
Supplementary Materials of Generative Category-level Object Pose Estimation via Diffusion Models

Anonymous Author(s)

Affiliation

Address

email

1	Contents	
2	1 Implementation Details	2
3	2 Qualitative Comparison on REAL275	2
4	3 More Results and Analysis	2
5	3.1 Per-category Results	2
6	3.2 Results on CAMERA	3
7	3.3 Real World Experiments	4
8	4 Ethics Statement and Boarder Impact	5

1 Implementation Details

Architecture of the Score Network The detailed architecture of the score network Φ_θ is illustrated in Figure 1. We utilize PointNet++ [1] to extract the global geometry feature \mathcal{F}_O of the partially observed point cloud O^* . And the sampled pose \mathbf{p} and timestep t features are embedded as \mathcal{F}_p and \mathcal{F}_t , respectively, using a Multi-Layer Perceptron (MLP). Then \mathcal{F}_O , \mathcal{F}_p and \mathcal{F}_t are concatenated to obtain the global feature \mathcal{F} , and three parallel branches are employed to predict the scores of R_x , R_y , and T individually, where $[R_x|R_y] \in \mathbb{R}^6$ and $T \in \mathbb{R}^3$ denote rotation and translation vectors, respectively. And $[R_x|R_y]$ is a continuous rotation representation proposed by [2] to address the discontinuity of quaternions and Euler angles in Euclidean space. As introduced in [2], the mapping from $\text{SO}(3)$ to the 6D representation of rotation is:

$$g_{GS}([\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_3]) = [\mathbf{a}_1 \ \mathbf{a}_2] \quad (1)$$

The mapping from the 6D representation to $\text{SO}(3)$ is:

$$f_{GS}([\mathbf{a}_1 \ \mathbf{a}_2]) = [\mathbf{b}_1 \ \mathbf{b}_2 \ \mathbf{b}_3] \quad (2)$$

$$b_i = \begin{cases} N(\mathbf{a}_1) & \text{if } i = 1 \\ N(\mathbf{a}_2 - (\mathbf{b}_1 \cdot \mathbf{a}_2)\mathbf{b}_1) & \text{if } i = 2 \\ \mathbf{b}_1 \times \mathbf{b}_2 & \text{if } i = 3 \end{cases} \quad (3)$$

Here $N(\cdot)$ denotes a normalization function.

Architecture of the Energy Network The energy network Ψ_ϕ shares exactly the same architecture with the score network Φ_θ . The inputs are first fed into Φ_θ to obtain a score-shaped vector $\Phi_\theta(\mathbf{p}, t|O) \in \mathbb{R}^{|\mathcal{P}|}$. Then, the output energy is calculated by the dot product between the input pose and the score-shaped vector $\Psi_\phi(\mathbf{p}, t|O) = \langle \mathbf{p}, \Phi_\theta(\mathbf{p}, t|O) \rangle \in \mathbb{R}^1$.

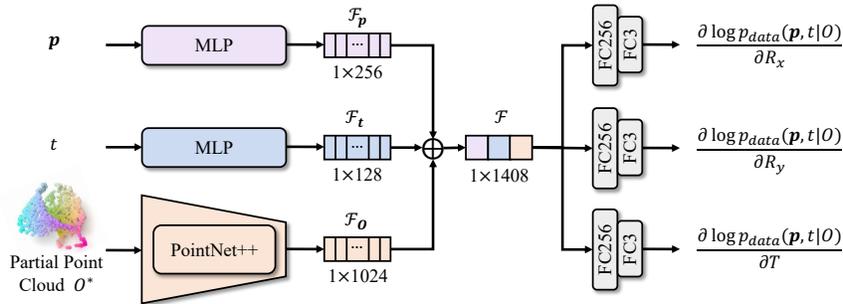


Figure 1: **Architecture of the score network Φ_θ .** \mathbf{p} denotes sampled 6D object poses. O^* denotes partially observed 3D point cloud condition. t denotes timestep. \oplus denotes the concatenation operator.

2 Qualitative Comparison on REAL275

Figure 2 illustrates the qualitative comparison results between our method and RBP-Pose [3] on the REAL275 dataset. The images are accompanied by red boxes highlighting objects that exhibit noticeable differences in the predicted results. Additionally, the bottom-right corner of each image provides an enlarged view of the highlighted object, showing the ground truth pose as well as the poses estimated by RBP-Pose and our approach. Our method demonstrates a significant performance improvement compared to RBP-Pose, particularly in the case of objects such as mugs. Notably, in the fourth row of the figure, it can be observed that our method achieves highly accurate poses even when only a small portion of the mug handle is visible. This success can be attributed to the fact that, during the training process, a unique pose exists when the mug handle is visible. However, when the mug handle becomes occluded, a multi-hypothesis problem arises, which our generative formulation effectively handles.

