapeNet dataset (Chang et al. 2015) to train our network and compare the performance with that of the state-of-the-art approaches. We follow the (23023,575,1152) train-val-test split on ShapeNet. Following the evaluation protocol of (Zhang, Nießner, and Wonka 2022), we include two metrics, the Chamfer-L1 distance, and F-Score. In our experiment, we set the threshold equal to 0.02. In Table 2, we present the Chamfer Distance and the F-Score for the proposed approach and the state-of-the-art approaches. We observe that, on average, the proposed approach achieves a marginally better Chamfer distance and marginally lesser F-Score than that of the 3DILG (Zhang, Nießner, and Wonka 2022) and IF-Net (Chibane, Alldieck, and Pons-Moll 2020). In Figure 4, we show the results of the proposed approach for three objects along with the ground-truth 3D surfaces. We observe that the reconstructed implicit surface captures topology information but in the case of the chair, it could not capture one of the spikes in the support of the chair.

Network and Training Detail, Run-time : Our architecture consists of two sub-networks: Basis Prediction Network (BPN) and Coefficient Prediction Network (CPN). It is essential for the CPN to be permutation invariant since there exists a correspondence between the basis and its coefficient. We ensure this through the use of shared MLPs. CPN takes kernel points as input and passes them through a shared-MLP (6 layers with 3, 64, 129, 256, 512, 1024 neurons, resp.) with a max-pooling after the last layer to generate 1024 global latent which is then concatenated with each of the kernel points and then passed through two shared MLPs (5 layers with 1024+3, 512, 512, 512, 512 neurons, resp.) and

Class	OccNet (Mescheder et al. 2019)		ConvNet (Peng et al. 2020)		IF-Net (Chibane et al. 2020)		3DILG (Zhang et al. 2023)		LIRS (Proposed)	
	CD↓	F-Score↑	CD↓	F-Score↑	CD↓	F-Score↑	CD↓	F-Score↑	CD↓	F-Score↑
Chair	0.058	0.890	0.044	0.934	0.031	0.990	0.029	0.992	0.027	0.993
Aeroplane	0.037	0.948	0.028	0.982	0.020	0.994	0.019	0.993	0.022	0.997
Lamp	0.090	0.820	0.050	0.945	0.038	0.970	0.036	0.971	0.028	0.980
Sofa	0.051	0.918	0.042	0.967	0.032	0.988	0.030	0.986	0.027	0.991
Table	0.041	0.961	0.036	0.982	0.029	0.998	0.026	0.999	0.033	0.974
Mean	0.046	0.907	0.040	0.962	0.030	0.988	0.028	0.988	0.027	0.987

Table 2: Comparison of the performance of the proposed approach (LIRS) with the performance of the state-of-the-art approaches OccNet (Mescheder et al. 2019), ConvNet (Peng et al. 2020), IF-Net (Chibane, Alldieck, and Pons-Moll 2020), and 3DILG (Zhang, Nießner, and Wonka 2022). We have used the standard evaluation metrics for this task Chamfer Distance (CD) and the F-Score and five classes (Chair, Aeroplane, Lamp, Sofa, and Table) from the ShapeNet dataset (Chang et al. 2015).

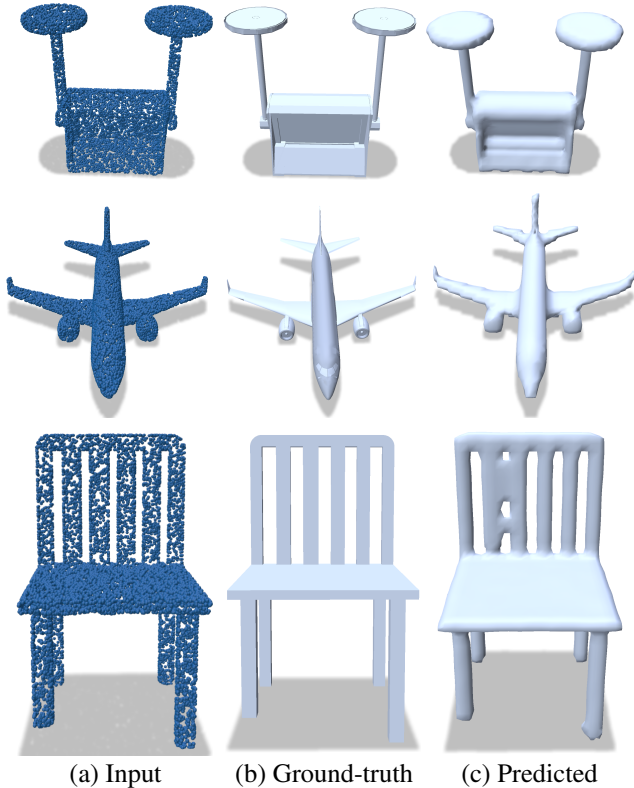


Figure 4: Visual results of 3D implicit surface reconstruction on three input point clouds.

