

Mending Neural Implicit Modeling for 3D Vehicle Reconstruction in the Wild

Shivam Duggal^{*1} Zihao Wang^{‡*1} Wei-Chiu Ma^{1,3} Sivabalan Manivasagam^{1,2}

Justin Liang¹ Shenlong Wang^{1,2} Raquel Urtasun^{1,2}

¹Uber ATG, ²University of Toronto, ³Massachusetts Institute of Technology

Abstract

Reconstructing high-quality 3D objects from sparse, partial observations from a single view is of crucial importance for various applications in computer vision, robotics, and graphics. While recent neural implicit modeling methods show promising results on synthetic or dense data, they perform poorly on sparse and noisy real-world data. We discover that the limitations of a popular neural implicit model are due to lack of robust shape priors and lack of proper regularization. In this work, we demonstrate high-quality in-the-wild shape reconstruction using: (i) a deep encoder as a robust-initializer of the shape latent-code; (ii) regularized test-time optimization of the latent-code; (iii) a deep discriminator as a learned high-dimensional shape prior; (iv) a novel curriculum learning strategy that allows the model to learn shape priors on synthetic data and smoothly transfer them to sparse real world data. Our approach better captures the global structure, performs well on occluded and sparse observations, and registers well with the ground-truth shape. We demonstrate superior performance over state-of-the-art 3D object reconstruction methods on two real-world datasets.

1. Introduction

Consider the street view image and the partial LiDAR scan in Fig. 1. As humans, we can effortlessly identify the complete vehicle in the scene and have a rough grasp of its 3D geometry. This is because human visual systems have accumulated hours of observations that help us develop mental models for these objects [24, 23]. While we may have never seen this particular car before, we know that cars shall be symmetric, they shall lie on the ground, and sedans shall have similar shapes and sizes. We can thus exploit these knowledge to infer the 3D structure of any object, given its sparse observations. The goal of this work is to equip computational visual machines with similar capabilities. One recently emerging class of such works is neural implicit shape modeling [34, 30, 31, 44, 45]. Neural implicit shape modeling involves reasoning about the surface

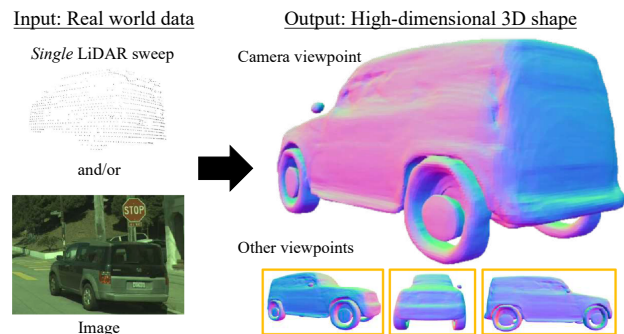


Figure 1: **3D Object Reconstruction in the Wild.** Our model takes as input a single LiDAR sweep and/or a RGB image, and outputs a high-quality 3D shape.

of a 3D shape as the level-set of a function (represented using a neural network) [42], for example: an inside-outside function or a signed-distance function. In this work, we exploit the neural implicit shape representation for reconstructing high-quality real-world objects.

Existing neural implicit reconstruction methods operate by learning a low-dimensional mapping for the corresponding high-dimensional 3D shape. Such a low-dimensional mapping is usually represented as a 1D (or multi-dimensional) latent-code, which is implicitly decoded back to the 3D shape. Previous approaches usually take one of two routes to generate the low-dimensional latent-code: (i) learn a direct mapping from observations to the latent-code, or (ii) model this problem as a structured optimization task and incorporate human knowledge into the model. Specifically, former methods capitalize on machine learning algorithms to directly learn the statistical priors from data. While they are more robust to noise in the observation, the decoded high-dimensional 3D shapes are not guaranteed to match with the input observations. Additionally, many feed-forward methods [30, 8] are learned with ground-truth (GT) supervision from synthetic data, which is often unavailable for real-world data. This leads to poor generalization performance on unseen real-world partial inputs. In contrast, optimization based approaches [34, 28] can produce coherent shape representations by incorporating the observations into carefully designed objective functions. In practice, however, designing the optimization objective function is cumbersome and hand-crafted priors alone can-

*Equal Contribution.

‡Work done as part of internship.




Figure 2: **Challenges of DeepSDF.** (i) *Instability*: while the input observation remains intact, different initializations may lead to distinct 3D reconstructions. See the red arrows. (ii) *Low-fidelity*: the model may easily overfit to the noise in the data, and fail to capture the global structure of the input observation. Here, the output shape is a hybrid of a van and a truck.

not include all possible phenomena. Despite hand-crafted regularization, optimization-based methods struggle with noise and can generate unusual artifacts. These limitations are more pronounced when reconstructing accurate and high-quality 3D shapes from real-world observations, because of potential noises, occlusions, and environment variability present in the real-world data.

In this work, we address the shortcomings of the prior works by highlighting the key components necessary for reconstructing high-quality 3D shapes in the wild. In order to reconstruct high-fidelity shapes from sparse and noisy real-world observations, it is necessary for the proposed approach to leverage (1) *strong shape priors* to extract the true 3D shape from input observations and be (2) *robust to the variable noise* in the input. To satisfy these properties, we begin with using the deep encoder network (adopted from feed-forward approaches [30, 8]) as a *robust initializer of the shape latent-code*. To guarantee that the predicted latent-code is faithful to the input observations, we further incorporate the *test-time optimization of the latent-code* [34] into our shape reconstruction method. However, we still have not addressed the hand-crafted objective function limitation. While it is necessary to carefully design an objective function for effective training and inference of the shape reconstruction system, the overall process is cumbersome and hand-designed. Moreover, test-time optimization is generally susceptible to noise in the observations. To overcome these limitations, we introduce a learned shape prior into our objective function, which serves as a *high-dimensional structural shape regularizer*. We do so by taking inspiration from the recent GAN inversion theory for 2D image generation [58, 19] and incorporate a 3D shape discriminator network into our approach. The discriminator judges the “naturalness” of the reconstructed shape both during training and test time. Finally, to effectively train our full model we introduce an *adversarial curriculum learning strategy*. This allows us to effectively learn several shape priors from synthetic data, and generalize them to real-world data.

