

3D Shape Reconstruction from Free-Hand Sketches

Jiayun Wang¹[0000-0002-8356-2775] Jierui Lin¹ Qian Yu^{1,2}
Runtao Liu¹ Yubei Chen¹ Stella X. Yu¹[0000-0002-3507-5761]

¹ UC Berkeley / ICSI

{peterwg, jerrylin0928, runtao_liu, yubeic, stellayu}@berkeley.edu

² Beihang Univerisity qianyu@buaa.edu.cn

Abstract. Sketches are arguably the most abstract 2D representations of real-world objects. Although a sketch usually has geometrical distortion and lacks visual cues, humans can effortlessly envision a 3D object from it. This suggests that sketches encode the information necessary for reconstructing 3D shapes. Despite great progress achieved in 3D reconstruction from distortion-free line drawings, such as CAD and edge maps, little effort has been made to reconstruct 3D shapes from free-hand sketches. We study this task and aim to enhance the power of sketches in 3D-related applications such as interactive design and VR/AR games.

Unlike previous works, which mostly study distortion-free line drawings, our 3D shape reconstruction is based on free-hand sketches. A major challenge for free-hand sketch 3D reconstruction comes from the insufficient training data and free-hand sketch diversity, e.g. individualized sketching styles. We thus propose data generation and standardization mechanisms. Instead of distortion-free line drawings, synthesized sketches are adopted as input training data. Additionally, we propose a sketch standardization module to handle different sketch distortions and styles. Extensive experiments demonstrate the effectiveness of our model and its strong generalizability to various free-hand sketches. Our code is available.

Keywords: Sketch, interactive design, shape reconstruction, data insufficiency, 3D reconstruction

1 Introduction

Human free-hand sketches are the most abstract 2D representations for 3D visual perception. Although a sketch may consist of only a few colorless strokes and exhibit various deformation and abstractions, humans can effortlessly envision the corresponding real-world 3D object from it. It is of interest to develop a computer vision model that can replicate this ability. Although sketches and 3D representations have drawn great interest from researchers in recent years, these two modalities have been studied relatively independently. We explore the plausibility of bridging the gap between sketches and 3D, and build a computer vision model to recover 3D shapes from sketches. Such a model will unleash many applications, like interactive CAD design and VR/AR games.




Fig. 1: **Left:** We study 3D reconstruction from a single-view free-hand sketch, differing from previous works [7, 30, 49] which use multi-view distortion-free line-drawings as training data. **Center:** While previous works [7, 49] employ distortion-free line drawings (e.g. edge-maps) as proxies for sketches, our model trained on synthesized sketches can generalize better to free-hand sketches. Further, the proposed sketch standardization module makes the method generalizes well to free-hand sketches by standardizing different sketching styles and distortion levels. **Right:** Our model unleashes many practical applications such as real-time 3D modeling with sketches. A demo is here.

With the development of new devices and sensors, sketches and 3D shapes, as representations of real-world objects beyond natural images, become increasingly important. The popularity of touch-screen devices makes sketching not a privilege to professionals anymore and increasingly popular. Researchers have applied sketch in tasks like image retrieval [12, 26, 45, 51, 58, 62] and image synthesis [12, 13, 29, 40, 52, 63] to leverage its power in expression. Furthermore, as depth sensors, such as structured light device, LiDAR, and TOF cameras, become more ubiquitous, 3D data become an emerging modality in computer vision. 3D reconstruction, the process of capturing the shape and appearance of real objects, is an essential topic in 3D computer vision. 3D reconstruction from multi-view images has been studied for many years [1, 4, 11, 38]. Recent works [8, 10, 43] have further explored 3D reconstruction from a single image.

Despite these trends and progress, there are limited works connecting 3D and sketches. We argue that sketches are abstract 2D representations of 3D perception, and it is of great significance to study sketches in a 3D-aware perspective and build connections between two modalities. Researchers have explored the potential of *distortion-free* line drawings (e.g. edge maps) for 3D modeling [27, 28, 57]. These works are based on *distortion-free* line drawings and generalize poorly to free-hand sketches (Fig.1L). Furthermore, the role of line drawings in such works is to provide geometrical information for the subsequent 3D modeling. Some other works [7, 30] employ neural networks to reconstruct 3D shapes directly from line drawings. However, their decent reconstructions come with two major limitations: **a)** they use distortion-free line drawings as training data, which makes such models hard to generalize to free-hand sketches; **b)** they usually require inputs depicting the object from multi-views to achieve satisfactory outcomes. Therefore, such methods cannot reconstruct the 3D shape from a single-view free-hand sketch well, as we show later in the experiment section. Other works such as [19, 49] tackle 3D retrieval instead of 3D shape reconstruction from

sketches. Retrieved shapes come from the pre-collected gallery set and may not resemble novel sketches well. Overall, reconstructing a 3D shape from a single *free-hand* sketch is still left not well explored.

We explore single-view free-hand sketch-based 3D reconstruction (Fig.1C). A free-hand sketch is defined as a line drawing created without any additional tool. As an abstract and concise representation, it is different from distortion-free line drawings (e.g. edge maps) since it commonly has some spatial distortions, but it can still reflect the essential geometric shape. 3D reconstruction from sketch is challenging due to the following reasons: **a)** Data insufficiency. Paired sketch-3D datasets are rare although there exist several large-scale sketch datasets and 3D shape datasets, respectively. Furthermore, collecting sketch-3D pairs can be very time-consuming and expensive than collecting sketch-image pairs, as each 3D shape could be sketched from various viewing angles. **b)** Misalignment between two representations. A sketch depicts an object from a certain view while a 3D shape can be viewed from multiple angles due to the encoded depth information. **c)** Due to the nature of hand drawing, a sketch is usually geometrically imprecise with a individual style compared to the real object. Thus a sketch can only provide suggestive shape and structural information. In contrast, a 3D shape is faithful to its corresponding real-world object with no geometric deformation.

To address these challenges, we propose a single-view sketch-to-3D shape reconstruction framework. Specifically, it takes a sketch from an *arbitrary* angle as input and reconstructs a 3D point cloud. Our model cascades a sketch standardization module U and a reconstruction module G . U handles various sketching styles/distortions and transfers inputs to standardized sketches while G takes a standardized sketch to reconstruct the 3D shape (point cloud) *regardless of the object category*. The key novelty lies in the mechanisms we propose to tackle the data insufficiency issue. Specifically, we first train an photo-to-sketch model on unpaired large-scale datasets. Based on the model, sketch-3D pairs can be automatically generated from 2D renderings of 3D shapes. Together with the standardization module U which unifies input sketch styles, the synthesized sketches provide sufficient information to train the reconstruction model G . We conduct extensive experiments on a composed sketch-3D dataset, spanning 13 classes, where sketches are synthesized and 3D objects come from the ShapeNet dataset [3]. Furthermore, we collect an evaluation set, which consists of 390 real sketch-3D pairs. Results demonstrate that our model can reconstruct 3D shapes with certain geometric details from real sketches under different styles, stroke line-widths, and object categories. Our model also enables practical applications such as real-time 3D modeling with sketches (Fig.1R).

