

CoordAR: One-Reference 6D Pose Estimation of Novel Objects via Autoregressive Coordinate Map Generation

Dexin Zuo¹, Ang Li¹, Wei Wang^{2*}, Wenxian Yu¹, Danping Zou^{1*}

¹Shanghai Jiao Tong University

²Corporate Research Center, State Key Laboratory of High-end Heavy-load Robots, Midea Group.
{dixin95, liang_sjtu, wxyu, dpzou}@sjtu.edu.cn, wangwei232@midea.com

Abstract

Object 6D pose estimation, a crucial task for robotics and augmented reality applications, becomes particularly challenging when dealing with novel objects whose 3D models are not readily available. To reduce dependency on 3D models, recent studies have explored one-reference-based pose estimation, which requires only a single reference view instead of a complete 3D model. However, existing methods that rely on real-valued coordinate regression suffer from limited global consistency due to the local nature of convolutional architectures and face challenges in symmetric or occluded scenarios owing to a lack of uncertainty modeling. We present CoordAR, a novel autoregressive framework for one-reference 6D pose estimation of unseen objects. CoordAR formulates 3D-3D correspondences between the reference and query views as a map of discrete tokens, which is obtained in an autoregressive and probabilistic manner. To enable accurate correspondence regression, CoordAR introduces 1) a novel coordinate map tokenization that enables probabilistic prediction over discretized 3D space; 2) a modality-decoupled encoding strategy that separately encodes RGB appearance and coordinate cues; and 3) an autoregressive transformer decoder conditioned on both position-aligned query features and the partially generated token sequence. With these novel mechanisms, CoordAR significantly outperforms existing methods on multiple benchmarks and demonstrates strong robustness to symmetry, occlusion, and other challenges in real-world tests.

Project Page — <https://sjtu-visys-team.github.io/CoordAR>

Introduction

Object 6-DoF (Degrees of Freedom) pose estimation, which recovers the rotation and translation of a rigid object from observations, is a fundamental task in computer vision and robotics, with extensive applications in augmented reality, robotic manipulation, and industrial automation. Despite its importance, real-world deployment remains challenging due to factors such as texture-less object surfaces, occlusion, and lighting variations.

Learning-based approaches have made significant progress but often rely on known 3D models during training

*indicates corresponding authors.

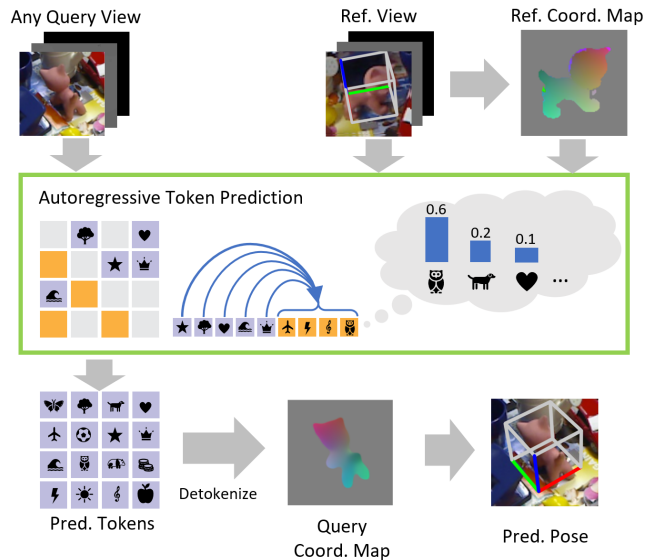


Figure 1: Given a reference view with known pose and depth-derived coordinate map, CoordAR predicts the corresponding coordinates in the query view and subsequently obtain the relative pose from the correspondences provided by the coordinate maps. Instead of inferring continuous coordinate values in parallel, our model autoregressively operate on patch-level token space with a pretrained tokenizer.

or inference and struggle to generalize to novel objects. For instance, instance-level methods (Su et al. 2022; Liu et al. 2025b) train a dedicated network per object using only synthetic data, achieving strong performance in the real world; however, they are costly and inflexible when it comes to novel objects. Category-level methods (Wang et al. 2019a; Cai et al. 2024) improve generalization across intra-class variation but still struggle with out-of-distribution objects.

An alternative is to learn correspondences between image observations and a given 3D model (Nguyen et al. 2024b; Caraffa et al. 2024), enabling zero-shot pose estimation for novel objects. However, these methods typically assume access to textured CAD models at inference, which is usually an unrealistic assumption in many real-world scenarios involving unknown objects.

