Shape Prior Deformation for Categorical 6D Object Pose and Size Estimation

 $\begin{array}{c} \text{Meng Tian}^{[0000-0001-9937-8975]}, \text{ Marcelo H Ang Jr}^{[0000-0001-8277-6408]}, \\ \text{Gim Hee Lee}^{[0000-0002-1583-0475]} \end{array}$

National University of Singapore, Singapore tianmeng@u.nus.edu, {mpeangh, gimhee.lee}@nus.edu.sg

Abstract. We present a novel learning approach to recover the 6D poses and sizes of unseen object instances from an RGB-D image. To handle the intra-class shape variation, we propose a deep network to reconstruct the 3D object model by explicitly modeling the deformation from a pre-learned categorical shape prior. Additionally, our network infers the dense correspondences between the depth observation of the object instance and the reconstructed 3D model to jointly estimate the 6D object pose and size. We design an autoencoder that trains on a collection of object models and compute the mean latent embedding for each category to learn the categorical shape priors. Extensive experiments on both synthetic and real-world datasets demonstrate that our approach significantly outperforms the state of the art. Our code is available at https://github.com/mentian/object-deformnet.

Keywords: category-level pose estimation, 3D object detection, shape generation, scene understanding

1 Introduction

Accurate 6D object pose estimation plays an important role in a variety of tasks, such as augmented reality, robotic manipulation, scene understanding, etc. In recent years, substantial progress has been made for instance-level 6D object pose estimation, where the exact 3D object models for pose estimation are given. Unfortunately, these methods [18, 32, 42] cannot be directly generalized to category-level 6D object pose estimation on new object instances with unknown 3D models. Consequently, the category, 6D pose and size of the objects have to be concurrently estimated. Although some other object instances from each category are provided as priors, the high variation of object shapes within the same category makes generalization to new object instances extremely challenging.

To the best of our knowledge, [24] is the first work to address the 6D object pose estimation problem at category-level. This approach defines 6D pose on semantically selected centers and trains a part-based random forest to recover the pose. However, building part representations upon 3D skeleton structures limits the generalization capability across unseen object instances. Another work

2 Meng Tian, Marcelo H Ang Jr, and Gim Hee Lee

[33] proposes the first data-driven solution and creates a benchmark dataset for this task. They introduce the Normalized Object Coordinate Space (NOCS) to represent different object instances within a category in a unified manner. A region-based network is trained to infer correspondences from object pixels to the points in NOCS. Class label and instance mask of each object are also obtained at the same time. These predictions are used together with the depth map to estimate the 6D pose and size of the object via point matching. However, the lack of explicit representation of shape variations limits their performance.

In this work, we propose to reconstruct the complete object models in the NOCS to capture the intra-class shape variation. More specifically, we first learn the categorical shape priors from the given object instances, and then train a network to estimate the deformation field of the shape prior (that is used to get the reconstructed object model) and the correspondences between object observation and the reconstructed model. The shape prior serves as prior knowledge of the category and encodes geometrical characteristics that are shared by objects of a given category. The deformation predicted by our network captures the instance-specific shape details, i.e. shape variation of that particular instance. We present a method which is applicable across different object categories and data representations to learn the shape priors. In particular, an autoencoder is trained on a collection of object models from various categories. For each category, we compute the mean latent embedding over all instances in the respective category. The categorical shape prior is constructed by passing the mean embedding through a decoder. Note that there is no restriction on the data representation (point cloud, mesh, or voxel) of shape priors or collected models as long as we choose a proper architecture for the encoder and decoder.

We use the Umeyama algorithm [29] to recover the 6D pose and metric size of the object from the correspondences estimated by our network that maps the point cloud obtained from the observed depth map to the points in NOCS. We evaluate our method on two standard benchmarks. Extensive experiments demonstrate the advantage of our network and prove the effectiveness of explicitly modeling the deformation. In summary, the main contributions of this work are:

- We propose a novel deep network for category-level 6D object pose and size estimation; our network explicitly models the deformation from the categorical shape prior to the object model.
- We present a learning-based method which utilizes the latent embeddings to construct the shape prior; our method is applicable across different categories and data representations.
- Our network achieves significantly higher mean average precisions on both synthetic and real-world benchmark datasets.

2 Related Work

Instance-Level Pose Estimation. Existing instance-level pose estimation approaches broadly fall into three categories. The first category casts votes in the

pose space and further refines coarse pose with algorithms such as iterative closest point. LINEMOD [9] uses holistic template matching to find the nearest viewpoint. [26] generates a latent code for the input image and search for its nearest neighbor in the codebook. [11, 27] aggregate the 6D votes cast by locally-sampled RGB-D patches. The second category directly maps input image to object pose. [10,37] extend 2D object detection network such that it can predict orientation as an add-on to the identity and 2D bounding box of the object. [15,32] regress 6D pose from RGB-D images in an end-to-end framework. The third category relies on establishing point correspondences. [2,12,17] regress the corresponding 3D object coordinate for each foreground pixel. [18,23,28] detect the keypoints of the object on image and then solve a Perspective-n-Point problem. [42] estimates a dense 2D-3D correspondence map between the input image and object model. Although our approach follows the approach from the third category, our task focuses on a more general setting where the object models are not available during inference.

