

# DeepJoin: Learning a Joint Occupancy, Signed Distance, and Normal Field Function for Shape Repair

NIKOLAS LAMB, Clarkson University, USA

SEAN BANERJEE, Clarkson University, USA

NATASHA KHOLGADE BANERJEE, Clarkson University, USA

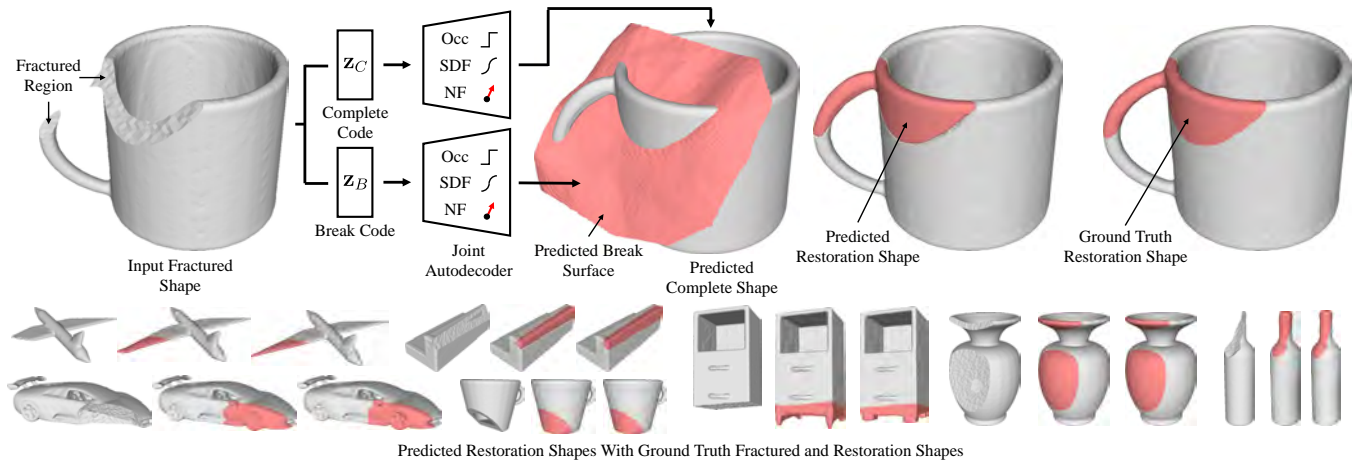


Fig. 1. From an input fractured shape, DeepJoin estimates latent codes for a complete shape and break surface using occupancy (Occ), signed distance function (SDF), and normal field (NF) samples, and predicts a restoration shape to repair the input fractured shape using the estimated latent codes.

We introduce DeepJoin, an automated approach to generate high-resolution repairs for fractured shapes using deep neural networks. Existing approaches to perform automated shape repair operate exclusively on symmetric objects, require a complete proxy shape, or predict restoration shapes using low-resolution voxels which are too coarse for physical repair. We generate a high-resolution restoration shape by inferring a corresponding complete shape and a break surface from an input fractured shape. We present a novel implicit shape representation for fractured shape repair that combines the occupancy function, signed distance function, and normal field. We demonstrate repairs using our approach for synthetically fractured objects from ShapeNet, 3D scans from the Google Scanned Objects dataset, objects in the style of ancient Greek pottery from the QP Cultural Heritage dataset, and real fractured objects. We outperform three baseline approaches in terms of chamfer distance and normal consistency. Unlike existing approaches and restorations using subtraction, DeepJoin restorations do not exhibit surface artifacts and join closely to the fractured region of the fractured shape. Our code is available at: <https://github.com/Terascale-All-sensing-Research-Studio/DeepJoin>.

CCS Concepts: • **Computing methodologies** → **Shape analysis; Neural networks**.

Additional Key Words and Phrases: Shape Representation, Implicit, Repair, Fracture, Deep Learning

Authors' addresses: Nikolas Lamb, Clarkson University, Potsdam, USA, [lambne@clarkson.edu](mailto:lambne@clarkson.edu); Sean Banerjee, Clarkson University, Potsdam, USA, [sbanerje@clarkson.edu](mailto:sbanerje@clarkson.edu); Natasha Kholgade Banerjee, Clarkson University, Potsdam, USA, [nbanerje@clarkson.edu](mailto:nbanerje@clarkson.edu).

## 1 INTRODUCTION

Household objects often undergo damage, e.g. denting, part loss, weathering, or fracturing. Fractured objects may be reassembled if all of their parts can be found. However, in some cases parts are destroyed during the fracture, e.g. the sugar jar in Figure 2(a), preventing reassembly of the object. Other damage, such as weathering, cannot be repaired using reassembly. Instead the object must be rebuilt using other materials. Users are likely to want to repair an object if the object has some sentimental or functional value. Repair is also necessary if the object is no longer manufactured. With the proliferation of 3D printing, the field of computational fabrication is ideally positioned to enable the repair of damaged objects by generating 3D printable repair parts.

Existing approaches that repair a single damaged object [Antle et al. 2011; Rengier et al. 2010; Schilling et al. 2014; Scopigno et al. 2011; Seixas et al. 2018] require modeling a repair part by hand using 3D design tools. While users may attempt to manually repair an object if the process is straightforward, repairing large or complex fractures, e.g. the cup in Figure 2(a), is outside the scope of an average user. Most automated repair approaches, e.g. for aerospace components [Gao et al. 2008; Zheng et al. 2006] and medical implants [Harrysson et al. 2007; Witek et al. 2016], are domain-specific and unlikely to generalize. Other automated approaches require fractures to be symmetric to existing object parts [Gregor et al. 2014; Papaioannou et al. 2017], or need to be fed a complete 3D proxy object [Lamb et al. 2019]. Though 3D-ORGAN [Hermoza and Sipiran



Fig. 2. (a) A fractured mug and a fractured jar with a destroyed part. (b) A complete, partial, and fractured shape.

2018] operates directly on fractured objects without geometric constraints, it encodes objects in low-resolution voxel space, producing restorations that are too coarse for physical repair.

We present DeepJoin, an approach rooted in deep learning that infers a restoration shape from an input fractured shape by deconstructing the fractured shape into a corresponding complete shape and break surface, shown in Figure 1. Our approach is related to approaches that input a partial shape and perform shape completion by inferring a latent code to fit the partial shape observation [Chen and Zhang 2019; Duggal et al. 2022; Hao et al. 2020; Park et al. 2019; Sitzmann et al. 2020; Tretschk et al. 2020; Yan et al. 2022a; Zheng et al. 2021]. However, different from approaches that perform partial shape completion, our approach addresses the challenge that, unlike a partial shape that is a strict subset of the complete shape, as shown in Figure 2(b), a fractured shape includes novel geometry at the fractured region that is not present in the complete shape.

To enable shape repair, our first contribution is a novel shape representation using a joint function that predicts three features—the occupancy function, the signed distance function (SDF), and the normal field. We find that when representing fractured shapes exclusively using the SDF, the most common implicit shape representation, deep networks struggle to accurately deconstruct the fractured shape into a complete shape and break surface. Our approach learns the occupancy function jointly with the SDF and partitions the fractured shape using the break occupancy, effectively treating estimation of the complete shape and break surface SDF as two partial completion tasks. The normal field (NF) represents surface normals for a shape. Our work uses the NF to capture the difference in surface properties between fractured and intact regions, in objects of materials such as ceramic or earthenware. Our second contribution is a set of loss functions for learning a representation of the complete shape and break surface using neural networks. Our third contribution is to use operations for constructive solid geometry (CSG) in occupancy and SDF space to formalize the dependence of fractured and restoration shapes on complete shapes and break surfaces in occupancy, SDF, and NF space. To generate a restoration mesh, we perform optimization over a fractured shape to obtain complete and break codes, predict the restoration SDF from the codes, and extract a mesh as the 0-level set.

