

SA-ConvONet: Sign-Agnostic Optimization of Convolutional Occupancy Networks

Jiapeng Tang^{1,4}, Jiabao Lei¹, Dan Xu², Feiying Ma⁴, Kui Jia^{1,5,6}, and Lei Zhang^{3,4}

¹School of Electronic and Information Engineering, South China University of Technology

²Department of Computer Science and Engineering, HKUST, HK

³Department of Computing, The Hong Kong Polytechnic University, HK

⁴DAMO Academy, Alibaba Group

⁵Pazhou Lab, Guangzhou, China

⁶Peng Cheng Laboratory, Shenzhen, China

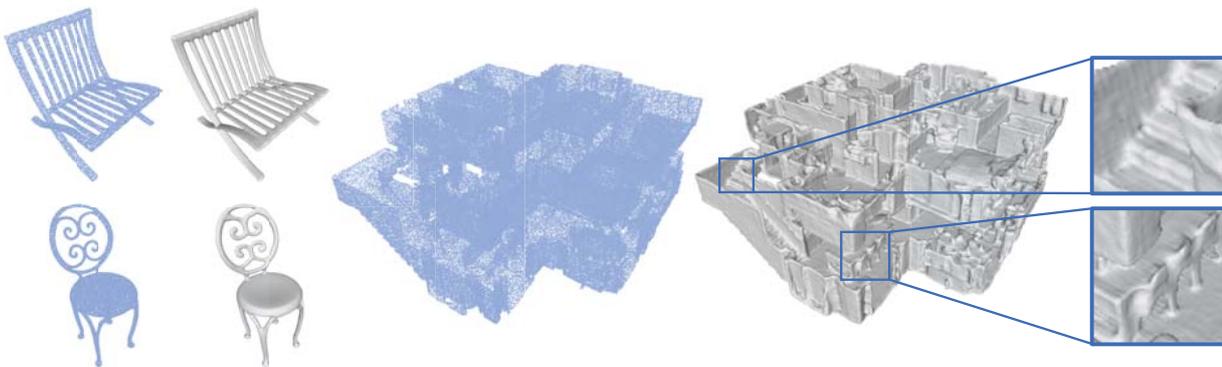


Figure 1: Given an un-oriented point cloud of a complex object (left) or a large-scale scene (right), our method can reconstruct an accurate surface mesh *without the use of oriented normals*.

Abstract

Surface reconstruction from point clouds is a fundamental problem in the computer vision and graphics community. Recent state-of-the-arts solve this problem by individually optimizing each local implicit field during inference. Without considering the geometric relationships between local fields, they typically require accurate normals to avoid the sign conflict problem in overlapped regions of local fields, which severely limits their applicability to raw scans where surface normals could be unavailable. Although SAL breaks this limitation via sign-agnostic learning, further works still need to explore how to extend this technique for local shape modeling. To this end, we propose to learn implicit surface reconstruction by sign-agnostic optimization of convolutional occupancy networks, to simultaneously achieve advanced scalability to large-scale scenes, generality to novel shapes, and applicability to raw scans in

a unified framework. Concretely, we achieve this goal by a simple yet effective design, which further optimizes the pre-trained occupancy prediction networks with an unsigned cross-entropy loss during inference. The learning of occupancy fields is conditioned on convolutional features from an hourglass network architecture. Extensive experimental comparisons with previous state-of-the-arts on both object-level and scene-level datasets demonstrate the superior accuracy of our approach for surface reconstruction from un-orientated point clouds. The code is available at <https://github.com/tangjiapeng/SA-ConvONet>.

1. Introduction

Surface reconstruction from point clouds is of significance to perceive and understand surrounding 3D worlds for intelligent systems, which plays a fundamental role in numerous practical applications, such as computer-aided design, 3D printing, and robotics grasping. Recently, this

Correspondence to Dan Xu and Kui Jia.

problem has attracted wide attention as inexpensive and portable commodity scanners such as the Microsoft Kinect make it much easier to acquire 3D point clouds. Classical methods [1, 5, 26, 24, 25] tackle this problem by mathematical optimization according to pre-defined geometric priors, while learning-based methods [17, 11, 28, 31] choose to learn geometric priors from large-scale 3D datasets in a data-driven manner. Recently, representing 3D surface as an implicit field has gained large popularity [10, 28, 30, 34, 46, 11, 22, 6, 14, 41, 31, 37]. Compared to other shape representations such as voxel [12, 44], octree [32, 39, 43, 20], point cloud [15] and mesh [17, 42, 23, 36, 29, 37], continuous implicit fields can enable surface reconstruction with infinite resolution and arbitrary topology.

A lot of methods have been proposed to advance the development of implicit surface reconstruction from various respects in terms of improving scalability, generality, and applicability. However, there is still not an approach in the literature to simultaneously achieve all these objectives with satisfactory performance. Targeting better scalability to large-scale scenes, several approaches [22, 6, 41, 31, 14] learn local implicit fields and model a global shape as a composition of local surface geometries, rather than conducting global shape reasoning from a latent code. Towards better generality to novel shapes, some works [30, 16, 22, 6, 41, 47] attempt to optimize the pre-trained priors at test time to obtain a better solution for each given input, instead of strictly respecting the learned priors. Existing state-of-the-art methods [22, 6, 41] improve both scalability and generality via individual optimization of each local implicit field during inference. However, without explicitly considering geometric relationships between local fields, they heavily rely on accurate normals to avoid the sign conflict problems in the overlapped regions of local fields. Although SAL [2] breaks this limitation via sign-agnostic learning that improves the applicability to real-world scans where surface normals are unavailable, it can only perform global shape modeling. Further works still need to explore how to extend this technique for local shape modeling.

