




Occlusion Resistant Object Rotation Regression from Point Cloud Segments

Ge Gao^(✉) , Mikko Lauri, Jianwei Zhang, and Simone Frintrop

Department of Informatics, University of Hamburg, Hamburg, Germany
{gao,lauri,zhang,frintrop}@informatik.uni-hamburg.de

Abstract. Rotation estimation of known rigid objects is important for robotic applications such as dexterous manipulation. Most existing methods for rotation estimation use intermediate representations such as templates, global or local feature descriptors, or object coordinates, which require multiple steps in order to infer the object pose. We propose to directly regress a pose vector from point cloud segments using a convolutional neural network. Experimental results show that our method achieves competitive performance compared to a state-of-the-art method, while also showing more robustness against occlusion. Our method does not require any post processing such as refinement with the iterative closest point algorithm.

Keywords: 6D pose estimation · Convolutional neural network · Point cloud · Lie algebra

1 Introduction

The 6D pose of an object is composed of 3D location and 3D orientation. The pose describes the transformation from a local coordinate system of the object to a reference coordinate system (e.g. camera or robot coordinate) [20], as shown in Fig. 1. Knowing the accurate 6D pose of an object is necessary for robotic applications such as dexterous grasping and manipulation. This problem is challenging due to occlusion, clutter and varying lighting conditions.

Many methods for pose estimation using only color information have been proposed [17, 21, 25, 32]. Since depth cameras are commonly used, there have been many methods using both color and depth information [1, 15, 18]. Recently, there are also many CNN based methods [15, 18]. In general, methods that use depth information can handle both textured and texture-less objects, and they are more robust to occlusion compared to methods using only color information [15, 18].

The 6D pose of an object is an inherently continuous quantity. Some works discretize the continuous pose space [8, 9], and formulate the problem as classification. Others avoid discretization by representing the pose using, e.g., quaternions [34], or the axis-angle representation [4, 22]. Work outside the domain of pose estimation has also considered rotation matrices [24], or in a more general case parametric representations of affine transformations [14]. In these cases the




Fig. 1. The goal of 6D pose estimation is to find the translation and rotation from the object coordinate frame O to the camera coordinate frame C .

problem is often formulated as regression. The choice of rotation representation has a major impact on the performance of the estimation method.

In this work, we propose a deep learning based pose estimation method that uses point clouds as an input. To the best of our knowledge, this is the first attempt at applying deep learning for directly estimating 3D rotation using point cloud segments. We formulate the problem of estimating the rotation of a rigid object as regression from a point cloud segment to the axis-angle representation of the rotation. This representation is constraint-free and thus well-suited for application in supervised learning.

Our experimental results show that our method reaches state-of-the-art performance. We also show that our method exceeds the state-of-the-art in pose estimation tasks with moderate amounts of occlusion. Our approach does not require any post-processing, such as pose refinement by the iterative closest point (ICP) algorithm [3]. In practice, we adapt PointNet [24] for the rotation regression task. Our input is a point cloud with spatial and color information. We use the geodesic distance between rotations as the loss function.

The remainder of the paper is organized as follows. Section 2 reviews related work in pose estimation. In Sect. 3, we argue why the axis-angle representation is suitable for supervised learning. We present our system architecture and network details in Sect. 4. Section 5 presents our experimental results. In Sect. 6 we provide concluding remarks and discuss future work.

2 Related Work

6D pose estimation using only RGB information has been widely studied [17, 21, 25, 32]. Since this work concentrates on using point cloud inputs, which contain depth information, we mainly review works that also consider depth information. We also review how depth information can be represented.

2.1 Pose Estimation

RGB-D Methods. A template matching method which integrates color and depth information is proposed by Hinterstoisser et al. [8,9]. Templates are built with quantized image gradients on object contour from RGB information and surface normals on object interior from depth information, and annotated with viewpoint information. The effectiveness of template matching is also shown in [12,19]. However, template matching methods are sensitive to occlusions [18].

Voting-based methods attempt to infer the pose of an object by accumulating evidence from local or global features of image patches. One example is the Latent-Class Hough Forest [30,31] which adapts the template feature from [8] for generating training data. During inference stage, a random set of patches is sampled from the input image. The patches are used in Hough voting to obtain pose hypotheses for verification.