As no dataset of real fractured shapes currently exists, we synthetically fracture 24,208 meshes from the ShapeNet [Chang et al. 2015] dataset and use them to validate our approach. We compare our approach to 3 baselines using existing shape completion approaches, i.e. DeepSDF [Park et al. 2019], ONet [Mescheder et al. 2019], and

ESSC [Zhang et al. 2018]. To demonstrate generalizability to cultural heritage objects and real 3D scans, we synthetically fracture and repair objects from the QP Cultural Heritage dataset [Koutsoudis et al. 2009] and the Google Scanned Objects dataset [GoogleResearch 2022]. We also show restorations for physically fractured objects.

## 2 RELATED WORK

*Fractured Shape Restoration.* Most existing approaches require a user to generate a repair part manually [Antlej et al. 2011; Rengier et al. 2010; Schilling et al. 2014; Scopigno et al. 2011; Seixas et al. 2018], which is out of the scope of the average user. Some approaches repair symmetric objects using reflection followed by subtraction [Gregor et al. 2014; Papaioannou et al. 2017]. These approaches cannot repair asymmetrical objects or objects that have undergone symmetric damage. The approach of Lamb et al. [2019] requires a ground truth complete proxy shape as input, which may be unobtainable e.g. in the case of a rare or specialized object. 3D-ORGAN [Hermosa and Sipiran 2018] performs shape completion from a fractured shape without requiring the shape to be symmetric. However, as 3D-ORGAN encodes shapes as  $32^3$  resolution voxels, their restorations cannot accurately represent the fractured surface and cannot be closely joined to the fractured shape. DeepJoin infers a restoration shape directly without geometric constraints, and generates high-resolution restorations that fit closely to the fractured shape.

*Partial Shape Completion.* Though not directly related to our work, a large body of prior work has focused on performing shape completion from partial inputs using deep neural networks. Many approaches use point clouds [Achlioptas et al. 2018; Dai et al. 2017; Han et al. 2017; Liu et al. 2020; Pan et al. 2021; Sarmad et al. 2019; Son and Kim 2020; Yuan et al. 2018] due to their compactness. However, point clouds cannot intrinsically represent closed surfaces, which are necessary to generate repair parts that may be 3D printed. Approaches that predict meshes directly [Groueix et al. 2018; Yu et al. 2022] struggle to reconstruct complex shapes [Mescheder et al. 2019], and cannot represent shapes of arbitrary topology. Voxel-based approaches [Brock et al. 2016; Sharma et al. 2016; Smith and Meger 2017; Wu et al. 2016; Zhang et al. 2018], become computationally intractable at high resolutions. Though approaches have reduced the memory footprint of voxels using hierarchical models [Dai et al. 2020, 2018] and sparse convolutions [Dai et al. 2020; Yi et al. 2021], these approaches discretize the output space, rendering them incapable of representing high frequency geometry.

Many recent shape completion approaches encode shapes implicitly using the signed distance function (SDF) [Chabra et al. 2020; Duggal et al. 2022; Hao et al. 2020; Lin et al. 2020; Ma et al. 2020; Park et al. 2019; Sitzmann et al. 2020; Tretschk et al. 2020; Xu et al. 2020; Yang et al. 2021; Zheng et al. 2021], the occupancy function [Chen and Zhang 2019; Chibane et al. 2020a; Genova et al. 2020; Jia and Kyan 2020; Liao et al. 2018; Lionar et al. 2021; Mescheder et al. 2019; Peng et al. 2020; Poursaeed et al. 2020; Sulzer et al. 2022; Yan et al. 2022b,a], or the unsigned distance function [Chibane et al. 2020b; Tang et al. 2021; Venkatesh et al. 2020]. DeepSDF [Park et al. 2019] introduced an autoencoder architecture that uses maximum *a posteriori* estimation to perform shape completion by estimating a latent code to fit a set of SDF samples. Approaches have extended the

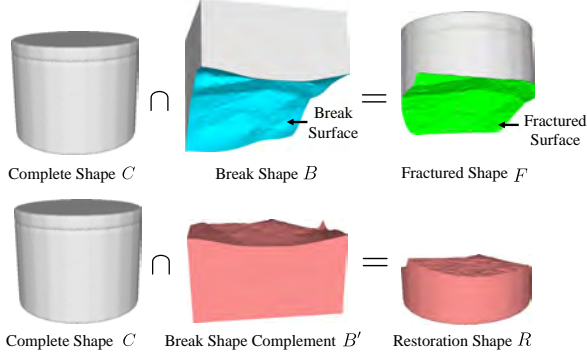


Fig. 3. We express the fractured shape  $F$  and the restoration shape  $R$  as the intersection of the complete shape  $C$  with the break shape  $B$  and break shape complement  $B'$  respectively.  $B$  is separated from  $B'$  by the break surface (blue). The fracture surface (green) is the surface shared by  $B$  and  $F$ .

autoencoder architecture by incorporating meta-learning [Sitzmann et al. 2020], encoding shapes at multiple resolutions [Hao et al. 2020], and proposing incremental loss functions [Duan et al. 2020]. Venkatesh et al. [2020] propose a joint shape function using the unsigned distance and NF, and do not learn the occupancy function. We demonstrate that excluding the occupancy function results in inaccurate restoration shapes with a high chamfer distance.

One could imagine an approach that identifies and removes the fractured region of the fractured shape to create a partial shape, applies a shape completion approach, and performs subtraction of the fractured shape from the complete shape to obtain a restoration shape. As we demonstrate in Section 7.2, approaches based on subtraction produce artifacts on the surface of the fractured shape. Attempts to remove artifacts by discarding components that have a volume less than a threshold cannot eliminate connected artifacts. Our approach automatically generates restoration shapes that join closely to the fractured shape without producing artifacts.

### 3 REPRESENTING FRACTURED SHAPES

We represent complete, fractured, and restoration shapes as point sets  $C$ ,  $F$ , and  $R$  respectively. For a given shape  $S \in \{C, F, R\}$ , we represent the occupancy of a point  $\mathbf{x}$  as  $o_S(\mathbf{x}) \in \{0, 1\}$ , where  $o_S(\mathbf{x}) = 1$  if  $\mathbf{x}$  is inside  $S$  and 0 if  $\mathbf{x}$  is outside or on the boundary of  $S$ . We exclude the boundary of  $S$  to prevent a point from being inside multiple shapes simultaneously, e.g. inside  $F$  and  $R$ . We define the break surface, shown in Figure 3 in blue, as a 2D surface that intersects  $F$  at the fractured region. As shown in Figure 3, we define the break shape,  $B$ , as the set of points on the same side of the break surface as the fractured shape. Theoretically the break shape  $B$  has an infinite size. In practice, we limit the break shape to occupy a unit cube. We define the fractured surface, shown in Figure 3 in green, as the surface shared by the fractured shape and the break surface. We express the fractured shape as the intersection of the complete shape and the break shape, i.e.  $F = C \cap B$ . Similarly, we express the restoration shape as the intersection of the complete shape and the complement of the break shape, i.e.  $R = C \cap B'$ .

For a shape  $S \in \{C, F, R\}$ , we define the SDF value of a point  $\mathbf{x}$  as  $s_S(\mathbf{x}) \in \mathbb{R}$ , the signed distance from  $\mathbf{x}$  to the surface of  $S$ . The value of  $s_S(\mathbf{x})$  is negative inside the shape and positive outside. We define the NF value of a point  $\mathbf{x}$  as  $\mathbf{n}_S(\mathbf{x}) \in \mathbb{S}^3$ , where  $\mathbf{n}_S(\mathbf{x})$  is the normal vector of the closest point to  $\mathbf{x}$  on the surface of  $S$ , and  $\mathbb{S}^3$  is the unit sphere. We define the SDF  $s_B(\mathbf{x})$  and the NF  $\mathbf{n}_B(\mathbf{x})$  for the break shape as the signed distance to the fractured surface, shown in Figure 3 in green, and the normal vector of the closest point to  $\mathbf{x}$  on the fractured surface respectively. The value of  $s_B(\mathbf{x})$  is negative on the fractured side and positive on the restoration side.