To this end, we propose to learn implicit surface reconstructions by sign-agnostic optimization of convolutional occupancy networks [31], to simultaneously achieve the three important reconstruction objectives, *i.e.* advanced generality, specialty, and applicability in a unified framework. We achieve this goal by a simple yet effective solution that further optimizes the pre-trained occupancy prediction networks via sign-agnostic learning. The learning of occupancy fields is conditioned on convolutional features from an hourglass network (*e.g.* U-Net [33]). Our solution is motivated by two key characteristics. The first characteristic is that, after being pre-trained on the accessible datasets with ground-truth signed fields, the occupancy decoder can provide a signed field as initialization for the test-time opti-

mization. Thus we can further apply unsigned objectives to optimize occupancy prediction networks, maximizing the consistency between the desired iso-surface with the observed un-oriented point cloud. The second characteristic is that, the U-Net [33] aggregates both local and global information in an hourglass convolutional manner. The use of local shape features not only preserves the fine-grained geometries, but also enables the surface recovery of large-scale indoor scenes. The integrated global shape features can enforce geometric consistency between learned local geometries and guarantee the assembly of local fields as a globally consistent one, although we do not utilize guidance from additional normal information. As shown in Figure 1, we can reconstruct surfaces with fine details directly from un-oriented point clouds *without the use of normals*, for both complicated objects to large-scale scenes.

Extensive experimental comparisons with state-of-the-arts on both object-level and scene-level datasets, including ShapeNet [8], synthetic indoor scene dataset [31], and real-world scene datasets (ScanNet [13] and Matterport3D [7]) demonstrate the superior performance of our approach for surface reconstruction from un-oriented point clouds.

2. Related Work

In this section, we briefly review existing methods for surface reconstruction from raw point clouds. Specifically, we only review those implicit reconstruction methods that find a field function (occupancy function or signed distance function) to approximate the given point cloud.

Classic Optimization-based Surface Reconstruction Computing a continuous surface from its discrete approximation is a severely ill-posed problem, since there could be infinitely possible solutions. Classical methods [1, 5, 26, 24, 25] formulate this task as a mathematical optimization problem and try to solve this problem utilizing pre-defined geometric priors such as local linearity and smoothness. There have been a number of representative reconstruction methods such as Radius Basis Function (RBF) [5], Moving Least Square (MLS) [1], and Poisson Surface Reconstruction (PSR) [24, 25]. The RBF [5] represents surface as a linear combination of a series of radial basis functions; the MLS [1] fits observed points via finding those constituent spatially-varying polynomials; the PSR [24, 25] models the surface reconstruction as a Poisson’s equation.

Learning-based Surface Reconstruction More recently, driven by large-scale 3D datasets (*e.g.* ShapeNet [8]), neural networks have achieved notable successes in the field of implicit surface reconstruction from point clouds, ranging from global to local field modeling. The global models [28, 10] intrinsically perform shape retrieval in the latent space [40], leading to limited generality to represent unseen shapes and restricted power to capture complex details. These drawbacks can be resolved by the local mod-

els [11, 34, 46, 31, 14] that focus on local geometry modeling. Our method also adopts the manner of the local implicit field learning. Thus it possesses the capability of representing large-scale scenes. Another advantage is the better generality to unseen shapes, which can provide relatively good initialization of signed fields for the test-time optimization.

Combination of Data-Driven Priors and Optimization

The above-mentioned learning methods [11, 34, 46, 31] fix the learned priors during inference. Specifically, they directly obtain a 3D surface via a single feed-forward pass. As the pre-trained priors are fixed, they have difficulty in generalizing well to unseen shapes that are dissimilar to the training samples. Some existing approaches [30, 22, 6, 41, 45, 47, 38] try to combine the data-driven priors with optimization strategy at the test phase to acquire better results for each given input. Among them, the methods of local field optimization, including LIG [22], DeepLocalSDF [6], and PatchNet [41], can achieve state-of-the-art performance. However, they require *additional surface normals* to solve the sign flipping problem when assembling local fields into a globally consistent one, which seriously limits their applicability to raw scans that lack reliable and accurate surface normals. Our method also belongs to this line of research. However, in contrast to them, we optimize local implicit fields that are conditioned on the convolutional features learned in an hourglass manner. Since the global consistency between local fields can be effectively maintained in the process of hourglass convolutional feature learning, we can always guarantee the local field assemblies as a globally consistent one during optimization, although we do not have the guidance of surface normals.

Sign-Agnostic Surface Reconstruction The raw point clouds scanned by sensing devices usually lack oriented normals. Although we can approximate them via normal estimation methods [21, 18, 35, 19, 4], the normal estimation errors can cause degenerated surfaces. Thus it is more appealing to model surface directly from un-oriented points in a sign-agnostic manner [2, 3]. The SAL [2] avoids the use of surface normals by properly initializing the implicit decoder network, such that they can produce signed solutions of implicit functions only using unsigned objectives. Our key idea of sign-agnostic implicit field optimization is similar to SAL. With the assistance of auxiliary datasets with ground-truth signed implicit fields, the occupancy decoder can be trained to represent signed fields. Given signed implicit fields as initialization for the test-time optimization, we further adapt the pre-trained priors to the given input, by applying the unsigned cross-entropy loss to align the desired iso-surface with the observed un-oriented point cloud. A concurrent work of SAIL-S3 [48] extended the geometric initialization of SAL [2] for local signed field learning from un-oriented point clouds. However, it still requires a post-optimization stage to avoid the local sign flipping issue.

3. Approach

3.1. Overview

Given a set of observed points $\mathcal{P} = \{\mathbf{p}_i \in \mathbb{R}^3\}_{i=1}^n$, the goal of our method is to reconstruct a surface \mathcal{S} that is as similar as possible to the underlying surface $\hat{\mathcal{S}}$. We choose to approximate the signed implicit field representation $\hat{\mathbf{O}}$ of $\hat{\mathcal{S}}$ by predicting a neural implicit field \mathcal{O} due to its advantage of reconstructing surfaces with infinite resolution and unrestricted topology. Our goal is to simultaneously achieve advancements in all three respects, *i.e.* scalability to large-scale scenes, generality to novel shapes, and applicability to raw scans. Towards this goal, we propose a simple yet effective solution of learning implicit surface reconstructions by sign-agnostic optimization of convolutional occupancy networks.