Category-Level Object Detection. The task of 3D object detection aims to estimate 3D bounding boxes of objects in the scene. [25] runs sliding windows in 3D space and generates amodal proposals for objects. [6,14,21] first generate 2D object proposals and then lift the proposals to 3D space. [38,44] are single-stage detectors which directly detect objects from 3D data. Although the above mentioned methods address the problem at category-level, the considered objects are usually constrained to the ground surface, e.g. instances of typical furniture classes in indoor scenes, cars, pedestrians, and cyclists in outdoor scenes. Consequently, the assumption that rotation is constrained to be only along the gravity direction has to be made. On the contrary, our approach can recover the full 6D pose of objects.

Category-Level Pose Estimation. There are only a few pioneering works focusing on estimating 6D pose of unseen objects. [3] leverages a generative representation of 3D objects and produces a multimodal distribution over poses with mixture density network. However, only rotation is considered in their work. [24] introduces a part-based random forest which employs simple depth comparison features, but our approach deals with RGB-D images. [33] proposes a canonical representation for all instances within an object category. Our approach also makes use of this representation. Instead of directly regressing the coordinates in NOCS, we account for intra-class shape variation by explicitly modeling the deformation from shape prior to object model. [4] trains a variational autoencoder to generate the complete object model. However, the reconstructed shape is not utilized for pose estimation. In our network, shape reconstruction and pose estimation are integrated together. [31] proposes the first category-level pose tracker, while our approach performs pose estimation without using temporal information.

Shape Deformation. 3D deformation is commonly applied for object reconstruction from a single image. [13,20,41] use free-form deformation in conjunction with voxel, mesh and point cloud representations, respectively. [34,36] starts from a coarse shape and predicts a series of deformations to progressively improve the



Fig. 1. Overview of our approach. We first obtain a foreground mask for each object instance. Next our network reconstructs the instance (*bowl* as an example) and establishes the correspondences between the observed points and the reconstructed model. Finally, the 6D pose is recovered by estimating a similarity transformation.

geometry. Similar to us, [35] also supervise the deformation with global feature of the target. However, we circumvent the fixed topology assumption of mesh representation by using point clouds instead.

3 Our Method

Refer to Fig. 3 for the details of our network.

Background. Given an RGB-D image as the input, our goal is to detect and localize all visible object instances in the 3D space. The object instances are not seen previously, but must come from known categories. Each object is represented by a class label and an amodal 3D bounding box parameterized by its 6D pose and size. The 6D pose is defined to be the rigid transformation (i.e. rotation and translation) that transforms the object from the reference to the camera coordinate frame. It is common to choose the coordinate frame of the given 3D object models as reference in instance-level 6D object pose estimation. Unfortunately, this is not viable for our category-level task since the instances of the 3D models are not available. To mitigate this problem, we leverage on the Normalized Object Coordinate Space (NOCS) - a shared canonical representation for all possible object instances within a category proposed in [33]. The categorical 6D object pose and size estimation problem is then reduced to finding the similarity transformation between the observed depth map of each object instance and its corresponding points in the NOCS (i.e. NOCS coordinates). **Overview.** In contrast to [33] that directly outputs the NOCS coordinates from

a Convolutional Neural Network (CNN), we propose an intermediate step to

estimate the deformation of a pre-learned shape prior to improve the learning of intra-class shape variation. Our shape priors are learned from a collection of models spanning all categories (Section 3.1). As shown in Fig. 1, our approach consists of three stages. The first stage performs instance segmentation on color image using an off-the-shelf network (e.g. Mask R-CNN [7]). Next we convert the masked depth map into a point cloud with camera intrinsic parameters for each instance and crop an image patch according to the bounding box of the mask. Taking the point cloud, image patch, and the corresponding shape prior as inputs, our network outputs a deformation field that deforms the shape prior into the shape of the desired object instance (a.k.a. reconstructed model). Furthermore, our network outputs a set of correspondences that associates each point in the point cloud obtained from the observed depth map of the object instance with the points of the reconstructed model. This set of correspondences is used to mask the reconstructed model into the NOCS coordinates (Section 3.2). Finally, the 6D pose and size of the object can be estimated by registering the NOCS coordinates and the point cloud obtained from the observed depth map (Section 3.3).