To overcome this limitation, recent studies have turned to *one-reference* methods, a promising paradigm that estimates the pose of a novel object using only a single reference view. Early methods (Corsetti et al. 2024; Fan et al. 2024; Zhang, Ramanan, and Tulsiani 2022) relied on sparse correspondences between reference and target views; however, they struggled with texture-less surfaces, occlusions, and large viewpoint changes. More recently, One2Any (Liu et al. 2025a) improves robustness by regressing coordinate maps as dense correspondences. However, it uses a convolutional decoder for real-valued coordinate regression, which introduces two key limitations.

Firstly, the limited receptive fields of the convolutional decoder restrict their ability to capture long-range dependencies, leading to inconsistent global reasoning in complex scenes. Secondly, training a regression model directly using continuous, real-valued coordinates fails to handle the inherent ambiguity arising from object symmetries and occlusions. Specifically, for symmetrical objects (e.g., cylinders, cubes), direct coordinate regression forces the network to reconcile multiple valid ground truths, leading to a wrong averaged result (Hodan, Barath, and Matas 2020); for occluded objects, the model lacks an explicit mechanism to represent uncertainty in unobserved regions. These issues collectively reduce robustness in real-world applications.

In this paper, we propose CoordAR, a novel framework for one-reference 6D pose estimation. Our model generates 3D-3D dense correspondences, represented by tokenized coordinate maps, conditioned on both the reference and query views. Based on the generated correspondences, the object pose can be calculated efficiently using the Umeyama algorithm (Umeyama 1991). Our model consists of three main stages: an encoding stage, where the reference RGB image and its coordinate map are encoded separately; a subsequent feature fusion stage; and a decoding stage, where the tokens are autoregressively decoded, conditioned on embeddings from both the reference and query views.

Our approach introduces three key innovations: (1) replacing traditional continuous coordinate regression with a probability prediction on a discretized space of the coordinate map, (2) introducing a modality-decoupled encoding strategy, where the RGB images and the reference coordinate map are encoded separately to obtain better performance and flexibility, and (3) designing a network that autoregressively generates the coordinate map, with the training objective formulated as predicting the conditional probability distribution of the next-set-of tokens given the query view, reference view, and the previously generated token sequence. Our method can accurately estimate 6D poses for novel objects in complex real-world scenarios using only a single reference view. Our contributions are :

- To the best of our knowledge, we are the first to introduce autoregressive coordinate map generation for 6D pose estimation of novel objects. We further demonstrate the superiority of autoregressive generation compared to parallel real-valued regression.
- We propose modality-decoupled encoders and transformer-style fusion blocks, integrating them

into the framework, which effectively fuses information from both the query and reference views.

- We achieve state-of-the-art performance across multiple benchmark datasets, significantly outperforming existing one-reference methods.

Related Works

Model-based Methods Model-based object pose estimation methods leverage 3D models of target objects as prior knowledge. Existing model-based approaches can be broadly categorized into instance-level methods, category-level methods, and category-agnostic methods. Instance-level methods (Xiang et al. 2018; Wang et al. 2021; Li, Wang, and Ji 2019) operate on a closed set of known 3D models during training, which inherently restricts their application to previously seen objects. While category-level methods (Wang et al. 2019b; Cai et al. 2024; Chen et al. 2020; Chen and Dou 2021) demonstrate improved generalization to novel objects within trained categories, they remain constrained by their predefined taxonomic boundaries. Recently, some category-agnostic methods (Labbé et al. 2022; Caraffa et al. 2024; Nguyen et al. 2024b) estimate the relative pose based on the rendered anchor views of the 3D model. Despite their strong performance, reliance on CAD models limits their application to unseen real-world scenarios, where a novel object usually lacks a corresponding CAD model.

Model-free Methods To address scenarios where 3D models are unavailable, some methods (Sun et al. 2022; He et al. 2022a; Wen et al. 2024; Liu et al. 2022) first reconstruct the 3D model from multiple views with known poses and then estimate poses by comparing them with the images rendered from the reconstructed model. However, these methods rely on a sufficient number of views to build a high-quality 3D model, and performance drops significantly when views are sparse. Other methods directly compare the query view with sparse reference views. For example, FS6D (He et al. 2022b) establishes 3D-3D correspondence through prototype matching between the query view and the reference views. Recently, the one-reference pose estimation problem has gained attention, where only a single reference view of the object is available. This setting poses significant challenges, primarily due to the wide diversity of object appearances, severe occlusions in both the reference and query views, and substantial viewpoint variations. One2Any (Liu et al. 2025a) regresses a coordinate map that encodes the relative pose between the query view and a single reference view. However, it trains a convolutional decoder with a regression objective for continuous-valued target coordinate maps, which consequently struggles in challenging scenarios involving symmetries or heavy occlusion, often producing inaccurate coordinate predictions.