We evaluated our approach for 3D vehicle reconstruction in the wild on two challenging self-driving datasets. We compare against state-of-the-art models on three tasks: LiDAR-based shape completion, image-based shape recon-

struction and image + LiDAR shape completion. Our results show superior performance in terms of both reconstruction accuracy and visual quality. Overall, we tackled the less-explored task of high-fidelity reconstruction of diverse real-world vehicles. We showcased significant improvement compared to prior works, thanks to the learned robust shape priors, the curriculum learning strategy used to extract them and the regularized optimization strategy.

2. Related Work

Feed Forward Shape Completion: Deep feed-forward approaches encode images or depth sensor data into latent representations, which are subsequently decoded into complete meshes or point clouds. Early works used voxel-grid based representations [11, 21] which had lossy resolution and memory constraints. Other works instead leveraged point clouds directly to create latent codes [36, 37, 18] for point cloud completion or mesh construction, using a variety of decoder architectures [55, 56, 27, 48, 51], enhanced intermediate representations [52, 17], and GAN-based loss functions [41]. While successful on synthetic datasets [5] or dense 3D scans [9], prior works [56, 52, 18] have had limited success on real world noisy datasets such as KITTI [15]. Other works [16, 10] directly perform mesh prediction from images and 3D data. Recent works [8, 35, 30, 39, 53, 40, 49, 33] represent shapes using neural implicit functions and have shown promising results on synthetic datasets [5, 1] or dense scans [4], but few have shown large-scale results on real data [18]. [46, 32] adapted feed-forward networks trained on synthetic data to real data by training an encoder that transforms partial point clouds to latent codes that match the observations. Concurrently, [2] showcased the importance of hierarchical priors (local and global shape priors) in generalizing to (unseen) real-world shapes. However, to our knowledge, most previous works applied to real sparse data generate shapes that are often amorphous, overly smooth, or restricted to small vehicles.

Optimization-based Shape Completion: Another line of work frames shape completion as an optimization problem, where the objective is to ensure the predicted shape is consistent with real sensor observations. Such works re-

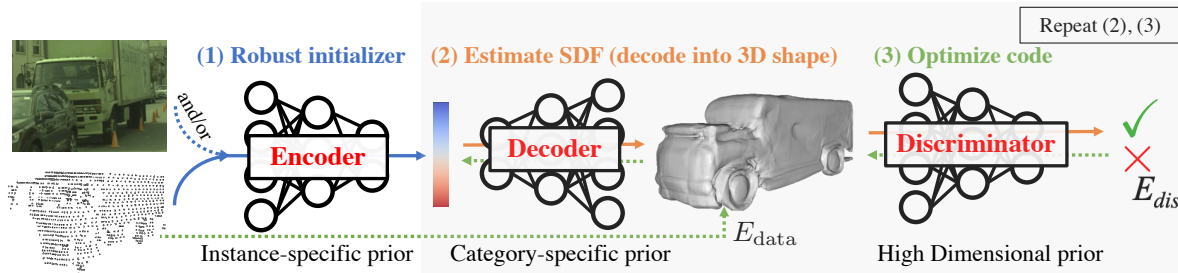


Figure 3: **Method Overview:** Given raw sensory data, our method first exploits a deep encoder to predict a robust initialization for the shape code. The latent shape-code is then optimized through the auto-decoder framework using a data-consistency energy term (E_{data}) and a discriminator-induced high-dimensional shape regularization term (E_{dis}).

quire strong shape priors, as the real-world observations are quite sparse, noisy, and have large holes. [13, 14, 50, 25] represent the shape prior as a PCA embedding of the volumetric signed distance fields, and optimize the shape and pose given sensor data. While showing promising performance on pose estimation and object tracking, the recovered shapes are coarse due to low dimensional linear embeddings, and generally represent sedans or smaller vehicles only. Recent works encode shape priors via neural implicit representations, and perform optimization supervised by input 3D observations [34, 12], with some using differentiable rendering [28, 22, 26]. While [34, 28, 22] optimized the low-dimensional latent-code, [26, 12, 54] directly optimized the neural network weights per object. Most of these works require dense supervision from the whole 3D space (GT occupancy or signed distance values) during training, preventing them from training on sparse real-world data. Thus, these works have focused only on synthetic shape completion. [57] have explored this approach for auto-labeling, but do not focus on shape quality.

Latent Space Optimization: Optimizing network weights or latent codes at test time has been quite popular for several tasks [3, 7, 58, 34]. CodeSLAM [3] optimizes geometric and motion codes during inference time for VisualSLAM, Chen *et al.* [7] optimizes geometric and appearance codes at inference time for pose estimation. Recently, latent code optimization has been explored in the field of GAN inversion [58, 43, 19] for 2D image generation. In this work, we take inspiration from the theory of GAN inversion and apply it to the task of high-fidelity 3D reconstruction. Another work related to ours is [20], which performs latent code optimization for point cloud completion and uses a GAN to ensure the optimized latent code is in the generator’s manifold. However, they demonstrated their approach on a controlled synthetic data environment and didn’t showcase any generalization to unseen data. Moreover, unlike our discriminator, their discriminator acts on a low-dimensional space. By directly passing the high-dimensional 3D shape to the discriminator, we allow the gradients to directly backpropagate from the

reconstructed shape to the latent-code and are thus able to guide the optimization to generate realistic shapes.