3.1 Categorical Shape Prior

Although object shape varies among different instances, an investigation over the 3D models reveals that objects of the same category (especially artificially generated objects) tend to have semantically and geometrically similar components. For example, cameras are usually made up of a nearly cuboid body and a cylindrical lens; and mugs are typically cylindrical with a handle. These categorical characteristics provide strong priors on the shape reconstruction of novel instances. We propose the learning of a mean shape to capture the high-level characteristics from all the available models for each respective category. To this end, we first train an autoencoder with all available object models and then compute the mean latent embedding of each object category with the encoder. These latent embeddings are passed into the decoder to get the mean shape priors for each object category. Unlike methods such as simple averaging [30] and principal component analysis (PCA) [3] that operate on voxel representations, our autoencoder framework can be easily altered to take any 3D representations.

Given a shape collection $\mathcal{M} = \{M_c^i \mid i = 1, 2, \cdots, N; c = 1, 2, \cdots, C\}$, where M_c^i is the 3D point cloud model of instance *i* from category *c*, we independently apply a similarity transformation to each model such that it is properly aligned in the NOCS. This step ensures that the learned shape prior has the same scale and orientation as the target shape to be reconstructed. The encoder Φ takes the point cloud and outputs a low-dimensional feature vector, i.e. the latent embedding $z_c^i \in \mathbb{R}^n$. The decoder Ψ takes this feature vector and outputs a point cloud that reconstructs the input:

$$\hat{M}_c^i = (\Psi \circ \Phi)(M_c^i) = \Psi(z_c^i). \tag{1}$$

Specifically, we adopt the PointNet-like encoder proposed in [40], and a threelayer fully-connected decoder as shown in Fig. 2a. The reconstruction error is



Fig. 2. (a) Architecture of the autoencoder. (b) The latent embeddings of all instances are mapped to \mathcal{R}^2 with T-SNE for visualization. These instances are from 6 categories *bottle, bowl, camera, can, laptop* and *mug.* (c) Shape priors are reconstructed by passing mean latent embedding of each category through the decoder.

measured by the Chamfer distance:

$$d_{\rm CD}(M_c^i, \hat{M}_c^i) = \sum_{x \in M_c^i} \min_{y \in \hat{M}_c^i} \|x - y\|_2^2 + \sum_{y \in \hat{M}_c^i} \min_{x \in M_c^i} \|x - y\|_2^2.$$
(2)

The autoencoder is trained on a shape collection by minimizing the reconstruction error. Once the training is converged, we obtain the latent embeddings $\{z_c^i\}$ of all instances in \mathcal{M} . Although not explicitly enforced during training, these latent embeddings form clusters in the latent space according to their categories. Fig. 2b visualizes the clustering effect of the embeddings. We use T-SNE [16] to further embed these features in \mathcal{R}^2 for visualization. Similar clustering results are also observed on a different set of models [39]. Based on this observation, we compute the mean latent embedding for each category and then pass it through the decoder to construct the shape prior:

$$M_c = \Psi(z_c) = \Psi\left(\frac{1}{N_c}\sum_i z_c^i\right).$$
(3)

The resulting categorical shape priors $\{M_c\}$ are shown in Fig. 2c.

3.2 Our Network Architecture

We denote the observation of an object instance as (V, I), where $V \in \mathcal{R}^{N_v \times 3}$ is the point cloud and $I \in \mathcal{R}^{H \times W \times 3}$ is the image patch. N_v denotes the number



Fig. 3. Our Network Architecture. The upper-left and lower-left branches extract point and global features from the instance and the shape prior respectively. The upper-right branch estimates the correspondence matrix, and the lower-right branch predicts the deformation field. The exchange of global features is the key part of our network.

of foreground pixels with a valid depth value. The corresponding shape prior is $M_c \in \mathcal{R}^{N_c \times 3}$, where N_c is the number of points in M_c . Our network takes V, I and M_c as inputs, and outputs the per-point deformation field $D \in \mathcal{R}^{N_c \times 3}$ and a correspondence matrix $A \in \mathcal{R}^{N_v \times N_c}$. The final reconstructed model is $M = M_c + D$. Each row of A sums to 1 since it represents the soft correspondences between a point in V and all points in M. As shown in Fig. 3, our network is composed of four parts: (1) extracts features from the object instance; (2) extracts features from the shape prior; (3) regresses the deformation field D; and (4) estimates the correspondence matrix A.

On the consideration that the depth and color are two different modalities, we follow the pixel-wise dense fusion approach proposed in [32] to effectively extract RGB-D features from the observation. For point cloud V, we use an embedding network similar to PointNet [22] to generate per-point geometric features by mapping each point in V to the d_g -dimensional feature space. The image patch I is processed with a fully convolutional network which follows an encoder-decoder architecture and maps I to $\mathcal{R}^{H \times W \times d_c}$. Next we associate the geometric feature of each point with its corresponding color feature and concatenate the feature pairs. Since each point in V has a corresponding pixel in I, not vice versa, redundant color features are discarded. The concatenated features are termed "instance point features" and fed to another shared multilayer perceptron. An average pooling layer is used to generate the "instance global feature". The categorical shape prior M_c is a point cloud with purely geometric information. We apply a simpler embedding network to extract the "category point features" and the "category global feature".