In this work, we explore autoregressive models (Jiang and Huang 2024; Xiong et al. 2024) to improve coordinate map regression, leveraging tokenization strategies that have demonstrated strong performance in tasks such as image generation (Esser, Rombach, and Ommer 2021), video generation (Yu et al. 2023) and embodied AI (Micheli,

Alonso, and Fleuret 2022). Specifically, we introduce 1) a novel coordinate map tokenization scheme enabling probabilistic prediction over discretized 3D space, 2) a modality-decoupled encoding strategy that separately models RGB appearance and coordinate cues, and 3) an autoregressive transformer decoder conditioned on pixel-aligned query features and the partially generated coordinate sequence. Together, these novel mechanisms lead to significant improvements over state-of-the-art methods.

Problem Statement

Estimating the 6D pose of an object is a challenging yet practically valuable task, especially in scenarios where full 3D models are unavailable or object instances appear in dynamic or unstructured environments. Unlike traditional pose estimation settings that rely on known 3D CAD models or extensive multi-view observations, we focus on a one-reference setting, where the model must infer the pose of a novel object using only one annotated view as prior knowledge. To achieve this, we aim to estimate the relative transformation $\mathbf{T}_{RQ} \in SE(3)$ that transforms points from the query view to the reference view using the following inputs:

- A reference RGB-D image $\mathcal{I}_R = (\mathcal{C}_R, \mathcal{D}_R)$ where \mathcal{C}_R denotes the color image and \mathcal{D}_R is the depth map.
- A query RGB-D image $\mathcal{I}_Q = (\mathcal{C}_Q, \mathcal{D}_Q)$ from an unknown viewpoint.
- The object’s binary mask \mathcal{M}_R in the reference image and \mathcal{M}_Q in the query image.

For evaluation, the absolute pose of the reference view \mathbf{T}_{RO} is assumed to be known, allowing us to derive the absolute pose of the query view \mathbf{T}_{QO} .

Reference Object Coordinates (ROC) Map In our method, the relative object pose between the reference and query views is represented by Reference Object Coordinates (ROC) maps, an effective representation introduced by (Liu et al. 2025a). The ROC map of the reference view $\mathbf{X}^R \in \mathbb{R}^{H \times W \times 3}$ is obtained by backprojecting the depth within the reference mask and applying normalization:

$$\mathbf{X}^R = \mathbf{S}\mathbf{\Pi}^{-1}(\mathcal{D}_R)[\mathcal{M}_R = 1], \quad (1)$$

where $\mathbf{\Pi}^{-1}(\cdot)$ denotes the backprojection operator, and $\mathbf{S} \in \mathbb{R}^{4 \times 4}$ is a normalization matrix that centers and scales the object point cloud. The normalization is derived from the inputs; further details are provided in the appendix. Likewise, the ROC map of the query view $\mathbf{X}^Q \in \mathbb{R}^{H \times W \times 3}$ is calculated as:

$$\mathbf{X}^Q = \mathbf{S}\mathbf{T}_{RQ}\mathbf{\Pi}^{-1}(\mathcal{D}_Q)[\mathcal{M}_Q = 1], \quad (2)$$

where \mathbf{T}_{RQ} is the relative transformation from the query to the reference view. Both \mathbf{X}^R and \mathbf{X}^Q represent 3D points in the reference object frame, thereby providing pixel-wise 3D–3D correspondences between the query and reference images. Since \mathbf{X}^R is known in advance, the task of 6D object pose estimation reduces to estimating the ROC map \mathbf{X}^Q of the query image.

The Proposed Method

We introduce CoordAR, a neural network for one-reference 6D object pose estimation. The overview of our method is shown in Fig. 2. Our network consists of three major stages: a modality-decoupled encoding stage, a subsequent fusion stage, and finally an autoregressive decoding stage, which we detail in the following sections. The output of our network is a pixel-aligned ROC map $\hat{\mathbf{X}}^Q \in \mathbb{R}^{H \times W \times 3}$ that directly corresponds to the object’s coordinates in the reference image, as described in Eq. (2).

Modality-decoupled Encoding An encoding stage for both the query and the reference is a prerequisite for visual understanding. Existing work (Liu et al. 2025a) leverages an encoder for the query and another for the reference, which we refer to as *role-specific encoding*. To encode the reference information, they concatenate the RGB image and ROC map channel-wise as input, ignoring the distinctness in structural patterns between the two modalities. In contrast to them, we assign separate encoders for different input modalities (RGB vs. ROC), allowing each encoder to specialize in its respective domain. As shown in Fig. 3, the modality-decoupled encoding we employ includes: 1) A *shared* RGB encoder that processes both the reference image \mathcal{C}_R and the query image \mathcal{C}_Q , and 2) Another encoder that handles the reference’s ROC map \mathbf{X}^R . Details about each encoder are presented in the appendix.