The set definitions of  $F$  and  $R$  as  $F = C \cap B$  and  $R = C \cap B'$  imply that the point occupancies  $o_F(\mathbf{x})$  and  $o_R(\mathbf{x})$  for  $F$  and  $R$  can be expressed as a logical conjunction of occupancy in the complete shape, and occupancy in the break shape and the break shape complement, i.e. as  $o_F(\mathbf{x}) = o_C(\mathbf{x}) \wedge o_B(\mathbf{x})$  and  $o_R(\mathbf{x}) = o_C(\mathbf{x}) \wedge \neg o_B(\mathbf{x})$ . The symbols  $\wedge$  and  $\neg$  represent the logical and and not operators. We relax these logical relationships to work with continuous values using the product T-norm [Gupta and Qi 1991], as

$$o_F(\mathbf{x}) = o_C(\mathbf{x})o_B(\mathbf{x}) \text{ and} \quad (1)$$

$$o_R(\mathbf{x}) = o_C(\mathbf{x})(1 - o_B(\mathbf{x})). \quad (2)$$

Figure 4 demonstrates the dependency of  $o_F(\mathbf{x})$  and  $o_R(\mathbf{x})$  on  $o_C(\mathbf{x})$  and  $o_B(\mathbf{x})$  given by Equations (1) and (2). In the figure,  $o_F(\mathbf{x})$  and  $o_R(\mathbf{x})$  can be seen as using the break shape and the inverse break shape to carve out the complete shape from a CSG perspective.

To compute the SDF value  $s_F(\mathbf{x})$  for the fractured shape from the SDF values for the complete and break shapes, we express  $s_F(\mathbf{x})$  as

$$s_F(\mathbf{x}) = \begin{cases} s_B(\mathbf{x}), & \text{if } o_B(\mathbf{x}) = 0 \text{ or } s_B(\mathbf{x}) > s_C(\mathbf{x}) \\ s_C(\mathbf{x}), & \text{otherwise.} \end{cases} \quad (3)$$

Visually, Figure 4 demonstrates how the SDF values for  $s_F(\mathbf{x})$  are obtained from the break shape for points outside the break shape and close to the break surface, i.e. when  $s_B(\mathbf{x}) > s_C(\mathbf{x})$ , following a CSG perspective in SDF space [Breen et al. 2000], similar to the perspective for occupancy. We express the SDF value  $s_R(\mathbf{x})$  for the restoration shape in terms of the complete and break shapes as

$$s_R(\mathbf{x}) = \begin{cases} -s_B(\mathbf{x}), & \text{if } o_B(\mathbf{x}) = 1 \text{ or } -s_B(\mathbf{x}) > s_C(\mathbf{x}) \\ s_C(\mathbf{x}), & \text{otherwise.} \end{cases} \quad (4)$$

We negate values for the break shape in Equation (4) compared to Equation (3) as the restoration shape is on the opposite side of the break surface from the fractured shape. Figure 4 illustrates how the SDF values for  $s_R(\mathbf{x})$  are obtained from the inverted break shape for points outside the break shape and near the break surface where the SDF for the break shape is more positive than the complete shape, i.e. when  $-s_B(\mathbf{x}) > s_C(\mathbf{x})$  holds. Computing the SDF for the fractured and restoration shapes as a function of the occupancy for the break shape allows our approach to treat the estimation of the complete and break SDF as two partial shape completion problems, where the partial shapes are defined by the break occupancy.

Similar to the SDF, the NF value at a point  $\mathbf{x}$  for a shape  $S$  is given by the closest point on the surface of  $S$  to  $\mathbf{x}$ . We use the same approach as in Equations (3) and (4) to define the fractured NF  $\mathbf{n}_F(\mathbf{x})$

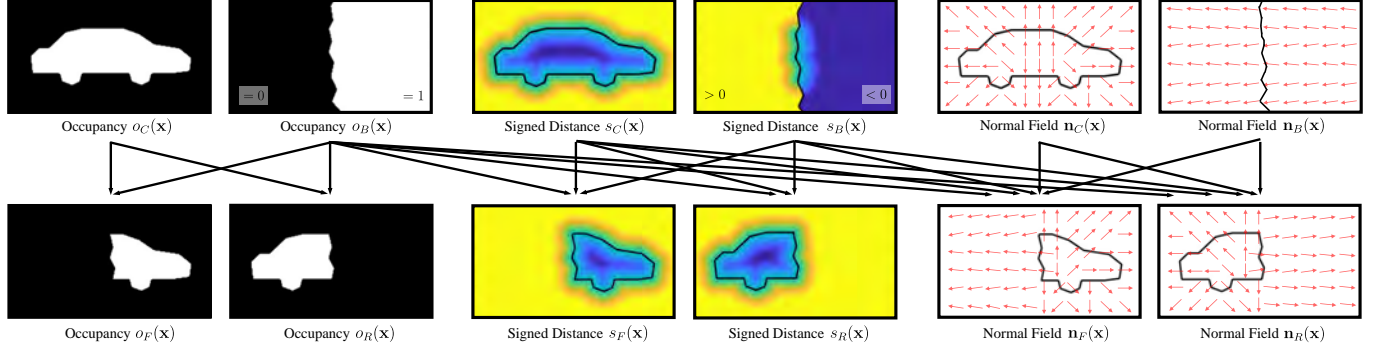


Fig. 4. The fractured  $o_F(\mathbf{x})$  and restoration occupancy  $o_R(\mathbf{x})$  are obtained by carving out the complete shape with the break shape. For the SDF and NF, the fractured shape values  $s_F(\mathbf{x})$  and  $\mathbf{n}_F(\mathbf{x})$  are the same as the break shape for points outside and near the surface of the break shape, i.e., the same as  $s_B(\mathbf{x})$  and  $\mathbf{n}_B(\mathbf{x})$ . The restoration shape values,  $s_R(\mathbf{x})$  and  $\mathbf{n}_R(\mathbf{x})$ , are the opposite of the break shape for points inside and near the surface of the break shape.

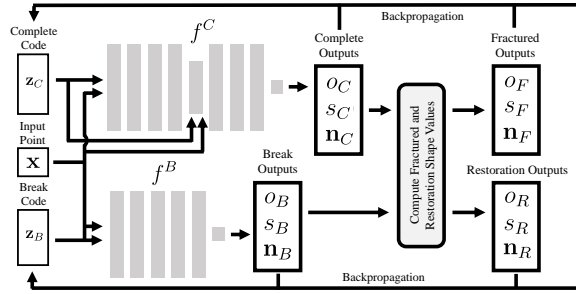


Fig. 5. Networks  $f^C$  and  $f^B$  predict the occupancy, SDF, and NF values for  $C$  and  $B$  given an input point  $\mathbf{x}$  and latent codes  $\mathbf{z}_C$  and  $\mathbf{z}_B$ . We compute the occupancy, SDF, and NF values for  $F$  and  $R$  from the values for  $C$  and  $B$ .