The shape prior M_c provides the prior knowledge of the category, i.e. the coarse shape geometry and canonical pose. Although the observation (V, I) is

8 Meng Tian, Marcelo H Ang Jr, and Gim Hee Lee

partial, it provides instance-specific details of the target shape. A natural way to reconstruct the object in NOCS is to deform M_c under the guidance of (V, I). Consequently, we concatenate the category and instance global features, and enrich the category point features with the concatenated features. The obtained feature vectors are successively convolved with 1×1 kernels to generate the deformation field D. Similar intuition and feature concatenation strategy also apply to the estimation of A. We combine the instance point features and global feature to aggregate both local and global information for each point. Each point in V gets mapped to the points of the reconstructed model through concatenation with the category global feature. We obtain the NOCS coordinates, denoted as P, of the points in V by multiplying A and M, i.e.

$$P = A \times M = A(M_c + D) \in \mathcal{R}^{N_v \times 3}.$$
(4)

3.3 6D Pose Estimation

Our goal is to estimate the 6D pose and size of the object instance. Given depth observation V and its NOCS coordinates P, the optimal similarity transformation parameters (rotation, translation, and scaling) can be computed by solving the absolute orientation problem using Umeyama algorithm [29]. We also implement the RANSAC algorithm [5] for robust estimation.

3.4 Loss Functions

In this section, we define the loss functions used to train our network, and explain how we handle object symmetry during training.

Reconstruction Loss. We assume that ground-truth model M_{gt} is available during training. The deformation field D is supervised indirectly by minimizing the Chamfer distance (c.f. Eq. 2) between M and M_{gt} , i.e. $L_{cd} = d_{CD}(M, M_{gt}) = d_{CD}(M_c + D, M_{gt})$.

Correspondence Loss. It is impractical and unnecessary to pre-compute the ground-truth value for A. Alternatively, we supervise A indirectly via the NOCS coordinates P (which is a result of applying the correspondence matrix on the reconstructed model) since the ground-truth NOCS coordinates P_{gt} can be obtained easily from the object model and its 6D pose through image rendering. We use the smooth L_1 loss function:

$$L_{\rm corr}(P, P_{gt}) = \frac{1}{N_v} \sum_{\mathbf{x} \in P} \sum_{i=1,2,3} \begin{cases} 5(x_i - y_i)^2, & \text{if } |x_i - y_i| \le 0.1\\ |x_i - y_i| - 0.05, & \text{otherwise} \end{cases},$$
(5)

where $\mathbf{x} = (x_1, x_2, x_3) \in P$, and $\mathbf{y} = (y_1, y_2, y_3) \in P_{gt}$.

Object symmetry is an inevitable problem for pose estimation algorithms, especially for those that require supervised training. For symmetrical objects, there exists at least one rotation such that appearance of the object is preserved under this rotation. In other words, two observations of a symmetric object can be very similar but with different rotation labels. We follow the solution proposed by [19] to map ambiguous rotations to a canonical one. More specifically, the Map operator for an arbitrary rotation $R \in SO(3)$ is defined as:

$$\operatorname{Map}(R) = R\hat{S}, \text{ with } \hat{S} = \underset{S \in \mathcal{S}(M_{c}^{i})}{\operatorname{arg\,min}} \|RS - I_{3}\|_{F}, \tag{6}$$

where the proper symmetry group $\mathcal{S}(M_c^i)$ is the set of rotations which preserve the appearance of a given object M_c^i . The experimental datasets assume continuous symmetry and the axis of symmetry is the y-axis of the NOCS. Hence, \hat{S} takes the following form:

$$\hat{S} = \begin{bmatrix} \cos\hat{\theta} & 0 & -\sin\hat{\theta} \\ 0 & 1 & 0 \\ \sin\hat{\theta} & 0 & \cos\hat{\theta} \end{bmatrix}, \text{ with } \hat{\theta} = \arctan 2(R_{13} - R_{31}, R_{11} + R_{33}), \quad (7)$$

where R_{11} , R_{13} , R_{31} , and R_{33} are the elements of R. During training, we apply the Map operator to the rotation label: $R_{gt} \leftarrow R_{gt}\hat{S}$ to eliminate the rotation ambiguity of any symmetric object with the ground-truth pose (R_{gt}, T_{gt}) . In practice, our network is supervised by ground-truth NOCS coordinates P_{gt} . Equivalently, we transform P_{gt} by \hat{S}^T : $P_{gt} \leftarrow \hat{S}^T P_{gt}$.

