

ROCA: Robust CAD Model Retrieval and Alignment from a Single Image

Can Gümeli

Angela Dai

Matthias Nießner

Technical University of Munich

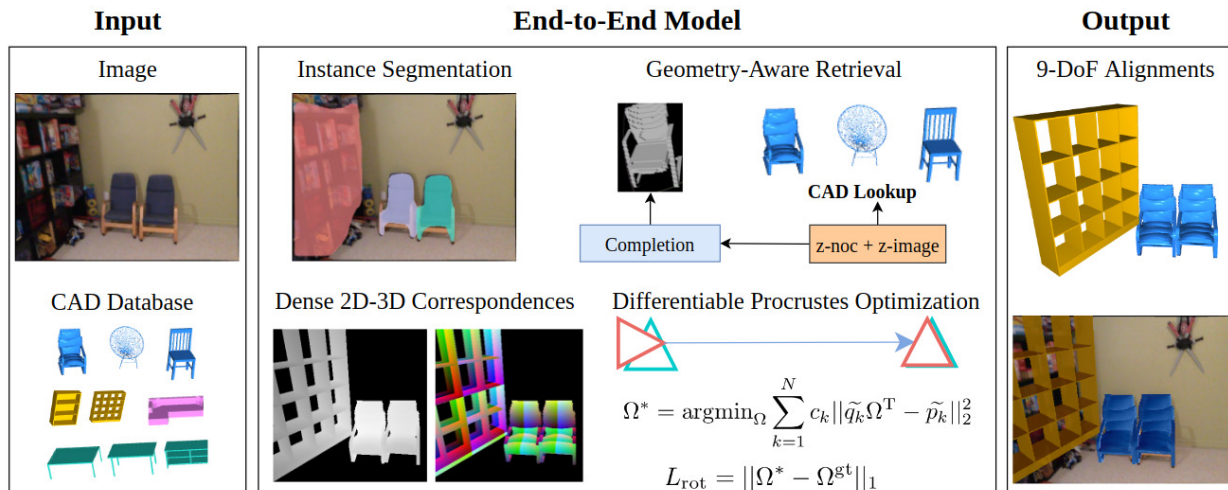


Figure 1. We present ROCA, a new end-to-end model that robustly retrieves and aligns 3D CAD models to a single image. From an RGB image and a database of CAD models, ROCA retrieves and aligns CAD models for each object in the image. In comparison to previous methods that perform direct pose regression, our approach leverages differentiable Procrustes optimization by predicting dense 2D-3D correspondences in the form of depths and normalized object coordinates (NOCs). In addition, predicted 3D correspondences help to learn retrieval of geometrically similar CAD models while simultaneously improving object alignments.

Abstract

We present ROCA¹, a novel end-to-end approach that retrieves and aligns 3D CAD models from a shape database to a single input image. This enables 3D perception of an observed scene from a 2D RGB observation, characterized as a lightweight, compact, clean CAD representation. Core to our approach is our differentiable alignment optimization based on dense 2D-3D object correspondences and Procrustes alignment. ROCA can thus provide a robust CAD alignment while simultaneously informing CAD retrieval by leveraging the 2D-3D correspondences to learn geometrically similar CAD models. Experiments on challenging, real-world imagery from ScanNet show that ROCA significantly improves on state of the art, from 9.5% to 17.6% in retrieval-aware CAD alignment accuracy.

¹The code is made available at <https://github.com/cangumeli/ROCA>.

1. Introduction

2D perception systems have seen remarkable advances in object recognition from 2D images in recent years, enabling widespread adoption of systems that can perform accurate 2D object localization, classification, and segmentation from an image [20, 31, 42]. Such advances have spurred forward developments in many fields, from classical image understanding to robotics and autonomous driving. However, unlike the human perception of 2D images, these systems tend to perform object recognition purely in 2D, whereas from a single RGB image, a human can perceive geometric shape, structure, and pose of the objects in the scene. In fact, such 3D understanding is crucial for many applications, enabling possible exploration and interaction with an observed environment.

At the same time, there has been notable progress in estimating 3D object geometry from visual data [9, 16, 19, 36, 44, 45]. In particular, Mesh R-CNN [18] introduced a formative approach to 3D object estimation from real-world images, bridging state-of-the-art 2D object detec-

tion with voxel-to-mesh estimation of the shape of each detected object in an image. In contrast to this generative approach, Mask2CAD proposed to retrieve and align CAD models from a database to produce a lightweight object reconstruction with high fidelity given by the CAD database [27, 35]. With significant availability of synthetic CAD models [5, 17, 47], such CAD model reconstruction of an observed scene shows strong promise in perceiving 3D from an image, as it enables geometric estimation in a clean, compact fashion more akin to artist-crafted 3D models, and easily consumed for downstream applications.

However, current approaches to estimate 3D object structure from an RGB image have largely focused on shape representation and generation, either without any explicit pose estimation [18], or simply regressing the object pose directly from the 2D features of the object [15, 27, 28, 35, 37]. Thus, we propose a new CAD retrieval and alignment approach to estimate 3D perception from an image by formulating a differentiable 9-DoF pose optimization directly coupled to the object retrieval; this enables a more robust, geometry-aware, end-to-end CAD alignment to the image.