3D object coordinates and object instance probabilities are learned using a Decision Forest in [1]. The 6D pose estimation is then formulated as an energy optimization problem which compares synthetic images rendered with the estimated pose with observed depth values. 3D object coordinates are also used in [18,23]. However, those approaches tend to be very computationally intensive due to generation and verification of hypotheses [18].

Most recent approaches rely on convolutional neural networks (CNNs). In [20], the work in [1] is extended by adding a CNN to describe the posterior density of an object pose. A combination of using a CNN for object segmentation and geometry-based pose estimation is proposed in [16]. PoseCNN [34] uses a similar two-stage network, in which the first stage extracts feature maps from RGB input and the second stage uses the generated maps for object segmentation, 3D translation estimation and 3D rotation regression in quaternion format. Depth data and ICP are used for pose refinement. Jafari et al. [15] propose a three-stage, instance-aware approach for 6D object pose estimation. An instance segmentation network is first applied, followed by an encoder-decoder network which estimates the 3D object coordinates for each segment. The 6D pose is recovered with a geometric pose optimization step similar to [1]. The approaches [15,20,34] do not directly use CNN to predict the pose. Instead, they provide segmentation and other intermediate information, which are used to infer the object pose.

Point Cloud-Based. Drost et al. [5] propose to extract a global model description from oriented point pair features. With the global description, scene data are matched with models using a voting scheme. This approach is further improved by [10] to be more robust against sensor noise and background clutter. Compared to [5,10], our approach uses a CNN to learn the global description.

2.2 Depth Representation

Depth information in deep learning systems can be represented with, e.g., voxel grids [26,28], truncated signed distance functions (TSDF) [29], or point clouds [24]. Voxel grids are simple to generate and use. Because of their regular

grid structure, voxel grids can be directly used as inputs to 3D CNNs. However, voxel grids are inefficient since they also have to explicitly represent empty space. They also suffer from discretization artifacts. TSDF tries to alleviate these problems by storing the shortest distance to the surface represented in each voxel. This allows a more faithful representation of the 3D information. In comparison to other depth data representations, a point cloud has a simple representation without redundancy, yet contains rich geometric information. Recently, PointNet [24] has allowed to use raw point clouds directly as an input of a CNN.

3 Supervised Learning for Rotation Regression

The aim of object pose estimation is to find the translation and rotation that describe the transformation from the object coordinate system \mathcal{O} to the camera coordinate system \mathcal{C} (Fig. 1). The translation consists of the displacements along the three coordinate axes, and the rotation specifies the rotation around the three coordinate axes. Here we concentrate on the problem of estimating rotation.

For supervised learning, we require a loss function that measures the difference between the predicted rotation and the ground truth rotation. To find a suitable loss function, we begin by considering a suitable representation for a rotation. We argue that the axis-angle representation is the best suited for a learning task. We then review the connection of the axis-angle representation to the Lie algebra of rotation matrices. The Lie algebra provides us with tools needed to define our loss function as the geodesic distance of rotation matrices. These steps allow our network to directly make predictions in the axis-angle format.

Notation. In the following, we denote by $(\cdot)^T$ vector or matrix transpose. By $\|\cdot\|_2$, we denote the Euclidean or 2-norm. We write $\mathbf{I}_{3 \times 3}$ for the 3-by-3 identity matrix.

3.1 Axis-Angle Representation of Rotations

A rotation can be represented, e.g., as Euler angles, a rotation matrix, a quaternion, or with the axis-angle representation. Euler angles are known to suffer from gimbal lock discontinuity [11]. Rotation matrices and quaternions have orthogonality and unit norm constraints, respectively. Such constraints may be problematic in an optimization-based approach such as supervised learning, since they restrict the range of valid predictions. To avoid these issues, we adopt the axis-angle representation. In the axis-angle representation, a vector $\mathbf{r} \in \mathbb{R}^3$ represents a rotation of $\theta = \|\mathbf{r}\|_2$ radians around the unit vector $\frac{\mathbf{r}}{\|\mathbf{r}\|_2}$ [7].

