

3D-AWARE HYPOTHESIS & VERIFICATION FOR GENERALIZABLE RELATIVE OBJECT POSE ESTIMATION

Chen Zhao

EPFL-CVLab

chen.zhao@epfl.ch

Tong Zhang*

EPFL-IVRL

tong.zhang@epfl.ch

Mathieu Salzmann

EPFL-CVLab, ClearSpace SA

mathieu.salzmann@epfl.ch

ABSTRACT

Prior methods that tackle the problem of generalizable object pose estimation highly rely on having dense views of the unseen object. By contrast, we address the scenario where only a single reference view of the object is available. Our goal then is to estimate the relative object pose between this reference view and a query image that depicts the object in a different pose. In this scenario, robust generalization is imperative due to the presence of unseen objects during testing and the large-scale object pose variation between the reference and the query. To this end, we present a new hypothesis-and-verification framework, in which we generate and evaluate multiple pose hypotheses, ultimately selecting the most reliable one as the relative object pose. To measure reliability, we introduce a 3D-aware verification that explicitly applies 3D transformations to the 3D object representations learned from the two input images. Our comprehensive experiments on the Objaverse, LINEMOD, and CO3D datasets evidence the superior accuracy of our approach in relative pose estimation and its robustness in large-scale pose variations, when dealing with unseen objects. Our project website is at: https://sailor-z.github.io/projects/ICLR2024_3DAHV.html.

1 INTRODUCTION

Object pose estimation is crucial in many computer vision and robotics tasks, such as VR/AR (Azuma, 1997), scene understanding (Geiger et al., 2012; Chen et al., 2017; Xu et al., 2018; Marchand et al., 2015), and robotic manipulation (Collet et al., 2011; Zhu et al., 2014; Tremblay et al., 2018; Pitteri et al., 2019). Much effort has been made toward estimating object pose parameters either by direct regression (Xiang et al., 2017; Wang et al., 2019a; Hu et al., 2020) or by establishing correspondences (Peng et al., 2019; Wang et al., 2021; Su et al., 2022) which act as input to a PnP algorithm (Lepetit et al., 2009). These methods have achieved promising results in the closed-set scenario, where the training and testing data contain the same object instances. However, this assumption restricts their applicability to the real world, where unseen objects from new categories often exist. Therefore, there has been growing interest in generalizable object pose estimation, aiming to develop models that generalize to unseen objects in the testing phase.

In this context, some approaches (Zhao et al., 2022b; Shugurov et al., 2022) follow a template-matching strategy, matching a query object image with reference images generated by rendering the 3D textured object mesh from various viewpoints. To address the scenario where the object mesh is unavailable, as illustrated in Fig. 1(a), some methods take real dense-view images as references. The object pose in the query image is estimated either by utilizing a template-matching mechanism (Liu et al., 2022) or by building 2D-3D correspondences (Sun et al., 2022). A computationally expensive 3D reconstruction (Schonberger & Frahm, 2016) is involved to either calibrate the reference images or reconstruct the 3D object point cloud. In any event, the requirement of dense-view references precludes the use of these methods for individual or sparse images, e.g., downloaded from the Internet. Intuitively, with sufficiently diverse training data, one could think of learning to regress the object pose parameters directly from a single query image. However, without access to a canonical object frame, the predicted object pose would be ill-defined as it represents the relative transformation between the camera frame and the object frame.

*Corresponding author.

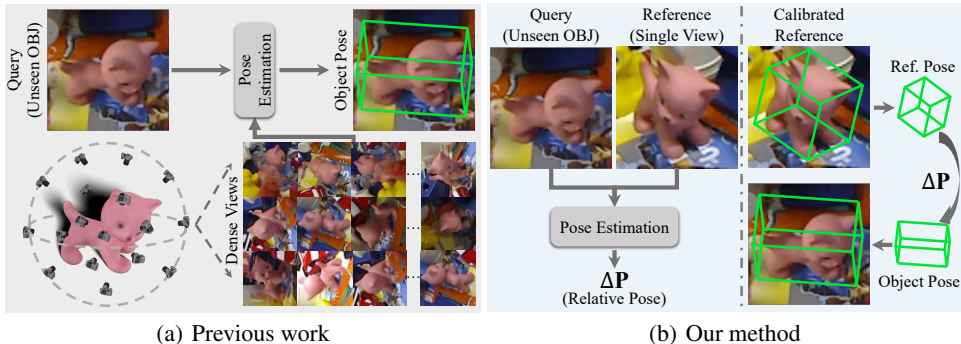


Figure 1: **Difference between previous work and our method.** Previous approaches (a) estimate the pose of an unseen object building upon either template matching or 2D-3D correspondences, which requires dense views of the object as references. By contrast, our method (b) takes only one reference as input and predicts the relative object pose between the reference and query. The object pose in the query can be derived when the pose of the reference is available.

To bypass this issue, we assume the availability of a single reference image that contains the novel object. As shown in Fig. 1(b), we take this reference to be the canonical view and estimate the relative object pose between the reference view and the query view, which is thus well-defined. If the object pose in the reference is provided, the object pose in the query can be derived. In this scenario, one plausible solution is to compute the relative object pose based on pixel-level correspondences (Lowe, 2004; Rublee et al., 2011). However, the two views may depict a large-scale object pose variation, and our experiments will evidence that even the state-of-the-art feature-matching approaches (Sarlin et al., 2020; Sun et al., 2021; Goodwin et al., 2022) cannot generate reliable correspondences in this case, which thus results in inaccurate relative object pose estimates. As an alternative, Zhang et al. (2022); Lin et al. (2023) predict the likelihood of pose parameters leveraging an energy-based model, which, however, lacks the ability to capture 3D information when learning 2D feature embeddings.

By contrast, we adopt a *hypothesis-and-verification* paradigm, drawing inspiration from its remarkable success in robust estimation (Fischler & Bolles, 1981). We randomly sample pose parameter hypotheses and verify the reliability of these hypotheses. The relative object pose is determined as the most reliable hypothesis. Since relative pose denotes a 3D transformation, achieving robust verification from two 2D images is non-trivial. Our innovation lies in a 3D-aware verification mechanism. Specifically, we develop a 3D reasoning module over 2D feature maps, which infers 3D structural features represented as 3D volumes. This lets us explicitly apply the pose hypothesis as a 3D transformation to the reference volume. Intuitively, the transformed reference volume should be aligned with the query one if the sampled hypothesis is correct. We thus propose to verify the hypothesis by comparing the feature similarities of the reference and the query. To boost robustness, we aggregate the 3D features into orthogonal 2D plane embeddings and compare these embeddings to obtain a similarity score that indicates the reliability of the hypothesis.