Regularization Losses. Row A_i of the matrix A represents the distribution over the correspondences between *i*-th point of V and the points in M. A_i can be understood as a relaxed one-hot vector, since each point of V usually can be well approximated by at most three points of M. We encourage A_i to be a peaked distribution by minimizing the average cross entropy: $L_{\text{entropy}} = \frac{1}{N_v} \sum_i \sum_j -A_{i,j} \log A_{i,j}$. We also regularize D to discourage large deformations: $L_{\text{def}} = \frac{1}{N_C} \sum_{\mathbf{d}_i \in D} ||\mathbf{d}_i||_2$. Minimal deformation preserves the semantic consistency between shape prior and the reconstructed model. For example, we want that the point belongs to the handle of the "mug" prior remains in the handle after deformation. This consistency loss is beneficial for the prediction of the correspondence matrix A.

In summary, the overall objective is a weighted sum of all four losses:

$$L = \lambda_1 L_{\rm cd} + \lambda_2 L_{\rm corr} + \lambda_3 L_{\rm entropy} + \lambda_4 L_{\rm def}.$$
(8)

4 Experiments

4.1 Experimental Setup

Datasets. The CAMERA [33] dataset is generated by rendering and compositing synthetic objects into real scenes in a context-aware manner. In total, there are 300K composite images, where 25K are set aside for evaluation. The training set contains 1085 object instances selected from 6 different categories - *bottle*, *bowl, camera, can, laptop* and *mug.* The evaluation set contains 184 different instances. The REAL [33] dataset is complementary to the CAMERA. It captures 4300 real-world images of 7 scenes for training, and 2750 real-world images of 6

10 Meng Tian, Marcelo H Ang Jr, and Gim Hee Lee

scenes for evaluation. Each set contains 18 real objects spanning the 6 categories. The two evaluation sets are referred to as CAMERA25 and REAL275.

Evaluation Metric. Following [33], we independently evaluate the performance of 3D object detection and 6D pose estimation. We report the average precision at different Intersection-Over-Union (IoU) thresholds for object detection. For 6D pose evaluation, the average precision is computed at $n^{\circ} mcm$. We ignore the rotational error around the axis of symmetry for symmetrical object categories (e.g. *bottle, bowl, and can*). Specially, we treat *mug* as symmetric object in the absence of the handle, and asymmetric object otherwise.

Baseline. Wang et al. [33] is currently the only publicly available code and datasets for the 6D object pose and size estimation task. Futhermore, it is also the state-of-the-art performance on the task. Hence, we choose [33] as our baseline for comparison.

4.2 Implementation Details

We collect all the instances in the CAMERA training dataset to train the autoencoder. Shape priors are learned from this collection and used in all experiments. Each prior consists of 1024 points. We use the Mask R-CNN implemented by matterport [1] for instance segmentation. For each detected instance, we resize the image patch to 192×192 , and randomly sample 1024 points by repetition (if insufficient point count) or downsampling (if sufficient point count). To extract instance color features, we choose the PSPNet [43] with ResNet-18 [8] as backbone. We randomly select 5 point-pairs to generate a hypothesis for the RANSAC-based pose fitting. The maximum number of iteration is 128 and inlier threshold is set to 10% of the object diameter. For the hyperparameters of the total loss, we empirically find that $\lambda_1 = 5.0$, $\lambda_2 = 1.0$, $\lambda_3 = 1e - 4$, and $\lambda_4 = 0.01$ are good choices.

4.3 Comparison to Baseline

We compare our approach to the Baseline [33] on CAMERA25 and REAL275. Quantitative results are summarized in Table 1.

CAMERA25. In the setting of estimating 6D object pose and size from an RGB-D image, we achieve a mAP of 83.1% for 3D IoU at 0.75, and a mAP of 54.3% for 6D pose at 5° 2cm. Our results are 14% and 22% higher than the Baseline [33], respectively. We naively remove the depth input and related sub-networks in our network (i.e. RGB image as the only input) to make fair comparisons with the Baseline [33], which takes an RGB image as its input. As shown in Table 1, our results without depth input are still significantly better than the Baseline [33] (i.e. +15.5% and +17.9%). On one hand, this experiment shows the advantage of explicit handling of the intra-class shape variation, and the effectiveness of our method which reconstructs the object via deformation. On the other hand, it also shows that adding depth to the network does help to improve overall performance, although our improved performance does not rely on it solely. Given that depth image is required to uniquely determine the

Data	Method	mAP						
		$3D_{50}$	$3D_{75}$	$5^{\circ} 2 \mathrm{cm}$	$5^{\circ} 5 \mathrm{cm}$	$10^{\circ} 2 \mathrm{cm}$	$10^{\circ} 5 \mathrm{cm}$	
CAMERA25	Baseline [33]	83.9	69.5	32.3	40.9	48.2	64.6	
	Ours (RGB)	93.1	84.6	50.2	54.5	70.4	78.6	
	Ours (RGB-D)	93.2	83.1	54.3	59.0	73.3	81.5	
REAL275	Baseline [33]	78.0	30.1	7.2	10.0	13.8	25.2	
	Ours (RGB)	75.2	46.5	15.7	18.8	33.7	47.4	
	Ours (RGB-D)	77.3	53.2	19.3	21.4	43.2	54.1	

Table 1. Comparisons on CAMERA25 and REAL275. We report the mAP w.r.t. different thresholds on 3D IoU, and rotation and translation errors.