3.2 The Lie Group $SO(3)$

The special orthogonal group $SO(3) = \{R \in \mathbb{R}^{3 \times 3} \mid RR^T = \mathbf{I}_{3 \times 3}, \det R = 1\}$ is a compact Lie group that contains the 3-by-3 orthogonal matrices with determinant one, i.e., all rotation matrices [6]. Associated with $SO(3)$ is the Lie algebra $so(3)$, consisting of the set of skew-symmetric 3-by-3 matrices.

Let $\mathbf{r} = [r_1 \ r_2 \ r_3]^T \in \mathbb{R}^3$ be an axis-angle representation of a rotation. The corresponding element of $so(3)$ is the skew-symmetric matrix

$$\mathbf{r}_\times = \begin{bmatrix} 0 & -r_3 & r_2 \\ r_3 & 0 & -r_1 \\ -r_2 & r_1 & 0 \end{bmatrix}. \quad (1)$$

The *exponential map* $\exp : so(3) \rightarrow SO(3)$ connects the Lie algebra with the Lie group by

$$\exp(\mathbf{r}_\times) = \mathbf{I}_{3 \times 3} + \frac{\sin \theta}{\theta} \mathbf{r}_\times + \frac{1 - \cos \theta}{\theta^2} \mathbf{r}_\times^2, \quad (2)$$

where $\theta = \mathbf{r}^T \mathbf{r} = \|\mathbf{r}\|_2$ as above¹.

Now let R be a rotation matrix in the Lie group $SO(3)$. The *logarithmic map* $\log : SO(3) \rightarrow so(3)$ connects R with an element in the Lie algebra by

$$\log(R) = \frac{\phi(R)}{2 \sin(\phi(R))} (R - R^T), \quad (3)$$

where

$$\phi(R) = \arccos\left(\frac{\text{trace}(R) - 1}{2}\right) \quad (4)$$

can be interpreted as the magnitude of rotation related to R in radians. If desired, we can now obtain an axis-angle representation of R by first extracting from $\log(R)$ the corresponding elements indicated in Eq. (1), and then setting the norm of the resulting vector to $\phi(R)$.

3.3 Loss Function for Rotation Regression

We regress to a predicted rotation $\hat{\mathbf{r}}$ represented in the axis-angle form. The prediction is compared against the ground truth rotation \mathbf{r} via a loss function $l : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_{\geq 0}$. Let \hat{R} and R denote the two rotation matrices corresponding to $\hat{\mathbf{r}}$ and \mathbf{r} , respectively. We use as loss function the geodesic distance $d(\hat{R}, R)$ of \hat{R} and R [7, 13], i.e.,

$$l(\hat{\mathbf{r}}, \mathbf{r}) = d(\hat{R}, R) = \phi(\hat{R}R^T), \quad (5)$$

where we first obtain \hat{R} and R via the exponential map, and then calculate $\phi(\hat{R}R^T)$ to obtain the loss value. This loss function directly measures the magnitude of rotation between \hat{R} and R , making it convenient to interpret. Furthermore, using the axis-angle representation allows to make predictions free of constraints such as the unit norm requirement of quaternions. This makes the loss function also convenient to implement in a supervised learning approach.

¹ In a practical implementation, the Taylor expansions of $\frac{\sin \theta}{\theta}$ and $\frac{1 - \cos \theta}{\theta^2}$ should be used for small θ for numerical stability.

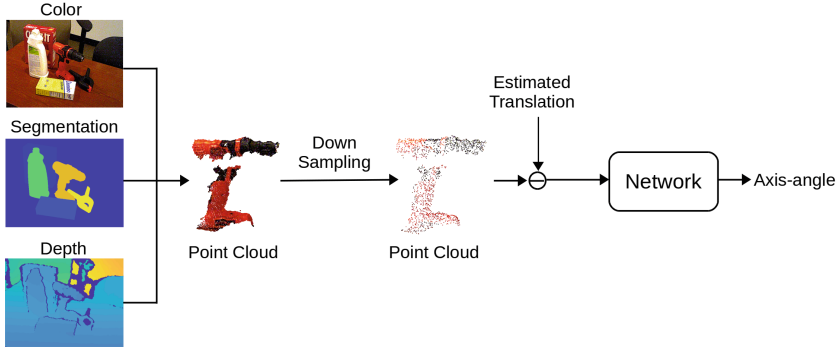


Fig. 2. System overview. The color and depth images together with a segmentation of the target object are used to create a point cloud. The segment is randomly down-sampled, and the estimated translation of the down-sampled segment is removed. The normalized segment is fed into a network for rotation prediction.