Our method achieves state-of-the-art performance on an existing benchmark of Lin et al. (2023). Moreover, we extend the experiments to a new benchmark for **generalizable relative object pose** estimation, which we refer to as GROP. Our benchmark contains over 10,000 testing image pairs, exploiting objects from Objaverse (Deitke et al., 2023) and LINEMOD (Hinterstoisser et al., 2012) datasets, thus encompassing both synthetic and real images with diverse object poses. In the context of previously unseen objects, our method outperforms the feature-matching and energy-based techniques by a large margin in terms of both relative object pose estimation accuracy and robustness. We summarize our contributions as follows:

- We highlight the importance of relative pose estimation for novel objects in scenarios where only one reference image is available for each object.
- We present a new hypothesis-and-verification paradigm where verification is made aware of 3D by acting on a learnable 3D object representation.

- We develop a new benchmark called GROP, where the evaluation of relative object pose estimation is conducted on both synthetic and real images with diverse object poses.

2 RELATED WORK

Instance-Specific Object Pose Estimation. The advancements in deep learning have revolutionized the field of object pose estimation. Most existing studies have focused on instance-level object pose estimation (Xiang et al., 2017; Peng et al., 2019; Wang et al., 2021; Su et al., 2022; Wang et al., 2019a), aiming to determine the pose of specific object instances. These methods have achieved remarkable performance in the closed-set setting, which means that the training data and testing data contain the same object instances. However, such an instance-level assumption restricts the applications in the real world where previously unseen objects widely exist. The studies of Zhao et al. (2022b); Liu et al. (2022) have revealed the limited generalization ability of the instance-level approaches when confronted with unseen objects. Some approaches (Wang et al., 2019b; Chen et al., 2020a; Lin et al., 2022) relaxed the instance-level constraint and introduced category-level object pose estimation. More concretely, the testing and training datasets consist of different object instances but the same object categories. As different instances belonging to the same category depict similar visual patterns, the category-level object pose estimation methods are capable of generalizing well to new instances. However, these approaches still face challenges in generalizing to objects from novel categories, since the object appearance could vary significantly.

Generalizable Object Pose Estimation. Recently, some effort has been made toward generalizable object pose estimation. The testing data may include objects from categories that have not been encountered during training. The objective is to estimate the pose of these unseen objects without retraining the network. In such a context, the existing approaches can be categorized into two groups, i.e., template-matching methods (Sundermeyer et al., 2020; Labbé et al., 2022; Zhao et al., 2022b; Liu et al., 2022; Shugurov et al., 2022) and feature-matching methods (Sun et al., 2022; He et al., 2022b). Given a query image of the object, the template-matching methods retrieve the most similar reference image from a pre-generated database. The object pose is taken as that in the retrieved reference. The database is created by either rendering the 3D object model or capturing images from various viewpoints. The feature-matching methods reconstruct the 3D object point cloud by performing SFM (Schonberger & Frahm, 2016) over a sequence of images. The 2D-3D matches are then built over the query image and the reconstructed point cloud, from which the object pose is estimated by using the PnP algorithm. Notably, these two groups both require dense-view reference images to be available. Therefore, they cannot be applied in scenarios where only sparse images are accessible.

Relative Object Pose Estimation. Some existing methods could nonetheless be applied for relative object pose estimation, even though they were designed for a different purpose. For example, one could use traditional (Lowe, 2004) or learning-based (Sarlin et al., 2020; Sun et al., 2021; Goodwin et al., 2022) methods to build pixel-pixel correspondences and compute the relative pose by using multi-view geometry (Hartley & Zisserman, 2003). However, as only two views (one query and one reference) are available, large-scale object pose variations are inevitable, posing challenges to the correspondence-based approaches. Moreover, RelPose (Zhang et al., 2022) and RelPose++ (Lin et al., 2023) build upon an energy-based model, which combines the pose parameters with the two-view images as the input and predicts the likelihood of the relative camera pose. However, RelPose and RelPose++ exhibit a limitation in their ability to reason about 3D information, which we found crucial for inferring the 3D transformation between 2D images. By contrast, we propose to explicitly utilize 3D information in a new hypothesis-and-verification paradigm, achieving considerably better performance in our experiments.

3 METHOD

3.1 PROBLEM FORMULATION

We train a network on *RGB* images depicting specific object instances from a set \mathcal{O}_{train} . During testing, we aim for the network to generalize to new objects in the set \mathcal{O}_{test} , with $\mathcal{O}_{test} \cap \mathcal{O}_{train} = \emptyset$. In contrast to some previous methods which assume that \mathcal{O}_{train} and \mathcal{O}_{test} contain the same cate-

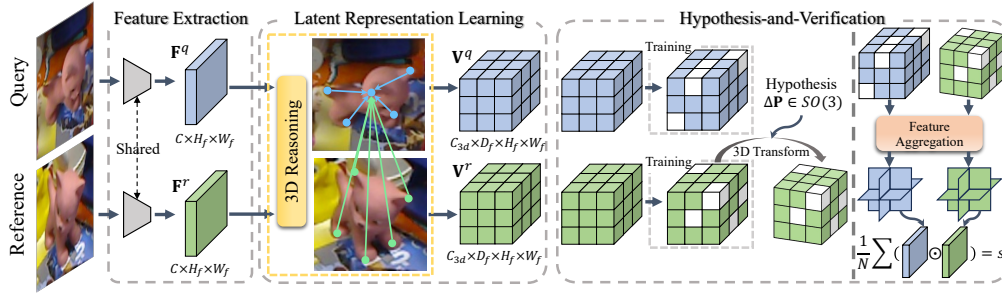


Figure 2: **Overview of our framework.** Our method estimates the relative pose of previously unseen objects given two images, building upon a new hypothesis-and-verification paradigm. A hypothesis $\Delta\mathbf{P}$ is randomly sampled and its accuracy is measured as a score s . To explicitly integrate 3D information, we perform the verification over a 3D object representation indicated as a learnable 3D volume. The sampled hypothesis is coupled with the learned representation via a 3D transformation over the reference 3D volume. We learn the 3D volumes from the 2D feature maps extracted from the RGB images by introducing a 3D reasoning module. To improve robustness, we randomly mask out some blocks colored in white during training.