Table 2. Evaluation of shape reconstruction with CD metric $(\times 10^{-3})$.

Data	Model	Bottle	Bowl	Camera	Can	Laptop	Mug	Average
CAMERA25	Reconstruction Shape Prior	$1.81 \\ 3.41$	$\begin{array}{c} 1.63 \\ 2.20 \end{array}$	$\begin{array}{c} 4.02\\ 9.01 \end{array}$	$\begin{array}{c} 0.97\\ 2.21 \end{array}$	$1.98 \\ 3.27$	$\begin{array}{c} 1.42 \\ 2.10 \end{array}$	$1.97 \\ 3.70$
REAL275	Reconstruction Shape Prior	$3.44 \\ 4.99$	$\begin{array}{c} 1.21 \\ 1.16 \end{array}$	$8.89 \\ 9.85$	$\begin{array}{c} 1.56 \\ 2.38 \end{array}$	$2.91 \\ 7.14$	$\begin{array}{c} 1.02\\ 0.97\end{array}$	$3.17 \\ 4.41$

scale of the object, we recommend it in practical applications. The top row of Fig. 4 shows the average precision at different error thresholds for all 6 object categories. It provides independent analysis for 3D IoU, rotation, and translation error.

REAL275. The REAL training set only contains 3 object instances per category, we enlarge this training set such that the network can generalize well to unseen objects. Following the Baseline [33], we randomly select data from CAM-ERA and REAL training set according to a ratio of 3 : 1. In fair comparison to the Baseline [33], our approach improves the mAP by 23.1% for 3D IoU at 0.75 and 12.1% for 6D pose at 5° 2cm. In strict comparison, we can still outperform by 16.4% and 8.5%, respectively. These results provide further evidence to support our approach. Fig. 4 (bottom) shows more detailed analysis of the errors.

4.4 Evaluation of Shape Reconstruction

To evaluate the quality of the reconstruction, we compute the CD metric (c.f. Eq. 2) of the reconstructed model from our method with the ground truth model in the NOCS. We get a CD metric of 1.97 on CAMERA25 and 3.17 on REAL275. In comparison, the CD metrics are 3.70 and 4.41 on the respective dataset for the shape priors from our autoencoder. The better CD metrics of the reconstructed models compared to the shape priors show that the deformation estimation in our framework improves the quality of the 3D model reconstruction. Table 2



Fig. 4. Average precision vs. error thresholds on CAMERA25 (top row) and REAL275 (bottom row).

shows the CD metric of our reconstructed models and the shape priors for each category.

4.5 Ablation Studies

Different shape priors. We first evaluate how different shape priors influence the performance. All settings are kept the same in this experiment except for the choices of the priors. Results are summarized in Table 3 and 4. "Embedding" refers to the priors obtained from decoding the mean latent embeddings. We also try the instance whose latent embedding has the minimum L_2 distance to the mean latent embedding (denoted as "NN"). In addition, we explore random selection of one instance per category from the shape collection to compose our priors (denoted as "Random"). In general, our approach remains stable under different priors. Our network can adapt to different shape priors because the deformation is explicitly estimated. We achieve the best result for accurate pose (i.e. $5^{\circ} 2$ cm) estimation when the learned categorical shape prior is used. Since our main target is to recover the 6D pose, we choose "Embedding" as our best model. To validate whether the priors are necessary, we use a point cloud uniformly sampled from a sphere of diameter one as our prior (denoted as "None"). The mAP decreases by 3.7% on real dataset when there are no priors, but the best result is achieved for object size estimation. Although shape priors are beneficial for estimating 6D pose, they sometimes bias shape reconstruction. Directly regress the NOCS Coordinates? As indicated by Eq. 4, our approach decouples the NOCS coordinates P to shape reconstruction M and dense

	Notwork	mAP						
	Network	$3D_{50}$	$3D_{75}$	$5^{\circ} 2 \mathrm{cm}$	$5^{\circ} 5 \mathrm{cm}$	$10^{\circ} 2 \mathrm{cm}$	$10^{\circ} 5 \mathrm{cm}$	
Shape Priors	Embedding NN	93.2	83.1	54.3 52.7	59.0 57.3	73.3	81.5 81.3	
	Random	93.3 93.3	85.7 85.8	53.4 54.0	58.0 58.8	72.8	81.0 81.6	
		90.0	05.0	54.0	50.0	73.1	01.0	
NOCS Coords	Regression	93.3	85.3	51.2	55.6	73.8	82.1	
Regularization	w/o Def. w/o Entropy	93.2 93.2	$\begin{array}{c} 85.1\\ 85.1 \end{array}$	$53.9 \\ 53.2$	$58.7 \\ 57.9$	$73.1 \\ 73.2$	$\begin{array}{c} 81.4\\ 81.8\end{array}$	

Table 3. Ablation studies on CAMERA25. Refer to text for more details.