In this work, we propose ROCA, a new method that jointly detects object regions in a given input image while simultaneously estimating depth and dense correspondences between each 2D object region and its location in its canonical object space. From the dense depth and correspondence estimates, we then formulate a differentiable Procrustes optimization and produce a final set of retrieved CAD models and their 9-DoF alignments. This geometry-aware differentiable optimization for the CAD alignment of each object enables more robust and accurate CAD alignment. In addition, our method learns geometry-aware embeddings for CAD retrieval. Our retrieval embedding module utilizes the canonical coordinates used in pose optimization and an auxiliary shape completion objective. We show learning geometry-aware embeddings improves both retrieval and alignment accuracy. Overall, our method significantly outperforms state of the art, improving by 8.1% and 9.5% in retrieval-aware alignment accuracy and alignment accuracy from a single RGB image.

In summary, we propose an end-to-end architecture for CAD model alignment to an RGB image with:

- a new differentiable pose optimization enabling geometry-aware 2D-to-3D CAD alignment to an RGB image,
- improved CAD retrieval by leveraging the dense object pose correspondences and proxy CAD completion objective to inform the construction of a joint embedding space between detected objects and CAD models,
- an interactive runtime of 53 milliseconds per image, facilitating its use in real-time applications.

2. Related Work

2D Object Recognition. Single-image 3D understanding requires strong 2D recognition capabilities. With recent advances in deep learning, we now have methods that have achieved remarkable success in image classification [21, 42], 2D object detection [31, 32], and 2D instance segmentation [20, 29]. We build our approach from the success of 2D recognition to extend to an end-to-end 3D object reasoning. In particular, to extract image features, we use a Mask-RCNN [20] recognition backbone that detects and segments objects in the image, from which we then estimate 3D CAD alignment and retrieval.

Monocular Depth Estimation. As we aim to predict 9-DoF object alignments, we must reason about absolute object depths. Monocular depth estimation from a single RGB image has been extensively studied to predict absolute per-pixel depths from an image. Modern monocular depth estimation frameworks typically train from pre-trained deep convolutional networks to predict per-pixel depth prediction in a fully convolutional fashion [14], leveraging learned semantic features to help resolve scale-depth ambiguities. Recent methods have proposed using deeper networks, multi-scale feature extractors, and alternative loss functions [22, 30]. To extract strong geometric features, we also predict dense depth for each detected object, leveraging state-of-the-art depth estimation techniques [22, 30].

Single-Image Shape Reconstruction. In recent years, many deep learning-based approaches have been developed to reconstruct 3D shapes from a 2D image. These approaches typically employ large synthetic 3D shape datasets [5], and render the synthetic objects for image input. Many different shape representations have been explored, including voxel grids [9, 44], point clouds [16], polygonal meshes [12, 19, 45], and neural implicit functions [36, 38]. Such approaches have seen significant success in the single-object scenario, inspiring Mesh-RCNN [18] to pioneer an approach built on 2D object recognition that generates 3D shapes for each detected object in an RGB image, as a first step to 3D perception in real-world images. More recent work has also focused on estimating scene layouts and inter-object relations [26, 37]. In contrast to these approaches that perform object-based reconstruction from an image, we propose to optimize for object poses by establishing dense geometric correspondences, rather than a direct regression, enabling more robust alignment optimization.

CAD Model Retrieval and Alignment. 3D reconstruction using CAD model priors has a long history in computer vision [3, 7, 41]. Recently, with the availability of large-scale 3D shape datasets [5], several approaches have been introduced to perform CAD model retrieval and alignment to an image based on analysis-by-synthesis reconstruction [23, 25]. While promising, these methods have

extensive computational costs (minutes to hours), making them ill-suited for many potential image-based reconstruction applications, such as real-time mobile or robotic applications. More recently, Mask2CAD [27] proposed a more lightweight approach that learns to simultaneously retrieve and 5-DoF align 3D CAD models to detected objects in an image, building on top of a state-of-the-art 2D recognition backbone [29, 32]. Its use of a CAD representation for lightweight object-based reconstruction and perception has inspired a patch-based approach to improve CAD retrieval [28], as well as extension to video input in Vid2CAD with full 9-DoF alignment for each object (estimating scale and depth, not originally predicted by Mask2CAD) [35]. We also aim to align and retrieve CAD models to reconstruct the objects in an RGB image, but rather than directly regressing the object pose, we formulate a differentiable optimization based on learned, dense object correspondences, enabling more accurate CAD alignment. Our established dense object correspondences also enable explicit geometry-aware retrieval, as we establish a joint embedding space between 3D CAD and 3D object correspondences.

Several recent approaches also solve for 9-DoF object alignment in an image. For instance, Points2Objects [15] builds from 2D object detection [13] to directly regress 9-DoF alignments and treats object retrieval as a classification problem, primarily demonstrating results on synthetic renderings. In contrast to their direct bounding box regression for alignment, we formulate a differentiable alignment based on dense geometric correspondence and perform nearest neighbor object retrieval from a CAD database.

Learned Differentiable Pose Optimization. Traditional pose optimization leverages correspondence finding to solve for pose alignment, with separate stages for establishing correspondences and the optimization for the best pose under the correspondences. With the introduction of deep learning, several methods have been proposed to bridge such pose optimization into an end-to-end learned pipeline, by establishing learned correspondences to inform various pose optimization tasks, such as point cloud registration [8], PnP optimization [6], or non-rigid tracking [4]. Most inspirational to our work is the differentiable Procrustes optimization of [2] for CAD model alignment to RGB-D scans. In contrast to their work, we must reason without 3D input data, and thus formulate a notably different optimization to focus on image-based object pose estimation with per-pixel correspondences and a weighted formulation for robust alignment estimation.

3. Method

3.1. Overview

From an input RGB image I , its camera intrinsics π , and a database of 3D CAD models \mathcal{S} , we aim to represent each

object in the image as a similar CAD model in its 9-DoF alignment to the image in metric camera space, thus providing a comprehensive lightweight, geometric reconstruction of the observed scene. An overview of our approach is shown in Figure 2.