gories, i.e., $\mathcal{C}_{train} = \mathcal{C}_{test}$, we work on generalizable object pose estimation. The testing objects in \mathcal{O}_{test} may belong to previously unseen categories, i.e., $\mathcal{C}_{test} \neq \mathcal{C}_{train}$. In such a context, we propose to estimate the relative pose $\Delta\mathbf{P}$ of the object depicted in two images \mathbf{I}_q and \mathbf{I}_r . As the 3D object translation can be derived by utilizing 2D detection (Saito et al., 2022; Wang et al., 2023; Kirillov et al., 2023), we focus on the estimation of 3D object rotation $\Delta\mathbf{R} \in SO(3)$, which is more challenging. As illustrated in Fig. 2, our method builds upon a hypothesis-and-verification mechanism (Fischler & Bolles, 1981). Concretely, we randomly sample an orientation hypothesis $\Delta\mathbf{R}_i$, utilizing the 6D continuous representation of Zhou et al. (2019). We then verify the correctness of $\Delta\mathbf{R}_i$ using a verification score $s_i = f(\mathbf{I}_q, \mathbf{I}_r | \Delta\mathbf{R}_i, \Theta)$, where f indicates a network with learnable parameters Θ . The expected $\Delta\mathbf{R}^*$ is determined as the hypothesis with the highest verification score, i.e.,

$$\Delta\mathbf{R}^* = \arg \max_{\Delta\mathbf{R}_i \in SO(3)} f(\mathbf{I}_q, \mathbf{I}_r | \Delta\mathbf{R}_i, \Theta). \quad (1)$$

To facilitate the verification, we develop a 3D transforming layer over a learnable 3D object representation. The details will be introduced in this section.

3.2 3D OBJECT REPRESENTATION LEARNING

Predicting 3D transformations from 2D images is inherently challenging, as it necessitates the capability of 3D reasoning. Furthermore, the requirement of generalization ability to unseen objects makes the problem even harder. Existing methods (Zhang et al., 2022; Lin et al., 2023) tackle this challenge by deriving 3D information from global feature embeddings, which are obtained through global pooling over 2D feature maps. However, this design exhibits two key drawbacks: First, the low-level structural features which are crucial for reasoning about 3D transformations are lost; Second, the global pooling process incorporates high-level semantic information (Zhao et al., 2022b), which is coupled with the object category. Therefore, these approaches encounter difficulties in accurately estimating the relative pose of previously unseen objects.

To address this, we introduce a 3D object representation learning module that is capable of reasoning about 3D information from 2D structural features. Concretely, the process begins by feeding the query and reference images into a pretrained encoder (Ranftl et al., 2020), yielding two 2D feature maps $\mathbf{F}^q, \mathbf{F}^r \in \mathbb{R}^{C \times H_f \times W_f}$. As no global pooling layer is involved, \mathbf{F}^q and \mathbf{F}^r contain more structural information than the global feature embeddings of (Zhang et al., 2022; Lin et al., 2023). Subsequently, \mathbf{F}^q and \mathbf{F}^r serve as inputs to a 3D reasoning module. Since each RGB image depicts the object from a particular viewpoint, inferring 3D features from a single 2D feature map is intractable. To address this issue, we combine the query and reference views and utilize the transformer (Vaswani et al., 2017; Dosovitskiy et al., 2020), renowned for its ability to capture relationships among local patches.

Our 3D reasoning block comprises a self-attention layer and a cross-attention layer, which account for the intra-view and inter-view relationships, respectively. Notably, unlike the existing method of Lin et al. (2023) that utilizes transformers at an image level, i.e., treating a global feature embedding as a token, our module takes \mathbf{F}^q and \mathbf{F}^r as input, thereby preserving more structural information throughout the process. Specifically, we compute

$$\mathbf{F}_{l+1}^q = g(\mathbf{F}_l^q, \mathbf{F}_l^r | \Omega_{\text{self}}^q, \Omega_{\text{cross}}^q), \quad (2)$$

$$\mathbf{F}_{l+1}^r = g(\mathbf{F}_l^r, \mathbf{F}_l^q | \Omega_{\text{self}}^r, \Omega_{\text{cross}}^r), \quad (3)$$

where g denotes the 3D reasoning block with learnable parameters $\{\Omega_{\text{self}}^q, \Omega_{\text{cross}}^q, \Omega_{\text{self}}^r, \Omega_{\text{cross}}^r\}$. Let us take \mathbf{F}^q as an example as the process over \mathbf{F}^r is symmetric. We tokenize \mathbf{F}^q by flattening it from $\mathbb{R}^{C \times H_f \times W_f}$ to $\mathbb{R}^{N \times C}$, where $N = H_f \times W_f$. A position embedding (Dosovitskiy et al., 2020) is added to the sequence of tokens, which accounts for positional information. To ensure a broader receptive field that covers the entire object, the tokens are fed into a self-attention layer, which is formulated as $\tilde{\mathbf{F}}_l^q = t(\mathbf{F}_l^q, \mathbf{F}_l^q | \Omega_{\text{self}}^q)$, where t denotes the attention layer. As aforementioned, $\tilde{\mathbf{F}}_l^q$ only describes the object in \mathbf{I}_q , which is captured from a single viewpoint. We thus develop a cross-attention layer, incorporating information from the other view \mathbf{I}_r into $\tilde{\mathbf{F}}_l^q$. We denote the cross attention as $\mathbf{F}_{l+1}^q = t(\tilde{\mathbf{F}}_l^q, \tilde{\mathbf{F}}_l^r | \Omega_{\text{cross}}^q)$, where $\tilde{\mathbf{F}}_l^r$ serves as the input of the next 3D reasoning block.

We denote the output of the last 3D reasoning block as $\hat{\mathbf{F}}^q, \hat{\mathbf{F}}^r \in \mathbb{R}^{C \times H_f \times W_f}$. $\hat{\mathbf{F}}^q$ and $\hat{\mathbf{F}}^r$ comprise both intra-view and inter-view object-related information. Nevertheless, it is still non-trivial to couple the 3D transformation with the 2D feature maps, which is crucial in the following hypothesis-and-verification module. To handle this, we derive a 3D object representation from the 2D feature map in a simple yet effective manner. We lift $\hat{\mathbf{F}}^q$ and $\hat{\mathbf{F}}^r$ from 2D space to 3D space, i.e., $\mathbb{R}^{C \times H_f \times W_f} \rightarrow \mathbb{R}^{C_{3d} \times D_f \times H_f \times W_f}$, where $C = C_{3d} \times D_f$. The 3D representations are thus encoded as 3D volumes \mathbf{V}^q and \mathbf{V}^r . Since the spatial dimensionality of \mathbf{V}^q and \mathbf{V}^r matches that of the 3D transformation, such a lifting process enables the subsequent 3D-aware verification.