(5 layers with 1027+512, 512, 512, 512, 3 neurons, resp.) with a skip connection between the two to generate the beta values. The model contains 6670122 trainable parameters. In essence, the basis kernel points provide a fixed foundation, and the coefficients control the influence of each basis function. The Voronoi cells of kernel points are used as their local support for the basis functions while training. We use ADAM with learning rate = 1e-8, weight_decay = 1e-7 along with OneCycleLR scheduler with max_lr=1e-3. We train the proposed network on a 48GB RAM GPU on a Linux machine and have used PyTorch (Paszke et al. 2019). The av-

erage time for a forward pass on test point clouds having 10,000 points is around 0.0139 Seconds.

3D Shape Completion To show the effectiveness of the proposed approach, we trained the proposed architecture for the 3D shape completion task on the MVP dataset (Pan, Cai, and Liu 2021; Pan et al. 2021) that consists of eight object classes. To evaluate our approach, we find the Chamfer-L1 distance between the predicted 3D surface and the ground-truth complete point clouds. The average Chamfer distance for the proposed approach for all 8 classes is 0.039.

Conclusion

In this work, we have addressed the problem of reconstructing the implicit 3D surface of an object from its partial 3D point cloud. We have proposed a radial basis functions-based linear representation of the underlying implicit surface. We further showed that the existing radial basis-based representation for implicit surface fail in a learning-based framework as the SDF predicted by these methods does not satisfy the full rank matrix criterion on the set of query points required for finding the mean squared loss between the predicted and the ground-truth SDFs. We have proposed an RBF-based approach where we learn linear radial basis functions with compact local supports. With this representation, along with the proposed query point selection, we were able to learn the coefficients of the proposed RBF-based representation. The proposed parametrization of the implicit surface is linear in the coefficients of RBF and has a compact surface representation. We achieved the best performance with respect to the Chamfer distance metric and comparable performance with the F-score metric.

Limitations and Future Scope: Our approach mostly supports single object-based tasks as it was trained on the ShapeNet dataset, and also it requires supervision from the ground-truth signed distance field. In future work, we would like to extend the proposed framework for scene-level tasks and solve it in semi-supervised or unsupervised approaches.

Acknowledgments

This research was supported by the start-up research grant (SRG), SERB, Government of India and TIH iHub Drishti, IIT Jodhpur.