The overall pipeline of the proposed approach is shown in Figure 2. Our approach consists of two stages, namely convolutional occupancy field pretraining and designed sign-agnostic, test-time implicit surface optimization. The former stage is responsible for learning the local shape priors with global consistency constraints, and provides relatively reasonable signed fields as initialization for the latter stage, which further optimizes the whole network using the unsigned cross-entropy loss to improve the accuracy of \mathbf{O} . We present the details of the proposed approach in the following two sections, *i.e.* Section 3.2 and Section 3.3.

3.2. Convolutional Occupancy Fields Pre-training

3.2.1 Convolutional Feature Learning

As shown in Figure 2, we first process the given input \mathcal{P} by a shallow PointNet [9] to obtain point-wise features. Then, we convert them to volumetric features with a dimension of $H \times W \times D$, by encapsulating local neighborhood information within a cell. Specifically, we integrate all point features belonging to the same voxel cell using the average pooling. To integrate both global and local information, we use a 3D-UNet to process \mathbf{V}_0 to obtain \mathbf{V} . Due to the issue of memory overhead of 3D-CNN, we set $H = W = D = 64$, and the depth of the 3D U-Net is set to 4 such that the receptive field is equal to the size of \mathbf{V}_0 . Due to the translation equivariance of convolution operations and rich shape features integrated by the hourglass network architecture, U-Net, we can enable scalable surface reconstruction for large-scale scenes.

3.2.2 Occupancy Field Predicting

Based on the obtained volumetric features \mathbf{V} , we can predict the occupancy probability $\mathbf{O}(\mathbf{q}) \in \mathbb{R}^3$ of a point \mathbf{q} randomly sampled in 3D space. To do this, we first perform trilinear interpolation to query the feature vector \mathbf{f}_q from \mathbf{V} according to the coordinate of \mathbf{q} , and then feed \mathbf{q} and

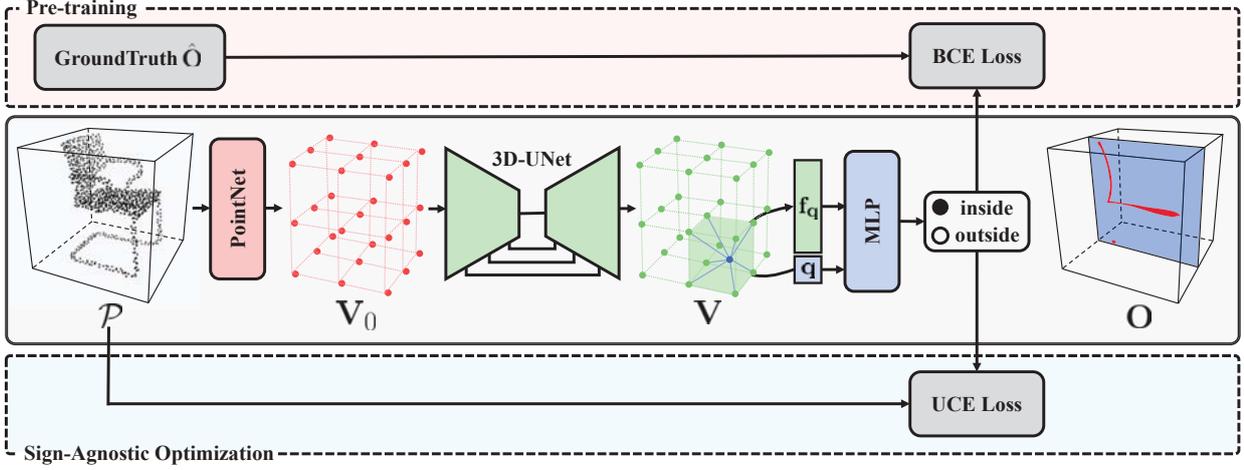


Figure 2: **Method Overview.** Our approach is built upon the convolutional occupancy networks (CONet) [31] (middle) that predicts an occupancy field \mathbf{O} based on convolutional features \mathbf{V} extracted from the input point cloud \mathcal{P} via a cascaded network of PointNet and 3D U-Net. We first pre-train the CONet [31] on the accessible datasets with ground-truth $\hat{\mathbf{O}}$ using the standard binary cross-entropy (BCE) loss (top). During inference, the proposed sign-agnostic optimization further fine-tunes the whole network parameters via unsigned cross-entropy (UCE) loss to improve the accuracy of \mathbf{O} (bottom).

\mathbf{f}_q into the occupancy decoder g that is implemented as a light-weight network of multi-layer perceptron (MLP):

$$\mathbf{O}(\mathbf{q}) = \text{sigmoid}(g(\mathbf{q}, \mathbf{f}_q)) \in (0, 1), \quad (1)$$

where the occupancy probability of \mathbf{q} is the sigmoid activation of final output logit $g(\mathbf{q}, \mathbf{f}_q)$.

3.2.3 Loss Function

During training, we uniformly sample some points \mathcal{Q} within the bounding volume of watertight mesh and compute their ground truth of occupancy values. And we punish the discrepancy between the predicted and the true occupancy values by a loss function written as:

$$\mathcal{L}(\mathbf{O}, \hat{\mathbf{O}}) = \sum_{\mathbf{q} \in \mathcal{Q}} \text{BCE}(\mathbf{O}(\mathbf{q}), \hat{\mathbf{O}}(\mathbf{q})), \quad (2)$$

where $\text{BCE}(x, y) = -y \log x - (1 - y) \log(1 - x)$ denotes the standard binary cross-entropy.