From the RGB image, we first detect and segment objects in 2D with a Mask-RCNN [20] backbone, while simultaneously estimating the object depths with a multi-scale FPN [31]. We then estimate the 9-DoF alignment for a detected object by using its estimated depths, 2D features, and instance mask to regress scale and an initial translation. Moreover, we also establish learned, dense correspondences to the object’s canonical space. We then formulate a weighted Procrustes optimization to solve for the rotation and refined translation, simultaneously estimating the optimization weights for robust alignment optimization. Our approach is trained end-to-end with supervision given by CAD models that have been aligned to RGB images.

In addition to alignment, our method learns to retrieve geometrically similar CAD models to represent each detected object. We first use the predicted Mask-RCNN object categories to determine the object class from which to search. We establish geometric similarity by learning geometry-aware joint embeddings between the detected objects and CAD models, leveraging the detected objects’ canonical correspondences and 2D features; this helps to significantly bridge the domain gap of 2D observations and 3D CADs. To additionally learn shape-guided features, our embedding spaces are also trained jointly with a proxy shape completion objective. This retrieval is trained end-to-end along with the object alignment predictions and enables CAD retrieval as a nearest-neighbor lookup from the input CAD database at inference time.

3.2. Object Recognition and Depth Estimation

From an input image I , we leverage a strong 2D backbone to estimate both 2D object recognition and dense depth. We use a ResNet-50-FPN [21, 31] to form the 2D backbone, which produces a feature map F that then informs both the 2D instance segmentation following Mask-RCNN [20], as well as the depth estimation using a Multi-scale Feature Fusion (MFF) [22] module. The 2D object recognition detects objects as a set of object bounding boxes $\{b_i\}$ and instance masks $\{m_i\}$, and the depth estimation produces a features map F^d from the MFF layer before outputting the estimated depth map D . For 2D object recognition, we use a Mask-RCNN [20] backbone, and for depth estimation, we use a ResNet-50-FPN [21, 31] with a Multi-scale Feature Fusion [22] module that up-projects [30] and concatenates four levels of FPN features. To up-sample depth to full image resolution, we use a pixel-shuffle layer [43]. We use a masked reverse Huber (berHu) [30] loss to optimize for depth map D .

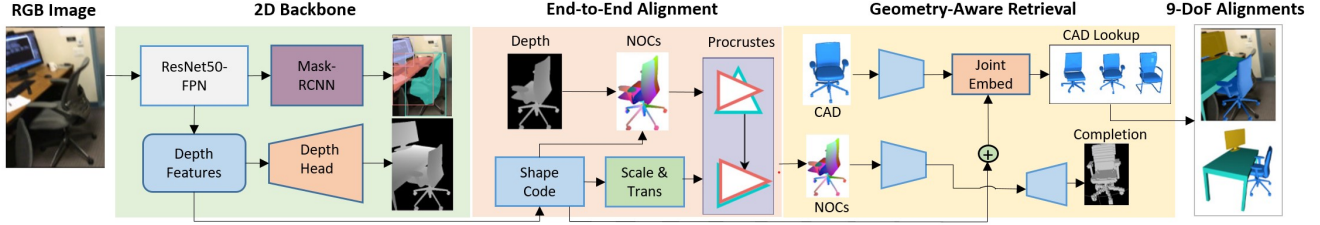


Figure 2. Method overview. From an input RGB image, we first estimate 2D instance segmentation and dense depth for each object. For each detected object, we then estimate dense correspondences to its canonical space as NOCs, which inform our differentiable Procrustes optimization for geometrically-informed object alignment. We additionally leverage the NOCs to construct a joint embedding space between estimated NOCs and voxelized CAD models to enable CAD retrieval, leveraging a proxy geometric completion loss. This enables geometrically informed CAD retrieval, and moreover, robust CAD alignment to the image.

We additionally refine the predicted depths in the regions of detected objects. For each detected object i with predicted 2D bounding box b_i , we consider the d_i cropped from D using b_i and resized with nearest-neighbor interpolation to an ROI size of 32×32 . We then apply an additional loss to focus on object depth estimation, $\|d_i - d_i^{\text{gt}}\|_1 + \|\mu(d_i) - \mu(d_i^{\text{gt}})\|_1$, where μ represents the average function.

We can then project the estimated per-object depths d_i with the intrinsics π to obtain camera-space positions p_i .

3.3. Robust Differentiable CAD Alignment

For each detected object, we then estimate its 9-DoF alignment to the image as its translation \mathbf{t}_i , scale s_i , and rotation \mathbf{r}_i .

Aggregating region features. To estimate the alignment for a detected object i , with predicted 2D bounding box b_i and instance mask m_i , we first aggregate its corresponding features from F^d which encode strong geometric descriptors of the object. We crop the features from F^d using ROIAlign [20] in the box region b_i , resulting in f_i^d of dimension $32 \times 32 \times 128$. For each object region, we then use a robust feature aggregation scheme to obtain a shape descriptor \mathbf{e}_i :

$$\mathbf{e}_i = \text{MLP}(\text{MaxPool}(m_i \odot \text{PixelMLP}(f_i^d))) \quad (1)$$

where \odot denotes the element-wise product operator, and PixelMLP is a fully-connected network that operates per-pixel and is shared over the pixels. Both MLPs use two layers with hidden size 512, and MaxPool performs a max pooling over the pixels in the object region to a single feature, to help increase robustness against mask and detection errors. The final MLP output activation contains a ReLU as well as a 30% dropout regularization.