To summarize our contributions: **a)** We are one of the pioneers to study the plausibility of reconstructing 3D shapes from single-view free-hand sketches. **b)** We propose a novel framework for this task and explore various design choices. **c)** To handle data insufficiency, we propose to train on synthetic sketches. Moreover, sketch standardization is introduced to make the model generalize to free-hand sketches better. It is a general method for zero-shot domain translation, and we show applications on zero-shot image translation tasks.

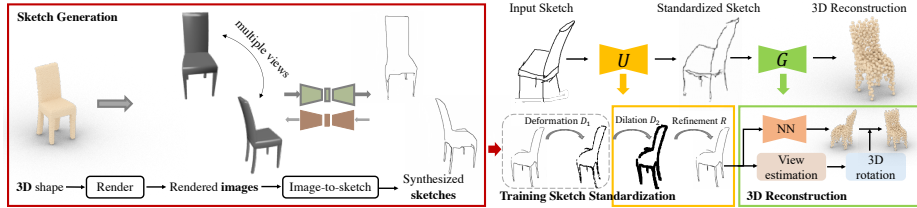


Fig. 2: Model overview. The model consists of three major components: **sketch generation**, **sketch standardization**, and **3D reconstruction**. To generate synthesized sketches, we first render 2D images for a 3D shape from multiple viewpoints. We then employ an image-to-sketch translation model to generate sketches of corresponding views. The standardization module standardize sketches with different distortions and deformations. *Deformation* D_2 is only used in training for augmentation such that the model is robust to geometric distortions of sketches. For inference, sketches are dilated (D_1) and refined (R) so their style matches that of training sketches. For 3D reconstruction, a view estimation module is adopted to align the output’s view and the ground-truth 3D shape.

2 Related Works

3D Reconstruction from Images. While SfM [38] and SLAM [11] achieve success in handling multi-view 3D reconstructions in various real-world scenarios, their reconstructions can be limited by insufficient input viewpoints and 3D scanning data. Deep-learning-based methods have been proposed to further improve reconstructions by completing 3D shapes with occluded or hollowed-out areas [4, 22, 59]. In general, recovering the 3D shape from a single-view image is an ill-posed problem. Attempts to tackle the problem include 3D shape reconstructions from silhouettes [8], shading [43], and texture [54]. However, these methods need strong presumptions and expertise in natural images [65], limiting their usage in real-world scenarios. Generative adversarial networks (GANs) [14] and variational autoencoders (VAEs) [24] have achieved success in image synthesis and enabled [55] 3D shape reconstruction from a single-view image. Fan *et al.* [10] further adopt point clouds as 3D representation, enabling models to reconstruct certain geometric details from an image. They may not directly work on sketches as many visual cues are missing.

3D reconstruction networks are designed differently depending on the output 3D representation. 3D voxel reconstruction networks [18, 48, 56] benefit from many image processing networks as convolutions are appropriate for voxels. They are usually constrained to low resolution due to the computational overhead. Mesh reconstruction networks [25, 53] are able to directly learn from meshes, where they suffer from topology issues and heavy computation [39]. We adopt point cloud representation as it can capture certain 3D geometric details with low computational overhead. Reconstructing 3D point clouds from images has been shown to benefit from well-designed network architectures [10, 32], latent embedding matching [31], additional image supervision [35], etc.

Sketch-Based 3D Retrievals/Reconstructions. Free-hand sketches are used for 3D shape retrieval [19, 49] given their power in expression. However, retrieval methods are significantly constrained by the gallery dataset. Precise sketching is also studied in the computer graphics community for 3D shape modeling




Fig. 3: Synthesized sketches are visually more similar to free-hand sketches than edge maps as they contain distortions and emphasize perceptually significant contours. After standardization, the free-hand sketches share a uniform style similar to training data.

or procedural modeling [20, 27, 28]. These works are designed for professionals and require additional information for shape modeling, e.g., surface-normal, procedural model parameters. Delanoy *et al.* [7] first employ neural networks to learn 3D voxels from line-drawings. While it achieves impressive performance, this model has several limitations: **a)** The model uses distortion-free edge map as training data. While working on some sketches with small distortions, it cannot generalize to general free-hand sketches. **b)** The model requires multiple inputs from different viewpoints for a satisfactory result. These limitations prevent the model from generalizing to real free-hand sketches. Recent works also explore reconstructing 3D models from sketches with direct shape optimization [17], shape contours [16], differential renderer [64], and unsupervised learning [50]. Compared to existing works, the proposed method in this work reconstructs the 3D point cloud based on a single-view free-hand sketch. Our model may make 3D reconstruction and its applications more accessible to the general public.

3 3D Reconstruction from Sketches

The proposed framework has three modules (Fig.2). To deal with data insufficiency, we first synthesize sketches as the training set. The module U transfers an input sketch to a standardized sketch. Then, the module G takes the standardized sketch to reconstruct a 3D shape (point clouds). We also present details of a new sketch-3D dataset, which is collected for evaluating the proposed model.

3.1 Synthetic Sketch Generation

To the best of our knowledge, there exists no paired sketch-3D dataset. While it is possible to resort to edge maps [7], edge maps are different from sketches (as shown in the 3rd and 4th rows of Fig.3). We show that the reconstruction model trained on edge maps cannot generalize well to real free-hand sketches in Sec. 4.4. Thus it is crucial to find an efficient and reliable way to synthesize sketches for 3D shapes. Inspired by [29], we employ a generative model to synthesize sketches from rendered images of 3D shapes. Fig.2L depicts the procedure. Specifically,

we first render m images for each 3D shape, where each image corresponds to a particular view of a 3D shape. We then adopt the model introduced in [29] to synthesize gray-scale sketches images, denoted as $\{S_i | S_i \in \mathbb{R}^{W \times H}\}$, as our training data. W, H refer to the width and height of a sketch image.