3.3. Sign-Agnostic Implicit Surface Optimization

In the inference stage, we can directly produce the implicit field through a single feed-forward pass. But we may not get satisfactory results if the given inputs are out of pre-trained priors. In order to improve the generality to unseen shapes, one can further optimize the pre-trained model for the given input. But we cannot apply the loss function in Equation 2 to supervise the network finetuning, because surface normals associated with observed points are not available, which causes the unavailability of in-out fields. Although we can choose to estimate the normals from \mathcal{P} , the

normal estimation errors would increase the difficulty of recovering clean surfaces.

However, the requirements of normals can be avoided by sign-agnostic optimization of occupancy field learning from hourglass convolutional networks. According to SAL [2], we know that by properly initializing network parameters, the implicit decoder can represent the signed field of a unit sphere, which helps us obtain signed solutions by unsigned learning objectives. Similarly, the pre-trained occupancy decoder can produce signed fields as initialization for the test-time optimization. As such, we can directly employ the unsigned cross-entropy loss to obtain consistency constraints between the occupancy field and unsigned inputs, without the use of surface normals. Besides, global consistency among local geometries can always be enforced during the optimization stage, because the features from \mathbf{V} are decoded from the same global features. Thus, without the guidance of normals, we can still guarantee globally consistent local field assemblies. Specifically, the unsigned cross-entropy (UCE) loss is formulated as:

$$\mathcal{L}_{uce} = \sum_{\mathbf{q} \in \mathcal{Q}} \text{BCE}(\mathbf{O}^\dagger(\mathbf{q}), \hat{\mathbf{O}}^\dagger(\mathbf{q})), \quad (3)$$

where prediction $\mathbf{O}^\dagger(\mathbf{q})$ and target $\hat{\mathbf{O}}^\dagger(\mathbf{q})$ are given by

$$\mathbf{O}^\dagger(\mathbf{q}) = \text{sigmoid}(|g(\mathbf{q}, \mathbf{f}_q)|) \in [0.5, 1), \quad (4)$$

$$\hat{\mathbf{O}}^\dagger(\mathbf{q}) = \begin{cases} 0.5, & \text{for } \mathbf{q} \in \mathcal{Q}_{\mathcal{S}} \\ 1.0, & \text{for } \mathbf{q} \in \mathcal{Q}_{\setminus \mathcal{S}} \end{cases}, \quad (5)$$

where $Q_{\hat{S}}$ is a point set obtained from the ground-truth surface \hat{S} , and $Q_{\setminus\hat{S}}$ is a point set sampled from non-surface volume $Q \setminus \hat{S}$. As \hat{S} is unknown at the test phase, we consider the observed surface \mathcal{P} as an approximation of \hat{S} , and identify randomly sampled points in 3D space as non-surface points $Q_{\setminus\mathcal{P}}$. More specifically, we force the observed surface \mathcal{P} to align with the 0.5 level set of occupancy field, and the signed occupancy values of non-surface points to be either 0 or 1.

After the sign-agnostic optimization of the implicit field, we apply the Multiresolution IsoSurface Extraction (MISE) [28] and marching cubes [27] to extract surface meshes as the final reconstruction results.

4. Experiments

Datasets We validate the efficiency of our method on experiments of both object-level and scene-level surface reconstruction tasks. For the former task, we conduct comparison on the *chair* category of the ShapeNet [8] dataset. For the latter task, we use the synthetic indoor scene dataset [31]. The split of train/val/test sets follows the same setting in CONet [31]. For each dataset, we randomly select 50 models from the test set to conduct quantitative evaluations. We use point clouds of size 30,000 sampled from true surfaces as inputs. Finally, we compare the synthetic-to-real generality by conducting experiments on ScanNet-V2 [13] and Matterport3D [7] datasets.

Implementation Details We first pre-train the convolutional occupancy networks with a batch size of 32 and a learning rate of 1×10^{-4} for overall 300k iterations. During the sign-agnostic optimization, the whole network is further optimized by the objective described in Equation 3 using a batch size of 16 for 1000 iterations. The initial learning rate is set to 3×10^{-5} , and decays by 0.3 every 400 iterations. We set $|Q_{\hat{S}}| = 512$ and $|Q_{\setminus\hat{S}}| = 1,536$ in Equation 5.

Baselines We conduct comparison with three categories of existing methods, i.e. classic optimization-based methods such as Screened Poisson Surface Reconstruction (SPSR) [25], deep optimization-based methods such as Sign-Agnostic Learning (SAL) [2] and Implicit Geometric Regularization (IGR) [16], learning-based methods such as Occupancy Networks (ONet) [28] and Convolutional Occupancy Networks (CONet) [31], and methods focusing on optimizing data-driven priors such as Local Implicit Grid (LIG) [22]. For SAL and IGR, we directly fit the neural implicit field to the observed point cloud. For ONet, CONet, and LIG, the evaluations are based on their provided pre-trained models. Specifically, we summarize their respective characteristics in Table 1. For methods that require oriented surface normals, we follow [21] to estimate un-oriented normals and then reorient their directions.

Evaluation Metrics We consider Chamfer Distance (CD

Methods	Without normals	Optimization of network parameters	Local geometry modeling
SPSR [25]	×	✓	✓
ONet [28]	✓	×	×
SAL [2]	✓	×	×
IGR [16]	✓	✓	×
CONet [31]	✓	×	✓
LIG [22]	×	✓	✓
Ours	✓	✓	✓

Table 1: Working condition summary of different methods. Note that our method is the first to maximize the scalability to large-scale scenes, generality to novel shapes, and applicability to real-world scans in a unified framework by performing local geometry modeling, optimizing network parameters while not requiring normals during inference.