Scale regression. We directly use predicted shape codes to regress scale, $s_i = \text{MLP}(\mathbf{e}_i)$, and optimize it with an ℓ_1 loss:

$$L_{\text{scale}} = \|s_i - s_i^{\text{gt}}\|_1. \quad (2)$$

Note that the final affine layer produces scale estimates for each class category, from which the scale is selected based

on the class category prediction, enabling capture of any class-specific scale characteristics.

Initial translation estimation. We additionally predict an initial translation estimate, based on the back-projected predicted object depths, p_i , as a learned offset \mathbf{t}_i^o from the center of the p_i :

$$\begin{aligned} \mathbf{t}_i^{\text{init}} &= 0.5(\max(p_i) + \min(p_i)) + \mathbf{t}_i^o, \\ \mathbf{t}_i^o &= \text{MLP}([\mathbf{e}_i, (\max(p_i) - \min(p_i))]). \end{aligned}$$

where $[\cdot, \cdot]$ denotes concatenation, and the MLP contains a single hidden layer of size 256.

We then apply losses on the p_i and $\mathbf{t}_i^{\text{init}}$:

$$L_p = H([\max(p_i), \min(p_i)], [\max(p_i^{\text{gt}}), \min(p_i^{\text{gt}})]) \quad (3)$$

$$L_{\text{trans.initial}} = \|\mathbf{t}_i^{\text{init}} - \mathbf{t}_i^{\text{gt}}\|_2, \quad (4)$$

where H denotes the Huber loss. Similar to the scale regression, we output per-category estimates from which the final output is selected based on the predicted class category.

Normalized Object Coordinates as Correspondences. In order to predict the final translation \mathbf{t}_i and rotation \mathbf{r}_i of the object, we propose to predict dense correspondences to the canonical object space, characterized as normalized object coordinates, to inform the alignment optimization. In contrast to a direct regression, this enables the alignment prediction to be geometrically informed and to learn the correspondences that best inform the alignment energy.

That is, for each foreground pixel in m_i , we predict its corresponding 3D coordinates in the normalized object coordinate (NOC) space, where the object lies in its canonical orientation is normalized into a unit cube $[-0.5, 0.5]^3$. Since NOCs will represent the same geometric surfaces as the back-projected depths, we predict NOC correspondences q_i based on the inverse translated depth points, predicted scale, and shape code, using a lightweight MLP with two 256-unit hidden layers:

$$q_i = \text{MLP}([p_i - \mathbf{t}_i^{\text{init}}, s_i, \mathbf{e}_i]). \quad (5)$$

We optimize for predicted NOCs using an ℓ_1 loss w.r.t. ground-truth NOCs:

$$L_{\text{noc}} = \|q_i - q_i^{\text{gt}}\|_1. \quad (6)$$

Differentiable Robust Procrustes Optimization. At the core of our alignment prediction, we formulate a differentiable weighted Procrustes optimizer that estimates rotation and refined translation. For each object, we have the regressed scale and translation $\mathbf{s} \in \mathbb{R}^3$ and $\mathbf{t}^{\text{init}} \in \mathbb{R}^3$, NOC-depth correspondences $(q, p) \in \mathbb{R}^{N \times 3}$, and m contains the segmentation mask probability of each foreground NOC-depth correspondence. For notational simplicity, now we omit the index suffix i from the previous sections.

Our objective is to solve for an orthogonal rotation $\Omega \in \mathbb{S}\mathbb{O}_3$ such that $(q \odot \mathbf{s})\Omega^T = (p - \mathbf{t})$. We formulate this problem as a least squares optimization:

$$\Omega^* = \operatorname{argmin}_{\Omega} \sum_{k=1}^N c_k \|\tilde{q}_k \Omega^T - \tilde{p}_k\|_2^2, \quad (7)$$

where $\tilde{q} = q \odot \mathbf{s}$, $\tilde{p} = p - \mathbf{t}$, and $c \in \mathbb{R}^N$ is a set of learned, positive weights that allow for learning to upweight more reliable correspondences. We predict c using a network conditioned on shape code \mathbf{e} and network predictions: $c_k = \text{MLP}([\mathbf{e}, \mathbf{s}, m_k, q_k, \tilde{p}_k])$. This weight prediction MLP contains a single 256-unit hidden layer followed by a sigmoid. The optimization problem is then solved analytically by differentiable weighted Procrustes optimization [8]:

$$U\Sigma V = \text{SVD}(\tilde{p}^T \begin{bmatrix} c_1 & & \\ & \ddots & \\ & & c_N \end{bmatrix} \tilde{q}) \quad (8)$$

$$\Omega^* = U \begin{bmatrix} 1 & & \\ & 1 & \\ & & \det(VU^T) \end{bmatrix} V^T \quad (9)$$

We optimize for the resulting rotation with an ℓ_1 loss using the ground-truth rotation matrix:

$$L_{\text{rot}} = \|\Omega^* - \Omega^{\text{gt}}\|_1. \quad (10)$$

Procrustes optimization can also be used to estimate translation by first zero-centering points for rotation and then using the resulting rotation and point averages for translation [8]. We found that directly estimating translation this way resulted in unstable training due to large transformations. To this end, we instead solve for a refined translation \mathbf{t} based on the initial estimate \mathbf{t}^{init} :

$$\mathbf{t} = \mathbf{t}^{\text{init}} + \mu_c(\tilde{p}) - \Omega^* \mu_c(\tilde{q}), \quad (11)$$

where μ_c denotes the weighted average operation w.r.t. predicted optimization weights c . We optimize for the refined translation with the objective:

$$L_{\text{trans}} = \|\mathbf{t} - \mathbf{t}^{\text{gt}}\|_2. \quad (12)$$

Thus, our final alignment loss is:

$$L_{\text{align}} = w_{\text{rot}} L_{\text{rot}} + L_{\text{trans}} + L_{\text{scale}} + w_{\text{noc}} L_{\text{noc}} + L_{\text{trans,initial}}.$$

We can then train for alignment in an end-to-end fashion, directly informing the predicted geometric correspondences learned from the image input, resulting in more robust alignment estimates.