3 More Results and Analysis

3.1 Per-category Results

Figure 3 demonstrates a quantitative comparison between our method and the state-of-the-art depth-based approach, RBP-Pose [3], for various object categories at different thresholds. The results

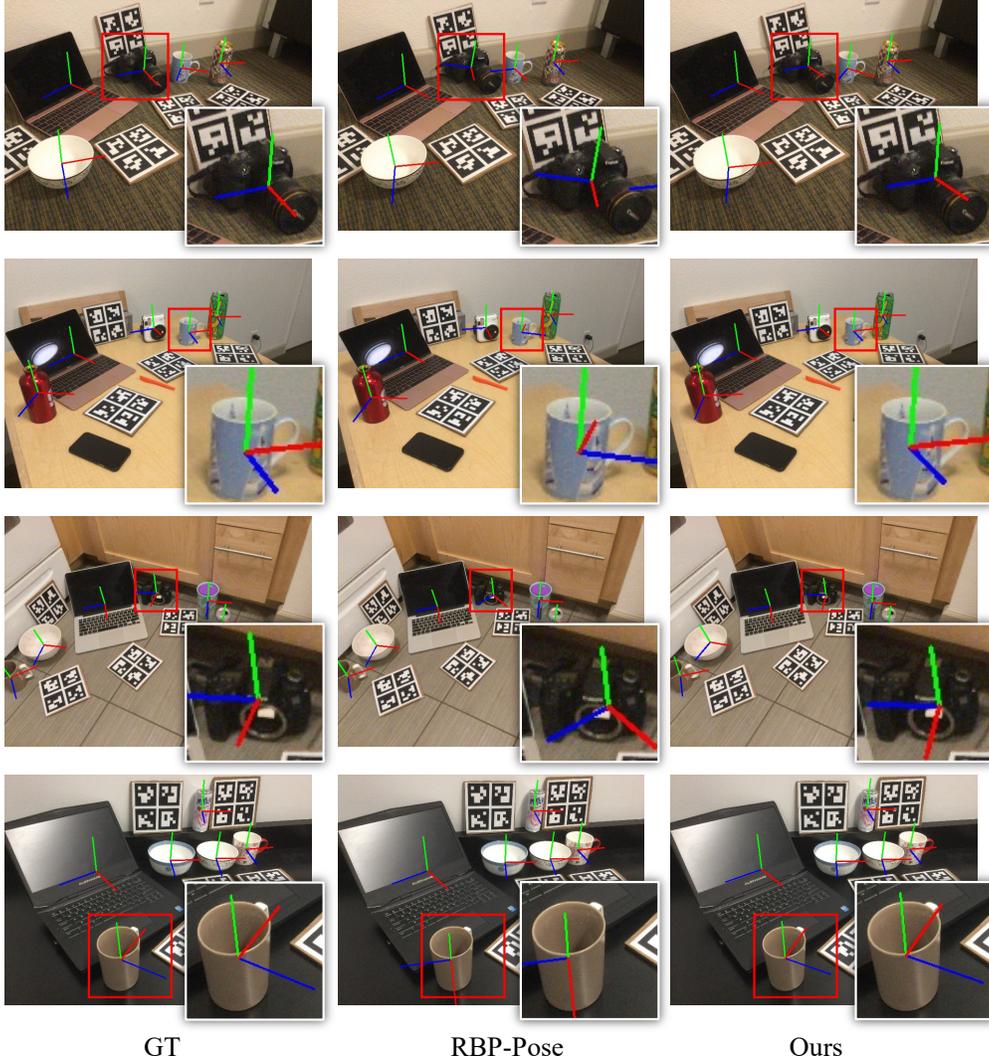


Figure 2: **Qualitative comparison with RBP-Pose [3] on REAL275.** The left column represents the ground truth pose, the middle column represents the results of RBP-Pose, the right column represents the results of our approach.

42 clearly indicate that our method outperforms RBP-Pose in all metrics, despite the fact that we do not
 43 incorporate augmentation specifically designed for symmetric objects during the training phase, unlike
 44 RBP-Pose. Our approach exhibits significant improvements, particularly in regions with stringent
 45 threshold requirements. This emphasizes the superior performance of our generative category-level
 46 object 6D pose estimation approach in effectively addressing the multi-hypothesis challenges posed
 47 by symmetric objects and partial observations, thereby enabling its successful application in robot
 48 manipulation tasks demanding precise object pose prediction. (*e.g.*, pouring liquids.)