3.3 3D-AWARE HYPOTHESIS AND VERIFICATION

The hypothesis-and-verification mechanism has achieved tremendous success as a robust estimator (Fischler & Bolles, 1981) for image matching (Yi et al., 2018; Zhao et al., 2021). The objective is to identify the most reliable hypothesis from multiple samplings. In such a context, an effective verification process is critical. Moreover, in the scenario of relative object pose estimation, we expect the verification to be differentiable and aware of the 3D transformation. We thus tailor the hypothesis-and-verification mechanism to meet these new requirements.

We develop a 3D masking approach in latent space before sampling hypotheses, drawing inspiration from the success of the masked visual modeling methods (He et al., 2022a; Xie et al., 2022). Instead of masking the RGB images, we propose to mask the learnable 3D volumes, which we empirically found more compact and effective. Specifically, we sample two binary masks $\mathbf{V}_b^q, \mathbf{V}_b^r \in \mathbb{R}^{C_{3d} \times D_f \times H_f \times W_f}$ during training, initialized as all ones. h of the elements in each mask are randomly set to 0. 3D masking is performed as $\tilde{\mathbf{V}}^q = \mathbf{V}^q \odot \mathbf{V}_b^q$, $\tilde{\mathbf{V}}^r = \mathbf{V}^r \odot \mathbf{V}_b^r$, where \odot denotes the Hadamard product. Note that the masking is asymmetric over \mathbf{V}_b^q and \mathbf{V}_b^r . Such a design enables the modeling of object motion between two images (Gupta et al., 2023), offering potential benefits to the task of relative object pose estimation.

The hypothesis-and-verification process begins by randomly sampling hypotheses, utilizing the 6D continuous representation of Zhou et al. (2019). Each hypothesis is then converted to a 3D rotation matrix $\Delta \mathbf{R}_i$, i.e., $\mathbb{R}^6 \rightarrow \mathbb{R}^{3 \times 3}$. During the verification, we explicitly couple the hypothesis with the learnable 3D representation by performing a 3D transformation. This is formulated as

$$\tilde{\mathbf{V}}^r = \varphi(\Delta \mathbf{R}_i \mathbf{X}^r), \mathbf{X}^r \in \mathbb{R}^{3 \times L}, \quad (4)$$

where \mathbf{X}^r denotes the 3D coordinates of the elements in $\tilde{\mathbf{V}}^r$ with $L = D_f \times H_f \times W_f$ and φ indicates trilinear interpolation. We keep the query 3D volume unchanged and only transform the reference 3D volume. Intuitively, the transformed $\tilde{\mathbf{V}}^r$ should be aligned with $\tilde{\mathbf{V}}^q$ if the sampled hypothesis is correct. Conversely, an incorrect 3D transformation is supposed to result in a noticeable disparity between the two 3D volumes. Therefore, our transformation-based approach facilitates the verification of $\Delta \mathbf{R}_i$, which could be implemented by assessing the similarity between $\tilde{\mathbf{V}}^q$ and

$\tilde{\mathbf{V}}^r$. However, the transformed $\tilde{\mathbf{V}}^r$ tends to be noisy in practice because of zero padding during the transformation and some nuisances such as the background. We thus introduce a feature aggregation module, aiming to distill meaningful information for robust verification. More concretely, we project $\tilde{\mathbf{V}}^q$ and $\tilde{\mathbf{V}}^r$ back to three orthogonal 2D planes, i.e., $\mathbb{R}^{C_{3d} \times D_f \times H_f \times W_f} \rightarrow \mathbb{R}^{3C \times H_f \times W_f}$ with $C = C_{3d} \times D_f$, and aggregate the projected features as $\mathbf{A}^q = g(\tilde{\mathbf{V}}^q | \Psi)$, $\mathbf{A}^r = g(\tilde{\mathbf{V}}^r | \Psi)$, where $\mathbf{A}^q, \mathbf{A}^r \in \mathbb{R}^{C_{2d} \times H_f \times W_f}$ represent the distilled feature embeddings and g is the aggregation module with learnable parameters Ψ . The verification score is then computed as

$$s_i = \frac{1}{N} \sum_{j,k} \frac{\mathbf{A}_{jk}^q \cdot \mathbf{A}_{jk}^r}{\|\mathbf{A}_{jk}^q\| \cdot \|\mathbf{A}_{jk}^r\|}, \quad \mathbf{A}_{jk}^q, \mathbf{A}_{jk}^r \in \mathbb{R}^{C_{2d}}. \quad (5)$$

We run the hypothesis and verification M times in parallel and the expected $\Delta \mathbf{R}^*$ is identified as

$$\Delta \mathbf{R}^* = \Delta \mathbf{R}_k, \quad k = \arg \max_i \{s_i, i = 1, 2, \dots, M\}. \quad (6)$$

Note that compared with the dynamic rendering method (Park et al., 2020) which optimizes the object pose by rendering and comparing depth images, our approach performs verification in the latent space. This eliminates the need for computationally intensive rendering and operates independently of depth information. An alternative to the hypothesis-and-verification mechanism consists of optimizing $\Delta \mathbf{R}$ via gradient descent. However, our empirical observations indicate that this alternative often gets trapped in local optima. Moreover, compared with the energy-based approaches (Zhang et al., 2022; Lin et al., 2023), our method achieves a 3D-aware verification. To highlight this, let us formulate the energy-based model with some abuse of notation as

$$\Delta \mathbf{R}^* = \arg \max_{\Delta \mathbf{R}_i \in SO(3)} s_i, \quad s_i = \text{FC}(f(\mathbf{I}_q, \mathbf{I}_r) + h(\Delta \mathbf{R}_i)), \quad (7)$$

where FC denotes fully connected layers. In this context, the 2D image embedding and the pose embedding are learned as $f(\mathbf{I}_q, \mathbf{I}_r)$ and $h(\Delta \mathbf{R}_i)$, separately. By contrast, in our framework, the volume features are conditioned on $\Delta \mathbf{R}_i$ via the 3D transformation, which thus facilitates the 3D-aware verification.

We train our network using an infoNCE loss (Chen et al., 2020b), which is defined as

$$\mathcal{L} = -\log \frac{\sum_{j=1}^P \exp(s_j^p / \tau)}{\sum_{i=1}^M \exp(s_i / \tau)}, \quad (8)$$

where s_j^p denotes the score of a positive hypothesis, and $\tau = 0.1$ is a predefined temperature. The positive samples are identified by computing the geodesic distance as

$$D = \arccos \left(\frac{\text{tr}(\Delta \mathbf{R}_i^T \Delta \mathbf{R}_{\text{gt}}) - 1}{2} \right) / \pi, \quad (9)$$

where $\Delta \mathbf{R}_{\text{gt}}$ is the ground truth. We then consider hypotheses with $D < \lambda$ as positive samples.

