

A Real World Dataset for Multi-view 3D Reconstruction

Rakesh Shrestha¹, Siqi Hu², Minghao Gou^{2,3}, Ziyuan Liu², Ping Tan^{1,2}
Simon Fraser University¹, Alibaba XR Lab²
Shanghai Jiao Tong University³
{rakeshs,pingtan}@sfu.ca, {husiqixiang, ziyuan-liu}@outlook.com, gmh2015@sjtu.edu.cn

Abstract. We present a dataset of 998 3D models of everyday tabletop objects along with their 847,000 real world RGB and depth images. Accurate annotation of camera pose and object pose for each image is performed in a semi-automated fashion to facilitate the use of the dataset in a myriad 3D applications like shape reconstruction, object pose estimation, shape retrieval *etc.* We primarily focus on learned multi-view 3D reconstruction due to the lack of appropriate real world benchmark for the task and demonstrate that our dataset can fill that gap. The entire annotated dataset along with the source code for the annotation tools and evaluation baselines is available at <http://www.ocrtoc.org/3d-reconstruction.html>.

Keywords: Dataset, Multi-view 3D reconstruction

1 Introduction

Deep learning has shown immense potential in the field of 3D vision in recent years, advancing challenging tasks such as 3D object reconstruction, pose estimation, shape retrieval, robotic grasping etc. But unlike for 2D tasks [10,28,23], large scale real world datasets for 3D object understanding is scarce. Hence, to allow for further advancement of state-of-the-art in 3D object understanding we introduce our dataset which consists of 998 high resolution, textured 3D models of everyday tabletop objects along with their 847K real world RGB-D images. Accurate annotation of camera pose and object pose is performed for each image. Figure 1 shows some sample data from our dataset.

We primarily focus on learned multi-view 3D reconstruction due to the lack of real world datasets for the task. 3D reconstruction methods [15,48,38,43,50] learn to predict 3D model of an object from its color images with known camera and object poses. They require large amount of training examples to be able to generalize to unseen images. While datasets like Pix3D [44], PASCAL3D+[52] and ObjectNet3D [51] provide 3D models and real world images, they are mostly limited to a single image per model.

Existing multi-view 3D reconstruction methods [8,21,38,43,50] rely heavily on synthetic datasets, especially ShapeNet [6], for training and evaluation. There are a few works [25,38] utilizing real world datasets [7], but only for qualitative evaluation purpose, not for training or quantitative evaluation. To remedy this,

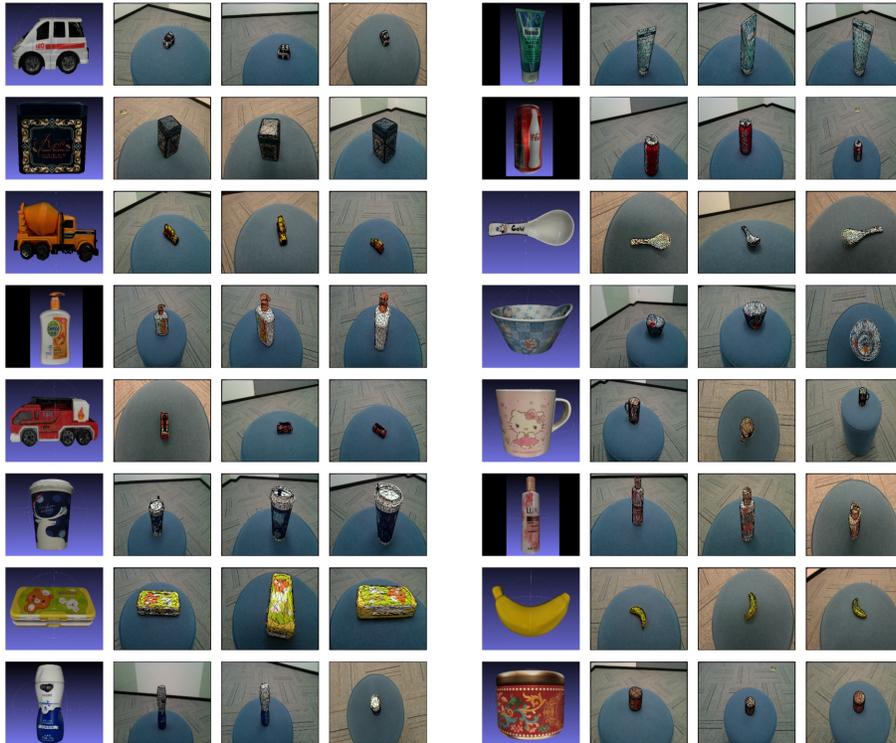


Fig. 1. Sample data from our dataset. From left to right, shown are visualization of textured 3D model, three sample multi-view images with wireframe object model superimposed based on annotated camera and object poses.

we present our dataset and validate its usefulness by performing training as well as qualitative/quantitative evaluation with various state-of-the-art multi-view 3D reconstruction baselines.

The contributions of our work are as follows:

1. To the best our knowledge, our dataset is the first real world dataset that can be used for training and quantitative evaluation of learning-based multi-view 3D reconstruction algorithms.
2. We present two novel methods for automatic/semi-automatic data annotation. We will make the annotation tools publicly available to allow future extensions to the dataset.

2 Related Work

3D Shapes Dataset: Datasets like Princeton shape benchmark [42], FAUST [2], ShapeNet [6] provide a large collection of 3D CAD models of diverse objects,

but without associated real world RGB images. PASCAL3D+[52] and ObjectNet3D [51] performed rough alignment between images from existing datasets and 3D models from online shape repositories. IKEA [27] also performed 2D-3D alignment between existing datasets but with finer alignment results on a smaller set of images and shapes (759 images and 90 shapes). Pix3D [44] extended IKEA to 10K images and 395 shapes through crowdsourcing and scanning some objects manually. These datasets mostly have single-view images associated with the shapes.