3. Background

The goal of 3D object shape reconstruction is to recover 3D geometry of an object given sensor observations (e.g: point cloud, images, etc). In this work, we parameterize a 3D shape implicitly using a signed-distance function, i.e. given a 3D point \mathbf{x} , we model the function $f(\mathbf{x}) = s$, where $|s|$ equals the distance of point \mathbf{x} from the surface, and $sign(s)$ represents inside / outside status of point \mathbf{x} w.r.t. the 3D object. We then represent shape reconstruction as:

$$f_{\theta}(\mathbf{x}, g_{\phi}(\mathbf{o})) = s,$$

where function g encodes the input observations \mathbf{o} into some latent representation. θ and ϕ are the parameters of the functions f and g respectively. To extract the final mesh, one can query the SDF values for a grid of points in the volume and then perform marching cubes [29].

Recent neural implicit modeling methods [34, 30, 8] generate high-fidelity shapes on synthetic or dense datasets. However, their shape quality degrades significantly on sparse and noisy real-world datasets. Towards our goal of reconstructing high quality shapes on real-world data, we first carefully analyze two standard approaches which prior works follow and highlight their limitations.

3.1. Feed-forward Approaches

Feed-forward networks [30, 8] formulate the above shape reconstruction task as an encoder-decoder task. The goal of the encoder (g_{ϕ}) is to extract discriminative latent cues, \mathbf{z} , from input observations (\mathbf{o}), and the role of the decoder (f_{θ}) is to map a 3D point (conditioned on the encoded latent features \mathbf{z}) to its corresponding signed-distance value.

During training, the encoder and decoder are jointly optimized with the GT signed distance field as supervision. The GT is basically a set of (\mathbf{x}, s) pairs, where the 3D points span over the whole volume (generally a normalized cube, enclosing the GT object). By generalizing the encoder over the training instances, feed-forward approaches make

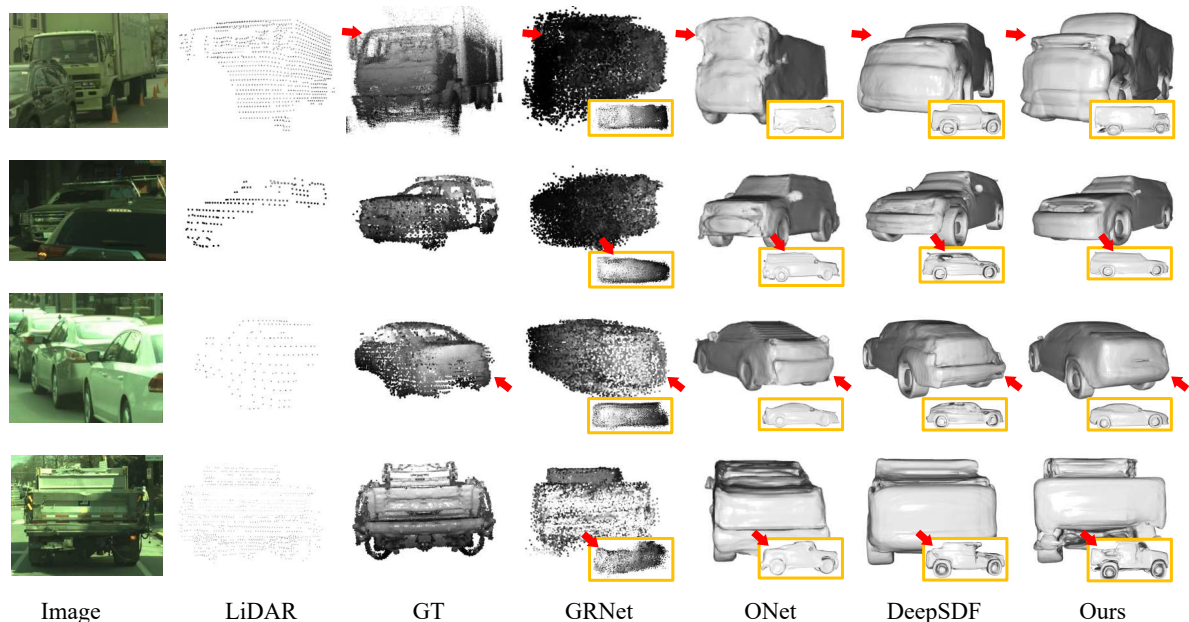


Figure 4: **LiDAR Completion on NorthAmerica:** Compared to baselines, our approach: (Row 1) captures the global structure, (Row 2) performs well in occluded scenarios, (Row 3) registers well with both input and GT points, (Row 4) maintains finer details (see cavity at the back of the car). Check supp. for additional comparisons (such as with Ours_{no-finetune}).

the encoder robust to input noise. However, this can cause over-smooth shapes at test-time, which are not realistic and do not register well with the observations (See ONet[30] results in Fig. 4). Low-quality shape predictions and artifacts are particularly pronounced for out-of-distribution inputs.

3.2. Neural Optimization Approaches

This line of neural implicit reconstruction works [34, 28, 12, 44, 46] focuses on combining the representation power of neural networks with structured optimization. One example is DeepSDF [34], which proposed an “auto-decoder” based optimization approach. Instead of inferring the latent code from observations ($\mathbf{z} = g_\phi(\mathbf{o})$), the auto-decoder framework directly optimizes the latent code, such that the predicted shape registers with the input observations.

During training, the DeepSDF model learns to both decode the spatial point \mathbf{x} into SDF value s (via the decoder f_θ) and assign a latent code \mathbf{z} to each training shape. The latent code learns to be a compact yet expressive representation of the 3D shape. Such a training procedure leads to the decoder emphasizing a *category-specific shape prior* (by generalizing over all training instances of a category) and the latent code encoding an *instance-specific shape prior*. During inference, the decoder is kept fixed, and the latent code (randomly initialized from a normal distribution) is optimized to minimize the following energy function:

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} E_{\text{data}}(\mathbf{o}, f_\theta(\mathbf{X}, \mathbf{z})) + \lambda E_{\text{reg}}(\mathbf{z}). \quad (1)$$

where \mathbf{X} typically represents a sparse set of input 3D points.