and restoration NF  $\mathbf{n}_R(\mathbf{x})$  in terms of  $C$  and  $B$ , as

$$\mathbf{n}_F(\mathbf{x}) = \begin{cases} \mathbf{n}_B(\mathbf{x}), & \text{if } o_B(\mathbf{x}) = 0 \text{ or } s_B(\mathbf{x}) > s_C(\mathbf{x}) \\ \mathbf{n}_C(\mathbf{x}), & \text{otherwise, and} \end{cases} \quad (5)$$

$$\mathbf{n}_R(\mathbf{x}) = \begin{cases} -\mathbf{n}_B(\mathbf{x}), & \text{if } o_B(\mathbf{x}) = 1 \text{ or } -s_B(\mathbf{x}) > s_C(\mathbf{x}) \\ \mathbf{n}_C(\mathbf{x}), & \text{otherwise.} \end{cases} \quad (6)$$

We negate the NF for the break shape in the definition of  $\mathbf{n}_R(\mathbf{x})$  in Equation (6), as normals on the restoration shape at the fracture are oriented in the opposite direction of normals on the break surface. Figure 4 shows that the normals for the fractured and break shape are the same outside of the break shape and near the break surface, i.e. where  $s_B(\mathbf{x}) > s_C(\mathbf{x})$  holds. Normals for the restoration shape are also the opposite of the break shape inside the break and near the break surface, i.e. where  $-s_B(\mathbf{x}) > s_C(\mathbf{x})$  holds.

We represent the joint occupancy, SDF, and NF function  $f^S$  for a shape  $S \in \{C, B\}$  using a neural network based on the autoencoder architecture of DeepSDF [Park et al. 2019] as shown in Figure 5. We condition the autoencoder for  $S$  by associating  $S$  with a latent code  $\mathbf{z}_S \in \mathbb{R}^p$  fed as input to the network, where  $p$  is the size of the code. We use  $p = 128$  and  $64$  for the complete and break shapes respectively. We use  $f_o^S(\mathbf{z}_S, \mathbf{x})$ ,  $f_s^S(\mathbf{z}_S, \mathbf{x})$ , and  $f_n^S(\mathbf{z}_S, \mathbf{x})$  to refer to the occupancy, SDF, and NF values predicted at point  $\mathbf{x}$ .

## 4 NETWORK OPTIMIZATION

To train our networks, we use a dataset of sample shapes where each sample is a tuple  $\{C, F, R, B\}$  corresponding to a complete, fractured, restoration, and break shape. We perform optimization over the network parameters, and over the complete and break shape codes. For a shape  $S$ , we use the notation  $o_S(\mathbf{x})$ ,  $s_S(\mathbf{x})$ , and  $\mathbf{n}_S(\mathbf{x})$  to represent the ground truth values for occupancy, the SDF, and NF respectively. During training, we optimize the loss

$$\mathcal{L}_{\text{train}} = \sum_{\mathbf{z}_C \in \mathcal{Z}_C, \mathbf{z}_B \in \mathcal{Z}_B} \mathcal{L}_{CB} + \mathcal{L}_F + \mathcal{L}_R + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (7)$$

where  $\mathcal{Z}_C$  and  $\mathcal{Z}_B$  are the sets of complete and break shape latent codes. We define the loss  $\mathcal{L}_{CB}$  for the complete and break shapes as

$$\mathcal{L}_{CB} = \frac{1}{|\mathcal{X}|} \sum_{S \in \{C, B\}} \sum_{\mathbf{x} \in \mathcal{X}} (BCE(f_o^S(\mathbf{z}_S, \mathbf{x}), o_S(\mathbf{x})) + \lambda_s \|f_s^S(\mathbf{z}_S, \mathbf{x}) - s_S(\mathbf{x})\|_1 + \lambda_n \|f_n^S(\mathbf{z}_S, \mathbf{x}) - \mathbf{n}_S(\mathbf{x})\|_2), \quad (8)$$

where  $BCE$  is the binary cross-entropy loss and  $\mathcal{X}$  is a set of probing sample points. We define the fractured shape loss  $\mathcal{L}_F$  as

$$\mathcal{L}_F = \frac{1}{|\mathcal{X}|} \sum_{\mathbf{x} \in \mathcal{X}} (BCE(f_o^C(\mathbf{z}_C, \mathbf{x}) f_o^B(\mathbf{z}_B, \mathbf{x}), o_F(\mathbf{x})) + \lambda_s \mathcal{L}_{F_s} + \lambda_n \mathcal{L}_{F_n}). \quad (9)$$

The first argument to  $BCE$ , corresponds to the definition for fractured occupancy from Equation (1). We represent the reconstruction error for the fractured shape SDF  $\mathcal{L}_{F_s}$ , as

$$\mathcal{L}_{F_s} = \begin{cases} \|f_s^B(\mathbf{z}_B, \mathbf{x}) - s_F(\mathbf{x})\|_1 & \text{if } f_o^B(\mathbf{z}_B, \mathbf{x}) \leq \mu \text{ or } f_s^B(\mathbf{z}_B, \mathbf{x}) > f_s^C(\mathbf{z}_C, \mathbf{x}), \\ \|f_s^C(\mathbf{z}_C, \mathbf{x}) - s_F(\mathbf{x})\|_1 & \text{otherwise,} \end{cases} \quad (10)$$

and the reconstruction error for the fractured shape NF  $\mathcal{L}_{F_n}$ , as

$$\mathcal{L}_{F_n} = \begin{cases} \|f_n^B(\mathbf{z}_B, \mathbf{x}) - \mathbf{n}_F(\mathbf{x})\|_2 & \text{if } f_o^B(\mathbf{z}_B, \mathbf{x}) \leq \mu \text{ or } f_s^B(\mathbf{z}_B, \mathbf{x}) > f_s^C(\mathbf{z}_C, \mathbf{x}), \\ \|f_n^C(\mathbf{z}_C, \mathbf{x}) - \mathbf{n}_F(\mathbf{x})\|_2 & \text{otherwise.} \end{cases} \quad (11)$$

Equations (10) and (11) use the definitions for the fractured shape SDF and NF given in Equations (3) and (5). We use  $\mu = 0.5$  as a



threshold to determine if a point is inside the predicted break shape.  $\mathcal{L}_R$ , the loss for the restoration shape, is given as

$$\mathcal{L}_R = \frac{1}{|\mathcal{X}|} (\sum_{\mathbf{x} \in \mathcal{X}} BCE(f_o^C(\mathbf{z}_C, \mathbf{x})(1 - f_o^B(\mathbf{z}_B, \mathbf{x})), o_R(\mathbf{x})) + \lambda_s \mathcal{L}_{R_s} + \lambda_n \mathcal{L}_{R_n}). \quad (12)$$

The first argument to  $BCE$  in Equation (12) corresponds to the definition for restoration shape occupancy given in Equation (2). We give the restoration SDF reconstruction error  $\mathcal{L}_{R_s}$  as

$$\mathcal{L}_{R_s} = \begin{cases} ||-f_s^B(\mathbf{z}_B, \mathbf{x}) - s_R(\mathbf{x})||_1 & \text{if } f_o^B(\mathbf{z}_B, \mathbf{x}) > \mu \text{ or} \\ & -f_s^B(\mathbf{z}_B, \mathbf{x}) > f_s^C(\mathbf{z}_C, \mathbf{x}), \\ ||f_s^C(\mathbf{z}_C, \mathbf{x}) - s_R(\mathbf{x})||_1 & \text{otherwise,} \end{cases} \quad (13)$$

and the error for the restoration shape NF  $\mathcal{L}_{R_n}$  as

$$\mathcal{L}_{R_n} = \begin{cases} ||-f_n^B(\mathbf{z}_B, \mathbf{x}) - \mathbf{n}_R(\mathbf{x})||_2 & \text{if } f_o^B(\mathbf{z}_B, \mathbf{x}) > \mu \text{ or} \\ & -f_n^B(\mathbf{z}_B, \mathbf{x}) > f_n^C(\mathbf{z}_C, \mathbf{x}), \\ ||f_n^C(\mathbf{z}_C, \mathbf{x}) - \mathbf{n}_R(\mathbf{x})||_2 & \text{otherwise,} \end{cases} \quad (14)$$

which use the definitions for the restoration shape SDF and NF from Equations (4) and (6). We negate the SDF and NF value for the predicted break shape in Equations (13) and (14) as the restoration lies on the opposite side of the break surface from the fractured shape. We define the regularization  $\mathcal{L}_{\text{reg}}$  loss as

$$\mathcal{L}_{\text{reg}} = ||\mathbf{z}_B||_1 + ||\mathbf{z}_C||_1. \quad (15)$$

$\mathcal{L}_{\text{reg}}$  imposes a zero-mean Laplacian prior on the complete and break codes. We use  $\lambda_s = 1.0$ ,  $\lambda_n = 1e - 1$ , and  $\lambda_{\text{reg}} = 1e - 4$ , i.e. the coefficients for the SDF, NF, and regularization losses respectively. We use the Adam optimizer [Kingma and Ba 2014].