Datasets like [24,4,19] have utilized RGB-D sensors to capture relatively small number of objects and are mostly geared towards robot manipulation tasks rather than 3D reconstruction. Knapitsch *et al.* [22] provided a small number of large scale scenes which are suitable for benchmarking traditional Structure-from-Motion (SfM) and Multi-view Stereo (MVS) algorithms rather than learned 3D reconstruction.

The dataset that is closest to ours is Redwood-OS [7]. It provides RGB-D videos of 398 objects and their 3D scene reconstructions. There are several crucial limitations that has prevented widespread adoption of this dataset for multi-view 3D reconstruction though. Firstly, the dataset is not annotated with camera and object pose information. While the camera pose can be obtained using Simultaneous Localization and Mapping (SLAM) or Structure-from-Motion (SfM) techniques [3,11,32,40,41], obtaining accurate object poses is relatively harder. Also, the 3D reconstructions were performed on scene level rather than object level, making it difficult to directly use it for supervision of object reconstruction.

More recently, Objectron [1] and CO3D [37] have provided large scale video sequences of real world objects along with point clouds and object poses but without precise dense 3D models. We aim to tackle the shortcomings of the existing datasets and create a dataset that can effectively serve as a real world benchmark for learning-based multi-view 3D reconstruction models. Table 1 shows the comparison between the relevant datasets.

	Ours	Objectron	CO3D	Redwood-OS	Pix3D	IKEA	PASCAL3D+	ObjectNet3D
Multi-view images	✓	✓	✓	✓	✗	✗	✗	✗
Dense 3D models	✓	✗	✗	✓	✓	✓	✓	✓
Scanned 3D models	✓	✓	✓	✓	*	✗	✗	✗
Object pose annotation	✓	✓	✓	✗	✓	✓	*	*
Textured 3D models	✓	✗	✗	✗	✗	✗	✗	✗

Table 1. Comparison between different datasets. Objectron and CO3D only provide point cloud models of the objects. Pix3D contains a mixture of scanned and CAD 3D models. PASCAL3D+ and ObjectNet3D only have rough object pose annotation, while the annotation is not provided in Redwood-OS. Only our dataset provides precisely scanned texture-mapped 3D models that are further registered to multi-view RGB images.

3D Reconstruction: The methods in [15,16,34,45,48,54] predict 3D models from single-view color images. Since a single-view image can only provide a limited coverage of a target object, multi-view input is preferred in many applications. SLAM and Structure-from-Motion methods [3,11,32,40,41] are popular ways of performing 3D reconstruction but they struggle with poorly textured and non-Lambertian surfaces and require careful input view selection. Deep learning has emerged as a potential solution to tackle these issues. Early works like [8,17,21] used Recurrent Neural Networks (RNN) to perform multi-view 3D reconstruction. Pixel2Mesh++ [50] introduced cross-view perceptual feature pooling and multi-view deformation reasoning to refine an initial shape. MeshMVS [43] predicted a coarse volume from Multi-view Stereo depths first and then applied deformations on it to get a finer shape. All of these works were trained and evaluated exclusively on synthetic datasets due to the lack of proper real world datasets.

Some recent works like DVR [33], IDR [55], Neus [49], Geo-Neus [13] have focused on unsupervised 3D reconstruction with expensive per-scene optimization for each object. These methods encode each scene into separate Multi-layer Perceptron (MLP) that implicitly represents the scene as Signed Distance Function (SDF) or Occupancy Field. These works have obtained impressive results on small scale datasets of real world objects [20,53]. Our dataset can be further applied to evaluate these methods quantitatively on a much larger scale dataset.

3 Data Acquisition

Our data acquisition takes place in two steps. First, a detailed and textured 3D model of an object is generated using Shining3D[®] EinScan-SE 3D scanner. The scanner uses a calibrated turntable, a 1.3 Megapixel camera and visible light sources to obtain the 3D model of an object. Then, an Intel[®] RealSense™ LiDAR Camera L515 is used to record a RGB-D video sequence of the object on a round ottoman chair, capturing 360° view around the object. The video is recorded at 30 frames per second in HD resolution (1280×720). Figure 1 shows a number of 3D models and some sample color images from our dataset.

Datasets like [7,24] perform 3D model generation and video recording in one step by reconstructing the 3D scene captured by the images. The quality of the 3D models generated this way depends heavily on the trajectory of the camera and requires some level of expertise for data collection. Furthermore, these datasets use consumer grade cameras which cannot reconstruct fine details in the 3D geometry. We therefore use specialized hardware designed for high quality 3D scanning.

Another approach is to utilize 3D CAD models from online repositories and match them with real world 2D images, which are also mostly collected online [9,27,51,52]. The downside of this approach is that it is difficult to ensure exact instance-level match between 3D models and 2D images. According to a survey conducted by Sun *et al.* [44], test subjects reported that only a small fraction of the images matched the corresponding shapes in datasets [51,52].

4 Data Annotation

The most challenging aspect of creating a large scale real world dataset for object reconstruction is generating ground truth annotations. Most learning-based 3D reconstruction methods require accurate camera poses as well as consistent object poses in the camera coordinate frame. While it is fairly easy to obtain the camera poses, obtaining accurate object poses is more challenging.

The methods in [44,52] perform object pose estimation by manually annotating corresponding keypoints in the 3D models and 2D images, and then performing 2D-3D alignment with the Perspective-n-Point (PnP) [14,26] and Levenberg-Marquardt algorithms [31]. Note that these datasets mostly contain a single image for each 3D model, which makes this kind of annotation feasible. In comparison, we aim to do this for video sequences with up to 1000 images, which could be manual intensive. Additionally, estimating object pose that is consistent over multi-view images will require keypoint matches at sub-pixel accuracy which is impossible by manual annotation.

On the other hand, the methods in [9,51] manually annotate the object pose directly by either trying to align the 3D model with the scene reconstruction [9] or the re-projected 3D model with 2D image[51]. We found these techniques to be inadequate for producing multi-view consistent object poses and therefore develop our own annotation systems.