3.4. Geometry-Aware CAD Retrieval

In addition to geometric object alignment, we propose a geometry-aware image-to-CAD retrieval that learns a joint embedding space of CAD models and object regions.

To represent a detected object i in the image, we use its shape code and predicted NOCs described in Section 3.3:

$$\mathbf{z}_i = \text{MLP}(\mathbf{e}_i) + 3\text{DCNN}(\text{Voxelize}(q_i)), \quad (13)$$

where $\mathbf{z}_i \in \mathbb{R}^{256}$, determined by the combination of an MLP that processes the shape code \mathbf{e}_i with a single hidden unit of size 1024 and a final ReLU activation, as well as a lightweight 3D convolutional network 3DCNN that operates on the 32^3 voxelized NOCs (the architecture is adopted from [10], we refer to the supplemental for more architecture details). To differentially voxelize the NOCs, we use the trilinear point cloud voxelization from PyTorch3D [40], and normalize the resulting density grid such that each occupancy probability is $\in [0, 1]$.

To encode the CAD models, we similarly represent them as 32^3 occupancy grids G , which are then encoded with a 3D CNN:

$$\tilde{\mathbf{z}} = 3\text{DCNN}(G). \quad (14)$$

Here, the 3DCNN is structured symmetrically to the NOC processing network, but they do not share weights, since NOC and CAD domains remain quite different.

Using the image and CAD representations above, our model learns a joint embedding space. For an image region embedding \mathbf{z}_i , its corresponding positive and negative CAD embeddings $\tilde{\mathbf{z}}_i^+$ and $\tilde{\mathbf{z}}_i^-$, we construct the space with a triplet loss:

$$L_{\text{ret}} = \max(\|\mathbf{z}_i - \tilde{\mathbf{z}}_i^+\|_2^2 - \|\mathbf{z}_i - \tilde{\mathbf{z}}_i^-\|_2^2 + 0.5, 0). \quad (15)$$

Since our predicted NOCs represent only the visible object geometry rather than the complete object geometry, whereas CAD models represent complete objects, we use an additional proxy completion loss on the NOC encoding:

$$L_{\text{comp}} = \text{BCE}(3\text{DU}p\text{CNN}(\mathbf{z}_i), G_i^+) \quad (16)$$

where BCE is the binary cross-entropy loss, G_i^+ is the positive CAD's occupancy grid, and 3DU_pCNN a 3D convolutional decoder that predicts the complete object geometry from the NOC encoding.

Our geometry-aware retrieval is trained end-to-end with the alignment prediction, allowing for retrieval to also inform the predicted NOC correspondences as well as the image features. At inference time, we pre-compute CAD embeddings and perform a nearest-neighbor lookup for each detected object region embedding prediction.

3.5. Implementation Details

Training. We train our approach end-to-end, using a momentum optimizer with momentum 0.9 and an initial learning rate of $1e-3$, starting with 1k warm-up iterations interpolating from $1e-4$ [20]. The learning rate is decayed by 0.1 at iteration 60k, and training takes 80k iterations for convergence. We use a batch size of 4, and 128 region proposals. To regularize our model, we use a weight decay of $1e-4$ and apply random contrast, brightness, and saturation augmentation, all sampled from the intensity range $[0.8, 1.25]$. We use loss weights $w_{rot} = 2, w_{noc} = 3$ to balance them with the translation and scale losses. To address class imbalances, we weight detection, segmentation, and alignment losses inversely proportional to the logarithm of training class frequencies. Note that we do not weight retrieval losses, as more frequent classes need better retrieval. Our 2D recognition backbone is pre-trained on ImageNet and then COCO [33,42]. In total, training takes ≈ 20 hours on a single RTX 3090 GPU.

Inference. For 2D object recognition at inference time, we follow Mask-RCNN [20], filtering detections using non-maxima suppression. Object detections with < 0.5 confidence are discarded, and we use a mask probability threshold of 0.7 to estimate foreground object regions for more robust feature aggregation and alignment. For retrieval, we use a 0.5 threshold for region segmentation. Although our model is not optimized for speed, running the full inference pipeline only takes 53 ms (≈ 19 frames per second), resulting in near real-time, interactive running speed.

Implementation. Our model implementation is based on PyTorch [39], Detectron2 [46], and PyTorch3D [40] frameworks. We use the weights of Detectron2 Mask-RCNN-R50-FPN-1x model to initialize our recognition backbone.

4. Results

4.1. Data and Evaluation

Dataset. We evaluate our approach on the ScanNet25k image data [11], following state of the art for CAD alignment to images [27, 28, 35]. This contains 20k training and 5k validation images, with training and validation images sampled from videos representing different scenes. We render Scan2CAD [1] CAD annotations into the corresponding image views to obtain object detection, segmentation, depth,




Figure 3. Qualitative evaluation on RGB images from ScanNet [11] with Scan2CAD [1] annotations. ROCA obtains more robust and accurate object alignments in many complex scenes.

and NOC targets for training. We consider input images at 360×480 resolution.