3.2 Sketch Standardization

Sketches usually have strong individual styles and geometric distortions. Due to the gap between the free-hand sketches and the synthesized sketches, directly using the synthesized sketches as training data would not lead to a robust model. The main issues are that the synthesized sketches have a uniform style and they do not contain enough geometric distortions. Rather, the synthesized sketches can be treated as an intermediate representation if we can find a way to project a free-hand sketch to the synthesized sketch domain. We propose a zero-shot domain translation technique, the sketch standardization module, to achieve this domain adaption goal without using the free-hand sketches as the training data. The training of the sketch standardization module only involves synthesized sketches. The general idea is to project a distorted synthesized sketch to the original synthesized sketch. The training consists of two parts: **a)** since the free-hand sketches usually have geometric distortions, we apply predefined distortion augmentation to the input synthesized sketches first. **b)** A geometrically distorted synthesized sketch still has a different style and line style compared to the free-hand sketches. Thus, the first stage of the standardization is to apply a dilation operation. The dilation operation would project distorted synthesized sketches and the free-hand sketches to the same domain. Then, a refinement network follows to project the dilated sketch back to the synthesized sketch domain.

In summary, as in Fig.2, the standardization module U first applies a dilation operator D_2 to the input sketch, which is followed by a refinement operator R to transfer to the standardized synthesized-sketch style (or training-sketch style) \tilde{S}_i , i.e. $U = R \circ D_2$. R is implemented as an image translation network. During training, a synthesized sketch \tilde{S}_i is first augmented by the deformation operator D_1 to mimic the drawing distortion, and then U aims to project it back to \tilde{S}_i . Please note that D_1 would not be used during the testing. We illustrate the standardization process in Fig.2R with more details in the following.

Deformation. When training U , each synthesized sketch is deformed with moving least squares [46] for random, local and rigid distortion. Specifically, we randomly sample a set of control points on sketch strokes and denote them as p , and denote the corresponding deformed point set as q . Following moving least squares, we solve for the best affine transformation $l_v(x)$ such that: $\min \sum_i w_i |l_v(p_i) - q_i|^2$, where p_i and q_i are row vectors and weights $w_i = \frac{1}{|p_i - v|^{2\alpha}}$. Affine transformation can be written as $l_v(p_i) = p_i M + T$. We add constraint $M^T M = I$ to make the deformation is rigid to avoid too much distortion. Details can be found in [46].

Style Translation. Adapting to unknown input free-hand sketch style during inference can be considered as a zero-shot domain translation problem, which is challenging. Inspired by [61], we first dilate the augmented training sketch strokes with 4 pixels and then use image-to-image translation network Pix2Pix [21] to translate the dilated sketches to the un-distorted synthesized sketches. During

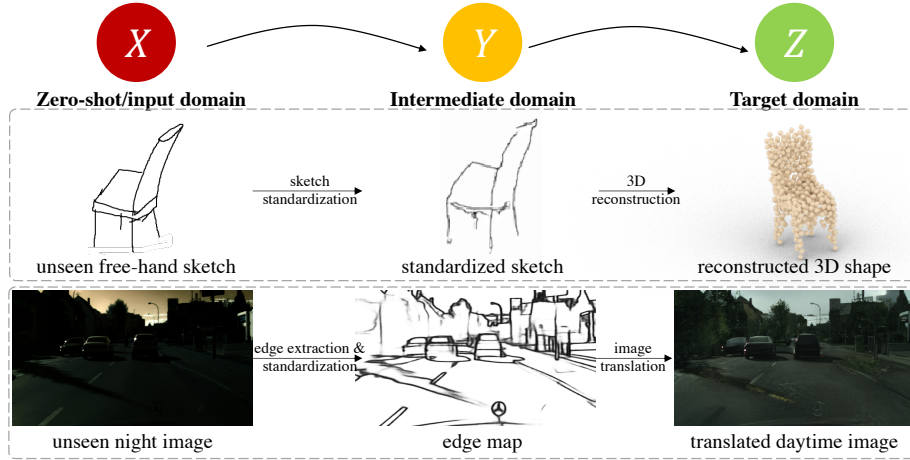


Fig. 4: Sketch standardization can be considered as a general zero-shot domain translation method. Given a sample from a zero-shot (input) domain X , we first translate it to a universal intermediate domain Y , and finally to the target domain Z . **1st example** (2nd row): the input domain is an unseen free-hand sketch. With sketch standardization, it is translated to an intermediate domain: standardized sketch, which shares similar style as synthesized sketch for training. With 3D reconstruction, the standardized sketch can be translated to the target domain: 3D point clouds. **2nd example** (last row): the input domain is an unseen nighttime image. With edge extraction, it gets translated to an intermediate domain: edge map. With the image-to-image translation model, the standardized edge map can be translated to the target domain: daytime image.

inference, we also dilate the free-hand sketches and apply the trained Pix2Pix model such that the style of an input free-hand sketch could be adapted to the synthesized sketch style during training. The dilation step can be considered as introducing uncertainty for the style adaption. Further, we show in Section 4.5 that the proposed style standardization module could be used as a general zero-shot domain translation technique, which generalizes to more applications such as sketch classification and zero-shot image-to-image translation.

A More General Message: Zero-Shot Domain Translation. We illustrate in Fig.4 a more general message of the standardization module: it can be considered as a general method for zero-shot domain translation. Consider the following problem: we would like to build a model to transfer domain X to domain Z but we do not have any training data from domain X . We propose a general idea to solve this problem is to build an intermediate domain Y as a bridge such that: **a)** we can translate data from domain X to domain Y and **2)** we can further translate data from domain Y to domain Z . We give two examples in the caption of Fig.4 and provided experimental results in Section 4.5.

3.3 Sketch-Based 3D Reconstruction

Our 3D reconstruction network G (pipeline in Fig.2R) consists of several components. Given a standardized sketch \tilde{S}_i , the view estimation module first estimates its viewpoint. \tilde{S}_i is then fed to the sketch-to-3D module to generate a point cloud

$P_{i,pre}$, whose pose aligns with the sketch viewpoint. A 3D rotation corresponding to the viewpoint is then applied to $P_{i,pre}$ to output the canonically-posed point cloud P_i . The objective of G is to minimize distances between reconstructed point cloud P_i and the ground-truth point cloud $P_{i,gt}$.

View Estimation Module. The view estimation module g_1 aims to determine the 3D pose from an input sketch \tilde{S} . Similar to the input transformation module of the PointNet [42], g_1 estimates a 3D rotation matrix A from a sketch \tilde{S} , i.e., $A = g_1(\tilde{S})$. A regularization loss $L_{orth} = \|I - AA^T\|_F^2$ is applied to ensure A is a rotation (orthogonal) matrix. The rotation matrix A rotates a point cloud from the viewpoint pose to a canonical pose, which matches the ground truth.