49 3.2 Results on CAMERA

50 Table 1 illustrates a quantitative comparison between our method and the baselines on the CAM-
 51 ERA [4] dataset. The results clearly demonstrate the remarkable performance enhancement achieved
 52 by our method. When compared to approaches that rely solely on depth data as network input, as well
 53 as those that utilize RGB-D and shape priors as network input, our method consistently outperforms
 54 them, surpassing the current state-of-the-art performance. Notably, our method exhibits a particularly
 55 pronounced advantage when stricter accuracy requirements are imposed, such as the $5^\circ 2cm$ metric.
 56 In this case, our method outperforms the current SOTA method, RBP-Pose, by an impressive margin
 57 of 6.4%. This significant improvement highlights the efficacy of our approach.

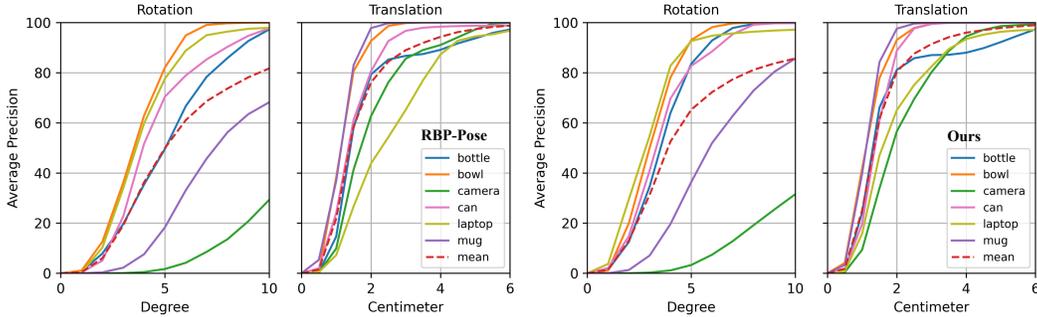


Figure 3: **Per-category quantitative comparison with RBP-Pose [3] on REAL275.** The left represents the results of RBP-Pose, while the right represents the results of our approach.

Table 1: **Quantitative comparison of category-level object pose estimation on CAMERA [4] dataset.** We summarize the results reported in the original paper for the baseline method. \uparrow represents a higher value indicating better performance, while \downarrow represents a lower value indicating better performance. **Data** refers to the format of the input data used by the method, and **Prior** indicates whether the method requires category prior information. ‘-’ indicates that the metrics are not reported in the original paper. **K** represents the number of hypotheses."

Method		Data	Prior	5°2cm \uparrow	5°5cm \uparrow	10°2cm \uparrow	10°5cm \uparrow	Parameters(M) \downarrow
Deterministic	NOCS [4]	RGB-D	\times	32.3	40.9	48.2	64.6	-
	DualPoseNet [5]	RGB-D	\times	64.7	70.7	77.2	84.7	67.9
	SPD [6]	RGB-D	\checkmark	54.3	59.0	73.3	81.5	18.3
	CR-Net [7]	RGB-D	\checkmark	72.0	76.4	81.0	87.7	-
	SGPA [8]	RGB-D	\checkmark	70.7	74.5	82.7	88.4	-
	GPV-Pose [9]	D	\times	72.1	79.1	-	89.0	-
	SAR-Net [10]	D	\checkmark	66.7	70.9	75.3	80.3	6.3
	SSP-Pose [11]	D	\checkmark	64.7	75.5	-	87.4	-
	RBP-Pose [3]	D	\checkmark	73.5	79.6	82.1	89.5	-
	Ours	D	\times	79.9	84.4	84.6	89.6	4.4
Probabilistic	Ours($K=10$)	D	\times	90.8	93.0	93.4	95.7	2.2
	Ours($K=50$)	D	\times	95.5	96.4	97.2	98.2	2.2

58 3.3 Real World Experiments

59 We have also successfully integrated our approach with
60 robot manipulation capabilities, as demonstrated through
61 various experiments conducted with the UFACTORY
62 xArm6 equipped with RealSense D435. **The demon-**
63 **strations can be found in the supplementary video**
64 **or on the project website.** As shown in Figure 4, we
65 illustrate the following three tasks:

66 **Pouring Task.** This task involves transferring the con-
67 tents (*e.g.*, water) from one container to another. The
68 demonstration highlights the potential of combining our
69 approach with heuristic strategies, enabling functional
70 robot operations.