Alignment Accuracy. To evaluate our 9-DoF alignment performance, we adopt the alignment accuracy metric of Vid2CAD [35]. For each scene, we first transform per-image 9-DoF alignment predictions to the ScanNet world space and apply a 3D clustering to predicted alignments, following the protocol described in [35]. Alignment accuracy is then computed following Scan2CAD [1]: an alignment is correct if the corresponding object classification is correct, translation error ≤ 20 cm, rotation error is $\leq 20^\circ$, and scale ratio is $\leq 20\%$. Alignments can be predicted up to the total number of ground truth alignments.

Retrieval-Aware Alignment Accuracy. The previous alignment accuracy metric, following [35], only takes into account object classification and 9-DoF alignment. Thus, to evaluate retrieval, we define a retrieval-aware alignment accuracy metric. In addition to alignment correctness, this metric enforces the correctness of the retrieved CAD model. Candidate CAD models are those that appear in the ScanNet scene, following retrieval in prior works [1, 2, 35].

4.2. Comparison to State of the Art

We compare our method against the state-of-the-art 9-DoF Mask2CAD-b5 architecture [27,35]. Since the original

Method	bathtub	bed	bin	bkshlf	cabinet	chair	display	sofa	table	class	instance
Total3D-ODN [37]	10.0	2.9	16.8	2.8	4.2	14.4	13.1	5.3	6.7	8.5	10.4
MDR-CN [34]	5.8	5.7	0.9	9.9	5.4	28.1	11.5	11.5	8.1	9.7	15.3
Mask2CAD-b5 [27]	8.3	2.9	25.9	3.8	5.4	30.9	17.3	5.3	7.1	11.9	17.9
Ours P	14.2	7.1	18.5	6.1	13.1	30.3	17.8	17.7	9.2	14.9 ^(+3.0)	19.3 ^(+1.4)
Ours P+E	21.7	4.3	27.2	13.7	15.0	42.9	22.0	13.3	14.5	19.4 ^(+7.5)	26.9 ^(+9.0)
Ours P+E+W	18.3	10.0	28.9	15.1	14.6	41.4	23.0	15.9	16.1	20.4 ^(+8.5)	27.0 ^(+9.1)
Ours P+E+W+R	22.5	12.9	29.7	11.3	13.8	40.6	28.3	12.4	14.8	20.7 ^(+8.8)	26.7 ^(+8.8)
Ours (final)	22.5	10.0	29.3	14.2	15.8	41.0	30.4	15.9	14.6	21.5 ^(+9.6)	27.4 ^(+9.5)

Table 1. **Alignment Accuracy on ScanNet [1, 11]** in comparison to the state of the art and ablations. Total3D-ODN and MDR-CN are the 3D object detectors of Total3D [37] and MDR-CenterNet [34], respectively. In both detectors, we provide ground-truth rotations in lieu of layout estimation. Mask2CAD-b5 is Mask2CAD that predicts full 9-DoF alignment [27, 35]. In our ablations, P denotes Procrustes alignment optimization, E end-to-end training, W learned weighted optimization, and R learned retrieval without any completion proxy. Both end-to-end training and weighted robust optimization improve alignment accuracy. Our learned retrieval and completion additionally improve the alignment performance.

Method	bathtub	bed	bin	bkshlf	cabinet	chair	display	sofa	table	class	instance
Mask2CAD-b5 [27]	7.5	2.9	23.3	2.8	4.2	23.0	12.0	3.5	6.0	9.5	13.8
Ours P+E+W	14.2	10.0	25.4	10.8	9.2	27.3	17.3	15.9	13.0	15.9 ^(+6.4)	19.4 ^(+5.6)
Ours P+E+W+R	21.7	12.9	27.6	5.7	10.8	30.6	19.9	9.7	11.2	16.7 ^(+7.2)	20.5 ^(+6.7)
Ours (final)	20.8	10.0	26.7	8.5	11.9	32.1	22.5	14.2	11.8	17.6 ^(+8.1)	21.7 ^(+7.9)

Table 2. **Retrieval-Aware Alignment Accuracy on ScanNet [1, 11]**. Mask2CAD-b5 is Mask2CAD that predicts full 9-DoF alignment [27, 35]. We additionally evaluate different retrieval strategies: P+E+W does not use learned retrieval and instead uses Chamfer distance as a retrieval metric, and P+E+W+R introduces learned retrieval but without proxy completion. Our final approach with learned retrieval with proxy completion achieves the best retrieval-aware alignment performance.

Mask2CAD predicted only 5-DoF alignments, Mask2CAD-b5 also learns to predict object depth and scale to produce 9-DoF alignments, as adapted by [35].

We additionally compare our method against state-of-the-art single-image 3D object detectors, namely Total3D [37] and MDR [34]. Total3D [37] predicts from a single image the room layout, object poses and generates object meshes. To compare with their object pose estimation, we evaluate against their Object Detection Network (Total3D-ODN). MDR [34] performs joint 3D object detection and voxel-based coarse-to-fine object reconstruction. We compare against their 3D CenterNet-based [13] detector (MDR-CN). Since both methods depend on the room layout, we provide ground truth rotation. All baselines are trained on ScanNet25k data.

Table 1 and Table 2 evaluate alignment accuracy and retrieval-aware alignment accuracy, respectively, in comparison with state of the art. Our approach outperforms state of the art in alignment accuracy by 9.6% (81% relative) and by 9.5% (53% relative) in class and instance averages, due to our differentiable Procrustes-based alignment optimization. Additionally, when considering retrieval-aware alignment accuracy, ROCA improves by 8.1% (85% relative improvement) and by 7.9% (57% relative improvement) over state of the art in class and instance accuracies. Our geometrically-grounded alignment formulation enables significant performance improvements in categories that are more scarcely represented in the train set (particularly “bed”, “bathtub”).