4 EXPERIMENTS

4.1 IMPLEMENTATION DETAILS

In our experiments, we employ 4 3D reasoning blocks. We set the number of hypotheses during training and testing to $M = 9,000$ and $M = 50,000$, respectively. We define the masking threshold $h = 0.25$ and the geodesic distance threshold $\lambda = 15^\circ$ (Zhang et al., 2022; Lin et al., 2023). We train our network for 25 epochs using the AdamW (Loshchilov & Hutter, 2017) optimizer with a batch size of 48 and a learning rate of 10^{-4} , which is divided by 10 after 20 epochs. Training takes around 4 days on 4 NVIDIA Tesla V100s.

4.2 EXPERIMENTAL SETUP

We compare our method with several relevant competitors including feature-matching methods, i.e., SuperGlue (Sarlin et al., 2020), LoFTR (Sun et al., 2021), and ZSP (Goodwin et al., 2022),

	SuperGlue	LoFTR	ZSP	Regress	RelPose	RelPose++	Ours
Angular Error ↓	67.2	77.5	87.5	46.0	50.0	38.5	28.5
Acc @ 30° (%) ↑	45.2	37.9	25.7	60.6	64.2	77.0	83.5
Acc @ 15° (%) ↑	37.7	33.1	14.6	42.7	48.6	69.8	71.0

Table 1: **Experimental results on CO3D.**

	SuperGlue	LoFTR	ZSP	Regress	RelPose	RelPose++	Ours
Angular Error ↓	102.4	134.1	107.2	55.9	80.4	33.5	28.1
Acc @ 30° (%) ↑	15.1	9.6	4.2	39.2	20.8	72.3	78.6
Acc @ 15° (%) ↑	12.1	7.7	1.5	15.6	6.7	42.9	58.4

Table 2: **Experimental results on Objaverse.**

	SuperGlue	LoFTR	ZSP	Regress	RelPose	RelPose++	Ours
Angular Error ↓	64.8	84.5	78.6	52.1	58.3	46.6	41.7
Acc @ 30° (%) ↑	26.2	24.2	10.7	26.5	26.1	42.5	61.5
Acc @ 15° (%) ↑	14.3	13.5	2.7	7.6	7.0	15.8	29.9

Table 3: **Experimental results on LINEMOD.**

energy-based methods, i.e., RelPose (Zhang et al., 2022) and RelPose++ (Lin et al., 2023), and a regression method (Lin et al., 2023). We first perform an evaluation using the benchmark defined in (Lin et al., 2023), where the experiments are conducted on the CO3D (Reizenstein et al., 2021) dataset. We report the angular error between the predicted $\Delta\mathbf{R}$ and the ground truth, which is computed as in Eq. 9, and the accuracy with thresholds of 30° and 15° (Zhang et al., 2022; Lin et al., 2023). Furthermore, We extend the evaluation by introducing a new benchmark called GROP. To this end, we utilize the Objaverse (Deitke et al., 2023) and LINEMOD (Hinterstoisser et al., 2012) datasets, which include synthetic and real data, respectively. We retrain RelPose, RelPose++, and the regression method in our benchmark, and use the pretrained models for SuperGlue and LoFTR since retraining these two feature-matching approaches requires additional pixel-level annotations. For ZSP, as there is no training process involved, we evaluate it using the code released by the authors. We derive $\Delta\mathbf{R}$ from the estimated essential matrix (Hartley & Zisserman, 2003) for the feature-matching methods because we only have access to RGB images. We evaluate all methods on identical predefined query and reference pairs (8,304 on Objaverse and 5,000 on LINEMOD), which ensures a fair comparison. Given our emphasis on relative object rotation estimation, we crop the objects from the original RGB image utilizing the ground-truth object bounding boxes (Xiao et al., 2019; Zhao et al., 2022b; Park et al., 2020; Nguyen et al., 2022). In Sec. 4.5, we evaluate robustness against noise in the bounding boxes.

4.3 EXPERIMENTS ON CO3D

Let us first evaluate our approach in the benchmark used in (Zhang et al., 2022; Lin et al., 2023), which builds upon the CO3D dataset (Reizenstein et al., 2021). All testing objects here are previously unseen and the evaluation thus emphasizes the generalization ability. Table 1 reports the results in terms of angular error and accuracy. Note that the results of SuperGlue, Regress, RelPose, and RelPose++ shown here align closely with the ones reported in (Lin et al., 2023), lending credibility to the evaluation. More importantly, our method produces consistently more precise relative object poses, with improvements of at least **10%** in angular error. This evidences the generalization ability of our approach to unseen objects.

4.4 EXPERIMENTS ON GROP

Let us now develop the evaluation in our benchmark. Table 2 and Table 3 provide the experimental results on Objaverse and LINEMOD, respectively. Our method also achieves superior generalization ability to unseen objects, outperforming the previous methods by a substantial margin. For instance,




Figure 3: **Qualitative results on Objaverse and LINEMOD.** Here, we assume the reference to be calibrated and visualize the object pose in the query, which is derived from the estimated relative object pose. The predicted and ground-truth object poses are indicated by blue and green arrows, respectively.

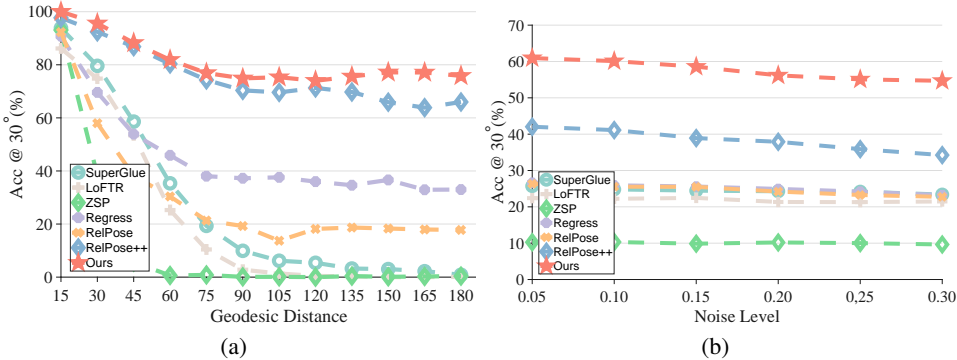


Figure 4: **Robustness.** (a) Acc @ 30° curves obtained with varying degrees of object pose variation between the reference and the query, measured by the geodesic distance. (b) Similar curves but for different levels of noise added to the object bounding boxes.

we achieve an improvement of at least **15.5%** on Objaverse and **14.1%** on LINEMOD, measured in terms of Acc @ 15°. Moreover, we illustrate some qualitative results in Fig. 3. To this end, we assume the object pose \mathbf{R}^r in the reference to be available, and the object pose \mathbf{R}^q in the query is computed as $\mathbf{R}^q = \Delta \mathbf{R} \mathbf{R}^r$. We represent the predicted and the ground-truth object poses as blue and green arrows, respectively. This evidences that our method consistently yields better predictions. In the scenario where there is a notable difference in object pose between the reference and query (as in the cat images in the third row), the previous methods struggle to accurately predict the pose for the unseen object, while our approach continues to deliver an accurate prediction.