The data term E_{data} ensures consistency between the estimated 3D shape and the observation, while the regularization term E_{reg} constrains the latent code. The final shape can be obtained by querying the SDF value (via $f_\theta(\mathbf{x}, \mathbf{z}^*)$).

In practice, for shape completion, DeepSDF requires GT (\mathbf{x}, s) pairs¹. The data term is implemented as clamped \mathcal{L}_1 distance between estimated SDF value $f_\theta(\mathbf{x}, \mathbf{z})$ and the corresponding GT SDF value, s . The regularization term equals ℓ_2 -norm of the latent code.

Limitations: (1) **Instability:** The auto-decoder framework can suffer from shape instability w.r.t. different latent code initializations. This is because the optimization landscape for Eq. 1 is highly non-convex. As shown in Fig. 2 (i), a minor perturbation in initialization may lead to completely different local minima and hence the final output. (2) **Low-quality:** The optimal code is not guaranteed to generate a high-quality and “natural” shape. This limitation (Fig. 2 (ii)) is due to lack of structural regularization in the objective. Specifically, the data term $E_{\text{data}}(\mathbf{o}, f_\theta(\mathbf{X}, \mathbf{z}))$ in Eq. 1 is typically decomposed into sum of independent terms per each individual point $\sum_{\mathbf{x} \in \mathbf{X}} E_{\text{data}}(\mathbf{o}, f_\theta(\mathbf{x}, \mathbf{z}))$. There is little constraint on the global SDF field — each data point operates individually. The model can thus easily overfit to noise it has never seen during training and the latent code may drop out of manifold during optimization, leading to artifacts. This issue is particularly severe for real-world scenarios, where observations are sparse and noisy.

¹During inference, DeepSDF not only uses on-surface partial point cloud (SDF=0), but also non-surface points with known GT SDF values.

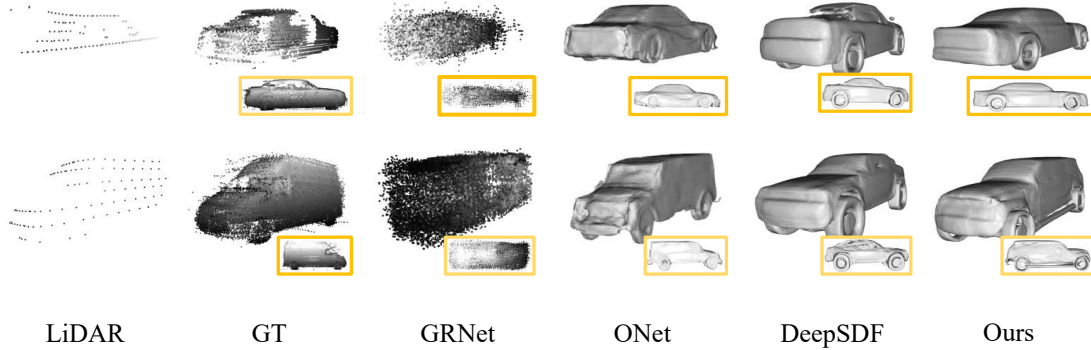


Figure 5: **LiDAR Completion on KITTI**: Compared to others, our approach maintains high fidelity with input points and generates finer shape details even with sparse input data. However, when input is significantly sparse, fidelity with GT drops.

4. Method

We now present our method for reconstructing accurate shapes from real-world observations. To address previous works’ limitations on real-world datasets, we combine components from both feed-forward and neural optimization approaches. From the feed-forward approaches, we leverage an encoder to robustly initialize a latent-code, given partial real world observations. Following neural-optimization, we utilize the auto-decoder framework to further optimize the latent-code, ensuring higher shape fidelity. However, the key challenge for shape reconstruction in the wild is that given (noisy, sparse) out-of-distribution observations, the reconstructed shapes no longer remain within the domain of realistic and naturally-looking shapes. To ensure that our reconstructed shapes always look realistic and maintain high-level properties of the shape category, we incorporate a high-dimensional shape regularizer in form of a discriminator, at both training and test-time optimization.

Thus, our shape reconstruction pipeline consists of encoder, decoder and discriminator as shown in Fig. 3. Finally, to ensure proper learning of all components, we propose a novel adversarial curriculum learning strategy. We now describe in detail our discriminator (Sec. 4.1), modified test-time optimization procedure (Sec. 4.2), and curriculum learning strategy (Sec. 4.3).

4.1. Discriminator as Shape Prior

Based on the observation that neural optimization based approaches lack a strong global shape prior, we introduce a discriminator to induce a learned prior over the predicted SDF. The discriminator function D_ψ (parameterized by ψ) evaluates the likelihood that the reconstructed shape is “natural” relative to a set of ground-truth synthetic shapes (e.g., ShapeNet). The benefits of a discriminator are two-fold: (1) During training, it serves as an objective function to improve the encoder through adversarial loss. (2) During inference, it regularizes the optimization process to ensure the latent code remains in the decoder’s domain. To capture the

overall geometry of the SDF, we randomly sample N points from the whole 3D space, query their SDF value, and pass all of them into the discriminator, which outputs a single scalar representing the “naturalness” of the shape.