Figure 3 shows a qualitative comparison of CAD retrieval and alignment on ScanNet images. ROCA obtains more robust and accurate object alignments across a diverse set of image views and object types.

4.3. Ablations

Effect of end-to-end optimization on alignment. Table 1 shows that our end-to-end Procrustes optimization notably improves class alignment accuracy, from 14.9% to 19.4%.

Effect of learned Procrustes weights on alignment. In Table 1, adding learned weights to end-to-end Procrustes additionally improves robustness in alignment estimation, resulting in an improvement of average class alignment accuracy from 19.4% to 20.4%.

Effect of learned retrieval on alignment. In fact, introducing learned retrieval from the predicted NOC correspondences helps to improve class alignment accuracy from 20.4% to 21.5%, as shown in Table 2. Here, the learned retrieval can provide additional signal to the correspondence learning, in our end-to-end formulation for joint retrieval and alignment.

Learned retrieval performance. In Table 2, we compare our method with a baseline retrieval approach of using a single-sided Chamfer distance from the NOC predictions to the database models (referred as “Ours P+E+W” in the tables). We achieve an improved performance when leveraging an end-to-end, learned approach that embeds NOC predictions into a shared space with CAD models. ROCA

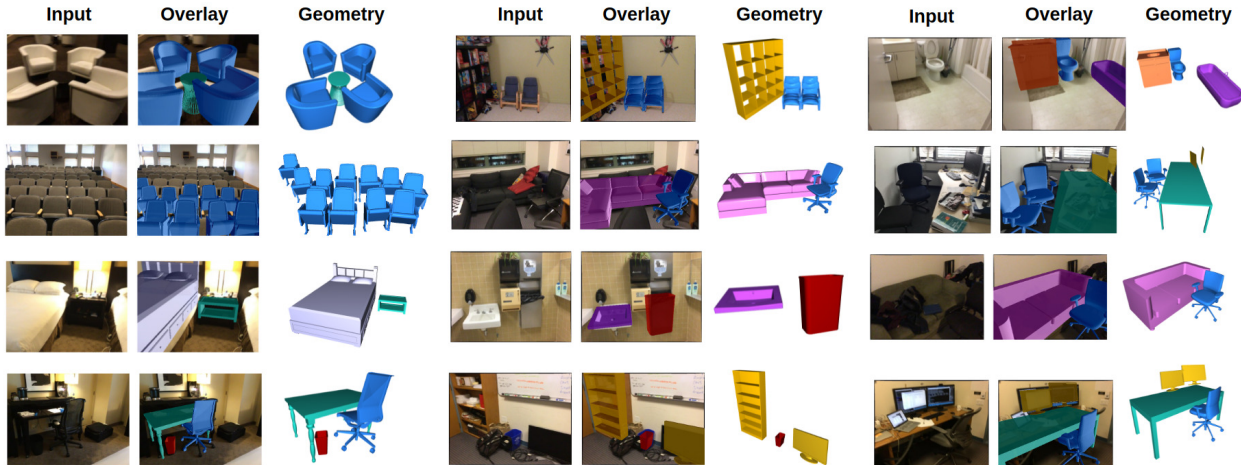


Figure 4. Sample predictions from our method. Our model shows promising 3D understanding ability in a wide variety of challenging real-world images [1, 11].

Data	Alignment		Retrieval	
	class	instance	class	instance
25k	21.5	27.4	17.6	21.7
400k	24.4	31.9	20.2	24.9

Table 3. Alignment and retrieval-aware alignment improvement due to augmenting the training image set, by sampling more frames from ScanNet [11] videos.

NOC noise σ	0.0	0.1	0.2	0.3
Accuracy	21.5	20.8	20.0	19.2

Table 4. Ablation of robustness to NOC noise under simulated Gaussian noise.

additionally leverages a completion proxy loss for the retrieval embedding, producing the best retrieval-aware alignment performance.

Scaling to augmented training data. While the ScanNet25k benchmark contains 20k train and 5k validation images obtained by sampling every 100 frames from ScanNet videos [11]. We consider augmenting training data by sampling images more densely. We prepare a 400k training dataset created by sampling every 5 frames, which helps to improve our alignment and retrieval-aware alignment performance, as shown in Table 3. Although such dense frame sampling leads to stronger performance, such augmentation can only be obtained from video datasets. Thus, we focus on the 25k scenario, as it is more representative for image-based reconstruction.

Robustness analysis. We show in Tab. 4 that under simulated Gaussian noise in NOC predictions, our method maintains robustness, recovering 90% of category average alignment accuracy even under strong noise of $\mathcal{N}(0, 0.3)$.

Limitations. While ROCA shows significantly more robust and accurate CAD alignment performance to images, various limitations remain. In particular, ROCA makes per-object predictions independently, while global scene context (e.g., object-object relationships) can provide additional important signal for object alignments. Addition-

ally, while CAD alignment and retrieval can provide an important semantic reconstruction of an observed scene, CAD databases cannot provide exact geometric matches in real-world scenarios, as is the case with Scan2CAD [1]. Thus, our approach is limited to approximate geometric representations; we believe that extending a CAD retrieval approach to couple with mesh deformation approaches (e.g., as in [24]) is a promising direction to address more exact geometric reconstruction.