Table 4. Ablation studies on REAL275. Refer to text for more details.

	Notwork	mAP						
	INCOM	$3D_{50}$	$3D_{75}$	$5^{\circ} 2 \mathrm{cm}$	$5^{\circ} 5 \mathrm{cm}$	$10^{\circ} 2 \mathrm{cm}$	$10^{\circ} 5 \mathrm{cm}$	
Shape Priors	Embedding	77.3	53.2	19.3	21.4	43.2	54.1	
	NN	75.9	52.6	17.0	19.0	42.0	51.6	
	Random	75.8	52.2	17.9	20.1	42.3	51.3	
	None	77.2	55.5	15.6	19.8	38.4	53.6	
NOCS Coords	Regression	78.7	54.9	13.7	14.9	42.5	51.4	
Regularization	w/o Def.	77.1	50.2	13.4	15.4	37.3	49.8	
	w/o Entropy	77.3	53.3	15.7	18.8	38.5	51.3	

correspondences A. However, both the network architecture and the training will be much simpler when we follow [33] to regress P directly (denoted as "Regression" in Table 3 and 4). For 6D pose estimation, the mAP of "Regression" at 5° 2cm is notably lower than "Embedding" on CAMERA25 (-3.1%) and REAL275 (-5.6%). This result further supports the benefit of handling shape variation via reconstruction over naive regression of the NOCS coordinates. "Regression" achieves slightly better mAP for object size estimation since it only finds the NOCS coordinates for the observed part, while "Embedding" needs to complete the unknown part of the object.

Regularization losses. To validate the necessity of the two regularization losses, we train the network without regularizing deformation or correspondence, while keeping all the rest same as "Embedding". The mean average precisions of both variants are still comparable to "Embedding" on synthetic dataset. However, mAP of 6D pose at 5° 2cm drops noticeably (-5.9% and -3.6%) on the more difficult real dataset.

4.6 Qualitative Results.

In Fig. 5, we provide several qualitative results from both synthetic and real instances. The 6D pose and object size can be reliably recovered from noisy



Fig. 5. Examples of qualitative results from CAMERA25 (top rows) and REAL275 (bottom rows). For each example, we visualize the results of pose estimation and the reconstructed model M. Our estimations are shown in red, while the ground truths are shown in green.

point correspondences using RANSAC-based pose fitting. Shape reconstruction can capture the variations between instances. The qualities of our predictions are generally better on synthetic data than real data, which is an indication that observation noise needs more attention in our future work. Out of the six categories, *camera* shows less accurate reconstruction due to its more complicated and varying geometry.

5 Conclusions

We present a novel approach for category-level 6D object pose estimation. Our network explicitly models intra-class shape variation by the estimation of the deformation from a shape prior to the object model. Shape priors are learned form a collection of object models and constructed in the latent space. Experiments on synthetic and real datasets demonstrate the advantage of our proposed approach.

Acknowledgements. This research is supported in parts by the Singapore MOE Tier 1 grant R-252-000-A65-114, and the Agency for Science, Technology and Research (A*STAR) under its AME Programmatic Funding Scheme (Project #A18A2b0046).

References

- 1. Abdulla, W.: Mask r-cnn for object detection and instance segmentation on keras and tensorflow. https://github.com/matterport/Mask_RCNN (2017)
- Brachmann, E., Krull, A., Michel, F., Gumhold, S., Shotton, J., Rother, C.: Learning 6d object pose estimation using 3d object coordinates. In: European conference on computer vision. pp. 536–551. Springer (2014)
- 3. Burchfiel, B., Konidaris, G.: Probabilistic category-level pose estimation via segmentation and predicted-shape priors. arXiv preprint arXiv:1905.12079 (2019)
- 4. Chen, D., Li, J., Wang, Z., Xu, K.: Learning canonical shape space for categorylevel 6d object pose and size estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24(6), 381–395 (1981)
- Gupta, S., Arbeláez, P., Girshick, R., Malik, J.: Aligning 3d models to rgb-d images of cluttered scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4731–4740 (2015)
- 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)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hinterstoisser, S., Lepetit, V., Ilic, S., Holzer, S., Bradski, G., Konolige, K., Navab, N.: Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In: Asian conference on computer vision. pp. 548–562. Springer (2012)
- Kehl, W., Manhardt, F., Tombari, F., Ilic, S., Navab, N.: Ssd-6d: Making rgbbased 3d detection and 6d pose estimation great again. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1521–1529 (2017)
- Kehl, W., Milletari, F., Tombari, F., Ilic, S., Navab, N.: Deep learning of local rgbd patches for 3d object detection and 6d pose estimation. In: European Conference on Computer Vision. pp. 205–220. Springer (2016)
- Krull, A., Brachmann, E., Michel, F., Ying Yang, M., Gumhold, S., Rother, C.: Learning analysis-by-synthesis for 6d pose estimation in rgb-d images. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 954–962 (2015)
- Kurenkov, A., Ji, J., Garg, A., Mehta, V., Gwak, J., Choy, C., Savarese, S.: Deformnet: Free-form deformation network for 3d shape reconstruction from a single image. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 858–866. IEEE (2018)
- Lahoud, J., Ghanem, B.: 2d-driven 3d object detection in rgb-d images. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4622–4630 (2017)
- Li, C., Bai, J., Hager, G.D.: A unified framework for multi-view multi-class object pose estimation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 254–269 (2018)
- Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(Nov), 2579–2605 (2008)