4.5 ABLATION STUDIES

To shed more light on the superiority of our method, we develop comprehensive ablation studies on Objaverse and LINEMOD. Most of the experiments are conducted on LINEMOD since it is a real dataset. As the two sparse views, i.e., a reference and a query, might result in a large-scale object pose variation, we start the ablations by analyzing the robustness in such a context. Specifically, we divide the Objaverse testing data into several groups based on the object pose variation between the reference and query, measured by geodesic distance. The task becomes progressively more

	w/o att.	w/o mask	w/ 2D mask	w/o agg.	RelPose*	Ours
Angular Error ↓	41.9	42.1	42.6	41.9	59.7	41.7
Acc @ 30° (%) ↑	60.0	59.6	60.1	59.4	26.4	61.5
Acc @ 15° (%) ↑	28.2	27.9	27.3	26.4	7.9	29.9

Table 4: Effectiveness of the key components in our pipeline.

challenging as the distance increases. We developed this experiment on Objaverse because of its wider range of pose variations compared to LINEMOD. Fig. 4(b) shows the Acc @ 30° curves as the distance varies from 0° to 180°. Note that all methods demonstrate satisfactory predictions when the distance is small, i.e., when the object orientations in the reference and query views are similar. However, the performance of feature-matching approaches, i.e., SuperGlue, LoFTR, and ZSP, dramatically drops as the distance increases. This observation supports our argument that the feature-matching methods are sensitive to the pose variations. By contrast, our method consistently surpasses all competitors, thus showing better robustness.

As the object bounding boxes obtained in practice are inevitably noisy, we evaluate the robustness against the noise in this context on LINEMOD. Concretely, we add noise to the ground-truth bounding boxes by applying jittering to both the object center and scale. The jittering magnitude varies from 0.05 to 0.30, which results in different levels of noise. The experimental results are shown in Fig. 4(b), where our method outperforms the competitors across all scenarios. This promising robustness underscores the possibility of integrating our method with existing unseen object detectors (Zhao et al., 2022a; Liu et al., 2022). To showcase this, we extend our method to 6D unseen object pose estimation by combining it with the detector introduced in (Liu et al., 2022) and provide some results in the appendix.

Furthermore, we evaluate the effectiveness of the key components in our framework. The results on LINEMOD are summarized in Table 4, where the evaluation of effectiveness encompasses four distinct aspects: First, we develop a counterpart by excluding self-attention and cross-attention layers (w/o att.) from the 3D reasoning blocks; Second, we modify the 3D masking by either omitting it (w/o mask) or substituting it with a 2D masking process over RGB images (w/ 2D mask); Third, we directly compute the similarity of 3D volumes without utilizing the 2D aggregation module (w/o agg.); Fourth, we replace our 3D-aware verification mechanism with the energy-based model (Zhang et al., 2022; Lin et al., 2023) (RelPose*), while retaining our feature extraction backbone unchanged. The modified versions, namely w/o att., w/o mask, w/ 2D mask, and w/o agg., exhibit worse performance, which thus demonstrates the effectiveness of the presented components, i.e., attention layers, 3D masking, and the feature aggregation module. Additionally, the inferior results yield by RelPose* highlight that the 3D-aware verification mechanism contributes to the high-accuracy predictions, instead of the feature extraction backbone in our framework. Consequently, this observation supports our claim that the proposed verification module facilitates the relative pose estimation for unseen objects by preserving the structural features and explicitly utilizing 3D information.

5 CONCLUSION

In this paper, we have tackled the problem of relative pose estimation for unseen objects. We assume the availability of only one object image as the reference and aim to estimate the relative object pose between the reference and a query image. In this context, we have tailored the hypothesis-and-verification paradigm by introducing a 3D-aware verification, where the 3D transformation is explicitly coupled with a learnable 3D object representation. We have developed comprehensive experiments on Objaverse, LINEMOD, and CO3D datasets, taking both synthetic and real data with diverse object poses into account. Our method remarkably outperforms the competitors across all scenarios and achieves better robustness against different levels of object pose variations and noise. Since our verification module incorporates local similarities when computing the verification scores, it could be affected by the occlusions. This stands as a potential limitation that we consider, and we intend to explore and address this issue in our future research endeavors.

ACKNOWLEDGMENT

This work was funded in part by the Swiss National Science Foundation via the Sinergia grant CRSII5-180359 and the Swiss Innovation Agency (Innosuisse) via the BRIDGE Discovery grant 40B2-0.194729.

REFERENCES

- Ronald T Azuma. A survey of augmented reality. *Presence: teleoperators & virtual environments*, 6(4):355–385, 1997.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Dengsheng Chen, Jun Li, Zheng Wang, and Kai Xu. Learning canonical shape space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11973–11982, 2020a.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pp. 1597–1607. PMLR, 2020b.
- Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 1907–1915, 2017.
- Alvaro Collet, Manuel Martinez, and Siddhartha S Srinivasa. The moped framework: Object recognition and pose estimation for manipulation. *The International Journal of Robotics Research*, 30(10):1284–1306, 2011.
- Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13142–13153, 2023.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Martin A Fischler and Robert C Bolles. 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.
- Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 3354–3361. IEEE, 2012.
- Walter Goodwin, Sagar Vaze, Ioannis Havoutis, and Ingmar Posner. Zero-shot category-level object pose estimation. In *Proceedings of the European Conference on Computer Vision*, pp. 516–532. Springer, 2022.
- Agrim Gupta, Jiajun Wu, Jia Deng, and Li Fei-Fei. Siamese masked autoencoders. *arXiv preprint arXiv:2305.14344*, 2023.
- Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16000–16009, 2022a.