5. Conclusion

We have presented ROCA, a robust end-to-end approach for single-image CAD model alignment and retrieval. We show that leveraging dense per-pixel depth and canonical point correspondences with our weighted differentiable Procrustes optimization leads to more robust and accurate pose predictions. Additionally, these correspondences can be leveraged for geometry-aware end-to-end retrieval to improve both retrieval and alignment performance. For challenging ScanNet/Scan2CAD image data [1, 11], our method significantly improves state-of-the-art retrieval-aware alignment accuracy from 9.5% to 17.6%. Our approach runs efficiently at test time, achieving interactive speeds of 53 milliseconds per image. We hope that this can further spur developments in 3D perception towards content creation, mixed reality, and domain transfer scenarios.

Acknowledgements. This project is funded by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt), the TUM Institute of Advanced Studies (TUM-IAS), the ERC Starting Grant Scan2CAD (804724), and the German Research Foundation (DFG) Grant Making Machine Learning on Static and Dynamic 3D Data Practical. We also want to thank Kevis-Kokitsi Maninis and Vanessa Wirth for helping with the evaluation and Yinyue Nie for helping with the Total3D baseline.

References

- [1] Armen Avetisyan, Manuel Dahnert, Angela Dai, Manolis Savva, Angel X Chang, and Matthias Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2614–2623, 2019. 6, 7, 8
- [2] Armen Avetisyan, Angela Dai, and Matthias Nießner. End-to-end cad model retrieval and 9dof alignment in 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2551–2560, 2019. 3, 6
- [3] Thomas O Binford. Survey of model-based image analysis systems. *The International Journal of Robotics Research*, 1(1):18–64, 1982. 2
- [4] Aljaz Bozic, Pablo Palafox, Michael Zollöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. 2020. 3
- [5] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 2
- [6] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8100–8109, 2020. 3
- [7] Roland T Chin and Charles R Dyer. Model-based recognition in robot vision. *ACM Computing Surveys (CSUR)*, 18(1):67–108, 1986. 2
- [8] Christopher Choy, Wei Dong, and Vladlen Koltun. Deep global registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2514–2523, 2020. 3, 5
- [9] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 1, 2
- [10] Manuel Dahnert, Angela Dai, Leonidas J Guibas, and Matthias Nießner. Joint embedding of 3d scan and cad objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8749–8758, 2019. 5
- [11] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 6, 7, 8
- [12] Angela Dai and Matthias Nießner. Scan2mesh: From unstructured range scans to 3d meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5574–5583, 2019. 2
- [13] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6569–6578, 2019. 3, 7
- [14] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep net work. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 2
- [15] Francis Engelmann, Konstantinos Rematas, Bastian Leibe, and Vittorio Ferrari. From points to multi-object 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4588–4597, 2021. 2, 3
- [16] Haoqiang Fan, Hao Su, and Leonidas J Guibas. 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*, pages 605–613, 2017. 1, 2
- [17] Huan Fu, Rongfei Jia, Lin Gao, Mingming Gong, Binqiang Zhao, Steve Maybank, and Dacheng Tao. 3d-future: 3d furniture shape with texture. *arXiv preprint arXiv:2009.09633*, 2020. 2
- [18] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9785–9795, 2019. 1, 2
- [19] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 216–224, 2018. 1, 2
- [20] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1, 2, 3, 4, 6
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 3
- [22] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1043–1051. IEEE, 2019. 2, 3
- [23] Siyuan Huang, Siyuan Qi, Yixin Zhu, Yinxue Xiao, Yuanlu Xu, and Song-Chun Zhu. Holistic 3d scene parsing and reconstruction from a single rgb image. In *Proceedings of the European conference on computer vision (ECCV)*, pages 187–203, 2018. 2
- [24] Vladislav Ishimtsev, Alexey Bokhovkin, Alexey Artemov, Savva Ignatyev, Matthias Nießner, Denis Zorin, and Burnaev Evgeny. Cad-deform: Deformable fitting of cad models to 3d scans. *ECCV 2020*, 2020. 8
- [25] Hamid Izadinia, Qi Shan, and Steven M Seitz. Im2cad. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5134–5143, 2017. 2
- [26] Nilesh Kulkarni, Ishan Misra, Shubham Tulsiani, and Abhinav Gupta. 3d-relnet: Joint object and relational network for 3d prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2212–2221, 2019. 2
- [27] Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. Mask2cad: 3d shape prediction by learning to segment

- and retrieve. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 260–277. Springer, 2020. 2, 3, 6, 7
- [28] Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. Patch2cad: Patchwise embedding learning for in-the-wild shape retrieval from a single image. *arXiv preprint arXiv:2108.09368*, 2021. 2, 3, 6
- [29] Weicheng Kuo, Anelia Angelova, Jitendra Malik, and Tsung-Yi Lin. Shapemask: Learning to segment novel objects by refining shape priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9207–9216, 2019. 2, 3
- [30] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, pages 239–248. IEEE, 2016. 2, 3
- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 1, 2, 3
- [32] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 2, 3
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [34] Feng Liu and Xiaoming Liu. Voxel-based 3d detection and reconstruction of multiple objects from a single image. *Advances in Neural Information Processing Systems*, 34, 2021. 7
- [35] Kevis-Kokitsi Maninis, Stefan Popov, Matthias Nießner, and Vittorio Ferrari. Vid2cad: Cad model alignment using multi-view constraints from videos. *arXiv preprint arXiv:2012.04641*, 2020. 2, 3, 6, 7
- [36] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 1, 2
- [37] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 55–64, 2020. 2, 7
- [38] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 2
- [39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. 6
- [40] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 5, 6
- [41] Lawrence G Roberts. *Machine perception of three-dimensional solids*. PhD thesis, Massachusetts Institute of Technology, 1963. 2
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 1, 2, 6
- [43] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016. 3
- [44] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. *CoRR*, abs/1703.09438, 2017. 1, 2
- [45] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018. 1, 2
- [46] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [47] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 2