4.2. Regularized Optimization for Real-World Inference

At inference time, given sparse and noisy real-world observations (\mathbf{o}), an encoder predicts a robust initialization for the shape latent code ($\mathbf{z} = g_\phi(\mathbf{o})$). Then, to ensure that the predicted shape aligns well with the observations, we optimize the latent code using the following objective function:

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} E_{\text{data}}(\mathbf{o}, f_\theta(\mathbf{X}, \mathbf{z})) + \lambda_{\text{reg}} E_{\text{reg}}(\mathbf{z}) + \lambda_{\text{dis}} E_{\text{dis}}(f_\theta(\mathbf{X}, \mathbf{z})). \quad (2)$$

Our data term, E_{data} , and latent-code regularization term, $E_{\text{reg}}(\mathbf{z})$, have the same format as DeepSDF (Eq. 1):

$$E_{\text{data}}(\mathbf{o}, f_\theta(\mathbf{X}, \mathbf{z})) = \sum_{\mathbf{x} \in X} \rho(\mathbf{s}, f_\theta(\mathbf{x}, \mathbf{z}))$$

where \mathbf{s} is the actual signed distance value at point \mathbf{x} , computed from the observation \mathbf{o} . ρ is a robust clamped \mathcal{L}_1 distance. The regularization term is the ℓ_2 -norm of the latent code (i.e., $E_{\text{reg}}(\mathbf{z}) = \|\mathbf{z}\|_2^2$). The additional discriminator prior term, $E_{\text{dis}}(f_\theta(\mathbf{X}, \mathbf{z})) = -\log D_\psi(\mathbf{X}, f_\theta(\mathbf{X}, \mathbf{z}))$, encodes the belief of the discriminator that the generated shape is natural. The 3D points (\mathbf{X}) used for E_{dis} are sampled from the whole 3D space, while the 3D points for E_{data} come directly from the observations.

4.3. Adversarial Curriculum Learning Strategy

Training neural implicit models require large amounts of dense 3D ground-truth data. While synthetic data provides this dense supervision and helps the model gain a rich prior over shapes, real-world in-the-wild data lacks dense GT. To effectively transfer the model trained on synthetic data to real-world data, its important to ensure that each component encodes complementary rich shape priors. Simply

Method	ACD (mm) ↓	Recall (%) ↑
ONet [30]	22.76	49.56
GRNet [52]	12.70	77.59
SAMP [14]	176.42	65.58
DIST [28]	19.55	71.54
DIST++ [28]	17.29	72.50
DeepSDF [34]	8.34	84.71
Our _{S^{no-finetune}}	7.02	86.48
Our _{S^{finetune}}	5.93	88.18

Table 1: **LiDAR Completion Results on NorthAmerica.**

jointly training the full model on synthetic data and directly applying to real data results in poor performance. We therefore introduce a curriculum learning strategy that allows our model to generalize to real-world data. We split training into two stages: Stage 1 and Stage 2. In Stage 1, we train the DeepSDF-based auto-decoder framework on clean and dense synthetic data. Then, to ensure our approach can handle sparse inputs at test time, in Stage 2 we perform adversarial training of the encoder using sparse synthetic observations. This additional Stage 2 allows using (dense and clean) synthetic GT data as strong supervision, while training the encoder to encode sparse observations.

Stage 1: We first train the decoder over synthetic training data, in the same manner as DeepSDF. Specifically, we jointly optimize the decoder weights and latent code to reconstruct the ground truth (GT) signed distances:

$$\min_{\{\mathbf{z}_i\}, \theta} \mathcal{L}^{\text{dec}} + \mathcal{L}^{\text{reg}}$$

where $\mathcal{L}^{\text{dec}} = \min_{\{\mathbf{z}_i\}, \theta} \sum_i^M \sum_j^N \rho(\mathbf{s}_{i,j}, f_\theta(\mathbf{x}_{i,j}, \mathbf{z}_i))$ and $\mathcal{L}^{\text{reg}} = \lambda_{\text{reg}} \|\mathbf{z}_i\|_2^2$. M is the total number of training shapes and N is the number of training point samples per shape. $\mathbf{s}_{i,j}$ is GT signed distance value for a training point sample $\mathbf{x}_{i,j}$. ρ refers to clamped \mathcal{L}_1 distance. Once the decoder (f_θ) is trained, we keep it fixed, preserving its ability to generate “natural” shapes. The optimized latent codes (\mathbf{z}') and the predicted signed distance fields ($\mathbf{s}' = f_\theta(\mathbf{x}, \mathbf{z}')$) are used as pseudo-ground-truth (pseudo-GT) in the next stage.

Stage 2: Next, given a synthetic *partial* point cloud or image observation, we adversarially train an encoder to predict a shape latent code ($\mathbf{z} = g_\phi(\mathbf{o})$) in the pre-trained decoder’s domain. For each synthetic training shape, we want to ensure that the predicted signed distance values match well with GT values (\mathcal{L}^{dec}), and that the encoder-predicted latent code matches the pseudo-GT latent code from Stage 1 (\mathcal{L}^z). In addition, we simultaneously train a discriminator to ensure the predicted signed-distance field matches the “naturalness” of the pseudo-GT signed distance field from Stage 1 (\mathcal{L}^{gan})². The overall loss function for training is:

²We observed that using Stage 1 predicted SDF (\mathbf{s}') instead of GT SDF as real samples for GAN loss results in better training and better results.

Method	ACD (mm) ↓	Recall (%) ↑
ONet [30]	18.10	59.10
GRNet [52]	13.66	78.21
DeepSDF [34]	14.81	79.76
Our _{S^{no-finetune}}	8.60	82.97

Table 2: **LiDAR Completion Results on KITTI.**

$$\mathcal{L} = \mathcal{L}^{\text{dec}} + \mathcal{L}^z + \mathcal{L}^{\text{gan}}$$

\mathcal{L}^{dec} measures the clamped \mathcal{L}_1 distance between the estimated and GT SDF values for each training point sample (same as \mathcal{L}^{dec} loss in Stage 1). \mathcal{L}^z measures the \mathcal{L}_2 distance between the encoder-predicted code and the pseudo-GT latent code ($\mathcal{L}^z = \|\mathbf{z} - \mathbf{z}'\|$). Additionally, we also enforce the encoder to embed the input sparse data into a realistic shape’s latent code (using $-\sum_{\mathbf{x} \in \mathbf{X}} \log(D_\psi(\mathbf{x}, f_\theta(\mathbf{x}, \mathbf{z})))$ as loss). \mathcal{L}^{gan} is the GAN loss between predicted SDF field and pseudo-GT SDF field for training the discriminator:

$$\mathcal{L}^{\text{gan}} = \sum_{\mathbf{x} \in \mathbf{X}} \log(D_\psi(\mathbf{x}, f_\theta(\mathbf{x}, \mathbf{z}))) + \log(1 - D_\psi(\mathbf{x}, \mathbf{s}'))$$

Such a training procedure disentangles the learning of the encoder, decoder and discriminator module. Through stage 1, the decoder learns to induce *category-specific shape prior*, by generalizing over all instances of the training set. By keeping the decoder fixed and training the encoder to generate an instance specific latent code, the encoder learns to induce *instance-specific shape prior*, while the discriminator acts as a *high-dimensional structural shape prior*.