## 5 INFERRING RESTORATION SHAPES

At inference time our approach generates a restoration shape for a novel fractured shape by performing optimization over occupancy, SDF, and NF samples from the fractured shape to obtain complete and break shape codes. We use the codes to predict the SDF value for the restoration shape. During inference we optimize the loss

$$\mathcal{L}_{\text{inf}} = \mathcal{L}_F + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (16)$$

We find that the surface for the break shape estimated in occupancy space may deviate slightly from the surface in SDF space, causing artifacts if Equation (4) is used to reconstruct the restoration mesh. To prevent artifacting, when generating the restoration mesh we obtain SDF values using the CSG equation for Boolean subtraction [Breen et al. 2000] in SDF space, i.e.  $s_R(\mathbf{x}) = \max(f_s^C(\mathbf{z}_C, \mathbf{x}), -f_s^B(\mathbf{z}_B, \mathbf{x}))$ , in place of Equation (4). To obtain a restoration mesh, we perform Marching Cubes [Lorensen and Cline 1987] on a  $256^3$  grid of points.

## 6 DATA PROCESSING AND DATASETS

We evaluate our approach on four datasets.

*ShapeNet*. We use 3 ShapeNet [Chang et al. 2015] classes corresponding to commonly fractured objects, i.e. jars, bottles, and mugs, and 5 classes with more complex geometry, i.e. airplanes, chairs, cars, tables, and sofas. We train one network per ShapeNet class.

*Google Scanned Objects dataset*. The dataset [GoogleResearch 2022] contains 3D scanned meshes of household objects such as shoes, pots, and plates. We train one network for the entire dataset.

*QP Cultural Heritage dataset*. The dataset [Koutsoudis et al. 2009] contains artist designed meshes in the style of Greek pottery. We use all meshes for testing on a network trained on ShapeNet jars.

*Real Fractured Objects*. We fracture and scan 3 mugs and use 2 items from Lamb et al. [2019]. We test them against networks trained on synthetically fractured ShapeNet mugs and jars respectively.

As the ShapeNet, Google, and QP datasets do not contain fractured meshes, we synthetically fracture meshes from these datasets using the fracturing approach described by Lamb et al. [2021]. To generate closed meshes we use the approach of Stutz and Geiger [2020]. We normalize meshes so they occupy a unit cube. We perform a fracture retention test by fracturing each mesh using a randomized geometric primitive 15 times. If between 5% and 20% of the vertices of the mesh are not removed after 15 attempts we discard the mesh. As the bottles, jars, and mugs ShapeNet classes have less than 600 samples we fracture meshes from these classes 3, 3, and 10 times respectively. We fracture all other meshes once. We retain 24,208 out of 26,166 meshes from ShapeNet, 1,042 out of 1,298 meshes from the Google dataset, and 333 out of 408 meshes from the QP dataset.

Though meshes from the Google dataset are oriented upright they are not facing in a uniform direction. We augment the Google dataset by randomly rotating meshes around the ground plane normal by 90 degrees. We partition the ShapeNet and Google datasets using a 70%/10%/20% train/validation/testing split. To obtain ground truth break surfaces we fit a thin-plate spline (TPS) [Duchon 1977] to the fractured region of each mesh such that the domain of the spline corresponds to a plane that is fitted to the fracture region vertices. We use the TPS to partition sample points into two sets, and denote the set that intersects with the fractured shape as the break shape. We discuss point our point sampling method in the supplementary.

## 7 RESULTS

We use the chamfer distance (CD), as defined by Park et al. [2019] and the normal consistency (NC), as defined by Mescheder et al. [2019] to evaluate the overall accuracy of predicted restoration shapes. In Section 7.2, we compare our approach to methods based on subtraction from completed shapes. Restorations generated using subtraction tend to exhibit physically implausible surface artifacts. We contribute the non-fractured region error (NFRE) metric to evaluate the degree of surface artifacting. To compute the NFRE we sample  $n$  points on the non-fractured region of the fractured shape and on the predicted and ground truth restoration shapes, and compute the percentage of points on the non-fractured region with a nearest neighbor on the predicted restoration that is closer than  $\eta$  and a nearest neighbor on the ground truth restoration farther than  $\eta$ . We use  $n = 30,000$  and  $\eta = 0.02$ . For success the NFRE and CD is low and the NC is high. In 1.2% of cases, our approach generates an empty restoration. Where applicable, we show the non-empty percentage, (NE%), i.e. the percentage of restorations generated.

We show restoration shapes generated using our approach on the ShapeNet dataset in Figure 6. Our approach generates restoration shapes that join closely to the fractured region of the fractured shape




Fig. 6. Predicted restoration shapes (red), joined to input fractured shapes (gray) and opened to show the fracture.

and restore complex structures such as the handle of the mug on the right, and the thin table leg on the left of Figure 6. Restoration shapes keep with the style of the input shape, e.g. the connected feet of the sofa in the bottom right. Unlike approaches based on symmetry [Gregor et al. 2014; Papaioannou et al. 2017], our approach repairs asymmetrical objects and objects with symmetrical fractures, such as the L-shaped sofa, airplane, beer bottle, vase, mugs, and both cars in Figure 6, and the mug, sofa, dresser, and bottle in Figure 1.

In Figure 7 we show restorations generated for meshes from the QP Cultural Heritage dataset, 3D scans from the Google Scanned Objects dataset, and real fractured objects. Our approach obtains a CD of 0.117 on the Google Scanned Objects dataset and a CD of 0.144 on the QP dataset. Our approach is able to generate plausible restoration shapes for Greek pottery, as shown in Figure 7(a), even having never been trained on objects from that time period. Though the Google Scanned Objects dataset is highly varied, our approach generates closely fitting repairs for simple household objects such as plates, bowls and cups, and for more complex objects, e.g. shoes, as shown in Figure 7(b). Figure 7(c) demonstrates that our approach generates feasible restoration shapes for real fractured objects, even when trained entirely on synthetic fractures. We 3D print a restoration part for the candlestick and the mug on the right. Though the base of the candlestick does not match the complete shape, the predicted restoration is physically plausible. For the mug, while small deviations in structure introduced by waterproofing and printer tolerances occur, the print provides a close fit, enabling repair.