4 System Architecture

Figure 2 shows the system overview. We train our system for a specific target object, in Fig. 2 the drill. The inputs to our system are the RGB color image, the depth image, and a segmentation mask indicating which pixels belong to the target object. We first create a point cloud segment of the target object based on the inputs. Each point has 6 dimensions: 3 dimensions for spatial coordinates and 3 dimensions for color information. We randomly sample n points from this point cloud segment to create a fixed-size downsampled point cloud. In all of our experiments, we use $n = 256$. We then remove the estimated translation from the point coordinates to normalize the data. The normalized point cloud segment is then fed into a network which outputs a rotation prediction in the axis-angle format. During training, we use the ground truth segmentation and translation. As we focus on the rotation estimation, during testing, we apply the segmentation and translation outputs of PoseCNN [34].

We consider two variants for our network presented in the following subsections. The first variant processes the point cloud as a set of independent points without regard to the local neighbourhoods of points. The second variant explicitly takes into account the local neighbourhoods of a point by considering its nearest neighbours.

4.1 PointNet (PN)

Our PN network is based on PointNet [24], as illustrated in Fig. 3. The PointNet architecture is invariant to all $n!$ possible permutations of the input point cloud, and hence an ideal structure for processing raw point clouds. The invariance is achieved by processing all points independently using multi-layer perceptrons (MLPs) with shared weights. The obtained feature vectors are finally max-pooled to create a global feature representation of the input point cloud. Finally, we attach a three-layer regression MLP on top of this global feature to predict the rotation.




Fig. 3. Network architecture. The numbers in parentheses indicate number of MLP layers, and numbers not in parentheses indicate intermediate vector dimensionality. A feature vector for each point is learned using shared weights. A max pooling layer then aggregates the individual features into a global feature vector. Finally, a regression network with 3 fully-connected layers outputs the rotation prediction.

4.2 Dynamic Nearest Neighbour Graph (DG)

In the PN architecture, all features are extracted based only on a single point. Hence it does not explicitly consider the local neighbourhoods of individual points. However, local neighbourhoods can contain useful geometric information for pose estimation [27]. The local neighbourhoods are considered by an alternative network structure based on the dynamic nearest-neighbour graph network proposed in [33]. For each point P_i in the point set, a k -nearest neighbor graph is calculated. In all our experiments, we use $k = 10$. The graph contains directed edges $(i, j_{i1}), \dots, (i, j_{ik})$, such that $P_{j_{i1}}, \dots, P_{j_{ik}}$ are the k closest points to P_i . For an edge e_{ij} , an edge feature $[P_i, (P_j - P_i)]^T$ is calculated. The edge features are then processed in a similar manner as in PointNet to preserve permutation invariance. This dynamic graph convolution can then be repeated, now calculating the nearest neighbour graph for the feature vectors of the first shared MLP layer, and so on for the subsequent layers. We use the implementation² provided by authors from [33], and call the resulting network DG for dynamic graph.

5 Experimental Results

This section shows experimental results of the proposed approach on the YCB video dataset [34], and compares the performance with state-of-the-art PoseCNN method [34]. Besides prediction accuracy, we investigate the effect of occlusions and the quality of the segmentation and translation estimates.

5.1 Experiment Setup

The YCB video dataset [34] is used for training and testing with the original train/test split. The dataset contains 133,827 frames of 21 objects selected from the YCB object set [2] with 6D pose annotation. 80,000 frames of synthetic data are also provided as an extension to the training set.

We select a set of four objects to test on, shown in Fig. 4. As our approach does not consider object symmetry, we use objects that have 1-fold rotational

² <https://github.com/WangYueFt/dgcnn>.



Fig. 4. Testing objects. From left to right: power drill, extra large clamp, banana, pitcher base.

symmetry (power drill, banana and pitcher base) or 2-fold rotational symmetry (extra large clamp).

We run all experiments using both the PointNet based (PN) and dynamic graph (DG) networks. During training, Adam optimizer is used with learning rate 0.008, and batch size of 128. Batch normalization is applied to all layers. No dropout is used.