Fusion Blocks To condition token generation on reference-view cues, we introduce several stacked fusion blocks that integrate reference-view information with the query-view features. Each fusion block has a similar structure to the decoder block in the transformer (Vaswani et al. 2017), which primarily consists of a self-attention layer, a following cross-attention layer, and a feedforward network (FFN). To accommodate modality-decoupled encoding, our cross-attention layer in the fusion blocks computes the affinity between encoded features within the same modality to mitigate the RGB-ROC domain gap, as demonstrated on the right side of Fig. 3. This decoupling improves both architectural clarity and performance, as validated by our ablation studies. More details about the fusion blocks can be found in the appendix. Finally, the output features of the fusion blocks are considered the condition features for token generation. Specifically, when decoding a token at a certain position, the position-aligned condition feature is selected and added to the intermediate output of the decoder.

Autoregressive ROC Map Generation Instead of directly regressing the ROC map, our network first generates patch-level tokens and then detokenizes them to obtain the pixel-level ROC map. To this end, we adopt a VQ-VAE (Van Den Oord, Vinyals et al. 2017) as our ROC map tokenizer. The VQ-VAE first encodes the ROC map into latent vectors, then quantizes them by replacing each vector with its nearest neighbor in a pre-trained codebook \mathcal{V} , yielding a discrete token sequence $\{s_1, \dots, s_{h \cdot w}\}$ where $s_* \in \mathcal{V}$. With the introduction of the VQ-VAE, the ROC map can be obtained indirectly by predicting discrete tokens, where a categorical distribution over \mathcal{V} can be established at each patch. Un-

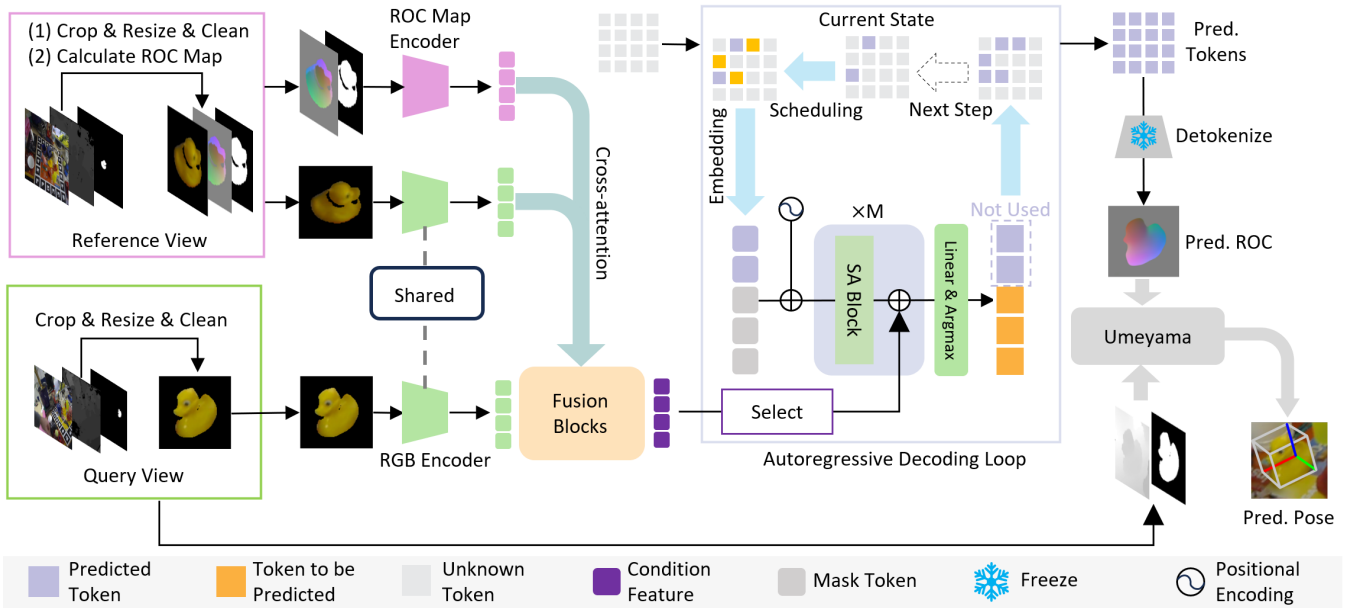


Figure 2: Overview of CoordAR framework. Images from both query view and reference view are cleaned by masks to remove the interference from background and occlusion. The inputs are encoded by modality-decoupled encoders, where encoders are shared among same modality. The reference features are then integrated with the query features to form the condition feature. Subsequently, the decoder, which consists of several self-attention (SA) blocks, autoregressively decode new tokens with learned mask tokens as input. Finally, all tokens are detokenized and combined with the query depth to compute the pose.