4.1 Notations

We represent an object pose by $\xi \in SE(3)$ where $SE(3)$ is the 3D Special Euclidean Lie group [47] of 4×4 rigid body transformation matrix:

$$\xi = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \quad (1)$$

where R is the 3×3 rotation matrix and t is the 3D translation vector.

We define object pose ${}_{\text{w}}\xi_{\text{obj}}$ as the transformation from canonical object frame (obj) to world frame (w). Similarly, the pose of the i^{th} camera ${}_{\text{w}}\xi_{\text{cam}_i}$ represents the transformation from camera to world frame. The canonical object frame is centered at the object with z-axis pointing upwards along the gravity direction while the world frame is arbitrary (*e.g.* pose of the first camera).

We use pinhole camera model with camera intrinsics matrix K :

$$K = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \quad (2)$$

where f_x and f_y are focal lengths and c_x and c_y principal points. These parameters are provided by the camera manufacturers.

The image coordinates p of a 3D point P_w in homogeneous world coordinate can be computed as:

$$p = K [R_i^T \quad -R_i^T t_i] P_w \quad (3)$$

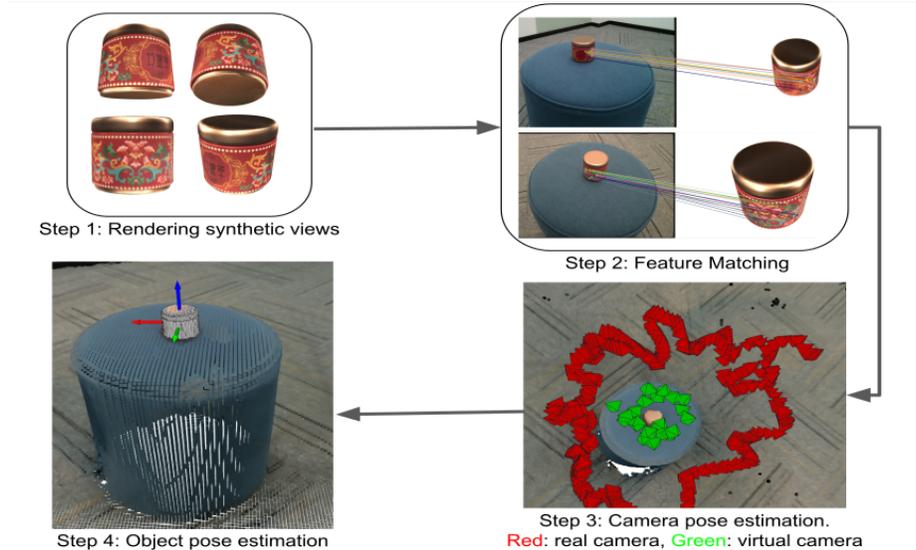


Fig. 2. Texture-rich Object Annotation. *Step 1:* Synthetic views of the 3D model are rendered. *Step 2:* Feature matching is performed between/ across real and synthetic images. *Step 3:* Pose of the real and virtual cameras are estimated. *Step 4:* Object pose is estimated by 7-DOF alignment between estimated and ground truth virtual camera poses.

where R_i and t_i are the rotation and translation components of the camera pose.

The images taken from our RGB-D camera suffer from radial and tangential distortion. But for the purpose of annotation, we undistort the images so that the pinhole camera model holds.

We now present two methods for annotating our dataset depending on the texture-richness of the object being scanned: **Texture-rich Object Annotation** and **Textureless Object Annotation**.

4.2 Texture-rich Object Annotation

Since our 3D models have high-fidelity textures from our 3D scanner, we can utilize it to annotate the object pose in the recorded video sequence. We perform joint camera and object pose estimation by matching keypoints between images and 3D model to ensure camera and object pose consistency over multiple views. Figure 2 illustrates the annotation process. Following are the steps involved:

i. Rendering synthetic views of a 3D model: Instead of directly matching keypoints between a 3D model and 2D images, we instead render synthetic views of the 3D model and perform 2D keypoint matching. We use the physically based rendering engine, Pyrender [29], to render synthetic views. This allows us to utilize robust keypoint matching algorithms developed for RGB images. The

virtual camera poses for rendering are randomly sampled around the object by varying the camera distance, and azimuth/elevation angles with respect to the object. We verify the quality of each rendered image by checking if there are sufficient keypoint matches against the real images. 150 images are rendered for each object model.

ii. Feature matching: We perform exhaustive feature matching across as well as within the real and synthetic images using neural network based feature matching technique SuperGlue [39].

iii. Camera pose estimation: Given the keypoint matches, we estimate the camera poses of both the real and virtual cameras in the same world coordinate frame using the SfM tool COLMAP [40,41].

iv. Object pose estimation: Let $\{\hat{\xi}_i \mid i = 1, \dots, 150\}$ be the ground truth poses of the virtual cameras in object frame (we keep track of the ground truth poses during the rendering step). Let $\{\xi_i \mid i = 1, \dots, 150\}$ be the corresponding poses estimated by COLMAP in world frame. By aligning $\{\xi_i\}$ and $\{\hat{\xi}_i\}$ we can estimate the object pose. We use the Kabsch-Umeyama algorithm [46] under Random Sample Consensus (RANSAC) [5] scheme to perform a 7-DOF (pose + scale) alignment. Since COLMAP only uses 2D image information, its poses have arbitrary scale; hence we perform a 7-DOF alignment instead of 6-DOF to obtain metric scale. After applying the Kabsch-Umeyama algorithm we get 7-DOF transformation S in $Sim(3)$ Lie Group parameterized as:

$$S = \begin{bmatrix} sR_s & t_s \\ 0 & 1 \end{bmatrix} \quad (4)$$

The camera poses from COLMAP can then be transformed to metric scale pose:

$${}_{\text{w}}\xi_{\text{cam}_i} = \begin{bmatrix} R_s R_i & sR_s t_i + t_s \\ 0 & 1 \end{bmatrix} \quad (5)$$

where R_i and t_i are the rotation and translation component of the camera poses from COLMAP.