- Xingyi He, Jiaming Sun, Yuang Wang, Di Huang, Hujun Bao, and Xiaowei Zhou. Onepose++: Keypoint-free one-shot object pose estimation without cad models. *Advances in Neural Information Processing Systems*, 35:35103–35115, 2022b.
- Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. 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.
- Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2930–2939, 2020.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
- Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. MegaPose: 6D Pose Estimation of Novel Objects via Render & Compare. In *CoRL*, 2022.
- Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate o(n) solution to the pnp problem. *International Journal of Computer Vision*, 81:155–166, 2009.
- Amy Lin, Jason Y Zhang, Deva Ramanan, and Shubham Tulsiani. Relpose++: Recovering 6d poses from sparse-view observations. *arXiv preprint arXiv:2305.04926*, 2023.
- Jiehong Lin, Zewei Wei, Changxing Ding, and Kui Jia. Category-level 6d object pose and size estimation using self-supervised deep prior deformation networks. In *Proceedings of the European Conference on Computer Vision*, pp. 19–34. Springer, 2022.
- 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 *Proceedings of the European Conference on Computer Vision*, pp. 740–755. Springer, 2014.
- Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. *arXiv preprint arXiv:2303.11328*, 2023.
- Yuan Liu, Yilin Wen, Sida Peng, Cheng Lin, Xiaoxiao Long, Taku Komura, and Wenping Wang. Gen6d: Generalizable model-free 6-dof object pose estimation from rgb images. *Proceedings of the European Conference on Computer Vision*, 2022.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose estimation for augmented reality: a hands-on survey. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2633–2651, 2015.
- Van Nguyen Nguyen, Yinlin Hu, Yang Xiao, Mathieu Salzmann, and Vincent Lepetit. Templates for 3d object pose estimation revisited: Generalization to new objects and robustness to occlusions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6771–6780, 2022.
- Keunhong Park, Arsalan Mousavian, Yu Xiang, and Dieter Fox. Latentfusion: End-to-end differentiable reconstruction and rendering for unseen object pose estimation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 10710–10719, 2020.
- Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4561–4570, 2019.

- Giorgia Pitteri, Slobodan Ilic, and Vincent Lepetit. Cornet: generic 3d corners for 6d pose estimation of new objects without retraining. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pp. 0–0, 2019.
- René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 44(3):1623–1637, 2020.
- Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10901–10911, 2021.
- Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2564–2571. Ieee, 2011.
- Kuniaki Saito, Ping Hu, Trevor Darrell, and Kate Saenko. Learning to detect every thing in an open world. In *Proceedings of the European Conference on Computer Vision*, pp. 268–284. Springer, 2022.
- Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4938–4947, 2020.
- Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4104–4113, 2016.
- Ivan Shugurov, Fu Li, Benjamin Busam, and Slobodan Ilic. Osop: A multi-stage one shot object pose estimation framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6835–6844, 2022.
- Yongzhi Su, Mahdi Saleh, Torben Fetzer, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Federico Tombari. Zebrapose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6738–6748, 2022.
- Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8922–8931, 2021.
- Jiaming Sun, Zihao Wang, Siyu Zhang, Xingyi He, Hongcheng Zhao, Guofeng Zhang, and Xiaowei Zhou. Onepose: One-shot object pose estimation without cad models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6825–6834, 2022.
- Martin Sundermeyer, Maximilian Durner, En Yen Puang, Zoltan-Csaba Marton, Narunas Vaskevicius, Kai O Arras, and Rudolph Triebel. Multi-path learning for object pose estimation across domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 13916–13925, 2020.
- Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *Conference on Robot Learning*, 2018. URL <https://arxiv.org/abs/1809.10790>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3343–3352, 2019a.

- Gu Wang, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 16611–16621, 2021.
- He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2642–2651, 2019b.
- Zhenyu Wang, Yali Li, Xi Chen, Ser-Nam Lim, Antonio Torralba, Hengshuang Zhao, and Shengjin Wang. Detecting everything in the open world: Towards universal object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11433–11443, 2023.
- Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017.
- Yang Xiao, Xuchong Qiu, Pierre-Alain Langlois, Mathieu Aubry, and Renaud Marlet. Pose from shape: Deep pose estimation for arbitrary 3D objects. In *British Machine Vision Conference (BMVC)*, 2019.
- Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9653–9663, 2022.
- Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 244–253, 2018.
- Kwang Moo Yi, Eduard Trulls, Yuki Ono, Vincent Lepetit, Mathieu Salzmann, and Pascal Fua. Learning to find good correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2666–2674, 2018.
- Jason Y Zhang, Deva Ramanan, and Shubham Tulsiani. Relpose: Predicting probabilistic relative rotation for single objects in the wild. In *Proceedings of the European Conference on Computer Vision*, pp. 592–611. Springer, 2022.
- Chen Zhao, Yixiao Ge, Feng Zhu, Rui Zhao, Hongsheng Li, and Mathieu Salzmann. Progressive correspondence pruning by consensus learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6464–6473, 2021.
- Chen Zhao, Yinlin Hu, and Mathieu Salzmann. Locposenet: Robust location prior for unseen object pose estimation. *arXiv preprint arXiv:2211.16290v2*, 2022a.
- Chen Zhao, Yinlin Hu, and Mathieu Salzmann. Fusing local similarities for retrieval-based 3d orientation estimation of unseen objects. In *Proceedings of the European Conference on Computer Vision*, pp. 106–122. Springer, 2022b.
- Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5745–5753, 2019.
- Menglong Zhu, Konstantinos G Derpanis, Yinfei Yang, Samarth Brahmabhatt, Mabel Zhang, Cody Phillips, Matthieu Lecce, and Kostas Daniilidis. Single image 3d object detection and pose estimation for grasping. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3936–3943. IEEE, 2014.

APPENDIX

A ARCHITECTURE OF THE 3D REASONING MODULE

We show the architecture of the 3D reasoning module in Fig. 5. Each 3D reasoning block consists of a self-attention layer and a cross-attention layer, which excel at capturing intra-view and inter-view relationships, respectively. The input 2D feature map is flattened from $\mathbb{R}^{C \times H_f \times W_f}$ to $\mathbb{R}^{N \times C}$, where $N = H_f \times W_f$. A position embedding, denoted as PE, is added to the flattened feature map. Fig. 5(b) illustrates the attention layer. The context refers to the input feature map itself in the self-attention layer and it represents the feature map of another view in the cross-attention layer. We use the standard multi-head attention (Vaswani et al., 2017) and layer normalization (Ba et al., 2016) in our attention layers.

B DATA CONFIGURATION

The synthetic images are generated by rendering objects of Objaverse from randomly sampled viewpoints (Liu et al., 2023). We attach these images to random backgrounds which are sampled from COCO (Lin et al., 2014). We randomly sample 128 objects from Objaverse and use 5 objects from LINEMOD sampled by Liu et al. (2022) as testing data, reserving the remaining objects for training. This design guarantees that all objects are previously unseen during the testing phase. We train the network on both synthetic and real data, alleviating the problem of domain gap.