### 7.1 Ablation Study: Joint Function Modalities

We evaluate the impact of learning a joint function for multiple features by training our approach to represent fractured shapes using joint functions for the SDF alone, for occupancy alone ('Occ'), SDF and NF ('SDF+NF'), occupancy and SDF ('Occ+SDF'), and for all three features ('Occ+SDF+NF'). We show the CD and NE% for

Method	bottles	cars	chairs	jars	mugs	planes	sofas	tables	Mean	NE%
Occ	0.047	0.089	0.159	0.092	0.048	0.057	0.135	0.170	0.099	89.4
SDF+NF	0.225	0.130	0.187	0.176	0.153	0.122	0.235	0.202	0.179	96.1
Occ+SDF	0.042	0.023	0.127	<b>0.071</b>	0.028	0.043	0.101	0.159	0.074	<b>98.9</b>
Occ+SDF+NF	<b>0.034</b>	<b>0.018</b>	<b>0.089</b>	0.090	<b>0.027</b>	<b>0.033</b>	<b>0.077</b>	<b>0.128</b>	<b>0.062</b>	98.8

Table 1. Chamfer (CD) and percentage of non-empty restorations (NE%), using DeepJoin with different features. Best values are bolded.




Fig. 7. Predicted and ground truth restorations for synthetically fractured (a) objects in the style of ancient Greek pottery and (b) 3D scans of common objects. (c) Objects with real fractures and with two 3D printed restorations. Complete objects from Lamb et al. [2019] are also shown in gray.

each variation of our approach in Table 1, over all restorations predicted by each approach. We do not show results for SDF as we find that the network is not able to learn a stable representation for the break shape and predicts no restorations. As shown in Table 1, using occupancy gives a low CD of 0.099. However, it often fails to converge, with a NE% of 89.4%. Though SDF+NF shows a higher CD than Occ, it predicts restorations more often. Using Occ+SDF shows lower CD compared to Occ, and predicts more restorations than SDF+NF, with a NE% of 98.9%. Adding NF gives lowest CD and a similar NE%, showing that NF improves restoration fidelity.

### 7.2 Comparison: Fracture Removal and Shape Completion

As no generalizable high-resolution shape repair approaches exist, we compare our approach to three baseline approaches that automatically remove the fractured region and perform shape completion. The baseline approaches use the existing shape completion approaches of DeepSDF [Park et al. 2019], ONet [Mescheder et al. 2019], and ESSC [Zhang et al. 2018]. For each shape completion approach we create a partial input shape by removing the fractured region. To generate partial shapes, we train a point cloud classifier based on PointNet++ [Qi et al. 2017] to classify the fractured region. Our classifier obtains a test accuracy of 97.1%.

We train DeepSDF to reconstruct complete shapes from SDF point samples. We create a partial shape for inference by removing input sample points that have a nearest neighbor in the fractured region of the fractured shape identified by our classifier. We train ONet to reconstruct complete shapes from complete point clouds. We create a partial point cloud for inference by removing points that are classified as belonging to the fractured region identified by our classifier. 3D-ORGAN [Hermosa and Sipiran 2018] generates low-resolution voxelized restoration shapes from fractured shapes. We find that the approach is unstable and does not converge during training. To provide a fair comparison to a voxel-based approach,

Method	Metric	bottles	cars	chairs	jars	mugs	planes	sofas	tables	Mean
DSDF	CD	<b>0.024</b>	0.025	0.109	0.134	0.043	0.034	<b>0.076</b>	0.129	0.072
	NC	<b>0.723</b>	0.636	0.471	0.506	0.679	0.562	<b>0.591</b>	0.495	0.583
	NFRE	0.084	0.158	0.316	0.217	0.092	0.125	0.215	0.238	0.181
ONet	CD	0.103	0.122	0.160	0.133	0.150	0.121	0.148	0.190	0.141
	NC	0.441	0.428	0.416	0.391	0.445	0.384	0.418	0.460	0.423
	NFRE	0.829	0.753	0.518	0.538	0.523	0.787	0.649	0.473	0.634
ESSC	CD	0.175	0.069	0.167	<b>0.056</b>	0.052	0.126	0.152	<b>0.122</b>	0.115
	NC	0.108	0.388	0.153	0.414	0.481	0.138	0.251	0.377	0.289
	NFRE	0.346	0.114	0.168	<b>0.051</b>	0.029	0.178	0.298	<b>0.030</b>	0.152
Ours	CD	0.034	<b>0.018</b>	<b>0.089</b>	0.090	<b>0.027</b>	<b>0.033</b>	0.077	0.128	<b>0.062</b>
	NC	0.687	<b>0.749</b>	<b>0.567</b>	<b>0.558</b>	<b>0.783</b>	<b>0.660</b>	0.537	<b>0.505</b>	<b>0.631</b>
	NFRE	<b>0.052</b>	<b>0.021</b>	<b>0.042</b>	0.054	<b>0.008</b>	<b>0.024</b>	<b>0.064</b>	0.040	<b>0.038</b>

Table 2. Chamfer distance (CD), normal consistency (NC) and non-fracture region error (NFRE) for baseline approaches and our approach. Best metric values are bolded. Mean is computed over class means.




Fig. 8. From left to right: input fractured shapes (gray) and restorations (red) from DeepSDF, ONet, ESSC, DeepJoin, and ground truth.

we use ESSC [Zhang et al. 2018] to perform partial shape completion. We train ESSC to reconstruct complete voxel grids from partial voxel grids at  $32^3$  spatial resolution to match the resolution of 3D-ORGAN. We use an input resolution of  $128^3$  encoded with the flipped-truncated SDF (FTSDF) [Song et al. 2017]. During inference, we create a partial input by computing the FTSDF with respect to a mesh with the fractured region, identified by our classifier, removed.

For each baseline approach we obtain a restoration shape by subtracting the predicted complete shape from the input fractured shape in occupancy space. For DeepSDF and ONet we obtain a mesh using Marching Cubes for  $256^3$  sample points. As restoration shapes must be closed meshes and not voxel grids, we generate a mesh for ESSC using Marching Cubes at  $32^3$  resolution. To mitigate surface artifacts, for each of the baseline approaches we automatically remove connected components from the predicted restoration shape that have a volume less than  $\delta$ , where we use  $\delta = 0.01$ .

Table 2 summarizes metrics over non-empty restorations predicted by our approach. Comparative approaches generate all non-empty restorations. DeepJoin outperforms baseline approaches in terms of the mean CD, NC, and NFRE. DeepSDF predicts restorations with a relatively low CD of 0.072 and high NC of 0.583 compared to the other baselines. However, as shown by the NFRE of 0.181 in Table 2, restorations predicted by DeepSDF exhibit artifacts on the surface of the fractured shape e.g. for the bottle, sofa, and pot in Figure 6. Complete shapes predicted by ONet are less accurate than DeepSDF, resulting in restoration shapes with large artifacts that




Fig. 9. Left: The complete and break shape may not intersect. Right: Restorations may be small predicted in the wrong location.

often cover most of the surface of the fractured shape, as shown in Figure 8, and by the NFRE of 0.634. ESSC shows a the lowest mean CD for 2 classes, and often predicts restoration shapes in the correct location and without surface artifacts. However, the restoration shapes do not precisely match the ground truth, as shown by the NC of 0.289, do not join to the fractured shape, e.g. the cutaway in Figure 8, and are too small to fully restore the fractured shape due to their coarse structure. DeepJoin generates accurate restoration shapes that do not exhibit artifacts, as shown in Figure 8.

## 8 CONCLUSION AND FUTURE WORK

In this paper we present DeepJoin, an approach to automatically generate restoration shapes from input fractured shapes by deconstructing the fractured shape into corresponding complete and break shapes. We contribute a novel method to encode a shape using a joint function for occupancy, signed distance, and NF and use this encoding to regress a high resolution restoration shape. Our approach overcomes the disadvantages of prior work using voxels for automated shape repair, and plays an important role in enabling consumer-driven object repair using additive manufacturing.