3D Reconstruction Module. The reconstruction network g_2 learns to reconstruct a 3D point cloud P_{pre} from a sketch \tilde{S} , i.e., $P_{pre} = g_2(\tilde{S})$. P_{pre} is further transformed by the corresponding rotation matrix A to P so that P aligns with the ground-truth 3D point cloud P_{gt} 's canonical pose. Overall, we have $P = g_1(\tilde{S}) \cdot g_2(\tilde{S})$. To train G , we penalize the distance between an output point cloud P and the ground-truth point cloud P_{gt} . We employ the Chamfer distance (CD) between two point clouds $P, P_{gt} \subset \mathbb{R}^3$:

$$d_{CD}(P||P_{gt}) = \sum_{\mathbf{p} \in P} \min_{\mathbf{q} \in P_{gt}} \|\mathbf{p} - \mathbf{q}\|_2^2 + \sum_{\mathbf{q} \in P_{gt}} \min_{\mathbf{p} \in P} \|\mathbf{p} - \mathbf{q}\|_2^2 \quad (1)$$

The final loss of the entire network is $L = \sum_i d_{CD}(G \circ U(S_i)||P_{i,gt}) + \lambda L_{orth} = \sum_i d_{CD}(A_i \cdot P_{i,pre}||P_{i,gt}) + \lambda L_{orth} = \sum_i d_{CD}(g_1(\tilde{S}_i) \cdot g_2(\tilde{S}_i)||P_{i,gt}) + \lambda L_{orth}$

where λ is the weight of the orthogonal regularization loss and $\tilde{S}_i = R \circ D_2 \circ D_1(S_i)$ is the standardized sketch from S_i . We employ CD rather than EMD (Section 4.2) to penalize the difference between the reconstruction and the ground-truth point clouds because CD emphasizes the geometric outline of point clouds and leads to reconstructions with better geometric details. EMD, however, emphasizes the point cloud distribution and may not preserve the geometric details well at locations with low point density.

4 Experimental Results

We first present the datasets, training and evaluation details, followed by qualitative and quantitative results. Then, we provide comparisons with some state-of-the-art methods. We also conduct ablation studies to understand each module.

4.1 3D Sketching Dataset

To evaluate the performance of our method, we collected a real-world evaluation set containing paired sketch-3D data. Specifically, we randomly choose ten 3D shapes from each of the 13 categories of the ShapeNet dataset [3]. Then we randomly render 3 images from different viewpoints for each 3D shape. Totally, there are 130 different 3D shapes and 390 rendered images. We recruited 11 volunteers to draw the sketches for the rendered images. Final sketches are reviewed for quality control. We present several examples in Fig.3.

4.2 Training Details and Evaluation Metrics

Training. The proposed model is trained on a subset of ShapeNet [3] dataset, following settings of [56]. The dataset consists of 43,783 3D shapes spanning 13

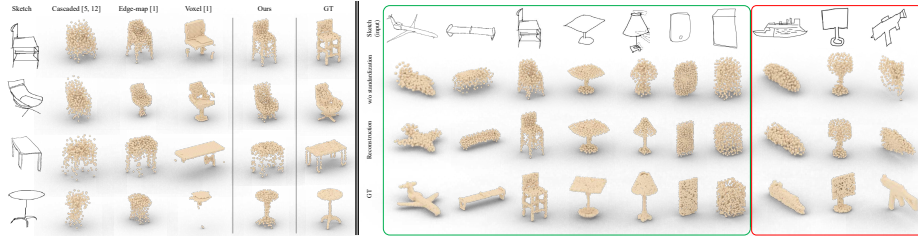


Fig. 5: **Left:** Performance on free-hand sketches with different design choices. The design pool includes the model with a cascaded two-stage structure (2nd column), the model trained on edge maps (3rd column), the model whose 3D output is represented by voxel (4th column), and the proposed model (5th column). Overall, the proposed method achieves better performance and keeps more fine-grained details, e.g., the legs of chairs. **Right:** 3D reconstructions on our newly-collected free-hand sketch evaluation dataset. *Outlined in green:* Examples of some good reconstruction results. Our model reconstructs 3D shapes with fine geometric fidelity of multiple categories *unconditionally*. *Outlined in red:* Examples of failure cases. Our model may not handle detailed structures well (e.g., *watercraft*), recognize the wrong category (e.g., *display* as a *lamp*) due to the ambiguity of the sketch, as well as not able to generate 3D shape from very abstract sketches where few geometric information is available (e.g., *rifle*).

categories, including car, chair, table, etc. For each category, we randomly select 80% 3D shapes for training and the rest for evaluation. As mentioned in Section 3.1, corresponding sketches of rendered images from 24 viewpoints of each 3D shape of ShapeNet are synthesized with our synthetic sketch generation module.

Evaluation. To evaluate our method’s 3D reconstruction performance on free-hand sketches, we use our proposed sketch-3D datasets (Section 4.1). To evaluate the generalizability of our model, we also evaluate on three additional free-hand sketch datasets, including the Sketchy dataset [45], the TU-Berlin dataset [9], and the QuickDraw dataset [15]. For these additional datasets, only sketches from categories that overlap with the ShapeNet dataset are considered.

Following the previous works [10, 31, 60], we adopt two evaluation metrics to measure the similarity between the reconstructed 3D point cloud P and the ground-truth point cloud P_{gt} . The first one is the Chamfer Distance (Eqn.1), and another one is the Earth Mover’s Distance (EMD): $d_{EMD}(P, P_{gt}) = \min_{\phi: P \rightarrow P_{gt}} \sum_{x \in P} \|x - \phi(x)\|$, where P, P_{gt} has the same size $|P| = |P_{gt}|$ and $\phi: P \rightarrow P_{gt}$ is a bijection. CD and EMD evaluate the similarity between two point clouds from two different perspectives (more details can be found in [10]).

4.3 Implementation Details

Sketch Generation. We utilize an off-the-shelf sketch-image translation model [29] to synthesize sketches for training. Given the appropriate quality of the generated sketches on the ShapeNet dataset (with some samples depicted in Fig.3), we directly use the model without any fine-tuning.