Recall that we assume we have access to only one reference image and the objective is to estimate the relative object pose between the reference and the query. Therefore, the selection of the reference image is a crucial aspect of our benchmark. As multi-view images are available in Objaverse and LINEMOD datasets, one could randomly sample a reference given a query. However, such a strategy may yield an inappropriate reference. As shown in Fig. 6, the object depicted in the reference image barely overlaps with the one in the query, which makes the relative object pose estimation too challenging. Therefore, we filter out the inappropriate references from the datasets during training and testing, which makes our evaluation more reasonable.

Specifically, we convert the object rotation matrices \mathbf{R}^r and \mathbf{R}^q to Euler angles $(\alpha_r, \beta_r, \gamma_r)$ and $(\alpha_q, \beta_q, \gamma_q)$, which indicate azimuth, elevation, and in-plane rotation, respectively. Note that only azimuth and elevation lead to viewpoint changes, which thus determine the co-visible regions between the reference and query. Consequently, we set the in-plane rotation to 0 and convert the Euler angle back to the rotation matrix, i.e., $\hat{\mathbf{R}} = h(\alpha, \beta, 0)$. We then measure the difference of the new rotation matrices $\hat{\mathbf{R}}^r$ and $\hat{\mathbf{R}}^q$ by computing the geodesic distance. We exclude the image pair with a distance larger than a predefined threshold (90° by default in our experiments). As illustrated in Fig. 4 in our main paper, the retained image pairs display acceptable variations in object pose. Moreover, we utilize the synthetic images on Objaverse generated by Liu et al. (2023). Each 3D object model is rendered from 10 randomly sampled viewpoints, which yields synthetic images without in-plane rotations. To introduce in-plane rotations, we rotate the reference and query images using randomly sampled 2D in-plane rotations during training and testing.

Fig. 7 shows the histograms of object pose variations between the reference and query images. We measure the variations based on the geodesic distance between the two object rotation matrices \mathbf{R}^r and \mathbf{R}^q . The histograms show that the image pairs we used in our experiments exhibit a diverse range of object pose variations, which makes our evaluation results convincing.

C QUALITATIVE RESULTS OF 6D OBJECT POSE ESTIMATION

We extend our method to 6D pose estimation for unseen objects by utilizing an off-the-shelf generalizable object detector (Liu et al., 2022). More concretely, instead of using dense-view reference images, we feed the one reference we have in our benchmark to the pretrained detection network, which predicts the object bounding box in the query image. We use the parameters of the object bounding box to compute 3D object translation, following the implementation in (Liu et al., 2022). Subsequently, we crop the object from the query and employ our approach to predict the relative 3D object rotation. The object rotation in the query is derived as $\mathbf{R}^q = \Delta \mathbf{R} \mathbf{R}^r$. Fig. 8 shows some qual-




Figure 5: Architecture of the 3D reasoning module.



Figure 6: Examples of inappropriate references.




Figure 7: Histograms of the object pose variation between the reference and query. We measure the object pose variation as the geodesic distance between the two object rotation matrices \mathbf{R}^r and \mathbf{R}^q . The histogram depicts the number of image pairs falling within different distance intervals.

itative results of 6D pose estimation for the unseen objects on LINEMOD. We draw the 3D object bounding boxes in blue and green, using the predicted 6D object pose and the ground truth, respectively. The promising results demonstrate the potential of our approach in terms of generalizable 6D object pose estimation.

D MORE DETAIL ABOUT THE ABLATION STUDIES

As we introduced in the main paper, we performed an ablation study, evaluating the robustness against the noise added to the 2D object bounding boxes. We simulate the bounding boxes in real-world applications by performing jittering to the ground truth with different levels of noise. We denote the object center and the size of the bounding box as c and s . We then randomly sample the perturbed parameters from the intervals $(c - 0.5 * n * s, c + 0.5 * n * s)$ and $(\frac{s}{1+n}, s * (1 + n))$, respectively, where n indicates the noise. We varied n from 0.05 to 0.3 in our experiments. Please refer to Fig. 5 (b) in our main paper for the experimental results.




Figure 8: **Qualitative results of 6D pose estimation for unseen objects on LINEMOD.** The blue and green 3D object bounding boxes are drawn using the predicted 6D object pose and the ground truth, respectively.

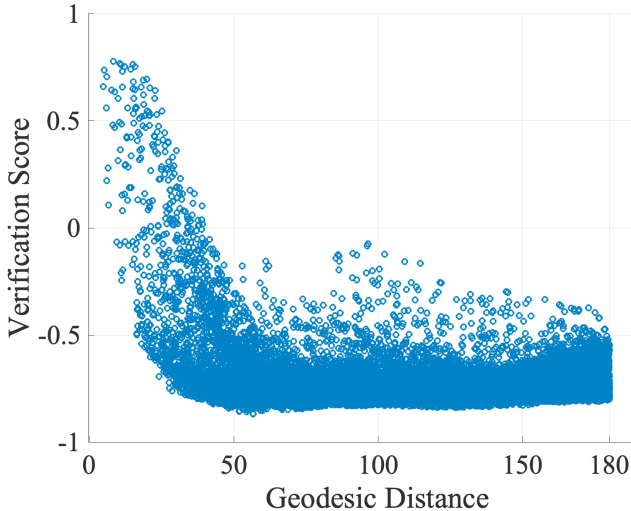


Figure 9: **Verification scores of all sampled pose hypotheses.** The x-axis and y-axis represent the geodesic distance between the pose samplings and the ground-truth relative object pose, and the verification scores, respectively.

Method	RelPose++	Ours	RelPose++-5000	Ours-5000
MACs	94.6	54.7	11.3	16.3
Angular Error	38.5	28.5	50.7	35.3

Table 5: **Efficiency.** Relpose++ uses 500,000 pose samples by default, while we sample 50,000 poses for our method in our experiments. RelPose++-5000 and Ours-5000 denote RelPose++ and our method with 5,000 pose samples, respectively. The multiply-accumulate operations (MACs) are used to measure the computation consumption.

E EFFICIENCY

It is worth noting that during testing, our method utilizes 50,000 pose samples, while RelPose++ uses 500,000. Despite processing fewer samples, our method achieves better accuracy in relative object pose estimation. To further evaluate the efficiency, we measure the computation cost in multiply-accumulate operations (MACs) and show the results in Table 5. All evaluated methods process the pose samples in parallel. “RelPose++-5000 and “Ours-5000” refer to RelPose++ and our method with 5,000 samples, respectively. The results clearly show that our method achieves a better trade-off between efficiency and accuracy in relative object pose estimation. Additionally, our method with only 5,000 samples still delivers more accurate results than RelPose++ with 500,000 samples.