References

- Berger, M.; Tagliasacchi, A.; Seversky, L. M.; Alliez, P.; Guennebaud, G.; Levine, J. A.; Sharf, A.; and Silva, C. T. 2017. A survey of surface reconstruction from point clouds. In *Computer graphics forum*, volume 36, 301–329. Wiley Online Library.
- Blane, M. M.; Lei, Z.; Civi, H.; and Cooper, D. B. 2000. The 3L algorithm for fitting implicit polynomial curves and surfaces to data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3): 298–313.
- Chabra, R.; Lenssen, J. E.; Ilg, E.; Schmidt, T.; Straub, J.; Lovegrove, S.; and Newcombe, R. 2020. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*, 608–625. Springer.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*.
- Chen, Z.; and Zhang, H. 2019. Learning implicit fields for generative shape modeling. In *CVPR*, 5939–5948.
- Chibane, J.; Alldieck, T.; and Pons-Moll, G. 2020. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6970–6981.
- Choy, C. B.; Xu, D.; Gwak, J.; Chen, K.; and Savarese, S. 2016. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 628–644. Springer.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 5828–5839.
- Dai, A.; Ruizhongtai Qi, C.; and Nießner, M. 2017. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *CVPR*, 5868–5877.
- Darmon, F.; Bascle, B.; Devaux, J.-C.; Monasse, P.; and Aubry, M. 2022. Improving neural implicit surfaces geometry with patch warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6260–6269.
- Deng, B.; Genova, K.; Yazdani, S.; Bouaziz, S.; Hinton, G.; and Tagliasacchi, A. 2020a. Cvxnet: Learnable convex decomposition. In *CVPR*, 31–44.
- Deng, B.; Lewis, J. P.; Jeruzalski, T.; Pons-Moll, G.; Hinton, G.; Norouzi, M.; and Tagliasacchi, A. 2020b. Nasa neural articulated shape approximation. In *ECCV*, 612–628. Springer.
- Erler, P.; Guerrero, P.; Ohrhallinger, S.; Mitra, N. J.; and Wimmer, M. 2020. Points2surf learning implicit surfaces from point clouds. In *European Conference on Computer Vision*, 108–124. Springer.
- Genova, K.; Cole, F.; Sud, A.; Sarna, A.; and Funkhouser, T. 2020. Local deep implicit functions for 3d shape. In *CVPR*, 4857–4866.
- Genova, K.; Cole, F.; Vlastic, D.; Sarna, A.; Freeman, W. T.; and Funkhouser, T. 2019. Learning shape templates with structured implicit functions. In *ICCV*, 7154–7164.
- Häne, C.; Tulsiani, S.; and Malik, J. 2017. Hierarchical surface prediction for 3d object reconstruction. In *3DV*, 412–420. IEEE.
- Huang, Z.; Carr, N.; and Ju, T. 2019. Variational implicit point set surfaces. *ACM Transactions on Graphics (TOG)*, 38(4): 1–13.
- Huang, Z.; Wen, Y.; Wang, Z.; Ren, J.; and Jia, K. 2022. Surface reconstruction from point clouds: A survey and a benchmark. *arXiv preprint arXiv:2205.02413*.
- Jiang, C.; Sud, A.; Makadia, A.; Huang, J.; Nießner, M.; Funkhouser, T.; et al. 2020. Local implicit grid representations for 3d scenes. In *CVPR*, 6001–6010.
- Keren, D.; Cooper, D.; and Subrahmonia, J. 1994. Describing complicated objects by implicit polynomials. *IEEE TPAMI*, 16(1): 38–53.
- Litany, O.; Bronstein, A.; Bronstein, M.; and Makadia, A. 2018. Deformable shape completion with graph convolutional autoencoders. In *CVPR*, 1886–1895.
- Liu, S.; Saito, S.; Chen, W.; and Li, H. 2019. Learning to infer implicit surfaces without 3d supervision. *Advances in Neural Information Processing Systems*, 32.
- Liu, S.; Wang, C. C.; Brunnett, G.; and Wang, J. 2016. A closed-form formulation of HRBF-based surface reconstruction by approximate solution. *Computer-Aided Design*, 78: 147–157.
- Liu, S.-L.; Guo, H.-X.; Pan, H.; Wang, P.-S.; Tong, X.; and Liu, Y. 2021. Deep implicit moving least-squares functions for 3D reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1788–1797.
- Macêdo, I.; Gois, J. P.; and Velho, L. 2011. Hermite radial basis functions implicit. In *Computer Graphics Forum*, volume 30, 27–42. Wiley Online Library.
- Mescheder, L.; Oechsle, M.; Niemeyer, M.; Nowozin, S.; and Geiger, A. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 4460–4470.
- Michalkiewicz, M.; Pontes, J. K.; Jack, D.; Baktashmotlagh, M.; and Eriksson, A. 2019a. Deep level sets: Implicit surface representations for 3d shape inference. *arXiv preprint arXiv:1901.06802*.
- Michalkiewicz, M.; Pontes, J. K.; Jack, D.; Baktashmotlagh, M.; and Eriksson, A. 2019b. Implicit surface representations as layers in neural networks. In *ICCV*, 4743–4752.
- Mittal, P.; Cheng, Y.-C.; Singh, M.; and Tulsiani, S. 2022. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *CVPR*, 306–315.
- Nan, L.; and Wonka, P. 2017. Polyfit: Polygonal surface reconstruction from point clouds. In *ICCV*, 2353–2361.
- Niemeyer, M.; Mescheder, L.; Oechsle, M.; and Geiger, A. 2020. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 3504–3515.
- Pan, L.; Cai, Z.; and Liu, Z. 2021. Robust Partial-to-Partial Point Cloud Registration in a Full Range. *arXiv preprint arXiv:2111.15606*.