For training, ground truth segmentations and translations are used as the corresponding inputs shown in Fig. 2. While evaluating 3D rotation estimation in Subsect. 5.3, the translation and segmentation predicted by PoseCNN are used.

5.2 Evaluation Metrics

For evaluating rotation estimation, we directly use geodesic distance described in Sect. 3 to quantify the rotation error. We evaluate 6D pose estimation using average distance of model points (ADD) proposed in [9]. For a 3D model \mathcal{M} represented as a set of points, with ground truth rotation R and translation \mathbf{t} , and estimated rotation \hat{R} and translation $\hat{\mathbf{t}}$, the ADD is defined as:

$$\text{ADD} = \frac{1}{m} \sum_{\mathbf{x} \in \mathcal{M}} \left\| (R\mathbf{x} + \mathbf{t}) - (\hat{R}\mathbf{x} + \hat{\mathbf{t}}) \right\|_2, \quad (6)$$

where m is the number of points. The 6D pose estimate is considered to be correct if ADD is smaller than a given threshold.

5.3 Rotation Estimation

Figure 5 shows the estimation accuracy as function of the rotation angle error threshold, i.e., the fraction of predictions that have an angle error smaller than the horizontal axis value. Results are shown for PoseCNN, PoseCNN with ICP refinement (PoseCNN + ICP), and our method with PointNet structure (PN), and with dynamic graph structure (DG). To determine the effect of the translation and segmentation input, we additionally test our methods while giving the ground truth translation and segmentation as input. The cases with ground truths provided are indicated by +gt, and shown with a dashed line.

With translation and segmentation results from PoseCNN, our methods show competitive or superior results compared to PoseCNN with ICP refinement. This

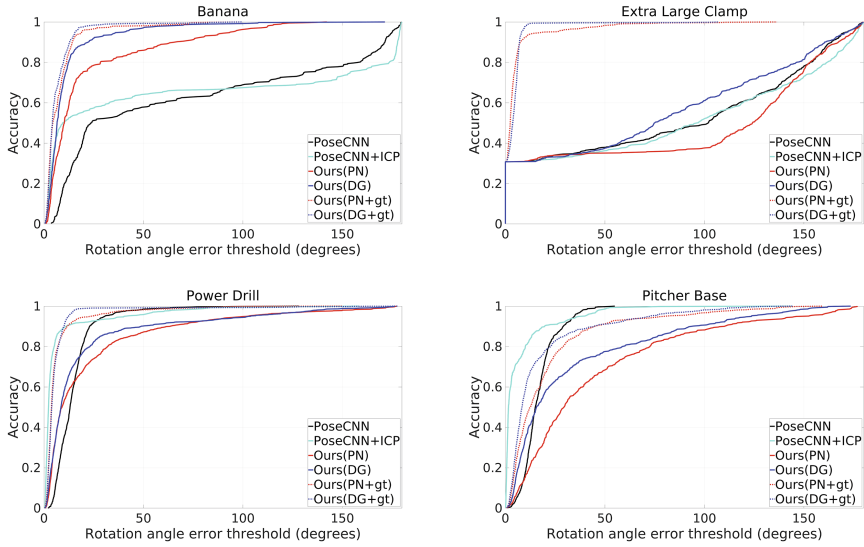


Fig. 5. Accuracy of rotation angle prediction shows the fraction of predictions with error smaller than the threshold. Results are shown for our method and PoseCNN [34]. The additional +gt denotes the variants where ground truth segmentation is provided.

demonstrates that our network is able to accurately predict rotation, and is able to do so without any post-processing or ICP-based pose refinement. We also note that in cases where our method does not work very well (e.g., extra large clamp), by providing the ground truth translation and segmentation (+gt), the results are greatly improved. This shows that good translation and segmentation are crucial for accurate rotation estimation. For pitcher base, our method does not perform well. One possible explanation is that information about the handle and water outlet parts of the pitcher, which are very discriminative for determining the pitcher’s rotation, may be lost in our downsampling step. In future work, we are planning to investigate other sampling methods such as farthest point sampling to ensure a full view of the object is preserved even with downsampling.