Data Augmentation. During training, to improve the model’s generalizability and robustness, we perform data augmentation for synthetic sketches before

Table 1: Ours outperforms baselines for 3D reconstruction. [7] uses edge-maps rather than sketches as input. [56] uses voxels rather than point clouds as output. [37] represents the 3D shape with multi-view depth maps. “cas.” refers to the two-stage cascaded training following [10, 29]. CD and EMD measure distances between reconstructions and ground-truths from different perspectives (see text for details). The lower, the better.

error	Chamfer Distance ($\times 10^{-4}$)							Earth Mover’s Distance ($\times 10^{-2}$)						
	points	edge [7]	voxel [56]	cas.	[37]	retrieval [49]	ours	points	edge [7]	voxel [56]	cas.	[37]	retrieval [49]	ours
airplane	11.4	7.8	35.1	71.7	8.0	11.2	6.1	8.5	7.3	10.8	12.7	8.5	11.9	6.5
bench	29.2	16.7	202.8	414.1	16.8	14.5	13.0	11.1	8.7	22.0	25.8	10.0	8.6	7.8
cabinet	61.7	50.4	59.1	354.5	51.5	45.3	39.2	17.6	17.8	17.0	29.6	18.4	17.2	16.0
car	20.8	13.3	173.2	114.2	14.1	14.2	10.4	8.9	20.0	25.2	20.0	21.6	21.2	18.0
chair	41.8	36.4	108.6	237.1	36.1	33.0	26.9	15.1	15.6	19.4	22.8	16.1	15.3	13.0
display	68.6	48.3	33.1	340.2	49.3	38.2	37.7	15.5	15.1	13.1	27.9	16.4	14.6	14.4
lamp	63.3	59.4	107.0	214.0	60.2	63.5	46.3	21.3	22.6	21.2	24.9	22.3	22.6	20.4
speaker	88.2	79.7	203.2	406.4	81.2	72.3	62.1	19.4	19.2	23.8	28.0	21.8	20.0	17.9
rifle	17.0	12.1	170.1	15.4	12.3	14.2	10.1	11.2	13.8	23.7	15.4	15.2	17.6	12.4
sofa	32.8	20.9	141.2	482.4	22.3	20.3	16.3	11.1	8.5	18.6	25.4	9.1	8.6	7.7
table	55.2	49.4	134.7	469.5	50.5	49.1	40.7	19.1	17.7	18.5	26.5	18.2	18.2	17.3
telephone	30.7	27.3	26.9	259.8	27.1	27.4	21.3	13.4	13.6	15.1	27.2	15.1	15.3	12.3
watercraft	32.9	26.0	129.1	53.8	26.0	27.3	20.3	12.5	11.1	23.1	17.8	12.2	12.7	10.6
avg.	42.6	34.4	117.2	264.1	35.0	33.1	26.9	14.2	14.7	19.3	23.4	15.8	15.7	13.4
free-hand sketch	87.1	89.0	162.5	334.2	91.8	89.2	86.1	18.6	16.4	22.9	26.1	17.0	16.8	16.0

feeding them to the standardization module. Specifically, we apply image spatial translation (up to ± 10 pixels) and rotation (up to $\pm 10^\circ$) on each input sketch.

Sketch Standardization. Each input sketch S_i is first randomly deformed with moving least squares [46] both globally and locally (D_1), and then binarized and dilated five times iteratively (D_2) to obtain a rough sketch S_r . The rough sketch S_r is then used to train a Pix2Pix model [21], R , to reconstruct the input sketch S_i . The network is trained for 100 epochs with an initial learning rate of $2e-4$. Adam optimizer [23] is used for the parameter optimization. During evaluation, random deformation D_1 is discarded.

3D Reconstruction. The 3D reconstruction network is based on [10]’s framework with hourglass network architecture [36]. We compare several different network architectures (simple encoder-decoder architecture, two-prediction-branch architecture, etc.) and find that hourglass network architecture gives the best performance. This may be due to its ability to extract key points from images [2, 36]. We train the network for 260 epochs with an initial learning rate of $3e-5$. The weight λ of the orthogonal loss is $1e-3$. To enhance the performance on every category, all categories of 3D shapes are trained together. The class-aware mini-batch sampling [47] is adopted to ensure a balanced category-wise distribution for each mini-batch. We choose Adam optimizer [23] for the parameter optimization. 3D point clouds are visualized with the rendering tool from [33].

4.4 Results and Comparisons

We first present our model’s 3D shape reconstruction performance, along with the comparisons with various baseline methods. Then we present the results on sketches from different viewpoints and of different categories, as well the results on other free-hand sketch datasets. Note that unless specifically mentioned, all evaluations are on the free-hand sketches rather than synthesized sketches.

Baseline Methods. Our 3D reconstruction network is a one-stage model where the input sketch is treated as an image, and point clouds represent the output 3D shape. As conducting the first work for single-view sketch-based 3D reconstruction, we explore different design options adopted by previous works on distortion-free line drawings and/or 3D reconstruction, including architectures, representation of sketches and 3D shapes. We compare with different variants to demonstrate the effectiveness of each choice of our model.

1) Model design: end-to-end vs. two-stage. Although the task of reconstructing 3D shapes from free-hand sketches is new, sketch-to-image synthesis and 3D shape reconstruction from images have been studied before [10, 29, 56]. Is a straight combination of the two models, instead of an end-to-end model, enough to perform well for the task? To compare these two architectures’ performance, We implement a cascaded model by composing a sketch-to-image model [66] and an image-to-3D model [10] to reconstruct 3D shapes.

2) Sketch: point-based vs. image-based. Considering a sketch is relatively sparse in pixel space and consists of colorless strokes, we can employ 2D point clouds to represent a sketch. Specifically, 512 points are randomly sampled from strokes of each binarized sketch, and we use a point-to-point network architecture (adapted from PointNet [42]) to reconstruct 3D shapes from the 2D point clouds.

3) Sketch: Using edge maps as proxy. We compare with a previous work [7]. Our proposed model uses synthetic sketch for training. However, an alternative option is using edge maps as a proxy of the free-hand sketch. As edge maps can be generated automatically (we use the Canny edge detector in implementation), the comparison helps us understand if our proposed synthesizing method is necessary.

4) 3D shape: voxel vs. point cloud. We compare with a previous work [56]. In this variant, we follow their settings and represent a 3D shape with voxels. As the voxel representation is adopted from the previous method, the comparison helps to understand if representing 3D shapes with point clouds has benefits.

5) 3D shape: depth map vs. point cloud. In this variant, we exactly follow a previous work [37] and represent the 3D shape with multi-view depth maps.