like One2Any (Liu et al. 2025a), where coordinates are generated in parallel, our decoder explicitly learns the dependencies between coordinates, which we find to be critical in our experiments. Mathematically, the distribution of tokens is represented as a masked autoregressive model (Li et al. 2024) with the query and reference images as conditions:

$$p(s) = \prod_{k=1}^K p(S_k | S_{<k}, C_F), \quad (3)$$

where $S_k = \{s_i, s_{i+1}, \dots, s_j\}$ are the tokens generated at the k -th step, and $s = \bigcup_{k=1}^K S_k$. Here C_F is the position-

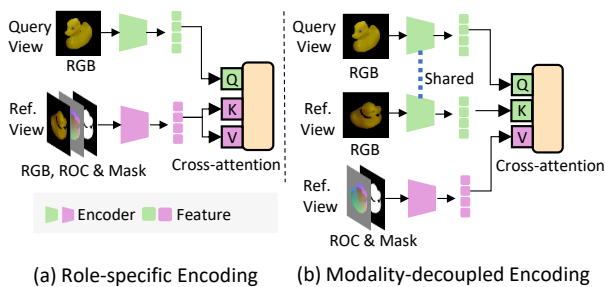


Figure 3: Different encoding schemes and cross-attention mechanisms in the fusion blocks. The modality-decoupled encoding improves architectural clarity and avoid affinity computation across modalities. To distinguish depth-absent regions from the background and invisible areas, we additionally concatenate the object mask to the input ROC map.

aligned condition feature, which is adapted from the feature after the fusion blocks. The training objective of the autoregressive decoder is to minimize the negative log-likelihood loss:

$$\mathcal{L}_{AR} = - \sum_k^K [\log(p(S_k | S_{<k}, C_F))]. \quad (4)$$

During inference, the previously generated tokens and position-aligned condition features both serve as conditioning information for predicting subsequent tokens. After all tokens have been generated, we leverage the decoder of the tokenizer to detokenize the tokens, producing an estimated ROC map $\hat{\mathbf{X}}^Q$. More details about the autoregressive decoder can be found in the appendix.

Recovering Object Pose from ROC Map As described in Eq. (2), given the estimated ROC map $\hat{\mathbf{X}}^Q$, we recover the predicted 3D object points in the reference camera frame by applying the inverse of the normalization matrix:

$$\hat{\mathbf{P}}_R^Q = \mathbf{S}^{-1} \hat{\mathbf{X}}^Q, \quad (5)$$

where \mathbf{S} is the normalization matrix computed from the reference object points, as defined in Eq. (1). To obtain the observed 3D points in the query camera frame, we backproject the depth map \mathcal{D}_Q within the query mask:

$$\mathbf{P}_Q^Q = \mathbf{\Pi}^{-1}(\mathcal{D}_Q)[\mathcal{M}_Q = 1]. \quad (6)$$

where $\mathbf{\Pi}^{-1}(\cdot)$ is the camera backprojection operator. Since $\hat{\mathbf{P}}_R^Q$ and \mathbf{P}_Q^Q are pixel-aligned, we estimate the relative pose \mathbf{T}_{RQ} using the Umeyama algorithm (Umeyama 1991),

which computes the optimal rigid transformation in a least-squares sense:

$$\hat{\mathbf{T}}_{RQ} = \text{Umeyama}(\hat{\mathbf{P}}_R^Q, \mathbf{P}_Q^Q) \quad (7)$$

Finally, given the object pose in the reference view \mathbf{T}_{RO} , the object pose in the query frame is obtained as $\mathbf{T}_{QO} = \hat{\mathbf{T}}_{RQ}^{-1} \mathbf{T}_{RO}$.

Experiments

Benchmark Datasets To evaluate our method under various real-world scenarios, we consider four datasets: Real275 (Wang et al. 2019a), Toyota-Light (Hodaň et al. 2018), LINEMOD (Hinterstoisser et al. 2011) and YCB-V (Xiang et al. 2018). These datasets encompass common challenges in 6D pose estimation, including illumination changes, occlusion, and significant variations in objects (geometric properties, materials, and textures), enabling a comprehensive evaluation of the algorithm.

Training Datasets Consistent with the previous work (Liu et al. 2025a), we train our models on the FoundationPose dataset (Wen et al. 2024) and a subset of the OO3D-9D (Cai et al. 2024) dataset. See the appendix for more details.