5. Experimental Details

5.1. Datasets

ShapeNet: We exploit 2364 watertight cars from ShapeNet [5] as our synthetic training dataset. We follow DeepSDF [34] to generate the GT SDF samples for the first stage of training. For stage 2, we simulate image and sparse point clouds (as input observations) for each object from five different viewpoints, resulting in a total dataset of 11820 (2364×5) image-point cloud pairs.

NorthAmerica: We further build a novel, large-scale 3D vehicle reconstruction dataset using a self-driving platform that collects data over multiple metropolitan cities in North America and under various weather conditions and time of the day. Using the acquired data, we generate a test-set of 935 high-quality instances. Each consists of aggregated LiDAR sweeps used as GT shape, partial single LiDAR sweep and cropped camera image used as input observations.

KITTI: We also generate 209 diverse test-set objects from 21 sequences of the KITTI tracking dataset [15]. For each object, we construct a dense ground truth 3D shape by aggregating multiple LiDAR sweeps using the GT bounding boxes. KITTI dataset is used only during test time for evaluation. Please refer to supp. for more dataset details.

Method	ACD (mm) ↓	Recall (%) ↑
DIST [28]	62.97	48.82
Ours	8.89	84.32

Table 3: **Single Image Reconstruction on NorthAmerica**

5.2. Metrics

Unlike their synthetic counterpart, real-world datasets do not possess complete watertight 3D shapes. We cannot evaluate 3D reconstruction metrics like Chamfer Distance, Volumetric IoU [38], normal consistency [38], and robust F-score [47] on them. We thus adopt asymmetric Chamfer Distance (ACD) between the GT point cloud and the reconstructed shape to measure the shape fidelity. ACD is defined as the sum of squared distance of each ground truth 3D point to the closest surface point on the reconstructed shape:

$$\text{ACD}(\mathbf{X}, \mathbf{Y}) = \frac{1}{|\mathbf{X}|} \sum_{x \in \mathbf{X}} \min_{y \in \mathbf{Y}} \|x - y\|_2.$$

We also compute the recall of the ground truth points from the reconstructed shape as a robust alternative:

$$\text{Recall}(\mathbf{X}, \mathbf{Y}) = \frac{1}{|\mathbf{X}|} \sum_{x \in \mathbf{X}} \left[\min_{y \in \mathbf{Y}} \|x - y\|_2 \leq t \right].$$

We set true-positive threshold $t = 0.1$ m in the paper. All the reported metrics are averaged over the test set objects.

5.3. Baselines

We compare against several state-of-the-art 3D object reconstruction algorithms: (i) feed-forward methods, such as Occupancy Networks (ONet) [30] and GRNet [52]; (ii) (deep) optimization based methods, such as linear SAMP [14], DeepSDF [34], and DIST [28]. DeepSDF samples off-surface points for SDF optimization to improve robustness. We also augment the original DIST approach with such sampling procedure, referred to as DIST++.

6. Experimental Results

We now showcase our experimental results. We first compare our proposed approach against the baselines on the LiDAR-based 3D shape completion task (Sec. 6.1). Next, we extend our approach to other sensor modalities (images (Sec. 6.2) and image + LiDAR combined (Sec. 6.3)). Finally, we perform an ablation of our approach (Sec. 6.4).

6.1. LiDAR-based Shape Completion

LiDAR-based shape completion refers to the task of reconstructing a 3D shape, given its partial point-cloud as the input observation (\mathbf{o}). We report LiDAR-based shape reconstruction results on NorthAmerica and KITTI datasets in Tab. 1 and Tab. 2 respectively. “Ours_{no-finetune}” result was generated by solely training the proposed approach on the

Method	ACD (mm) ↓	Recall (%) ↑
DIST [28]	23.40	71.99
DIST++ [28]	17.52	72.65
Ours	5.36	89.05

Table 4: **Image+LiDAR Completion on NorthAmerica.**

ShapeNet dataset (as mentioned in Sec. 4.3), followed by inference-stage optimization. Ours_{no-finetune} approach significantly outperforms all the baselines. In particular, we reduce ACD error by 45% (37%) compared to the best feed-forward network GRNet and 16% (42%) compared to the best optimization method DeepSDF on NorthAmerica (KITTI) dataset. In contrast to NorthAmerica, GRNet performs better than optimization-based DeepSDF on KITTI. As KITTI point clouds are noisier, this demonstrates the value of feed-forward approaches being robust. We note that SAMP has particularly large error on NorthAmerica dataset due to its difficulty in handling larger objects of the dataset (such as trucks), as the embedding was trained mostly with smaller synthetic vehicles. Overall, the improvements suggest the effectiveness of our approach compared to feed-forward or optimization-only methods.

Fine-tuning on sparse NorthAmerica dataset: These improvements demonstrated that the shape regularization provided by the discriminator makes the encoder and decoder more robust for real-world shape reconstruction. Encouraged by these results, we explored whether we could further improve performance by fine-tuning our model’s encoder on sparse real-world data. Specifically, we supervise the encoder to generate a latent code such that the decoder-predicted SDF values for the input observation points (sparse on-surface LiDAR points) are 0. We also ensure that the shape latent-code generated by the encoder ($g_\phi(\mathbf{o}) = \mathbf{z}$) remains close to the shape latent-code generated by the pre-trained encoder without any fine-tuning ($\mathbf{z}_{\text{iter-0}}$). This serves as a regularization loss ($\mathcal{L}_{\text{finetune}}^z = \|\mathbf{z} - \mathbf{z}_{\text{iter-0}}\|$), preventing the encoder from generating an out-of-distribution latent-code. We also regularize the reconstructed shape using discriminator-guided $\mathcal{L}^{\text{dis}}(-\log D_\psi(\mathbf{X}, f_\theta(\mathbf{X}, \mathbf{z})))$ regularization term.