Comparison and Results. Table 1 and Fig.5L present quantitative and qualitative results of our method and different design variants. Specifically for quantitative comparisons (Table 1), we report 3D shape reconstruction performance on both synthesized (evaluation set) and free-hand sketches. This is due to that the collected free-hand sketch dataset is relatively small and together they provide a more comprehensive evaluation. We have the following observations:

a) Representing sketches as images outperforms representing them as 2D point clouds (points vs. ours). **b)** The model trained on synthesized sketches performs better on real free-hand sketches than the model trained on edge maps (89.0 vs. 86.1 on CD, 16.4 vs. 16.0 on EMD). Training with edge maps could reconstruct okay overall coarse shape. However, the unsatisfactory performance on geometric details reveals such methods are hard to generalize to free-hand sketches with distortions. It also shows the necessity of the proposed sketch generation and standardization modules. **c)** For model design, the end-to-end model outperforms the two-stage model by a large margin (cas. vs. ours). **d)** For 3D shape represen-

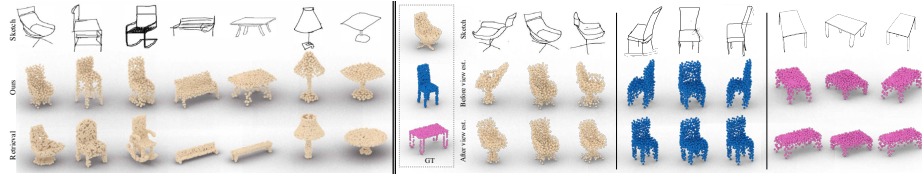


Fig. 6: **Left:** Ours (2nd row) versus nearest-neighbor retrieval results (last row) of given sketches. Our model generalizes to unseen 3D shapes better and has higher geometry fidelity. **Right:** 3D reconstructions of sketches from different viewpoints. Before the view estimation module, the reconstructed 3D shape aligns with the input sketch’s viewpoint. The module transforms the pose of the output 3D shape to align with canonical pose, i.e. the pose of the ground-truth 3D shape.

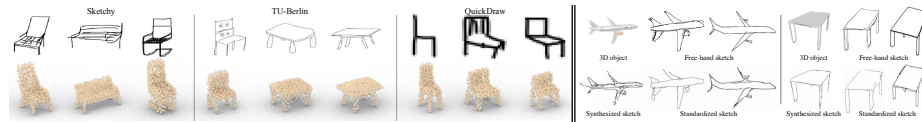


Fig. 7: **Left:** Our approach trained on ShapeNet can be directly applied to other unseen sketch datasets [9, 15, 45] and generalize well. Our model is able to reconstruct 3D shapes from sketches with different styles and line-widths, and even low-resolution data. **Right:** Standardized sketches converted from different individual styles (by different volunteers). For each rendered image of a 3D object, we show free-hand sketches from two volunteers and the standardized sketches from these free-hand sketches. Contents are preserved after the standardization process, and standardized sketches share the style similar to the synthesized ones.

tation, while the voxel representation can reconstruct the general shape well, the fine-grained details are frequently missing due to its low resolution ($32 \times 32 \times 32$). Thus, point clouds outperform voxels. The proposed method also outperforms a previous work that uses depth maps as 3D shape representation [37]. Note that the resolution of voxels can hardly improve much due to the complexity and computational overhead. However, we show that increasing the number of points improves the reconstruction quality (details in supplementary).

Retrieval Results. We compare with nearest-neighbor retrievals, following methods and settings of [49] (Fig.6L). We could generalize to unseen 3D shapes and reconstructs with higher geometry fidelity (e.g., stand of the lamp).

Reconstruction with Different Categories and Views. Fig.5R shows 3D reconstruction results with sketches from different object categories. Our model reconstructs 3D shapes of multiple categories *unconditionally*. There are some failure cases that the model may not handle well.

Fig.6R depicts reconstructions with sketches from different views. Our model can reconstruct 3D shapes from different views even if certain parts are occluded (e.g. legs of the table). Slight variations in details exist for different views.

Evaluation on Other Free-Hand Sketch Datasets. We also evaluate on three other free-hand sketch datasets [9, 15, 45]. Our model can reconstruct 3D shapes from sketches with different styles, line-widths, and levels of distortions even at low resolution (Fig.7L).

Table 2: Ablation studies of standardization and view estimation module. CD is enlarged by 10^4 , and EMD by 10^2 . **(a)** 3D shape reconstruction errors of ablation studies of standardization and view estimation module. Having both standardization and view estimation module gives the highest performance. The lower, the better. **(b)** Reconstruction error with different components of the standardization module: deformation and style translation. Having both parts gives the highest performance. **(c)** The sketch standardization module improves cross-dataset sketch classification accuracy. A ResNet-50 model is trained on TU-Berlin [9] and evaluated Sketchy dataset [45]. The sketch standardization module gives 3 percentage points gain.

(a)			(b)		(c)				
error	no standard.	no view est.	ours	deform.	trans.	CD	EMD	(%)	acc.
CD	92.6	86.8	86.1	×	×	92.6	18.2	w/o std.	75.1
EMD	18.2	16.2	16.0	×	✓	87.2	16.3	w/ std.	78.1
				✓	×	90.1	17.4		
				✓	✓	86.1	16.0		

4.5 Sketch Standardization Module

Visualization. The standardization module can be considered as a domain translation module designed for sketches. We show the standardized sketches of these free-hand sketches and compare them to the synthesized ones Fig.7R. With the standardization module, sketches share a style similar to synthesized sketches which are used as training data. Thus, standardization diminishes the domain gap of sketches with various styles and enhances the generalizability.

Ablation Studies of the Entire Module. The sketch standardization module is introduced to handle various drawing styles of humans. We thus verify this module’s effectiveness on real sketches, both quantitatively (Table 2a) and qualitatively (Fig. 5R). As shown in Table 2a, the reconstruction performance has a significant drop when removing the standardization module. Its effect is also proved in visualizations. In Fig. 5R, we can observe that our full model equipping with the standardization module can produce 3D shapes with higher quality, being more similar to GT shapes, e.g., the airplane and the lamp.

Ablation Studies of Different Components. The standardization module consists of two components: sketch deformation and style translation. We study each module’s performance and report in Table 2b. We observe that the style transformation part improves the reconstruction performance better compared with the deformation part, while having both parts gives the highest performance.

Additional Applications. We show the effectiveness of the proposed sketch standardization with two more applications. The first applications is on cross-dataset sketch classification. We identify the common 98 categories of TU-Berlin sketch dataset [9] and Sketchy dataset [45]. Then we train on TU-Berlin and evaluate on Sketchy. As reported in Table 2c, adding additional sketch standardization module, the classification accuracy improves 3 percentage points.