Evaluation Metrics To follow the baseline protocols for each setup, we evaluate pose estimation performance using the following metrics:

- Recall of the ADD(-S) error, which is within 0.1 of the object diameter, as used in (He et al. 2022b; Corsetti et al. 2024), shot for ADD(-S).
- Area under the curve (AUC) of ADD and ADD-S (Xiang et al. 2018).
- Average Recall of MSSD, MSPD, and VSD metrics defined in the BOP challenge (Hodaň et al. 2018), shot for AR.

Results on Pose Estimation

We primarily compare our method with model-free pose estimation approaches. For systematic comparison, our analysis includes model-free methods based on both single-view references and multi-view references. The single-view-based methods include Oryon (Corsetti et al. 2024), ObjMatch (Gümel, Dai, and Nießner 2023), NOPE (Nguyen

| Methods | Modality | Real275 | | Toyota-Light | |
|----------------|----------|-------------|-------------|--------------|-------------|
| | | AR | ADD(-S) | AR | ADD(-S) |
| LatentFusion | RGB | 22.6 | 9.6 | 28.2 | 10.2 |
| ObjectMatch | RGBD | 26.0 | 13.4 | 9.8 | 5.4 |
| Oryon | RGBD | 46.5 | 34.9 | 34.1 | 22.9 |
| Any6D | RGBD | 51.0 | – | 43.3 | – |
| One2Any | RGBD | 54.9 | 41.0 | 42.0 | 34.6 |
| CoordAR | RGBD | 71.0 | 82.2 | 62.5 | 82.6 |

Table 1: Comparison of 6D pose estimation methods on Real275 (Wang et al. 2019a) and Toyota-Light (Hodaň et al. 2018) datasets using AR and ADD(-S) metrics. Methods were evaluated on 2K reference-query image pairs.




Figure 4: Visual comparison on LINEMOD and YCB-V datasets. We present results from several state-of-the-art methods alongside our method (**CoordAR**). Ground-truth poses are visualized in the last column, represented by bounding boxes with three distinctly colored edges.

et al. 2024a), Any6D (Lee et al. 2025) and One2Any (Liu et al. 2025a); the multi-view-based methods include FoundationPose (Wen et al. 2024), LatentFusion (Park et al. 2020) FS6D (He et al. 2022b) OnePose (Sun et al. 2022) and OnePose++ (He et al. 2022a). Baseline results are adopted from One2Any (Liu et al. 2025a) with the assumption that ground-truth masks are available. Meanwhile, we follow the same evaluation protocols that they used. More specifically, on the LINEMOD (Hinterstoisser et al. 2011) and YCB-V (Xiang et al. 2018) datasets, the first view is chosen as the reference view for the entire test set. On the Real275 (Wang et al. 2019a) and Toyota-Light (Hodaň et al. 2018) datasets, 2K reference-query image pairs are randomly sampled for evaluation.

Generalization to Real-world Novel Objects We first evaluate our method on Real275 (Wang et al. 2019a) and Toyota-Light (Hodaň et al. 2018) for performance on real-world novel objects. Real275 contains 18 different real-world scenes comprising 42 unique objects across 6 categories, while Toyota-Light includes 21 rigid household objects under 5 different lighting conditions. As shown in Tab. 1, our method surpasses the existing methods in both AR and ADD(-S) metrics, demonstrating excellent generalization to real-world novel objects. Qualitative results of the two datasets can be found in the appendix.

| Methods Ref. Images | Predator | | FS6D | | FoundationPose | | FoundationPose | | NOPE | | One2Any | | CoordAR | |
|------------------------|----------|-------|------|-------|----------------|-------------|----------------|-------|--------------|-------------|-------------|-------------|-------------|-------------|
| | 16 | | 16 | | 16 - CAD | | 1 - CAD | | 1 + GT trans | | 1 | | 1 | |
| metrics of AUC | ADD | ADD-S | ADD | ADD-S | ADD | ADD-S | ADD | ADD-S | ADD | ADD-S | ADD | ADD-S | ADD | ADD-S |
| can | 29.23 | 73.6 | 50.0 | 91.9 | 85.2 | 97.2 | 81.5 | 90.8 | 32.9 | 95.6 | 75.5 | 86.8 | 60.9 | 93.0 |
| box | 21.33 | 62.58 | 45.0 | 93.1 | 94.4 | 98.0 | 81.2 | 91.1 | 20.3 | 85.3 | 90.2 | 78.7 | 92.5 | 98.5 |
| bottle | 23.62 | 73.1 | 39.1 | 87.7 | 90.5 | 97.0 | 73.3 | 90.0 | 26.7 | 89.3 | 90.7 | 93.0 | 87.5 | 99.7 |
| block | 22.75 | 74.85 | 36.8 | 95.2 | 94.1 | 97.8 | 36.8 | 96.7 | 35.7 | 95.0 | 84.9 | 91.1 | 84.6 | 97.3 |
| others | 24.1 | 71.72 | 40.2 | 83.2 | 94.9 | 97.5 | 40.2 | 93.4 | 32.2 | 86.4 | 81.3 | 95.2 | 75.4 | 93.1 |
| mean | 24.3 | 71.0 | 42.1 | 88.4 | 91.5 | 97.4 | 76.1 | 90.4 | 25.1 | 86.0 | 80.6 | 90.3 | 78.5 | 95.5 |