The results also confirm the fact that ICP based refinement usually only improves the estimation quality if the initial guess is already good enough. When the initial estimation is not accurate enough, the use of ICP can even decrease the accuracy, as shown by the PoseCNN + ICP curve falling below the PoseCNN curve for large angle thresholds.

Effect of Occlusion. We quantify the effect of occlusion on the rotation prediction accuracy. For a given frame and target object, we estimate the occlusion factor O of the object by

$$O = 1 - \frac{\lambda}{\mu}, \quad (7)$$

Table 1. Average rotation angle error in degrees with 95% confidence interval in frames with low ($O < 0.2$) and moderate (mod, $O \geq 0.2$) occlusion

Object	Banana		Power drill		Extra large clamp	
Occlusion	low	mod	low	mod	low	mod
PoseCNN [34]	$62.0^\circ \pm 3.1^\circ$	$8.2^\circ \pm 0.25^\circ$	$14.7^\circ \pm 0.3^\circ$	$37.4^\circ \pm 2.4^\circ$	$109.8^\circ \pm 2.0^\circ$	$151.0^\circ \pm 3.6^\circ$
PoseCNN+ICP	$56.5^\circ \pm 3.4^\circ$	$7.1^\circ \pm 0.9^\circ$	$6.9^\circ \pm 0.4^\circ$	$44.1^\circ \pm 3.5^\circ$	$115.5^\circ \pm 2.0^\circ$	$140.5^\circ \pm 6.0^\circ$
Ours (PN)	$25.1^\circ \pm 1.4^\circ$	$53.5^\circ \pm 3.8^\circ$	$23.2^\circ \pm 0.8^\circ$	$34.8^\circ \pm 6.2^\circ$	$125.7^\circ \pm 1.8^\circ$	$158.1^\circ \pm 2.3^\circ$
Ours (DG)	$12.6^\circ \pm 0.8^\circ$	$6.7^\circ \pm 0.4^\circ$	$18.6^\circ \pm 0.8^\circ$	$49.1^\circ \pm 6.7^\circ$	$96.9^\circ \pm 1.9^\circ$	$100.8^\circ \pm 6.1^\circ$
Ours (PN+gt)	$8.6^\circ \pm 0.5^\circ$	$10.6^\circ \pm 1.8^\circ$	$7.0^\circ \pm 0.3^\circ$	$3.8^\circ \pm 0.3^\circ$	$6.2^\circ \pm 0.4^\circ$	$36.4^\circ \pm 3.7^\circ$
Ours (DG+gt)	$7.4^\circ \pm 0.4^\circ$	$4.3^\circ \pm 0.2^\circ$	$5.8^\circ \pm 0.3^\circ$	$5.3^\circ \pm 0.6^\circ$	$4.7^\circ \pm 0.1^\circ$	$3.5^\circ \pm 0.1^\circ$

where λ is the number of pixels in the 2D ground truth segmentation, and μ is the number of pixels in the projection of the 3D model of the object onto the image plane using the camera intrinsic parameters and the ground truth 6D pose, when we assume that the object would be fully visible. We noted that for the test frames of the YCB-video dataset O is mostly below 0.5. We categorize $O < 0.2$ as low occlusion and $O \geq 0.2$ as moderate occlusion.

Table 1 shows the average rotation angle error (in degrees) and its 95% confidence interval³ for PoseCNN and our method in the low and moderate occlusion categories. We also investigated the effect of the translation and segmentation by considering variants of our methods that were provided with the ground truth translation and segmentation. These variants are shown in the table indicated by +gt. We observe that in the moderate occlusion category, our methods have significantly better performance than PoseCNN. We note that for the extra large clamp, the results are greatly improved if the ground truths are provided. This indicates that the failure of both PoseCNN and our method for extra large clamp is due to the poor quality translation and segmentation. Furthermore, with the dynamic graph architecture (DG), the average error tends to be lower. This shows the local neighbourhood information extracted by DG is useful for rotation estimation. One observation is that for banana, the rotation error in low occlusion is significantly higher than it is in the moderate case for PoseCNN. This is because near 25% of the test frames in low occlusion case present an rotation error in range of 160° to 180° .