The second application corresponds to the second example depicted in Fig.4. The target domain is CityScapes dataset [5], where the training data comes from. We extract corresponding edge maps with a deep learning approach [41]

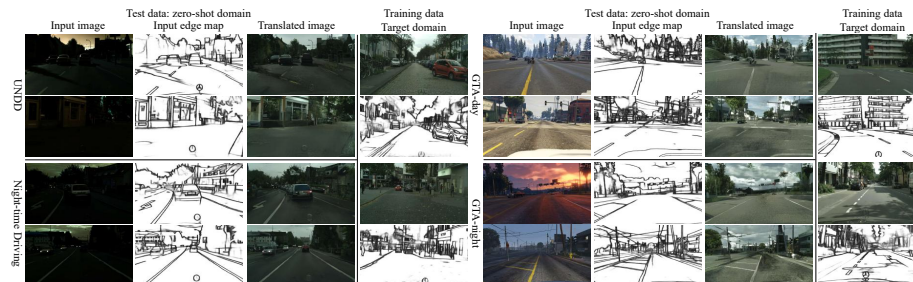


Fig. 8: Zero-shot domain translation results. We aim to translate zero-shot images (**Left**) to the training data domain (**Right**). Specifically, we evaluate the proposed zero-shot domain translation performance on three new datasets: UNDD [34], Night-Time Driving [6] and GTA [44]. The novel domains of night-time and simulated images can be translated to the target domain of daytime and real-world images by leveraging the synthetic edge map domain as a bridge. The target domain is CityScapes dataset [5]. We extract corresponding edge maps and train an image-to-image translation model [21] to translate edge maps to the corresponding RGB images. 1st, 3rd, 5th, 7th rows of column 1 depict some sample training RGB images and 2nd, 4th, 6th, 8th rows of column 1 depict the corresponding edge maps respectively.

and train an image-to-image translation model [21] to translate edge maps to the corresponding RGB images. We evaluate the zero-shot domain translation performance on three new datasets: UNDD [34] (night images), Night-Time Driving [6] (night images) and GTA [44] (synthetic images; screenshots taken from simulated environment). The novel domains of night and simulated images can be translated to the target domain of daytime and real-world images. We visualize the results in Fig.8.

4.6 View Estimation Module

Removing the view estimation module leads to a performance drop of CD and EMD (Table 2a). For qualitative results (Fig.6R), without the 3D rotation, the reconstructed 3D shape has the pose aligned with the input sketch. With the 3D rotation, the 3D shape is aligned to the ground truth’s canonical pose.

5 Summary

We study 3D shape reconstruction from a single-view free-hand sketch. The major novelty is that we use synthesized sketches as training data and introduce a sketch standardization module, in order to tackle the data insufficiency and sketch style variation issues. Extensive experimental results shows that the proposed method is able to successfully reconstruct 3D shapes from single-view free-hand sketches *unconditioned* on viewpoints and categories. The work may unleash more potentials of the sketch in applications such as sketch-based 3D design/games, making them more accessible to the general public.

Acknowledgement We thank Charles R. Qi and Kaichun Mo for the help with the point cloud renderer, as well as Yifei Xing, Haoran Liao, Suifang Mao, Mengyang Zhang, Jiazheng Zhao, Ziyun Zhao, Ruiqing Xiang, Yu Wang, Yan He and Mengyao Lu for providing free-hand sketches of 3D objects.

References

1. Alexiadis, D.S., Chatzitofis, A., Zioulis, N., Zoidi, O., Louizis, G., Zarpalas, D., Daras, P.: An integrated platform for live 3d human reconstruction and motion capturing. *IEEE Transactions on Circuits and Systems for Video Technology* **27**(4), 798–813 (2016)
2. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017)
3. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015)
4. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: *European conference on computer vision*. Springer (2016)
5. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
6. Dai, D., Van Gool, L.: Dark model adaptation: Semantic image segmentation from daytime to nighttime. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. pp. 3819–3824. IEEE (2018)
7. Delanoy, J., Aubry, M., Isola, P., Efros, A.A., Bousseau, A.: 3d sketching using multi-view deep volumetric prediction. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* **1**(1), 1–22 (2018)
8. Dibra, E., Jain, H., Oztireli, C., Ziegler, R., Gross, M.: Human shape from silhouettes using generative hks descriptors and cross-modal neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017)
9. Eitz, M., Hays, J., Alexa, M.: How do humans sketch objects? *ACM Transactions on Graphics* **31**(4), 44:1–44:10 (2012)
10. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017)
11. Fuentes-Pacheco, J., Ruiz-Ascencio, J., Rendón-Mancha, J.M.: Visual simultaneous localization and mapping: a survey. *Artificial intelligence review* **43**(1), 55–81 (2015)
12. Gao, X., Wang, N., Tao, D., Li, X.: Face sketch–photo synthesis and retrieval using sparse representation. *IEEE Transactions on circuits and systems for video technology* **22**(8), 1213–1226 (2012)
13. Ghosh, A., Zhang, R., Dokania, P.K., Wang, O., Efros, A.A., Torr, P.H., Shechtman, E.: Interactive sketch & fill: Multiclass sketch-to-image translation. In: *Proceedings of the IEEE International Conference on Computer Vision* (2019)
14. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems* (2014)
15. Google: The quick, draw! dataset (2017), <https://quickdraw.withgoogle.com/data>
16. Guillard, B., Remelli, E., Yvernay, P., Fua, P.: Sketch2mesh: Reconstructing and editing 3d shapes from sketches. *arXiv preprint arXiv:2104.00482* (2021)
17. Han, Z., Ma, B., Liu, Y.S., Zwicker, M.: Reconstructing 3d shapes from multiple sketches using direct shape optimization. *IEEE Transactions on Image Processing* **29**, 8721–8734 (2020)

18. Häne, C., Tulsiani, S., Malik, J.: Hierarchical surface prediction for 3d object reconstruction. In: 2017 International Conference on 3D Vision. IEEE (2017)
19. He, X., Zhou, Y., Zhou, Z., Bai, S., Bai, X.: Triplet-center loss for multi-view 3d object retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018)
20. Huang, H., Kalogerakis, E., Yumer, E., Mech, R.: Shape synthesis from sketches via procedural models and convolutional networks. *IEEE transactions on visualization and computer graphics* **23**(8), 2003–2013 (2016)
21. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition (2017)
22. Kar, A., Häne, C., Malik, J.: Learning a multi-view stereo machine. In: Advances in neural information processing systems (2017)
23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
24. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
25. Kolotouros, N., Pavlakos, G., Daniilidis, K.: Convolutional mesh regression for single-image human shape reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
26. Lei, J., Song, Y., Peng, B., Ma, Z., Shao, L., Song, Y.Z.: Semi-heterogeneous three-way joint embedding network for sketch-based image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology* **30**(9), 3226–3237 (2019)
27. Li, C., Pan, H., Liu, Y., Tong, X., Sheffer, A., Wang, W.: Bendsketch: modeling freeform surfaces through 2d sketching. *ACM Transactions on Graphics* **36**(4), 1–14 (2017)
28. Li, C., Pan, H., Liu, Y., Tong, X., Sheffer, A., Wang, W.: Robust flow-guided neural prediction for sketch-based freeform surface modeling. *ACM Transactions on Graphics (TOG)* **37**(6), 1–12 (2018)
29. Liu, R., Yu, Q., Yu, S.: An unpaired sketch-to-photo translation model. *arXiv preprint arXiv:1909.08313* (2019)
30. Lun, Z., Gadelha, M., Kalogerakis, E., Maji, S., Wang, R.: 3d shape reconstruction from sketches via multi-view convolutional networks. In: 2017 International Conference on 3D Vision (3DV). pp. 67–77. IEEE (2017)
31. Mandikal, P., Navaneet, K., Agarwal, M., Babu, R.V.: 3d-lmnet: Latent embedding matching for accurate and diverse 3d point cloud reconstruction from a single image. *arXiv preprint arXiv:1807.07796* (2018)
32. Mandikal, P., Radhakrishnan, V.B.: Dense 3d point cloud reconstruction using a deep pyramid network. In: 2019 IEEE Winter Conference on Applications of Computer Vision. IEEE (2019)
33. Mo, K., Guerrero, P., Yi, L., Su, H., Wonka, P., Mitra, N., Guibas, L.J.: Structurenet: Hierarchical graph networks for 3d shape generation. *ACM Transactions on Graphics* **38**(6) (2019)
34. Nag, S., Adak, S., Das, S.: What’s there in the dark. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 2996–3000. IEEE (2019)
35. Navaneet, K., Mandikal, P., Agarwal, M., Babu, R.V.: Capnet: Continuous approximation projection for 3d point cloud reconstruction using 2d supervision. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8819–8826 (2019)
36. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European conference on computer vision. Springer (2016)