*Limitations.* In 1.2% of cases our approach fails to generate the restoration shape, e.g. on the left of Figure 9, when the break surface is predicted away from the fractured shape. However, as our approach also predicts a complete shape, it enables fallback to generating a restoration using subtraction. As we represent break surfaces using thin-plate splines, our approach is unable to accurately represent concave breaks. In future work, we will explore using more generalizable representations for the break surface, such as non-uniform rational basis splines (NURBS) [Piegl and Tiller 1996]. Our approach may predict a restoration shape that is smaller than the ground truth, e.g. the cabinet in Figure 9, or may predict the restoration shape in the wrong location if the fractured region is small, e.g. the TV stand in Figure 9. However, these repair parts can still be used to partially repair the fracture, e.g. for the table.

*Future Work.* Though we focus on single component fractures, future work can explore simultaneous estimation of inter-part transformations to facilitate multi-fracture assembly and repair. In future work, we will investigate automated repair deformation near the fracture surface to enable fine-precision joins of 3D prints for physical water-tightness. Though our fractured shapes simulate the surface roughness found in ceramic or earthenware objects, proliferation of our approach requires large-scale analysis of fracture across diverse materials. Our future work will contribute a comprehensive fracture dataset of 3D scans of physically fractured objects with accurate waterproofing, to facilitate data-driven study of fracture.

## REFERENCES

- Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. 2018. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*. *International conference on machine learning* 80, 35, 40–49.
- Kaja Antleij, Miran Eric, Mojca Savnik, Bernarda Zupanek, Janja Slabe, and B Borut Battestin. 2011. Combining 3D Technologies in the Field of Cultural Heritage: Three Case Studies.. In *VAST (Short and Project Papers)*. The Eurographics Association, Geneva, Switzerland, 1–4.
- David E Breen, Sean Mauch, and Ross T Whitaker. 2000. 3d scan-conversion of CSG models into distance, closest-point and colour volumes. In *Volume Graphics*. Springer, London, UK, 135–158.
- Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. 2016. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236* 1, 1 (2016), 1–9.
- Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. 2020. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*. Springer, Berlin, Germany, 608–625.
- Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. 2015. *ShapeNet: An Information-Rich 3D Model Repository*. Technical Report arXiv:1512.03012 [cs.GR]. Stanford University — Princeton University — Toyota Technological Institute at Chicago.
- Zhiqin Chen and Hao Zhang. 2019. Learning implicit fields for generative shape modeling. In *Proc. CVPR*. IEEE, Piscataway, NJ, 5939–5948.
- Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. 2020a. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proc. CVPR*. IEEE, Piscataway, NJ, 6970–6981.
- Julian Chibane, Gerard Pons-Moll, et al. 2020b. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems* 33 (2020), 21638–21652.
- Angela Dai, Christian Diller, and Matthias Nießner. 2020. Sg-nn: Sparse generative neural networks for self-supervised scene completion of rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 849–858.
- Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, and Matthias Nießner. 2018. Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In *Proc. CVPR*. IEEE, Piscataway, NJ, 4578–4587.
- Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. 2017. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 5868–5877.
- Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J Guibas. 2020. Curriculum deepsf. In *European Conference on Computer Vision*. Springer, Berlin, Germany, 51–67.
- Jean Duchon. 1977. Splines minimizing rotation-invariant semi-norms in Sobolev spaces. In *Constructive theory of functions of several variables*. Springer, Berlin, Germany, 85–100.
- Shivam Duggal, Zihao Wang, Wei-Chiu Ma, Sivabalan Manivasagam, Justin Liang, Shenlong Wang, and Raquel Urtasun. 2022. Mending Neural Implicit Modeling for 3D Vehicle Reconstruction in the Wild. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. IEEE, Piscataway, NJ, 1900–1909.
- Jian Gao, Xin Chen, Oguzhan Yilmaz, and Nabil Gindy. 2008. An integrated adaptive repair solution for complex aerospace components through geometry reconstruction. *The International Journal of Advanced Manufacturing Technology* 36, 11–12 (2008), 1170–1179.
- Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. 2020. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 4857–4866.
- GoogleResearch. 2022. *Google Scanned Objects*. Open Robotics. <https://fuel.gazebosim.org/1.0/GoogleResearch/fuel/collections/Google%20Scanned%20Objects>
- Robert Gregor, Ivan Sipiran, Georgios Papaioannou, Tobias Schreck, Anthousis Andreadis, and Pavlos Mavridis. 2014. Towards Automated 3D Reconstruction of Defective Cultural Heritage Objects.. In *GCH. EUROGRAPHICS*, Geneva, Switzerland, 135–144.
- Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 2018. A papier-mâché approach to learning 3d surface generation. In *Proc. CVPR*. IEEE, Piscataway, NJ, 216–224.
- Madan M Gupta and J11043360726 Qi. 1991. Theory of T-norms and fuzzy inference methods. *Fuzzy sets and systems* 40, 3 (1991), 431–450.
- Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. 2017. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *Proceedings of the IEEE international conference on computer vision*. IEEE, Piscataway, NJ, 85–93.
- Zekun Hao, Hadar Averbuch-Elor, Noah Snavely, and Serge Belongie. 2020. Dualsf: Semantic shape manipulation using a two-level representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 7631–7641.
- Ola LA Harrysson, Yasser A Hosni, and Jamal F Nayfeh. 2007. Custom-designed orthopedic implants evaluated using finite element analysis of patient-specific computed tomography data: femoral-component case study. *BMC musculoskeletal disorders* 8, 1 (2007), 1–10.
- Renato Hermoza and Ivan Sipiran. 2018. 3D reconstruction of incomplete archaeological objects using a generative adversarial network. In *Proceedings of Computer Graphics International*. ACM, New York, NY, 5–11.
- Meng Jia and Matthew Kyan. 2020. Learning Occupancy Function from Point Clouds for Surface Reconstruction. *arXiv preprint arXiv:2010.11378* 1 (2020), 1–11.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *Proc. ICLR*. International Conference on Representation Learning, La Jolla, CA, 1–15.
- Anestis Koutsoudis, George Pavlidis, Fotis Arnaoutoglou, Despina Tsiafakis, and Christodoulos Chamzas. 2009. Qp: A tool for generating 3D models of ancient Greek pottery. *Journal of Cultural Heritage* 10, 2 (2009), 281–295.
- Nikolas Lamb, Sean Banerjee, and Natasha Kholgade Banerjee. 2019. Automated reconstruction of smoothly joining 3d printed restorations to fix broken objects. In *Proc. SCF*. ACM, New York, NY, 1–12.
- Nikolas Lamb, Noah Wiederhold, Benjamin Lamb, Sean Banerjee, and Natasha Kholgade Banerjee. 2021. Using Learned Visual and Geometric Features to Retrieve Complete 3D Proxies for Broken Objects. In *Proc. SCF*. ACM, New York, NY, 1–15.
- Yiyi Liao, Simon Donne, and Andreas Geiger. 2018. Deep marching cubes: Learning explicit surface representations. In *Proc. CVPR*. IEEE, Piscataway, NJ, 2916–2925.
- Chen-Hsuan Lin, Chaoyang Wang, and Simon Lucey. 2020. SDF-SRN: Learning Signed Distance 3D Object Reconstruction from Static Images. *arXiv preprint arXiv:2010.10505* 1, 1 (2020), 1–17.
- Stefan Lionar, Danil Emtsev, Dusan Svilarkovic, and Songyou Peng. 2021. Dynamic Plane Convolutional Occupancy Networks. In *Proc. WACV*. IEEE, Piscataway, NJ, 1829–1838.
- Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. 2020. Morphing and sampling network for dense point cloud completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. AAAI, Menlo Park, CA, 11596–11603.
- William E Lorensen and Harvey E Cline. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics* 21, 4 (1987), 163–169.
- Baorui Ma, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. 2020. Neural-Pull: Learning Signed Distance Functions from Point Clouds by Learning to Pull Space onto Surfaces. *arXiv preprint arXiv:2011.13495* 1, 1 (2020), 1–12.
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proc. CVPR*. IEEE, Piscataway, NJ, 4460–4470.
- Liang Pan, Xinyi Chen, Zhongang Cai, Junzhe Zhang, Haiyu Zhao, Shuai Yi, and Ziwei Liu. 2021. Variational relational point completion network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 8524–8533.
- Georgios Papaioannou, Tobias Schreck, Anthousis Andreadis, Pavlos Mavridis, Robert Gregor, Ivan Sipiran, and Konstantinos Vardis. 2017. From reassembly to object completion: A complete systems pipeline. *Journal on Computing and Cultural Heritage* 10, 2 (2017), 1–22.
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. Deepsf: Learning continuous signed distance functions for shape representation. In *Proc. CVPR*. IEEE, Piscataway, NJ, 165–174.
- Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. 2020. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III* 16. Springer, Berlin, Germany, 523–540.
- Les Pielg and Wayne Tiller. 1996. *The NURBS book*. Springer Science & Business Media, Berlin, Germany.
- Omid Poursaeed, Matthew Fisher, Noam Aigerman, and Vladimir G Kim. 2020. Coupling explicit and implicit surface representations for generative 3d modeling. In *European Conference on Computer Vision*. Springer, Berlin, Germany, 667–683.
- Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* 30 (2017), 1–10.
- Fabian Rengier, Amit Mehndiratta, Hendrik Von Tengg-Kobligh, Christian M Zechmann, Roland Unterhinninghofen, H-U Kauczor, and Frederik L Giesel. 2010. 3D printing based on imaging data: review of medical applications. *International journal of computer assisted radiology and surgery* 5, 4 (2010), 335–341.
- Muhammad Sarmad, Hyunjo Jenny Lee, and Young Min Kim. 2019. RL-gan-net: A reinforcement learning agent controlled gan network for real-time point cloud shape completion. In *Proc. CVPR*. IEEE, Piscataway, NJ, 5898–5907.
- René Schilling, Benjamin Jastram, Oliver Wings, Daniela Schwarz-Wings, and Ahi Sema Issever. 2014. Reviving the dinosaur: virtual reconstruction and three-dimensional printing of a dinosaur vertebra. *Radiology* 270, 3 (2014), 864–871.