Here we demonstrate shape completion results (“Ours_{finetune}”) obtained by fine-tuning our model’s encoder on sparse NorthAmerica dataset³. From Tab. 1, we see that fine-tuning reduces the ACD by 18% (Ours_{finetune} vs Ours_{no-finetune}). This showcase the effectiveness of our learned priors which allow us to fine-tune our model even using sparse on-surface LiDAR points as weak supervision. On the other hand, prior (auto-decoder based) optimization methods like DeepSDF [34] and DIST [28] cannot be trained or fine-tuned on real data easily. This is because

³For fine-tuning/ training on NorthAmerica, we generate 3100 instances, which are completely different from NorthAmerica test samples.

Syn. Training Stage			Test-time Opt. Stage			ACD ↓ (mm)	Recall ↑ (%)
Dec.	Enc.	Disc.	Dec.	Enc.	Disc.		
✓			✓			8.34	84.71
✓	✓		✓	✓		7.30	86.21
✓	✓	✓	✓	✓		6.96	86.55
✓	✓	✓	✓		✓	8.26	84.71
✓	✓	✓	✓	✓	✓	7.02	86.48

Table 5: Ablation study on encoder and discriminator.

(1) real-world data lacks accurate dense supervision, and (2) DeepSDF and DIST do not have a pre-trained representation of the latent code for real-world objects.

Qualitative Results: Fig 4 compares Ours_{finetune} reconstruction results with the prior works on NorthAmerica dataset. GRNet generates non-watertight shapes and fails to recover fine details. ONet produces overly-smooth shapes at times and doesn’t have high-fidelity to the observations. DeepSDF maintains high-quality local details for visible regions, but doesn’t predict the correct shape. Our approach produces results that are both visually appealing and have high-fidelity with the observations. Fig 5 compares Ours_{no-finetune} results with the prior works on KITTI, where we observe similar trends. However, when input is very sparse, the reconstruction fails to match with the GT shape.

6.2. Image-based 3D Reconstruction

Our proposed pipeline can be easily extended to reconstruct shapes using other sensor modalities. Here, instead of encoding point-clouds, we demonstrate 3D reconstruction results by encoding a single-view image observation on the NorthAmerica dataset. Since the synthetic dataset we used does not have realistic texture maps for the corresponding 3D CAD models, we directly train the image encoder and discriminator on the real-world dataset. The training procedure is the same as stage 2 of the proposed curriculum learning technique (Sec. 4.3), but using an image encoder and using the real world GT shapes (same as those used for lidar-based fine-tuning) as supervision. Please check supp. for detailed training procedure. We compare our image-reconstruction results with state-of-the-art DIST[28] approach on the NorthAmerica dataset in Tab. 3. Compared to DIST, our image-based approach produces significantly better results (85% reduction in ACD), especially for occluded regions. Moreover, our image-only method (Tab. 3 last-row) is competitive with or better than all LiDAR completion baselines (Tab. 1 baselines).

6.3. Image + LiDAR Shape Completion

Our approach is amenable to combining sensor modalities for shape reconstruction without any re-training. We simply use the two pre-trained encoder modules (point-cloud encoder and image-encoder) and generate two shape latent-codes, $\mathbf{z}_{\text{img}}^{\text{init}}$ and $\mathbf{z}_{\text{LiDAR}}^{\text{init}}$. Then, inspired by GAN inver-

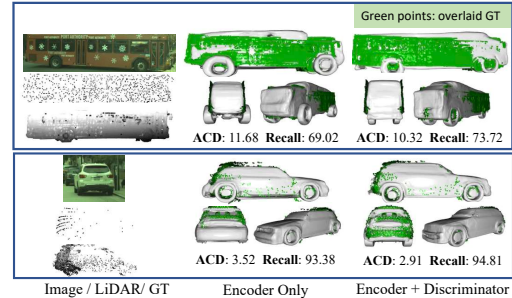


Figure 6: Analyzing the effectiveness of discriminator. The discriminator improves the fidelity of the reconstructed shapes, even without any real-world training.

sion [19] and photometric stereo [6], we fuse the two latent codes at the decoder features level. Please refer to supp. for more details. As shown in the Tab. 4, our proposed image + LiDAR shape completion technique performs significantly better than DIST and DIST++. Moreover, using both LiDAR and image observations reduces the ACD error by 10% and 40% compared to LiDAR and image only models.

6.4. Analysis

Ablation study on Encoder and Discriminator: We analyse the significance of the encoder (Enc.) and the discriminator (Disc.) in Tab. 5 and in Fig. 6. To showcase our “generalizability”, no real-world fine-tuning was performed. Adding both Enc. (row 1 vs. row 2) and Disc. (row 2 vs. row 3) to the reconstruction pipeline helps. Comparing row 3 with row 5 shows that adopting Disc. during optimization slightly affects quantitative performance, possibly due to noise in real-world aggregated LiDAR GT. However, row 1 of Fig. 6 demonstrates that Disc. doesn’t hinder generalization to unseen shapes, rather enhances the visual fidelity of the reconstructed shapes (eg: model was trained on ShapeNet cars and tested on NorthAmerica bus category).

7. Conclusion

In this paper, we present a simple yet effective solution for 3D object reconstruction in the wild. Unlike previous approaches that suffer from sparse, partial and potentially noisy observations, our method can generate accurate, high-quality 3D shapes from real-world sensory data. The key to success is to fix the inherent limitations of existing methods: lack of robust shape priors and lack of proper regularization. To address these limitations, we highlight the importance of (i) deep encoder as a robust initializer of the shape latent-code; (ii) regularized test-time optimization of the latent-code; (iii) learned high-dimensional shape regularization. Our proposed curriculum training strategy allows us to effectively learn disentangled shape priors. We significantly improve the shape reconstruction quality and achieve state-of-the-art performance on two real-world datasets.