Since the ground truth virtual camera poses $\{\hat{\xi}_i \mid i = 1, \dots, 150\}$ are in object frame, the transformation in Equation (5) will lead to camera poses in object frame *i.e.* ${}_{\text{w}}\xi_{\text{obj}} = \mathbb{I}$ where \mathbb{I} is the 4×4 identity matrix.

4.3 Textureless Object Annotation

While the pipeline outlined in Sub-section 4.2 can accurately annotate texture-rich objects, it will fail for textureless objects since correct feature matches among the images cannot be established. To tackle this problem we develop another annotation system shown in Figure 3 that can handle objects lacking good textures which consists of the following steps:

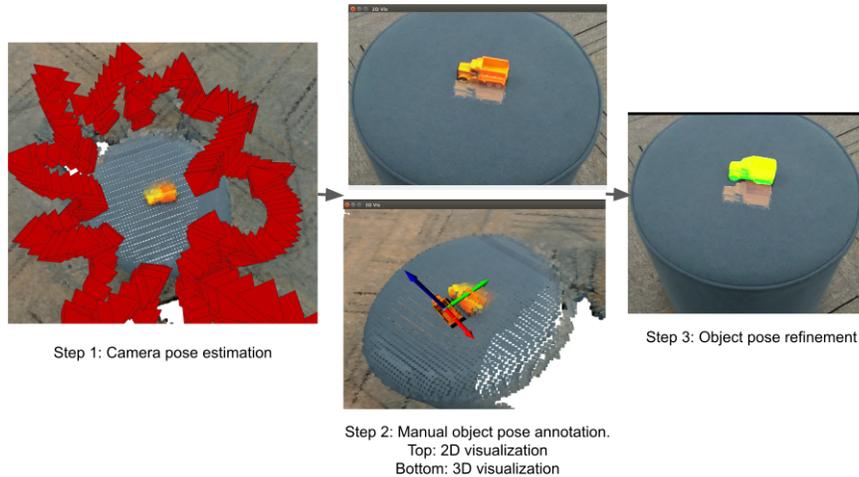


Fig. 3. Textureless Object Annotation. *Step 1:* Camera pose annotation (+ dense scene reconstruction). *Step 2:* Manual annotation of rough object pose where a transparent projection of the object model is superimposed over an RGB image for 2D visualization (top) and the 3D object is placed alongside the dense scene reconstruction for 3D visualization (bottom). *Step 3:* Object pose is refined such that the object projection overlaps with the ground truth mask (green).

i. Camera pose estimation: Even when the object being scanned is textureless, our background has sufficient textures to allow successful camera pose estimation. We therefore utilize the RGB-D version of ORB-SLAM2 [32] to obtain the camera poses $\{w\xi_{\text{cam}_i}\}$. Since it uses depth information alongside RGB, the poses are in metric scale.

ii. Manual annotation of rough object pose: We create an annotation interface as shown in Step 2 of Figure 3 to estimate the rough object pose. To facilitate the annotation, we reconstruct the 3D scene using the RGB-D images and camera poses estimated in the previous step by employing Truncated Signed Distance Function (TSDF) fusion [56]. The object pose $w\xi_{\text{obj}}$ is initialized to be a fixed distance in front of the first camera and the z-axis is aligned with the principle axis of the 3D scene found using Principal Component Analysis (PCA). An annotator can then update the 3 translation and 3 Euler angle (roll-pitch-yaw) components of the 6D object pose using keyboard to align the object model with the scene. In addition to the 3D scene, we also show the projection of the object model over an RGB image. The RGB image can be changed to verify the consistency of the object pose over multiple views.

iii. Object pose refinement: We find that obtaining accurate object pose through manual annotation is difficult, so we refine it further by aligning the projection of the 3D object model with ground truth object masks in different

images. The ground truth object masks are obtained from Cascade Mask R-CNN [18] with a 152-layer ResNetXt backbone pretrained on ImageNet.

Let ${}_{\text{w}}\xi_{\text{obj}}$ be the rough object pose from manual annotation and ${}_{\text{w}}\xi_{\text{cam}_i}$ be the pose of the i^{th} camera. The camera-centric object pose is represented as follows:

$$\xi = {}_{\text{cam}_i}\xi_{\text{obj}} = ({}_{\text{w}}\xi_{\text{cam}_i})^{-1} \times {}_{\text{w}}\xi_{\text{obj}} \quad (6)$$

The transformation $\xi \in SE(3)$ is used to differentiably render [36] the object model onto the image of camera i to obtain the rendered object mask by applying the projection model of Equation (3). Since direct optimization in the manifold space $SE(3)$ is not possible, we instead optimize the linearized increment of the manifold around ξ . This is a common technique in SLAM and Visual Odometry [11,32].

Let $\delta\xi \in \mathfrak{se}(3)$ represent the linearized increment of ξ belonging to the Lie algebra $\mathfrak{se}(3)$ corresponding to Lie Group $SE(3)$ [47]. The updated object pose is given by:

$$\xi' = \xi \times \exp(\delta\xi) \quad (7)$$

Here, \exp represents the exponential map that transforms $\mathfrak{se}(3)$ to $SE(3)$. The object pose w.r.t. world frame can also be updated by right multiplication of the initial pose with $\exp(\delta\xi)$.

We can optimize $\delta\xi$ in order to increase the overlap between the rendered mask M at ξ' and ground truth mask \hat{M} using least-squares minimization of the mask loss:

$$\mathcal{L}_{\text{mask}} = \text{mean}(\|M \ominus \hat{M}\|_2) \quad (8)$$

where \ominus represents element-wise subtraction.