- Pan, L.; Chen, X.; Cai, Z.; Zhang, J.; Zhao, H.; Yi, S.; and Liu, Z. 2021. Variational Relational Point Completion Network. In *CVPR*, 8524–8533.
- Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 165–174.
- Peng, S.; Niemeyer, M.; Mescheder, L.; Pollefeys, M.; and Geiger, A. 2020. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, 523–540. Springer.
- Riegler, G.; Osman Ulusoy, A.; and Geiger, A. 2017. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, 3577–3586.
- Rouhani, M.; and Sappa, A. D. 2010. Relaxing the 3L algorithm for an accurate implicit polynomial fitting. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 3066–3072. IEEE.
- Saito, S.; Huang, Z.; Natsume, R.; Morishima, S.; Kanazawa, A.; and Li, H. 2019. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2304–2314.
- Sitzmann, V.; Zollhöfer, M.; and Wetzstein, G. 2019. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32.
- Sun, B.; Kim, V. G.; Aigerman, N.; Huang, Q.; and Chaudhuri, S. 2022. PatchRD: detail-preserving shape completion by Learning patch retrieval and Deformation. In *ECCV*, 503–522. Springer.
- Tatarchenko, M.; Richter, S. R.; Ranftl, R.; Li, Z.; Koltun, V.; and Brox, T. 2019. What do single-view 3d reconstruction networks learn? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3405–3414.
- Tiwari, G.; Sarafianos, N.; Tung, T.; and Pons-Moll, G. 2021. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11708–11718.
- Turk, G.; and O’Brien, J. F. 2002. Modelling with implicit surfaces that interpolate. *ACM Transactions on Graphics (TOG)*, 21(4): 855–873.
- Venkatesh, R.; Karmali, T.; Sharma, S.; Ghosh, A.; Babu, R. V.; Jeni, L. A.; and Singh, M. 2021. Deep implicit surface point prediction networks. In *ICCV*, 12653–12662.
- Wang, Z.; Zhou, S.; Park, J. J.; Paschalidou, D.; You, S.; Wetzstein, G.; Guibas, L.; and Kadambi, A. 2023. Alto: Alternating latent topologies for implicit 3d reconstruction. In *CVPR*, 259–270.
- Wu, J.; Zhang, C.; Xue, T.; Freeman, B.; and Tenenbaum, J. 2016. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *Advances in neural information processing systems*, 29.
- Wu, Q.; Liu, X.; Chen, Y.; Li, K.; Zheng, C.; Cai, J.; and Zheng, J. 2022. Object-compositional neural implicit surfaces. In *European Conference on Computer Vision*, 197–213. Springer.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1912–1920.
- Xiang, P.; Wen, X.; Liu, Y.-S.; Cao, Y.-P.; Wan, P.; Zheng, W.; and Han, Z. 2022. Snowflake point deconvolution for point cloud completion and generation with skip-transformer. *IEEE TPAMI*, 45(5): 6320–6338.
- Xiu, Y.; Yang, J.; Tzionas, D.; and Black, M. J. 2022. Icon: Implicit clothed humans obtained from normals. In *CVPR*, 13286–13296. IEEE.
- Xu, Q.; Wang, W.; Ceylan, D.; Mech, R.; and Neumann, U. 2019. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *Advances in neural information processing systems*, 32.
- Xu, Y.; Nan, L.; Zhou, L.; Wang, J.; and Wang, C. C. 2022. Hrbf-fusion: Accurate 3d reconstruction from rgb-d data using on-the-fly implicit. *ACM Transactions on Graphics (TOG)*, 41(3): 1–19.
- Yan, X.; Lin, L.; Mitra, N. J.; Lischinski, D.; Cohen-Or, D.; and Huang, H. 2022a. Shapeformer: Transformer-based shape completion via sparse representation. In *CVPR*, 6239–6249.
- Yan, X.; Yan, H.; Wang, J.; Du, H.; Wu, Z.; Xie, D.; Pu, S.; and Lu, L. 2022b. Fbnet: Feedback network for point cloud completion. In *ECCV*, 676–693. Springer.
- Yavartanoo, M.; Chung, J.; Neshatavar, R.; and Lee, K. M. 2021. 3dias: 3d shape reconstruction with implicit algebraic surfaces. In *ICCV*, 12446–12455.
- Yew, Z. J.; and Lee, G. H. 2022. Regtr: End-to-end point cloud correspondences with transformers. In *CVPR*, 6677–6686.
- Yu, X.; Rao, Y.; Wang, Z.; Liu, Z.; Lu, J.; and Zhou, J. 2021. PointR: Diverse point cloud completion with geometry-aware transformers. In *ICCV*, 12498–12507.
- Yu, Z.; Peng, S.; Niemeyer, M.; Sattler, T.; and Geiger, A. 2022. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35: 25018–25032.
- Yuan, W.; Khot, T.; Held, D.; Mertz, C.; and Hebert, M. 2018. PCN: Point Completion Network. In *2018 3DV*, 728–737.
- Zhang, B.; Nießner, M.; and Wonka, P. 2022. 3dilg: Irregular latent grids for 3d generative modeling. *Advances in Neural Information Processing Systems*, 35: 21871–21885.
- Zhao, T.; Alliez, P.; Boubekur, T.; Buse, L.; and Thiery, J.-M. 2021. Progressive discrete domains for implicit surface reconstruction. In *Computer Graphics Forum*, volume 40, 143–156.