37. Nozawa, N., Shum, H.P., Ho, E.S., Morishima, S.: Single sketch image based 3d car shape reconstruction with deep learning and lazy learning. In: VISIGRAPP (1: GRAPP). pp. 179–190 (2020)
38. Özyeşil, O., Voroninski, V., Basri, R., Singer, A.: A survey of structure from motion*. *Acta Numerica* **26**, 305–364 (2017)
39. Pan, J., Han, X., Chen, W., Tang, J., Jia, K.: Deep mesh reconstruction from single rgb images via topology modification networks. In: Proceedings of the IEEE International Conference on Computer Vision (2019)
40. Peng, C., Gao, X., Wang, N., Li, J.: Superpixel-based face sketch–photo synthesis. *IEEE Transactions on Circuits and Systems for Video Technology* **27**(2), 288–299 (2015)
41. Poma, X.S., Riba, E., Sappa, A.: Dense extreme inception network: Towards a robust cnn model for edge detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1923–1932 (2020)
42. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition (2017)
43. Richter, S.R., Roth, S.: Discriminative shape from shading in uncalibrated illumination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2015)
44. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: European conference on computer vision. pp. 102–118. Springer (2016)
45. Sangkloy, P., Burnell, N., Ham, C., Hays, J.: The sketchy database: learning to retrieve badly drawn bunnies. *ACM Transactions on Graphics* **35**(4), 1–12 (2016)
46. Schaefer, S., McPhail, T., Warren, J.: Image deformation using moving least squares. In: ACM SIGGRAPH 2006 Papers, pp. 533–540 (2006)
47. Shen, L., Lin, Z., Huang, Q.: Relay backpropagation for effective learning of deep convolutional neural networks. In: European conference on computer vision. Springer (2016)
48. Tatarchenko, M., Dosovitskiy, A., Brox, T.: Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In: Proceedings of the IEEE International Conference on Computer Vision (2017)
49. Wang, F., Kang, L., Li, Y.: Sketch-based 3d shape retrieval using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2015)
50. Wang, L., Qian, C., Wang, J., Fang, Y.: Unsupervised learning of 3d model reconstruction from hand-drawn sketches. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 1820–1828 (2018)
51. Wang, L., Qian, X., Zhang, X., Hou, X.: Sketch-based image retrieval with multi-clustering re-ranking. *IEEE Transactions on Circuits and Systems for Video Technology* **30**(12), 4929–4943 (2019)
52. Wang, N., Gao, X., Sun, L., Li, J.: Anchored neighborhood index for face sketch synthesis. *IEEE Transactions on Circuits and Systems for Video Technology* **28**(9), 2154–2163 (2017)
53. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: Proceedings of the European Conference on Computer Vision (2018)
54. Witkin, A.P.: Recovering surface shape and orientation from texture. *Artificial intelligence* **17**(1-3), 17–45 (1981)

55. Wu, J., Zhang, C., Xue, T., Freeman, B., Tenenbaum, J.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In: *Advances in neural information processing systems* (2016)
56. Xie, H., Yao, H., Sun, X., Zhou, S., Zhang, S.: Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In: *Proceedings of the IEEE International Conference on Computer Vision* (2019)
57. Xu, B., Chang, W., Sheffer, A., Bousseau, A., McCrae, J., Singh, K.: True2form: 3d curve networks from 2d sketches via selective regularization. *Transactions on Graphics* **33**(4) (2014). <https://doi.org/2601097.2601128>
58. Xu, P., Liu, K., Xiang, T., Hospedales, T.M., Ma, Z., Guo, J., Song, Y.Z.: Fine-grained instance-level sketch-based video retrieval. *IEEE Transactions on Circuits and Systems for Video Technology* **31**(5), 1995–2007 (2020)
59. Yang, B., Rosa, S., Markham, A., Trigoni, N., Wen, H.: Dense 3d object reconstruction from a single depth view. *IEEE transactions on pattern analysis and machine intelligence* **41**(12), 2820–2834 (2018)
60. Yang, G., Huang, X., Hao, Z., Liu, M.Y., Belongie, S., Hariharan, B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In: *Proceedings of the IEEE International Conference on Computer Vision* (2019)
61. Yang, S., Wang, Z., Liu, J., Guo, Z.: Deep plastic surgery: Robust and controllable image editing with human-drawn sketches. *arXiv preprint arXiv:2001.02890* (2020)
62. Yu, Q., Liu, F., Song, Y.Z., Xiang, T., Hospedales, T.M., Loy, C.C.: Sketch me that shoe. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016)
63. Zhang, S., Gao, X., Wang, N., Li, J.: Face sketch synthesis from a single photo-sketch pair. *IEEE Transactions on Circuits and Systems for Video Technology* **27**(2), 275–287 (2015)
64. Zhang, S.H., Guo, Y.C., Gu, Q.W.: Sketch2model: View-aware 3d modeling from single free-hand sketches. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6012–6021 (2021)
65. Zhang, Y., Liu, Z., Liu, T., Peng, B., Li, X.: Realpoint3d: An efficient generation network for 3d object reconstruction from a single image. *IEEE Access* **7**, 57539–57549 (2019)
66. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE international conference on computer vision* (2017)