$\times 0.01$), Normal Consistency (NC $\times 0.01$), and F-score (FS $\times 0.01$) as primary evaluation metrics. The F-score is reported with thresholds of τ and 2τ ($\tau = 0.01$). The quantitative results between two point clouds are measured from randomly sampled ten thousand surface points. For the CD, the lower is better. For NC and FS, the higher is better.

5. Object-level Reconstruction

Methods	CD ↓	NC ↑	FS (τ) ↑	FS (2τ) ↑
SPSR [25]	1.923	81.54	80.86	85.13
ONet[28]	1.117	84.58	62.35	86.57
SAL [2]	2.418	78.67	54.33	73.70
IGR [16]	2.678	75.97	69.02	76.01
CONet [31]	0.821	91.12	74.73	96.85
LIG [22]	2.200	80.35	60.62	65.99
Ours	0.522	93.51	97.16	99.37

Table 2: Quantitative comparison for surface reconstruction from un-oriented point clouds on the ShapeNet-chair.

We first conduct the object-level reconstruction experiments. To simulate the influence of sensing noises during real scans acquisition, we perturb the input by gaussian noise with zero mean and standard deviation 0.05.

As shown in Figure 3, our approach has demonstrated superiority in terms of visual quality in the complex object reconstruction. Compared to the methods of only using global shape features such as ONet [28], SAL [2], and IGR [16]), ours is more capable of recovering complicated geometries because we utilize rich shape features, including both local and global information. Besides, instead of strictly respecting the learned priors like CONet [31], more faithful surface details (e.g. slender bars and tiny holes) can be preserved by breaking the barrier of pre-trained priors during inference. In addition, SPSR [25], IGR [16], and LIG [22] tend to produce degenerated meshes, caused by

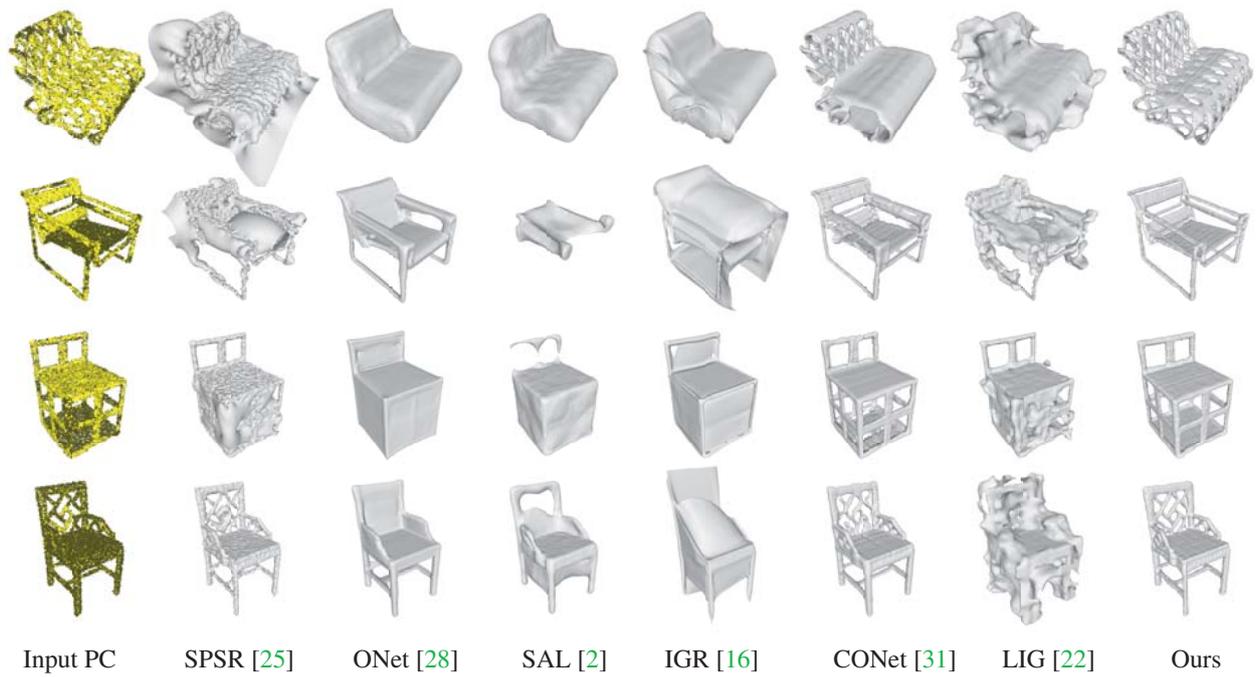


Figure 3: **Object-level Reconstruction on ShapeNet.** Qualitative comparison for surface reconstruction from un-orientated point clouds of ShapeNet-chair [13].