The optimization is performed using stochastic gradient descent for each camera for 30 iterations in PyTorch [35] library. Since $\delta\xi \in \mathfrak{se}(3)$ cannot represent large changes in pose, we update the pose $\xi \leftarrow \xi'$ every 30 iterations and relinearize $\delta\xi$ around the new ξ .

5 Dataset Statistics

We collected in total 998 objects. It typically takes about 20 minutes to scan the 3D model of an object and record a video, but about 2 hours to register the scanned 3D model to all the video frames. Table 2 shows the category distribution of objects in our dataset along with the method used to annotate the object (texture-rich vs textureless). Each category in our dataset contains 39-115 objects, with average 67 objects per category. A majority of the objects (89%) were annotated using texture-rich pipeline which requires no user input. Table 3 shows the distribution of images over the categories. We have on average 56K images for each category.

Category	Bottle	Bowl	Cleanser	Cup	Eating Utensils	Box	Plate	Toy Animal	Toy Car	Toy Fruit	Toy Aircraft	Toy Boat	Toy Food	Toy Figure	Misc	Total
Texture-rich Annotation	60	61	51	39	27	83	45	69	115	12	61	39	82	95	51	890
Textureless Annotation	0	11	0	12	14	0	6	32	0	33	0	0	0	0	0	108
Total	60	72	51	54	41	83	51	101	115	45	61	39	82	95	51	998

Table 2. Annotation statistics.

Bottle	Bowl	Cleanser	Cup	Eating Utensils	Box	Plate	Toy Animal	Toy Car	Toy Fruit	Toy Aircraft	Toy Boat	Toy Food	Toy Figure	Misc	Total
54K	61K	44K	45K	33K	68K	45K	82K	104K	38K	51K	32K	69K	78K	43K	849K

Table 3. Image distribution over the categories. Number of images in each category has been rounded to nearest 1000.

6 Evaluation

To verify the usefulness of our dataset, we train and evaluate state-of-the-art multi-view 3D reconstruction baselines exclusively on our dataset. From each object, we randomly sample 100 different 3-view image tuples as the multi-view inputs. To ensure fair evaluation and avoid overfitting we split our dataset into training, testing and validation sets in approximately 70%-20%-10% ratio. The train-test-validation split is performed such that the distribution in each object category is also 70%-20%-10%. Only the data in training set is used to fit the baseline models while validation set is used to decide when to save the model parameters during training (known as checkpointing). All the evaluation results presented here are on the test set entirely held out during the training process.

6.1 Experiments

We evaluate our datasets with several recent learning-based 3D reconstruction baseline methods, including Multi-view Pixel2Mesh (MVP2M) [50], Pixel2Mesh++ (P2M++) [50], Multi-view extension of Mesh R-CNN [15] (MV M-RCNN) provided by [43], MeshMVS [43], DVR [33], IDR [55] and COLMAP [40,41]. We use the ‘Sphere-Init’ version of Mesh R-CNN and ‘Back-projected depth’ version of MeshMVS.

MVP2M pools multi-view image features and uses it to deform an initial ellipsoid to the desired shape. Pixel2Mesh++ deforms the mesh predicted by MVP2M by taking the weighted sum of deformation hypothesis sampled near the MVP2M mesh vertices. MV M-RCNN improves on MVP2M with a deeper backbone, better training recipe and higher resolution initial shape.

MeshMVS first predicts depth images using Multi-view Stereo and uses the depths to obtain a coarse shape which is deformed using similar techniques as MVP2M and MV MR-CNN. To train the depth prediction network of MeshMVS, we use depths rendered from the 3D object models since the recorded

depth can be inaccurate or altogether missing at close distances. We also evaluate the baseline MeshMVS (RGB-D) which uses ground truth depths instead of predicted depths to obtain the coarse shape, essentially performing shape completion instead of prediction.

We also include per-scene optimized baselines DVR, IDR and COLMAP which do not require training generalizable priors with 3D supervision. DVR and IDR perform NeRF [30] like optimization to learn 3D models from images using implicit neural representation. COLMAP performs Structure-from-Motion (SfM) to first generate sparse point cloud which are further densified using Patch Match Stereo algorithm [41]. These methods require larger number of images to produce satisfactory results, hence we use 64 input images. Since the time required to reconstruct a scene is large for these methods, we evaluate these methods only on 30 scenes from the test set - 2 from each category.

All of the baselines require the object in the images to be segmented out of the background. We do this by rendering the 2D image masks of 3D object models using the annotated camera/object pose. Also, we transform the images to the size and intrinsics (Equation (2)) required by the baselines before training/testing.

Metrics: We follow recent works [15,43,50] and choose F1-score (harmonic mean of precision and recall) at a thresholds $\tau = 0.3$ as our evaluation metric. Precision in this context is defined as the fraction of points in predicted model within τ distance from the ground truth points while recall is the fraction of point in ground truth model within τ distance from the predicted points.

We also report Chamfer Distance between a predicted model P and ground truth model Q which measures the mean distance between the closest pairs of points $A_{P,Q} = \{(p, \arg \min_q \|p - q\|) : p \in P, q \in Q\}$ in the two models:

$$\mathcal{L}_{\text{chamfer}}(P, Q) = |P|^{-1} \sum_{(p,q) \in A_{P,Q}} \|p - q\|^2 + |Q|^{-1} \sum_{(q,p) \in A_{Q,P}} \|q - p\|^2 \quad (9)$$

We uniformly sample 10k points from predicted and ground truth meshes to evaluate these metrics. Following [12,15], we rescale the 3D models so that the longest edge of the ground truth mesh bounding box has length 10.

Results: The quantitative comparison results of different learning-based 3D reconstruction baselines on our dataset are presented in Table 4. Note that both training and testing set contain objects from all categories, but test F1-score on individual categories as well as over all categories are reported here. Figure 4 visualizes the shapes generated by different methods for qualitative evaluation.