Qualitative results for rotation estimation are shown in Fig. 6. In the leftmost column, the occlusion factor O of the target object is denoted. Then, from left to right, we show the ground truth, PoseCNN+ICP, and our method using DG and our method using DG with ground truth translation and segmentation (DG+gt) results. In all cases, the ground truth pose, or respectively, the pose estimate, are indicated by the green overlay on the figures. To focus on the difference in the rotation estimate, we use the ground truth translation for all methods for the visualization. The rotation predictions for Ours (DG) are still based on translation and segmentation from PoseCNN.

³ The results for pitcher base are not reported here since all samples in testing set for pitcher base have low occlusion.




Fig. 6. Qualitative results for rotation estimation. The number on the left indicates the occlusion factor O for the target object. Then, from left to right: ground truth, PoseCNN [34] with ICP refinement, our method using dynamic graph (DG) with PoseCNN segmentation, and dynamic graph with ground truth segmentation (DG+gt). The green overlay indicates the ground truth pose, or respectively, the predicted pose of the target object. Ground truth translation is used in all cases.

The first two rows of Fig. 6 show cases with moderate occlusion. When the discriminative part of the banana is occluded (top row), PoseCNN can not recover the rotation, while our method still produces a good estimate. The situation is similar in the second row for the drill. The third row illustrates that the quality of segmentation has a strong impact on the accuracy of rotation estimation. In this case the segmentation fails to detect the black clamp on the black background, which leads to a poor rotation estimate for both PoseCNN and our method. When we provide the ground truth segmentation (third row, last column), our method is able to recover the correct rotation. The final fourth row shows a failure case for the pitcher base. Our method fails while it loses information about the discriminative handle and water outlet parts of the pitcher in the subsampling phase.

6 Conclusion

We propose to directly predict the 3D rotation of a known rigid object from a point cloud segment. We use axis-angle representation of rotations as the regression target. Our network learns a global representation either from individual input points, or from point sets of nearest neighbors. Geodesic distance is used as the loss function to supervise the learning process. Without using ICP refinement, experiments shows that the proposed method can reach competitive and sometimes superior performance compared to PoseCNN.

Our results show that point cloud segments contain enough information for inferring object pose. The axis-angle representation does not have any constraints, making it a suitable regression target. Using Lie algebra as a tool provides a valid distance measure for rotations. This distance measure can be used as a loss function during training.

We discovered that the performance of our method is strongly affected by the quality of the target object translation and segmentation, which will be further investigated in future work. We will extend the proposed method to full 6D pose estimation by additionally predicting the object translations. We also plan to integrate object classification into our system, and study a wider range of target objects.

Acknowledgments. This work was partially funded by the German Science Foundation (DFG) in project Crossmodal Learning, TRR 169.

References

1. Brachmann, E., Krull, A., Michel, F., Gumhold, S., Shotton, J., Rother, C.: Learning 6D object pose estimation using 3D object coordinates. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8690, pp. 536–551. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10605-2_35
2. Calli, B., Walsman, A., Singh, A., Srinivasa, S., Abbeel, P.: Benchmarking in manipulation research using the Yale-CMU-Berkeley object and model set. *Robot. Autom. Mag. IEEE* **22**(3), 36–52 (2015)
3. Chen, Y., Medioni, G.: Object modelling by registration of multiple range images. *Image Vis. Comput.* **10**(3), 145–155 (1992)
4. Do, T., Cai, M., Pham, T., Reid, I.: Deep-6DPose: recovering 6D object pose from a single RGB image. arXiv preprint [arXiv:1802.10367](https://arxiv.org/abs/1802.10367) (2018)
5. Drost, B., Ulrich, M., Navab, N., Ilic, S.: Model globally, match locally: efficient and robust 3D object recognition. In: CVPR (2010)
6. Hall, B.C.: Lie Groups, Lie Algebras, and Representations. GTM, vol. 222. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-13467-3>
7. Hartley, R., Trunpf, J., Dai, Y., Li, H.: Rotation averaging. *Int. J. Comput. Vis.* **103**(3), 267–305 (2013)
8. Hinterstoisser, S., et al.: Multimodal templates for real-time detection of textureless objects in heavily cluttered scenes. In: ICCV (2011)