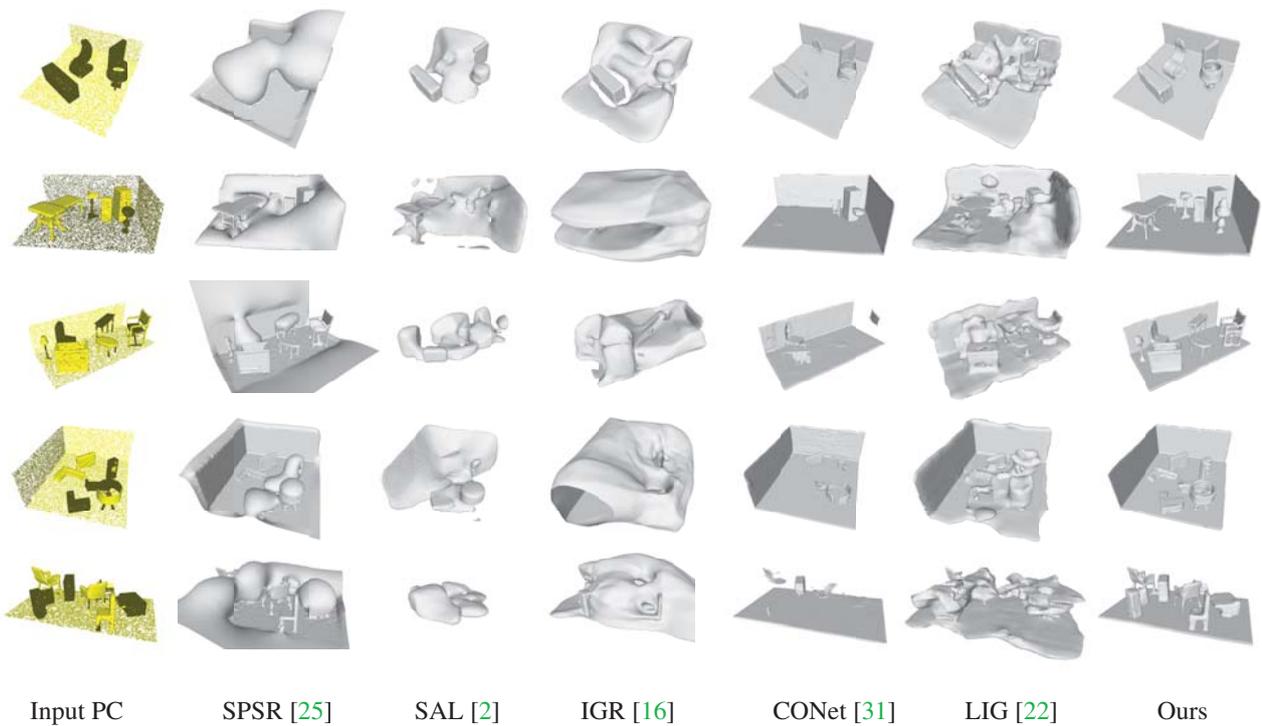


Figure 4: **Scene-level Reconstruction on synthetic rooms.** Qualitative comparison for surface reconstruction from un-orientated point clouds on the synthetic room indoor scene dataset provided by [31].

inaccurate normal orientation estimation. However, it has no negative effects on our results, as we do not use the additional information during the test-time optimization. Our superiority is also verified by numerical results reported in Table 2, where our method outperforms existing state-of-the-art methods by large margins.

6. Scene-Level Reconstruction

Methods	CD ↓	NC ↑	FS (τ) ↑	FS (2τ) ↑
SPSR [25]	2.083	78.21	76.17	81.22
SAL [2]	2.720	73.85	40.47	59.79
IGR [16]	1.923	77.94	74.02	81.23
CONet [31]	2.020	83.43	73.28	81.74
LIG [22]	1.953	79.82	62.46	70.96
Ours	0.495	90.04	93.85	98.82

Table 3: Quantitative results for surface reconstruction from un-oriented point clouds on the synthetic room dataset [31].

To investigate whether our method possesses the scalability to indoor scene reconstructions, we further conduct the experiments of 3D reconstruction from un-oriented point clouds on the synthetic indoor scene dataset [31]. From the qualitative comparison shown in Figure 4, we can observe that some subtle legs of chairs can also be recovered by our approach, while others cannot capture these details. This demonstrates that ours can scale well to large scenes as we adopt the strategy of local geometry reasoning, instead of global shape modeling as in SAL [2] and IGR [16]. Compared to CONet [31], more fine-grained surface recoveries demonstrate that our approach can achieve better generality to novel scenes, due to the effective sign-agnostic optimization that conforms the desired implicit surface to the observed un-oriented surfaces. Moreover, the bypassing of normal estimation enables more robust scene-level surface reconstructions. Again, better quantitative results presented in Table 3 consistently demonstrate the superiority of our approach.

Methods	CD ↓	NC ↑	FS (τ) ↑	FS (2τ) ↑
SPSR [25]	1.339	84.60	82.33	87.83
SAL [2]	2.026	81.24	61.54	80.90
IGR [16]	2.392	84.12	78.07	83.98
CONet [31]	1.559	82.05	59.55	80.76
LIG [22]	1.501	81.99	70.39	78.30
Ours	0.728	86.40	82.08	95.86

Table 4: Quantitative comparison for surface reconstruction from un-oriented point clouds on the real-world ScanNet dataset [31]. As watertight meshes are not provided by ScanNet, we directly evaluate all methods using the pre-trained models on the synthetic room dataset.

7. Real-World Scenes Generalization

To compare the generalization performance on real-world scans, we also evaluate our approach on the real-world datasets, including ScanNet [13] and Matterport3D [7]. Notably, all models are only trained on the same synthetic indoor scene dataset.

ScanNet-V2 The qualitative and quantitative comparisons are respectively shown in Figure 5 and Table 4. As can be seen, compared to other methods, our results achieve clearly better numerical scores and more fine-grained surface geometries, which effectively verifies the better generalization ability of the proposed method on real-world scans.

Matterport3D To evaluate its scalability to huge scenes that contain multiple rooms, we finally conduct experiments on the Matterport3D dataset [7]. Following the sliding-window strategy presented in [31], we separately apply the designed sign-agnostic optimization of convolutional occupancy networks to each room. The visualization comparison is presented in Figure 5 (d). Notably, the Matterport3D [7] is significantly different from the synthetic indoor room dataset that is used to pre-train our network. But our reconstruction results can still preserve rich details inside each room while adhering to the room layout, which fully demonstrates that our method can achieve better scalability to huge scenes and better robustness to noises from different sensing devices.

8. Conclusion

For the task of surface reconstruction from un-oriented point clouds, we have proposed a simple yet effective solution of learning implicit surface reconstructions by sign-agnostic optimization of convolutional occupancy networks, which achieves scalability to large scenes, generality to novel shapes, and applicability to real-world scans in a unified framework. The characteristics of implicit field learning from convolutional features of hourglass networks enable the test-time optimization without the use of surface normals. Extensive experiments on both object-level and scene-level datasets show that our method significantly outperforms the existing methods, both quantitatively and qualitatively. A limitation of our approach is the slow inference speed, which is also a common drawback of test-time optimization methods. We leave it as our future effort.