We can see that overall Pixel2Mesh++ performs the best (barring MeshMVS RGB-D). This is contrary to the results on ShapeNet reported in [43] where MeshMVS performs the best. This can be attributed to the high depth prediction error of MeshMVS (average depth error is $\sim 6\%$ of the total depth range). When predicted depth is replaced with ground truth depth, we indeed see a significant improvement in the performance of MeshMVS indicating that depth prediction is the main bottleneck in its performance.

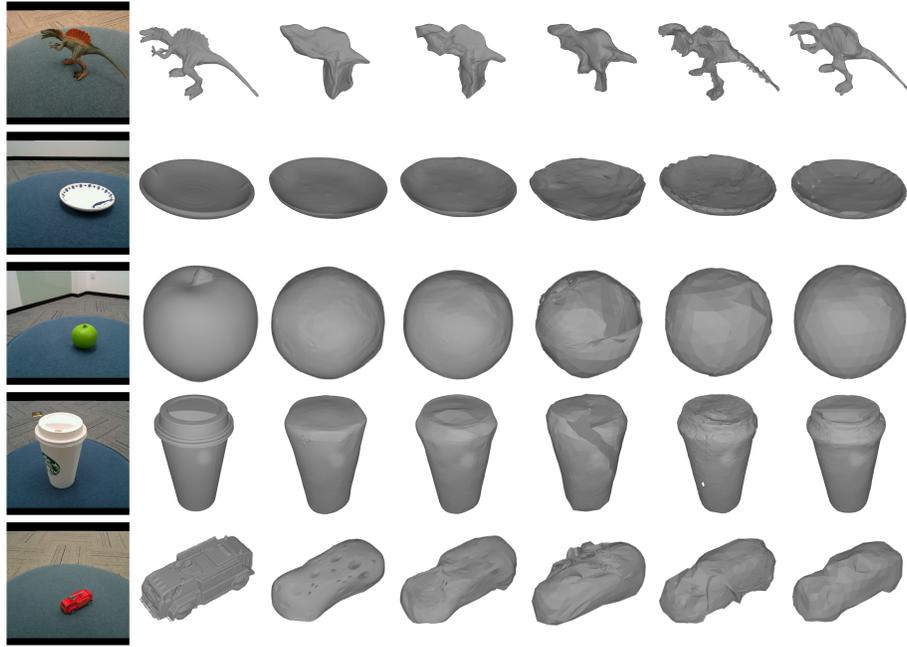


Fig. 4. Qualitative Evaluation. Left to right, shown are an input image, ground truth mesh, results from MVP2M, P2M++, MV M-RCNN, MeshMVS, and MeshMVS (RGB-D) respectively.

Table 5 shows the quantitative comparison between different unsupervised, per-scene optimized baselines. Here, IDR outperforms the other two baselines which is in line with the results presented in [55] on the DTU dataset. COLMAP performs worse than the rest because the textures on most of the objects are insufficient for dense reconstruction using Patch Match stereo leading to sparse and noisy results (Figure 5).

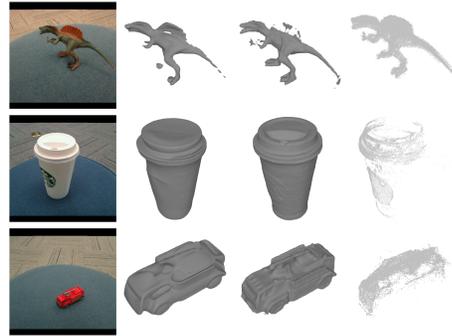


Fig. 5. Qualitative evaluation. Left to right: Input image, results from DVR, IDR, COLMAP.

Single category training: We compare the difference in the performance when each category is trained and evaluated separately. In this case, there will be a different set of model parameters for each category. For these experiments we sample 200 different 3-view images as inputs from each scene. Table 6 shows the results for MV M-RCNN baseline when each category is trained separately versus when all are trained together. We see that the performance

Category	F1@0.3 \uparrow					Chamfer \downarrow				
	MVP2M	P2M++	MV M-RCNN	MeshMVS	MeshMVS (RGB-D)	MVP2M	P2M++	MV M-RCNN	MeshMVS	MeshMVS (RGB-D)
Bottle	74.86	81.72	67.12	55.35	94.11	0.23	0.19	1.78	0.53	0.06
Bowl	74.49	80.60	70.94	59.69	91.88	0.31	0.23	1.41	0.46	0.09
Box	61.32	72.24	67.90	59.28	91.04	0.39	0.29	0.84	0.48	0.10
Cleanser	68.69	80.90	72.37	64.01	95.97	0.30	0.19	0.96	0.39	0.04
Cup	65.31	73.67	66.00	52.41	87.40	0.39	0.30	1.59	0.64	0.12
Eating Utensils	74.34	88.40	81.67	72.74	96.08	0.24	0.12	0.84	0.26	0.04
Plate	76.13	81.65	66.69	72.03	82.70	0.34	0.26	3.42	0.27	0.16
Toy Boat	55.77	65.57	57.95	59.58	88.19	0.46	0.37	6.54	5.22	5.79
Toy Aircraft	52.85	65.31	56.91	64.94	91.13	0.73	0.55	2.91	0.34	0.08
Toy Animals	49.46	68.12	64.07	59.89	93.53	0.81	0.51	0.94	0.42	0.06
Toy Food	60.92	71.24	61.08	49.12	90.16	0.35	0.25	2.78	0.59	0.10
Toy Fruit	56.78	72.85	59.84	41.28	88.86	0.55	0.37	3.77	0.87	0.11
Miscellaneous	56.54	69.40	66.34	61.57	91.40	0.63	0.46	0.96	0.47	0.11
Toy Car	59.32	71.33	63.90	57.65	88.28	0.37	0.25	1.47	0.47	0.10
Toy Figure	50.42	68.32	57.23	58.63	91.39	0.70	0.46	5.12	0.52	0.08
All	61.25	73.30	64.77	58.79	90.85	0.47	0.33	2.24	0.65	0.29

Table 4. Quantitative comparison of state-of-the-art learning-based multi-view 3D reconstruction methods on our dataset. We report F1-score and Chamfer Distance on each semantic category as well as over all categories. The baseline MeshMVS (RGB-D) is not considered for highlighting the best performance since it uses ground truth depth as additional input.

is generally better when using all categories, showing that 3D reconstruction models can learn to generalize over multiple categories in our dataset.