71 **Stacking Task.** In this task, we focused on piling up
72 objects of the same category, like organizing scattered
73 bowls on a tabletop. This demonstrates the precision
74 of the estimated object pose, as accurate knowledge of
75 object poses is crucial for completing this task.

76 **Handover Task.** This task involved either receiving
77 objects from human hands to perform tasks or passing ob-
78 jects to person. The demonstration exemplified one form
79 of human-robot interaction empowered by our method.

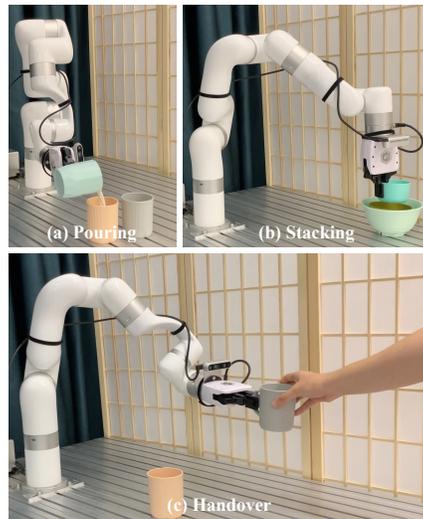


Figure 4: **Pose estimation for robot ma-**
nipulation tasks. We demonstrate three
types of tasks.

80 4 Ethics Statement and Boarder Impact

81 Our method has the potential to develop home-assisting robot, thus contributing to social welfare.
82 We evaluate our method in synthesized or human-collected datasets, which may introduce data bias.
83 However, similar studies also have such general concerns. We do not see any possible major harm in
84 our study.

85 References

- 86 [1] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical
87 feature learning on point sets in a metric space. *Advances in neural information processing*
88 *systems*, 30, 2017.
- 89 [2] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation
90 representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer*
91 *Vision and Pattern Recognition*, pages 5745–5753, 2019.
- 92 [3] Ruida Zhang, Yan Di, Zhiqiang Lou, Fabian Manhardt, Federico Tombari, and Xiangyang
93 Ji. Rbp-pose: Residual bounding box projection for category-level pose estimation. In *Com-*
94 *puter Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022,*
95 *Proceedings, Part I*, pages 655–672. Springer, 2022.
- 96 [4] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas.
97 Normalized object coordinate space for category-level 6d object pose and size estimation. In
98 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
99 2642–2651, 2019.
- 100 [5] Jiehong Lin, Zewei Wei, Zhihao Li, Songcen Xu, Kui Jia, and Yuanqing Li. Dualposenet:
101 Category-level 6d object pose and size estimation using dual pose network with refined learning
102 of pose consistency. In *Proceedings of the IEEE/CVF International Conference on Computer*
103 *Vision*, pages 3560–3569, 2021.
- 104 [6] Meng Tian, Marcelo H Ang, and Gim Hee Lee. Shape prior deformation for categorical 6d
105 object pose and size estimation. In *Computer Vision–ECCV 2020: 16th European Conference,*
106 *Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 530–546. Springer, 2020.
- 107 [7] Jiaze Wang, Kai Chen, and Qi Dou. Category-level 6d object pose estimation via cascaded
108 relation and recurrent reconstruction networks. In *2021 IEEE/RSJ International Conference on*
109 *Intelligent Robots and Systems (IROS)*, pages 4807–4814. IEEE, 2021.
- 110 [8] Kai Chen and Qi Dou. Sgpa: Structure-guided prior adaptation for category-level 6d object pose
111 estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*,
112 pages 2773–2782, 2021.
- 113 [9] Yan Di, Ruida Zhang, Zhiqiang Lou, Fabian Manhardt, Xiangyang Ji, Nassir Navab, and
114 Federico Tombari. Gpv-pose: Category-level object pose estimation via geometry-guided point-
115 wise voting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
116 *Recognition*, pages 6781–6791, 2022.
- 117 [10] Haitao Lin, Zichang Liu, Chilam Cheang, Yanwei Fu, Guodong Guo, and Xiangyang Xue.
118 Sar-net: shape alignment and recovery network for category-level 6d object pose and size
119 estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
120 *Recognition*, pages 6707–6717, 2022.
- 121 [11] Ruida Zhang, Yan Di, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Ssp-pose:
122 Symmetry-aware shape prior deformation for direct category-level object pose estimation. In
123 *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages
124 7452–7459. IEEE, 2022.