Acknowledgement. This work was partially supported by the Guangdong R&D key project of China (No.: 2019B010155001), the National Natural Science Foundation of China (No.: 61771201), the Program for Guangdong Introducing Innovative and Entrepreneurial Teams (No.: 2017ZT07X183), the Early Career Scheme of the Research Grants Council (RGC) of the Hong Kong SAR GRF (No.: 26202321), HKUST Startup Fund (No.: R9253), and Alibaba DAMO Academy.

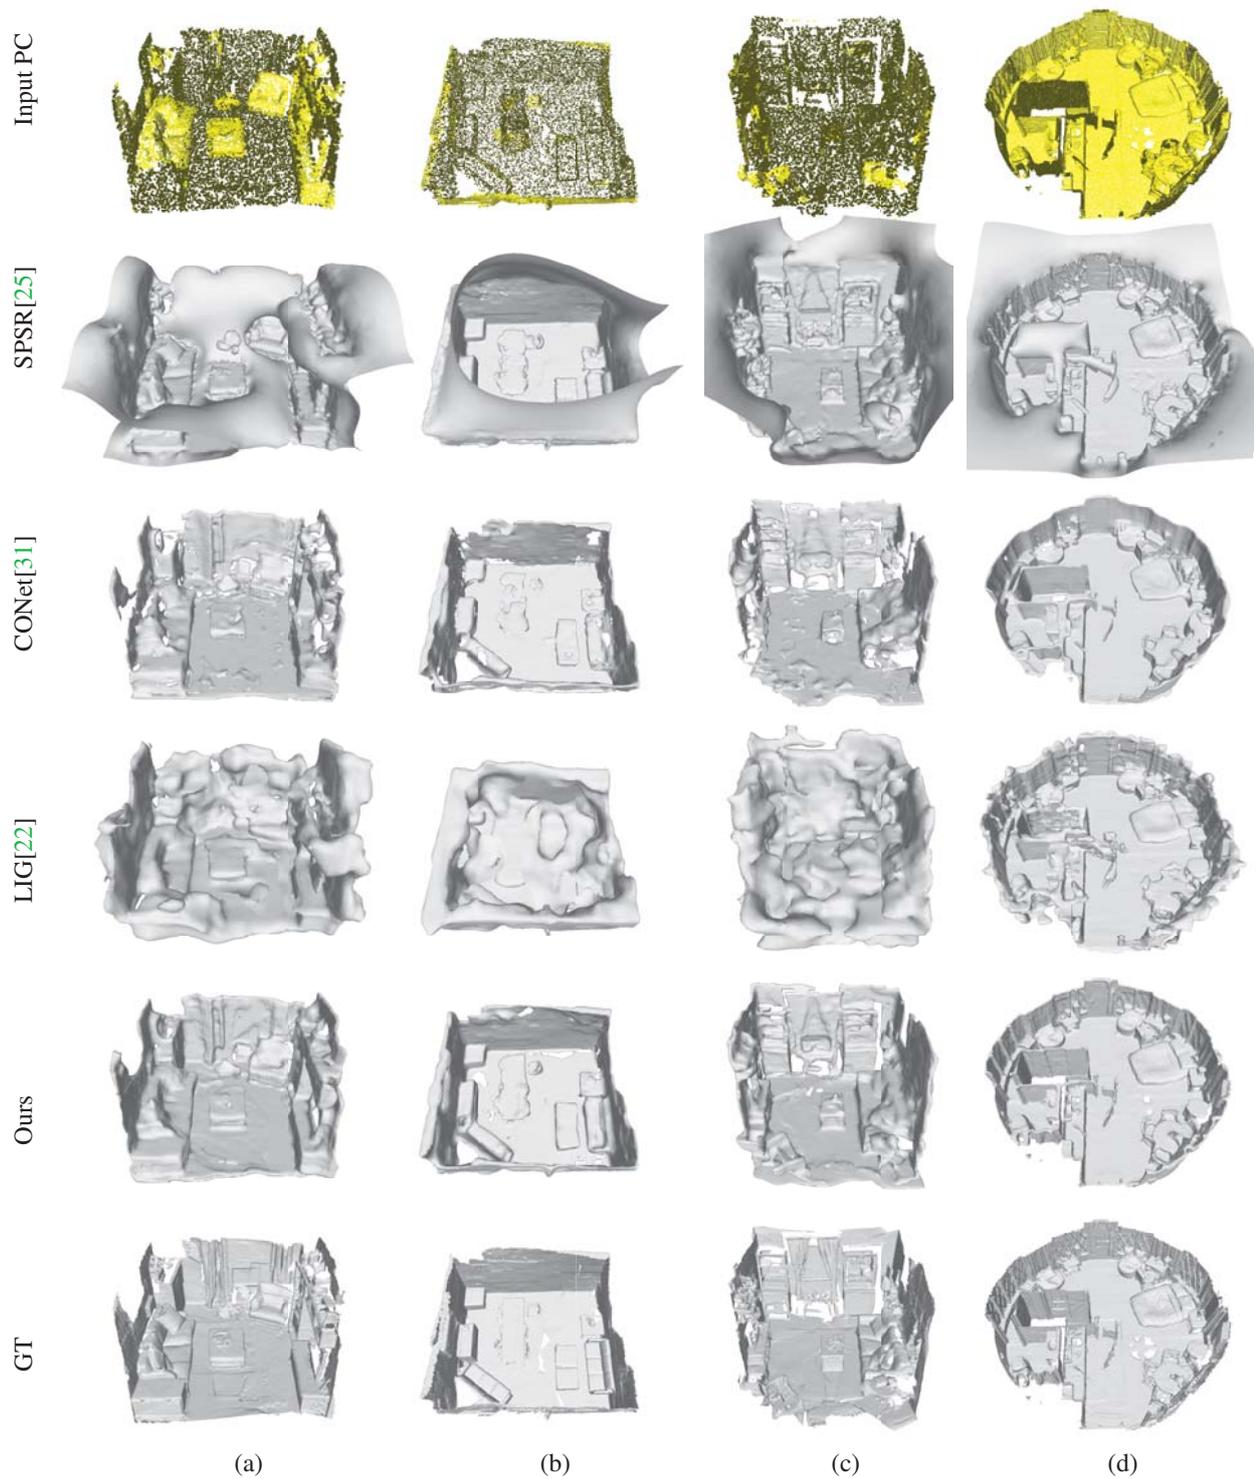


Figure 5: **Scene-level Reconstruction on ScanNet [13] and Matterport3D [7]**. Qualitative comparison for surface reconstruction from un-orientated scans of ScanNet (a, b, c) and Matterport3D (d). All methods except SPSR are trained on the synthetic room dataset and directly evaluated on ScanNet.

References

- [1] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T. Silva. Computing and rendering point set surfaces. *TVCG*, 2003.
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *CVPR*, 2020.
- [3] Matan Atzmon and Yaron Lipman. Sald: Sign agnostic learning with derivatives. *ICLR*, 2021.
- [4] Yizhak Ben-Shabat and Stephen Gould. Deepfit: 3d surface fitting via neural network weighted least squares. In *ECCV*. Springer, 2020.
- [5] Jonathan C Carr, Richard K Beatson, Jon B Cherrie, Tim J Mitchell, W Richard Fright, Bruce C McCallum, and Tim R Evans. Reconstruction and representation of 3d objects with radial basis functions. In *SIGGRAPH*, 2001.
- [6] Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*. Springer, 2020.
- [7] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niebner, M. Savva, S. Song, A. Zeng, and Y. Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *3DV*, 2017.
- [8] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [9] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019.
- [11] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *CVPR*, 2020.
- [12] Christopher Bongsoo Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 2016.
- [13] 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.
- [14] Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J Mitra, and Michael Wimmer. Points2surf learning implicit surfaces from point clouds. In *ECCV*. Springer, 2020.
- [15] Haoqiang Fan, Hao Su, and Leonidas J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, 2017.
- [16] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *ICML*, 2020.
- [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] Gaël Guennebaud and Markus Gross. Algebraic point set surfaces. In *SIGGRAPH*. 2007.
- [19] Paul Guerrero, Yanir Kleiman, Maks Ovsjanikov, and Niloy J Mitra. Pcpnet learning local shape properties from raw point clouds. In *Computer Graphics Forum*. Wiley Online Library, 2018.
- [20] Christian Häne, Shubham Tulsiani, and Jitendra Malik. Hierarchical surface prediction for 3d object reconstruction. In *3DV*, 2017.
- [21] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *SIGGRAPH*, 1992.