9. Hinterstoisser, S., et al.: Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes. In: Lee, K.M., Matsushita, Y., Rehg, J.M., Hu, Z. (eds.) ACCV 2012. LNCS, vol. 7724, pp. 548–562. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-37331-2_42
10. Hinterstoisser, S., Lepetit, V., Rajkumar, N., Konolige, K.: Going further with point pair features. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9907, pp. 834–848. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_51
11. Hoag, D.: Apollo guidance and navigation: considerations of Apollo IMU gimbal lock, pp. 1–64. MIT Instrumentation Laboratory, Cambridge (1963)
12. Hodaň, T., Zabulis, X., Lourakis, M., Obdržálek, S., Matas, J.: Detection and fine 3D pose estimation of texture-less objects in RGB-D images. In: IROS (2015)
13. Huynh, D.Q.: Metrics for 3D rotations: comparison and analysis. *J. Math. Imag. Vis.* **35**(2), 155–164 (2009)
14. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: NIPS (2015)
15. Jafari, H., Mustikovela, S.K., Pertsch, K., Brachmann, E., Rother, C.: iPose: instance-aware 6D pose estimation of partly occluded objects. arXiv preprint [arXiv:1712.01924](https://arxiv.org/abs/1712.01924) (2018)
16. Jafari, O.H., Mustikovela, S.K., Pertsch, K., Brachmann, E., Rother, C.: The best of both worlds: learning geometry-based 6D object pose estimation. arXiv preprint [arXiv:1712.01924](https://arxiv.org/abs/1712.01924) (2017)
17. Kehl, W., Manhardt, F., Tombari, F., Ilic, S., Navab, N.: SSD-6D: making RGB-based 3D detection and 6D pose estimation great again. In: ICCV (2017)
18. Kehl, W., Milletari, F., Tombari, F., Ilic, S., Navab, N.: Deep learning of local RGB-D patches for 3D object detection and 6D pose estimation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9907, pp. 205–220. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_13
19. Kehl, W., Tombari, F., Navab, N., Ilic, S., Lepetit, V.: Hashmod: a hashing method for scalable 3D object detection. In: BMVC (2015)
20. Krull, A., Brachmann, E., Michel, F., Yang, M.Y., Gumhold, S., Rother, C.: Learning analysis-by-synthesis for 6D pose estimation in RGB-D images. In: ICCV (2015)
21. Li, Y., Wang, G., Ji, X., Xiang, Y., Fox, D.: DeepIM: deep iterative matching for 6D pose estimation. arXiv preprint [arXiv:1804.00175](https://arxiv.org/abs/1804.00175) (2018)
22. Mahendran, S., Ali, H., Vidal, R.: 3D pose regression using convolutional neural networks. In: ICCV (2017)
23. Michel, F., et al.: Global hypothesis generation for 6D object pose estimation. In: CVPR (2017)
24. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: deep learning on point sets for 3D classification and segmentation. In: CVPR (2017)
25. 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: ICCV (2017)
26. Riegler, G., Ulusoy, A.O., Geiger, A.: OctNet: learning deep 3D representations at high resolutions. In: CVPR (2017)
27. Rusu, R., Bradski, G., Thibaux, R., Hsu, J.: Fast 3D recognition and pose using the viewpoint feature histogram. In: IROS (2010)
28. Sedaghat, N., Zolfaghari, M., Amiri, E., Brox, T.: Orientation-boosted voxel nets for 3D object recognition. In: BMVC (2017)
29. Song, S., Xiao, J.: Deep sliding shapes for Amodal 3D object detection in RGB-D images. In: CVPR (2016)

30. Tejani, A., Kouskouridas, R., Doumanoglou, A., Tang, D., Kim, T.: Latent-class Hough forests for 6 DoF object pose estimation. *PAMI* **40**(1), 119–132 (2018)
31. Tejani, A., Tang, D., Kouskouridas, R., Kim, T.-K.: Latent-class Hough forests for 3D object detection and pose estimation. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8694, pp. 462–477. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10599-4_30
32. Tekin, B., Sinha, S.N., Fua, P.: Real-time seamless single shot 6D object pose prediction. In: *CVPR* (2018)
33. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph CNN for learning on point clouds. arXiv preprint [arXiv:1801.07829](https://arxiv.org/abs/1801.07829) (2018)
34. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: PoseCNN: a convolutional neural network for 6D object pose estimation in cluttered scenes. In: *RSS* (2018)