- Roberto Scopigno, Marco Callieri, Paolo Cignoni, Massimiliano Corsini, Matteo Dellepiane, Federico Ponchio, and Guido Ranzuglia. 2011. 3D models for cultural heritage: Beyond plain visualization. *Computer* 44, 7 (2011), 48–55.
- Maria Luiza Seixas, Paulo Santos Assis, João Cura D’Ars Figueiredo, Maria Aparecida Pinto, and Daniella Gualberto Caldeira Paula. 2018. The use of rapid prototyping in the joining of fractured historical silver object. *Rapid Prototyping Journal* 24 (2018), 532–538.
- Abhishek Sharma, Oliver Grau, and Mario Fritz. 2016. Vconv-dae: Deep volumetric shape learning without object labels. In *ECCV*. Springer, Berlin, Germany, 236–250.
- Vincent Sitzmann, Eric R Chan, Richard Tucker, Noah Snavely, and Gordon Wetstein. 2020. Metasdf: Meta-learning signed distance functions. *arXiv preprint arXiv:2006.09662* 1, 1 (2020), 1–17.
- Edward J Smith and David Meger. 2017. Improved adversarial systems for 3d object generation and reconstruction. In *Conference on Robot Learning*. PMLR, Cambridge, UK, 87–96.
- Hyeontae Son and Young Min Kim. 2020. SAUM: Symmetry-Aware Upsampling Module for Consistent Point Cloud Completion. In *Proc. ACCV*. Springer, Berlin, Germany, 1–17.
- Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. 2017. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 1746–1754.
- David Stutz and Andreas Geiger. 2020. Learning 3d shape completion under weak supervision. *International Journal of Computer Vision* 128, 5 (2020), 1162–1181.
- Raphael Sulzer, Loic Landrieu, Alexandre Boulch, Renaud Marlet, and Bruno Vallet. 2022. Deep Surface Reconstruction from Point Clouds with Visibility Information. *arXiv preprint arXiv:2202.01810* 1, 1 (2022), 1–13.
- Jiapeng Tang, Jiabao Lei, Dan Xu, Feiying Ma, Kui Jia, and Lei Zhang. 2021. Sign-Agnostic CONet: Learning Implicit Surface Reconstructions by Sign-Agnostic Optimization of Convolutional Occupancy Networks. *arXiv preprint arXiv:2105.03582* 1, 1 (2021), 1–16.
- Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Carsten Stoll, and Christian Theobalt. 2020. PatchNets: Patch-based generalizable deep implicit 3D shape representations. In *Proc. ECCV*. Springer, Berlin, Germany, 293–309.
- Rahul Venkatesh, Sarthak Sharma, Aurobrata Ghosh, Laszlo Jeni, and Maneesh Singh. 2020. DUDE: Deep Unsigned Distance Embeddings for Hi-Fidelity Representation of Complex 3D Surfaces. *arXiv preprint arXiv:2011.02570* 1, 1 (2020), 1–9.
- Lukasz Witek, Kimberly S Khouri, Paulo G Coelho, and Roberto L Flores. 2016. Patient-specific 3D models for autogenous ear reconstruction. *Plastic and Reconstructive Surgery—Global Open* 4, 10 (2016), e1093.
- Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. 2016. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Proc. NeurIPS*. Neural Information Processing Systems, San Diego, CA, 82–90.
- Yifan Xu, Tianqi Fan, Yi Yuan, and Gurprit Singh. 2020. Ladybird: Quasi-monte carlo sampling for deep implicit field based 3d reconstruction with symmetry. In *European Conference on Computer Vision*. Springer, Berlin, Germany, 248–263.
- Siming Yan, Zhenpei Yang, Haoxiang Li, Li Guan, Hao Kang, Gang Hua, and Qixing Huang. 2022b. Implicit Autoencoder for Point Cloud Self-supervised Representation Learning. *arXiv preprint arXiv:2201.00785* 1, 1 (2022), 1–24.
- Xingguang Yan, Liqiang Lin, Niloy J Mitra, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2022a. Shapeformer: Transformer-based shape completion via sparse representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 6239–6249.
- Mingyue Yang, Yuxin Wen, Weikai Chen, Yongwei Chen, and Kui Jia. 2021. Deep optimized priors for 3d shape modeling and reconstruction. In *Proc. CVPR*. IEEE, Piscataway, NJ, 3269–3278.
- Li Yi, Boqing Gong, and Thomas Funkhouser. 2021. Complete & label: A domain adaptation approach to semantic segmentation of lidar point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 15363–15373.
- Qian Yu, Chengzhuan Yang, and Hui Wei. 2022. Part-Wise AtlasNet for 3D point cloud reconstruction from a single image. *Knowledge-Based Systems* 242 (2022), 108395.
- Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. 2018. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*. IEEE, Piscataway, NJ, 728–737.
- Jiahui Zhang, Hao Zhao, Anbang Yao, Yurong Chen, Li Zhang, and Hongen Liao. 2018. Efficient semantic scene completion network with spatial group convolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, Berlin, Germany, 733–749.
- Jianming Zheng, Zhongguo Li, and Xi Chen. 2006. Worn area modeling for automating the repair of turbine blades. *The International Journal of Advanced Manufacturing Technology* 29, 9 (2006), 1062–1067.
- Zerong Zheng, Tao Yu, Qionghai Dai, and Yebin Liu. 2021. Deep implicit templates for 3D shape representation. In *Proc. CVPR*. IEEE, Piscataway, NJ, 1429–1439.