Table 2: Performance on occluded YCB-V (Xiang et al. 2018) dataset. The Predator (Huang et al. 2021), originally proposed for point cloud registration, is additionally provided for reference. The methods are evaluated by AUC of ADD and AUC of ADD-S metrics. The baseline results are adopted from One2Any (Liu et al. 2025a) and reproduced by their released model, where objects are categorized into five groups. Results on each object can be found in the appendix.

| Methods | Modality | Ref. Images | ape | benchvise | cam | can | cat | driller | duck | eggbox | glue | holepuncher | iron | lamp | phone | avg. |
|----------------|----------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|
| OnePose | RGB | 200 | 11.8 | 92.6 | 88.1 | 77.2 | 47.9 | 74.5 | 34.2 | 71.3 | 37.5 | 54.9 | 89.2 | 87.6 | 60.6 | 63.6 |
| OnePose++ | RGB | 200 | 31.2 | 97.3 | 88.0 | 89.8 | 70.4 | 92.5 | 42.3 | 99.7 | 48.0 | 69.7 | 97.4 | 97.8 | 76.0 | 76.9 |
| LatentFusion | RGBD | 16 | 88.0 | 92.4 | 74.4 | 88.8 | 94.5 | 91.7 | 68.1 | 96.3 | 49.4 | 82.1 | 74.6 | 94.7 | 91.5 | 83.6 |
| FS6D + ICP | RGBD | 16 | 78.0 | 88.5 | 91.0 | 89.5 | 97.5 | 92.0 | 75.5 | 99.5 | 99.5 | 96.0 | 87.5 | 97.0 | 97.5 | 91.5 |
| FoundationPose | RGBD | 1-CAD | 36.5 | 55.5 | 84.2 | 71.7 | 65.3 | 16.3 | 49.8 | 42.6 | 64.8 | 52.7 | 20.7 | 15.8 | 51.7 | 48.3 |
| NOPE | RGB | 1 + GT trans | 2.0 | 4.5 | 2.5 | 2.2 | 0.7 | 4.7 | 0.5 | 100.0 | 79.4 | 2.9 | 4.5 | 4.2 | 3.9 | 16.3 |
| Oryon | RGBD | 1 | 1.2 | 1.3 | 3.9 | 0.8 | 12.7 | 8.5 | 0.8 | 63.2 | 18.4 | 1.6 | 0.6 | 2.9 | 11.7 | 9.8 |
| One2Any | RGBD | 1 | 33.1 | 15.7 | 72.7 | 37.0 | 66.2 | 68.2 | 35.8 | 100.0 | 99.9 | 42.0 | 28.2 | 31.9 | 53.2 | 52.6 |
| CoordAR | RGBD | 1 | 45.6 | 76.9 | 70.7 | 77.3 | 88.1 | 96.5 | 50.2 | 97.0 | 99.8 | 67.5 | 52.7 | 91.4 | 61.2 | 75.0 |

Table 3: Performance on LINEMOD (Hinterstoisser et al. 2011) dataset with large view variations. We report the recall of ADD(-S) metric. Baseline results of taken from One2Any (Liu et al. 2025a).

Occluded Scenes The YCB-V (Xiang et al. 2018) dataset contains numerous occluded scenes, including cases where even the first reference view is occluded. As displayed in rows 4 to 7 of Fig. 4, our method exhibits robustness when the query and reference are occluded. We also observe that our method performs well when the object frame is ill-defined (see row 6), suggesting that our method is independent of the definition of the canonical object frame. While we observe a slightly lower ADD AUC in Tab. 2 compared to One2Any (Liu et al. 2025a), our method achieves a significantly higher ADD-S AUC than existing single-view methods. We attribute this discrepancy to a bias in the YCB-V evaluation protocol. Note that YCB-V contains texture-rich food containers that are geometrically symmetric but have different texture-based symmetry definitions during evaluation, such as the *tomato_soup_can*. In such cases, ADD AUC can penalize predictions that are geometrically correct but differ in texture alignment, even though they achieve strong performance under the ADD-S metric (see the full YCB-V results in the appendix for details). Interestingly, when switching to the first frame of each scene as the reference, our method outperforms One2Any in both ADD AUC and ADD-S AUC, as demonstrated by the pose tracking results reported in the appendix.