References

- [1] Renderpeople, 2018.
- [2] Jan Bechtold, Tatarchenko Maxim, Fischer Volker, and Brox Thomas. Fostering generalization in single-view 3d reconstruction by learning a hierarchy of local and global shape priors. In *CVPR*, 2021.
- [3] Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J. Davison. Codeslam - learning a compact, optimisable representation for dense visual slam. In *CVPR*, 2018.
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments, 2017.
- [5] 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. Shapenet: An information-rich 3d model repository, 2015.
- [6] Guanying Chen, Kai Han, and Kwan-Yee K Wong. Ps-fcn: A flexible learning framework for photometric stereo. In *ECCV*, 2018.
- [7] Xu Chen, Zijian Dong, Jie Song, Andreas Geiger, and Otmar Hilliges. Category level object pose estimation via neural analysis-by-synthesis, 2020.
- [8] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019.
- [9] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017.
- [10] Angela Dai and Matthias Nießner. Scan2mesh: From unstructured range scans to 3d meshes. In *CVPR*, 2019.
- [11] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *CVPR*, 2017.
- [12] Thomas Davies, Derek Nowrouzezahrai, and Alec Jacobson. On the effectiveness of weight-encoded neural implicit 3d shapes, 2021.
- [13] Francis Engelmann, Jörg Stückler, and Bastian Leibe. Joint Object Pose Estimation and Shape Reconstruction in Urban Street Scenes Using 3D Shape Priors. In *GCP*, 2016.
- [14] Francis Engelmann, Jorg Stuckler, and Bastian Leibe. Samp: Shape and motion priors for 4d vehicle reconstruction. *WACV*, 2017.
- [15] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [16] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *ICCV*, 2019.
- [17] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. Atlasnet: A papier-mâché approach to learning 3d surface generation. In *CVPR*, 2018.
- [18] Jiayuan Gu, Wei-Chiu Ma, Sivabalan Manivasagam, Wenyuan Zeng, Zihao Wang, Yuwen Xiong, Hao Su, and Raquel Urtasun. Weakly-supervised 3d shape completion in the wild. In *ECCV*, 2020.
- [19] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *CVPR*, 2020.
- [20] Swaminathan Gurumurthy and Shubham Agrawal. High fidelity semantic shape completion for point clouds using latent optimization. In *WACV*, 2019.
- [21] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *ICCV*, 2017.
- [22] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *CVPR*, 2020.
- [23] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018.
- [24] Abhishek Kar, Shubham Tulsiani, Joao Carreira, and Jitendra Malik. Category-specific object reconstruction from a single image. In *CVPR*, 2015.
- [25] Abhijit Kundu, Yin Li, and James M. Rehg. 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. In *CVPR*, 2018.
- [26] Chen-Hsuan Lin, Chaoyang Wang, and Simon Lucey. Sdfsrm: Learning signed distance 3d object reconstruction from static images. In *NeurIPS*, 2020.
- [27] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion. In *AAAI*, 2020.
- [28] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In *CVPR*, 2020.
- [29] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, 1987.
- [30] Lars M. Mescheder, Michael Oechsle, M. Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019.
- [31] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [32] Mahyar Najibi, Guangda Lai, Abhijit Kundu, Zhichao Lu, Vivek Rathod, Thomas Funkhouser, Caroline Pantofaru, David Ross, Larry S. Davis, and Alireza Fathi. Dops: Learning to detect 3d objects and predict their 3d shapes. In *CVPR*, 2020.
- [33] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020.
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019.

- [35] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, 2020.
- [36] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017.
- [38] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, 2017.
- [39] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019.
- [40] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020.
- [41] Muhammad Sarmad, Hyunjoon Jenny Lee, and Young Min Kim. RL-gan-net: A reinforcement learning agent controlled gan network for real-time point cloud shape completion. In *CVPR*, 2019.
- [42] Stan Sclaroff and Alex Pentland. Generalized implicit functions for computer graphics. In *SIGGRAPH*, 1991.
- [43] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *TPAMI*, 2020.
- [44] Vincent Sitzmann, Eric R. Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. In *NeurIPS*, 2020.
- [45] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *NeurIPS*, 2019.
- [46] David Stutz and Andreas Geiger. Learning 3d shape completion under weak supervision. *IJCV*, 2018.
- [47] Maxim Tatarchenko, Stephan R. Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, 2019.
- [48] Lyne P Tchapmi, Vineet Kosaraju, Hamid Rezaatofghi, Ian Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *CVPR*, 2019.
- [49] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d scene representation and rendering, 2020.
- [50] Rui Wang, Nan Yang, Joerg Stueckler, and Daniel Cremers. Directshape: Direct photometric alignment of shape priors for visual vehicle pose and shape estimation. In *ICRA*, 2020.
- [51] Xin Wen, Tianyang Li, Zhizhong Han, and Yu-Shen Liu. Point cloud completion by skip-attention network with hierarchical folding. In *CVPR*, 2020.
- [52] Haozhe Xie, Hongxun Yao, Shangchen Zhou, Jiageng Mao, Shengping Zhang, and Wenxiu Sun. Grnet: Gridding residual network for dense point cloud completion. In *ECCV*, 2020.
- [53] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *NeurIPS*, 2019.
- [54] Mingyue Yang, Yuxin Wen, Weikai Chen, Yongwei Chen, and Kui Jia. Deep optimized priors for 3d shape modeling and reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [55] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *CVPR*, 2018.
- [56] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *3DV*, 2018.
- [57] Sergey Zakharov, Wadim Kehl, Arjun Bhargava, and Adrien Gaidon. Autolabeling 3d objects with differentiable rendering of sdf shape priors. In *CVPR*, 2020.
- [58] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020.