- [22] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *CVPR*, 2020.
- [23] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018.
- [24] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *SGP*, 2006.
- [25] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM ToG*, 2013.
- [26] David Levin. Mesh-independent surface interpolation. In *Geometric modeling for scientific visualization*. Springer, 2004.
- [27] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH*, 1987.
- [28] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019.
- [29] Junyi Pan, Xiaoguang Han, Weikai Chen, Jiapeng Tang, and Kui Jia. Deep mesh reconstruction from single rgb images via topology modification networks. In *ICCV*, 2019.
- [30] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019.
- [31] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, 2020.
- [32] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, 2017.
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*. Springer, 2015.
- [34] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. *ICCV*, 2019.
- [35] Lv Sheng, Dan Xu, Wanli Ouyang, and Xiaogang Wang. Un-supervised collaborative learning of keyframe detection and visual odometry towards monocular deep slam. In *ICCV*, 2019.
- [36] Jiapeng Tang, Xiaoguang Han, Junyi Pan, Kui Jia, and Xin Tong. A skeleton-bridged deep learning approach for gener-

- ating meshes of complex topologies from single rgb images. In *CVPR*, 2019.
- [37] Jiapeng Tang, Xiaoguang Han, Mingkui Tan, Xin Tong, and Kui Jia. Skeletonnet: A topology-preserving solution for learning mesh reconstruction of object surfaces from rgb images. *TPAMI*, 2021.
- [38] Jiapeng Tang, Dan Xu, Kui Jia, and Lei Zhang. Learning parallel dense correspondence from spatio-temporal descriptors for efficient and robust 4d reconstruction. In *CVPR*, 2021.
- [39] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *ICCV*, 2017.
- [40] 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.
- [41] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Carsten Stoll, and Christian Theobalt. Patchnets: Patch-based generalizable deep implicit 3d shape representations. In *ECCV*. Springer, 2020.
- [42] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018.
- [43] Peng-Shuai Wang, Chun-Yu Sun, Yang Liu, and Xin Tong. Adaptive o-cnn: a patch-based deep representation of 3d shapes. In *SIGGRAPH Asia*, 2018.
- [44] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NeurIPS*, 2016.
- [45] Dan Xu, Weidi Xie, and Andrew Zisserman. Geometry-aware video object detection for static cameras. In *BMVC*, 2019.
- [46] 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.
- [47] Mingyue Yang, Yuxin Wen, Weikai Chen, Yongwei Chen, and Kui Jia. Deep optimized priors for 3d shape modeling and reconstruction. *CVPR*, 2021.
- [48] Wenbin Zhao, Jiabao Lei, Yuxin Wen, Jianguo Zhang, and Kui Jia. Sign-agnostic implicit learning of surface self-similarities for shape modeling and reconstruction from raw point clouds. In *CVPR*, 2021.