- 16 Meng Tian, Marcelo H Ang Jr, and Gim Hee Lee
- Michel, F., Kirillov, A., Brachmann, E., Krull, A., Gumhold, S., Savchynskyy, B., Rother, C.: Global hypothesis generation for 6d object pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 462–471 (2017)
- Peng, S., Liu, Y., Huang, Q., Zhou, X., Bao, H.: Pvnet: Pixel-wise voting network for 6dof pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4561–4570 (2019)
- Pitteri, G., Ramamonjisoa, M., Ilic, S., Lepetit, V.: On object symmetries and 6d pose estimation from images. In: 2019 International Conference on 3D Vision (3DV). pp. 614–622. IEEE (2019)
- Pontes, J.K., Kong, C., Sridharan, S., Lucey, S., Eriksson, A., Fookes, C.: Image2mesh: A learning framework for single image 3d reconstruction. In: Asian Conference on Computer Vision. pp. 365–381. Springer (2018)
- Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum pointnets for 3d object detection from rgb-d data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 918–927 (2018)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 652–660 (2017)
- Rad, M., Lepetit, V.: Bb8: a scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3828–3836 (2017)
- Sahin, C., Kim, T.K.: Category-level 6d object pose recovery in depth images. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (September 2018)
- Song, S., Xiao, J.: Deep sliding shapes for amodal 3d object detection in rgb-d images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 808–816 (2016)
- Sundermeyer, M., Marton, Z.C., Durner, M., Brucker, M., Triebel, R.: Implicit 3d orientation learning for 6d object detection from rgb images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 699–715 (2018)
- Tejani, A., Tang, D., Kouskouridas, R., Kim, T.K.: Latent-class hough forests for 3d object detection and pose estimation. In: European Conference on Computer Vision. pp. 462–477. Springer (2014)
- Tekin, B., Sinha, S.N., Fua, P.: Real-time seamless single shot 6d object pose prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 292–301 (2018)
- Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. IEEE Transactions on Pattern Analysis & Machine Intelligence (4), 376–380 (1991)
- Wallace, B., Hariharan, B.: Few-shot generalization for single-image 3d reconstruction via priors. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3818–3827 (2019)
- Wang, C., Martín-Martín, R., Xu, D., Lv, J., Lu, C., Fei-Fei, L., Savarese, S., Zhu, Y.: 6-pack: Category-level 6d pose tracker with anchor-based keypoints. In: International Conference on Robotics and Automation (ICRA) (2020)
- Wang, C., Xu, D., Zhu, Y., Martín-Martín, R., Lu, C., Fei-Fei, L., Savarese, S.: Densefusion: 6d object pose estimation by iterative dense fusion. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3343–3352 (2019)

- 33. Wang, H., Sridhar, S., Huang, J., Valentin, J., Song, S., Guibas, L.J.: Normalized object coordinate space for category-level 6d object pose and size estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2642–2651 (2019)
- 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)
- Wang, W., Ceylan, D., Mech, R., Neumann, U.: 3dn: 3d deformation network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1038–1046 (2019)
- Wen, C., Zhang, Y., Li, Z., Fu, Y.: Pixel2mesh++: Multi-view 3d mesh generation via deformation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1042–1051 (2019)
- Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. Robotics: Science and Systems (RSS) (2018)
- Yang, B., Luo, W., Urtasun, R.: Pixor: Real-time 3d object detection from point clouds. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 7652–7660 (2018)
- Yang, Y., Feng, C., Shen, Y., Tian, D.: Foldingnet: Point cloud auto-encoder via deep grid deformation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 206–215 (2018)
- Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. In: 2018 International Conference on 3D Vision (3DV). pp. 728–737. IEEE (2018)
- Yumer, M.E., Mitra, N.J.: Learning semantic deformation flows with 3d convolutional networks. In: European Conference on Computer Vision. pp. 294–311. Springer (2016)
- Zakharov, S., Shugurov, I., Ilic, S.: Dpod: 6d pose object detector and refiner. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1941–1950 (2019)
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2881–2890 (2017)
- Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4490–4499 (2018)