7 Discussion

The results presented in Tables 4 and 6 as well as the qualitative evaluation of Figure 4 show that the problem of generalizable multi-view 3D reconstruction is far from solved. While works like Pixel2Mesh++, Mesh R-CNN and MeshMVS have offered promising avenues for advancement of the state-of-the-art, more research is still needed in this direction. Table 5 and Figure 5 shows the limitations of traditional 3D reconstruction methods like COLMAP. While more recent NeRF-based methods like DVR and IDR generates high quality reconstruction, their running time is at the order of 10 hours in general and requires a larger number of input images (64 in our case). We hope that our dataset can serve as a challenging benchmark for these problems; aiding and inspiring future work in 3D shape generation.

8 Conclusion

We present a large scale dataset of 3D models and their real world multi-view images. Two methods were developed for annotation of the dataset which can provide high accuracy camera and object poses. Experiments show that our dataset can be used for training and evaluating multi-view 3D reconstruction methods, something that has been lacking in existing real world datasets.

Category	F1@0.3 \uparrow			Chamfer \downarrow		
	DVR	IDR	COLMAP	DVR	IDR	COLMAP
Bottle	92.95	91.90	25.80	0.11	0.09	1.36
Bowl	69.20	83.50	14.72	0.60	0.62	1.81
Box	75.97	78.95	34.64	0.70	0.22	1.66
Cleanser	86.61	98.32	45.54	0.09	0.02	1.38
Cup	70.77	78.69	14.86	0.64	0.68	3.23
Eating utensils	87.93	95.12	41.34	0.08	0.04	1.10
Plate	62.46	75.01	76.98	0.50	0.45	0.14
Toy Boat	86.71	99.87	45.66	0.08	0.01	0.37
Toy Aerocraft	78.48	99.07	86.37	0.24	0.02	0.09
Toy Animals	85.26	90.61	53.31	0.19	0.11	0.67
Toy Food	86.06	89.59	34.74	0.12	0.08	0.96
Toy Fruit	63.27	90.25	0.55	0.75	0.09	40.56
Miscellaneous	89.21	89.31	45.96	0.10	0.09	0.66
Toy Car	76.25	94.26	27.24	0.24	0.06	1.32
Toy Figure	82.25	97.60	49.39	0.16	0.03	0.32
All	79.56	89.54	39.81	0.30	0.17	3.71

Table 5. Quantitative comparison of state-of-the-art NeRF-based 3D reconstruction methods along with COLMAP on our dataset. We report F1-score and Chamfer Distance on each semantic category as well as over all categories.

Category	F1@0.3 \uparrow		Chamfer \downarrow	
	All	Single	All	Single
Bottle	67.12	64.02	1.78	3.33
Bowl	70.94	53.41	1.41	27.37
Box	67.90	65.81	0.84	1.85
Cleanser	72.37	73.00	0.96	1.21
Cup	66.00	61.95	1.59	1.61
Eating utensils	81.67	77.06	0.84	0.80
Plate	66.69	62.15	3.42	50.72
Toy Boat	57.95	55.39	6.54	10.21
Toy Aerocraft	56.91	43.51	2.91	5.16
Toy Animals	64.07	62.32	0.94	1.16
Toy Food	61.08	62.02	2.78	4.92
Toy Fruit	59.84	20.54	3.77	67.82
Miscellaneous	66.34	44.29	0.96	3.71
Toy Car	63.90	65.18	1.47	1.26
Toy Figure	57.23	50.81	5.12	9.0
Mean	64.77	58.15	2.24	10.57

Table 6. Single Vs All Category Training evaluation on MV M-RCNN baseline.

References

1. Ahmadyan, A., Zhang, L., Ablavatski, A., Wei, J., Grundmann, M.: Objectron: A large scale dataset of object-centric videos in the wild with pose annotations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7822–7831 (2021)
2. Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: Dataset and evaluation for 3d mesh registration. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3794–3801 (2014)
3. Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., Leonard, J.J.: Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics* **32**(6), 1309–1332 (2016)
4. Calli, B., Walsman, A., Singh, A., Srinivasa, S., Abbeel, P., Dollar, A.M.: Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set. *IEEE Robotics & Automation Magazine* **22**(3), 36–52 (2015)
5. Cantzler, H.: Random sample consensus (ransac). Institute for Perception, Action and Behaviour, Division of Informatics, University of Edinburgh (1981)
6. 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)
7. Choi, S., Zhou, Q.Y., Miller, S., Koltun, V.: A large dataset of object scans. arXiv preprint arXiv:1602.02481 (2016)
8. 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. pp. 628–644. Springer (2016)
9. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
11. Engel, J., Schöps, T., Cremers, D.: Lsd-slam: Large-scale direct monocular slam. In: European conference on computer vision. pp. 834–849. Springer (2014)
12. Fouhey, D.F., Gupta, A., Hebert, M.: Data-driven 3d primitives for single image understanding. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3392–3399 (2013)
13. Fu, Q., Xu, Q., Ong, Y.S., Tao, W.: Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. arXiv preprint arXiv:2205.15848 (2022)
14. Gao, X.S., Hou, X.R., Tang, J., Cheng, H.F.: Complete solution classification for the perspective-three-point problem. *IEEE transactions on pattern analysis and machine intelligence* **25**(8), 930–943 (2003)
15. Gkioxari, G., Malik, J., Johnson, J.: Mesh r-cnn. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9785–9795 (2019)
16. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: A papier-mâché approach to learning 3d surface generation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 216–224 (2018)
17. Gwak, J., Choy, C.B., Chandraker, M., Garg, A., Savarese, S.: Weakly supervised 3d reconstruction with adversarial constraint. In: 2017 International Conference on 3D Vision (3DV). pp. 263–272. IEEE (2017)

18. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
19. Hodan, T., Haluza, P., Obdržálek, Š., Matas, J., Lourakis, M., Zabulis, X.: T-less: An rgb-d dataset for 6d pose estimation of texture-less objects. In: 2017 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 880–888. IEEE (2017)
20. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanæs, H.: Large scale multi-view stereopsis evaluation. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition. pp. 406–413. IEEE (2014)
21. Kar, A., Häne, C., Malik, J.: Learning a multi-view stereo machine. *Advances in neural information processing systems* **30** (2017)
22. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* **36**(4), 1–13 (2017)
23. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al.: The open images dataset v4. *International Journal of Computer Vision* **128**(7), 1956–1981 (2020)
24. Lai, K., Bo, L., Ren, X., Fox, D.: A large-scale hierarchical multi-view rgb-d object dataset. In: 2011 IEEE International Conference on Robotics and Automation. pp. 1817–1824 (2011). <https://doi.org/10.1109/ICRA.2011.5980382>
25. Lei, J., Sridhar, S., Guerrero, P., Sung, M., Mitra, N., Guibas, L.J.: Pix2surf: Learning parametric 3d surface models of objects from images. In: European Conference on Computer Vision. pp. 121–138. Springer (2020)
26. Lepetit, V., Moreno-Noguer, F., Fua, P.: Epn: An accurate o (n) solution to the pnp problem. *International journal of computer vision* **81**(2), 155–166 (2009)
27. Lim, J.J., Pirsivash, H., Torralba, A.: Parsing ikea objects: Fine pose estimation. In: 2013 IEEE International Conference on Computer Vision. pp. 2992–2999 (2013). <https://doi.org/10.1109/ICCV.2013.372>
28. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
29. Matl, M.: Pyrender. <https://github.com/mmatl/pyrender> (2019)
30. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European conference on computer vision. pp. 405–421. Springer (2020)
31. Moré, J.J.: The levenberg-marquardt algorithm: implementation and theory. In: *Numerical analysis*, pp. 105–116. Springer (1978)
32. Mur-Artal, R., Montiel, J.M.M., Tardos, J.D.: Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics* **31**(5), 1147–1163 (2015)
33. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3504–3515 (2020)
34. 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. pp. 9964–9973 (2019)
35. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: *Advances*

- in *Neural Information Processing Systems* 32, pp. 8024–8035. Curran Associates, Inc. (2019)
36. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv:2007.08501 (2020)
 37. Reizenstein, J., Shapovalov, R., Henzler, P., Sbordone, L., Labatut, P., Novotny, D.: Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In: *International Conference on Computer Vision* (2021)
 38. Runz, M., Li, K., Tang, M., Ma, L., Kong, C., Schmidt, T., Reid, I., Agapito, L., Straub, J., Lovegrove, S., et al.: Frodo: From detections to 3d objects. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14720–14729 (2020)
 39. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4938–4947 (2020)
 40. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
 41. Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: *European Conference on Computer Vision (ECCV)* (2016)
 42. Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The princeton shape benchmark. In: *Proceedings Shape Modeling Applications, 2004*. pp. 167–178. IEEE (2004)
 43. Shrestha, R., Fan, Z., Su, Q., Dai, Z., Zhu, S., Tan, P.: Meshmvs: Multi-view stereo guided mesh reconstruction. In: *2021 International Conference on 3D Vision (3DV)*. pp. 1290–1300. IEEE (2021)
 44. Sun, X., Wu, J., Zhang, X., Zhang, Z., Zhang, C., Xue, T., Tenenbaum, J.B., Freeman, W.T.: Pix3d: Dataset and methods for single-image 3d shape modeling. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2974–2983 (2018)
 45. Tang, J., Han, X., Pan, J., Jia, K., Tong, X.: A skeleton-bridged deep learning approach for generating meshes of complex topologies from single rgb images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4541–4550 (2019)
 46. Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(4), 376–380 (1991). <https://doi.org/10.1109/34.88573>
 47. Varadarajan, V.S.: *Lie groups, Lie algebras, and their representations*, vol. 102. Springer Science & Business Media (2013)
 48. 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 (ECCV)*. pp. 52–67 (2018)
 49. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems* **34**, 27171–27183 (2021)
 50. Wen, C., Zhang, Y., Li, Z., Fu, Y.: Pixel2mesh++: Multi-view 3d mesh generation via deformation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1042–1051 (2019)
 51. Xiang, Y., Kim, W., Chen, W., Ji, J., Choy, C., Su, H., Mottaghi, R., Guibas, L., Savarese, S.: Objectnet3d: A large scale database for 3d object recognition. In: *European conference on computer vision*. pp. 160–176. Springer (2016)

52. Xiang, Y., Mottaghi, R., Savarese, S.: Beyond pascal: A benchmark for 3d object detection in the wild. In: IEEE winter conference on applications of computer vision. pp. 75–82. IEEE (2014)
53. Yao, Y., Luo, Z., Li, S., Zhang, J., Ren, Y., Zhou, L., Fang, T., Quan, L.: Blended-mvs: A large-scale dataset for generalized multi-view stereo networks. *Computer Vision and Pattern Recognition (CVPR)* (2020)
54. Yao, Y., Schertler, N., Rosales, E., Rhodin, H., Sigal, L., Sheffer, A.: Front2back: Single view 3d shape reconstruction via front to back prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 531–540 (2020)
55. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* **33**, 2492–2502 (2020)
56. Zhou, Q.Y., Koltun, V.: Dense scene reconstruction with points of interest. *ACM Transactions on Graphics (ToG)* **32**(4), 1–8 (2013)