Large View Variations To evaluate robustness to large viewpoint variations, a comparison is conducted on the LINEMOD (Hinterstoisser et al. 2011) dataset. This dataset features multiple texture-less objects, such as a toy ape

and a hole-puncher. The images are captured by circling around each object, resulting in significant viewpoint variations. As displayed in Tab. 3, our approach achieves the highest performance across the majority of objects compared with single-view-based methods. Qualitative results are shown in rows 1 to 4 of Fig. 4. Notably, our method succeeds in estimating the pose of the top view (see row 2), the side view (see row 3), and even the back view (see row 1). While our method does not surpass the state-of-the-art multi-view approaches, it demonstrates overall superiority over OnePose (Sun et al. 2022) and achieves better performance on several objects compared to OnePose++ (He et al. 2022a).

Ablation Studies

To justify the key design choices, we conduct ablation experiments on the LINEMOD dataset. Due to computational limitations, models are trained with reduced iterations (see the appendix for more details).

Effect of Autoregressive Decoder We first study the overall effectiveness of the autoregressive decoder for ROC maps by comparing it against a convolutional regression decoder that has a similar architecture to the decoder in One2Any (Liu et al. 2025a). For a fair comparison, we reduce the number of parameters in our decoder to match those of the convolutional decoder. As shown in Tab. 4, both metrics decrease after replacing the autoregressive decoder with the convolutional decoder. For further understanding, we

| Component | Variations | #Params (B) | ADD(-S) | AR |
|----------------|--------------------|-------------|-------------|-------------|
| ROC decoding | convolutional | 0.28 | 70.9 | 59.7 |
| | autoregressive | 0.28 | 73.1 | 61.6 |
| Autoregression | w/o | 0.37 | 60.7 | 52.1 |
| | w/ | 0.37 | 73.6 | 61.9 |
| Test-time AR | w/o | 0.37 | 68.0 | 56.3 |
| | w/ | 0.37 | 73.6 | 61.9 |
| Tokenization | w/o | 0.37 | 56.4 | 48.7 |
| | w/ | 0.37 | 60.7 | 52.1 |
| Encoding | role-specific | 0.37 | 61.6 | 49.8 |
| | modality-decoupled | 0.37 | 73.6 | 61.9 |

Table 4: Ablation study on critical design choices. All evaluations are conducted on the full LINEMOD dataset using the AR and ADD(-S) metrics. Parameter counts (in billions, excluding the tokenizer) are provided for reference.

keep the tokenization and disable the autoregression: 1) during both training and testing by training a model that predicts tokens in parallel, or 2) only during testing by generating all tokens in a single step. According to Tab. 4, disabling autoregression notably degrades performance. Fig. 5 provides a reasonable explanation for this result in columns 2 and 3, where non-autoregressive predictions exhibit disrupted spatial coherence. For example, on the top-right of the *can* (row 2, column 3), dark-purple values are incorrectly predicted where light-green should appear. As demonstrated in Tab. 4, disabling test-time autoregression reduces both ADD(-S) and AR by 5.6%, decreasing ADD(-S) by 0.7% when training is prolonged (see Tab. 5), suggesting that previously generated tokens can serve as an effective conditional context.

Effect of Tokenization Afterward, we further remove the reliance on tokenization by regressing real-valued ROC maps from the features of the pre-final layer in the transformer decoder. As quantified in Tab. 4, replacing token prediction with real-value regression further decreases recall. This may be attributed to their handling of ambiguous cases. As depicted in Fig. 5, this model produces ambiguous ROC maps when dealing with symmetry ambiguity on the *bowl* (row 1, column 2) and occlusion ambiguity on the *cup* (row 3, column 2). A similar failure also occurs in the convolutional-head variant (see results on the *bowl* in row 1 and the *box* in row 5). Notably, discrete-token-based models (columns 3–4) demonstrate improved performance with clearer coordinate maps, suggesting that probabilistic modeling plays a key role in resolving such ambiguities.

Effect of Modality-decoupled Encoding To validate the effectiveness of the modality-decoupled encoding, we replace it with the role-specific encoding used in One2Any (Liu et al. 2025a). More specifically, we employ two DINOv2-structured encoders: one processing the query RGB image and another handling the channel-wise concatenation of the reference RGB image and its corresponding ROC map and mask. As shown in Tab. 4, the performance degrades significantly after switching to role-specific encoding, indicating that allocating encoders by modality is critical for visual understanding.

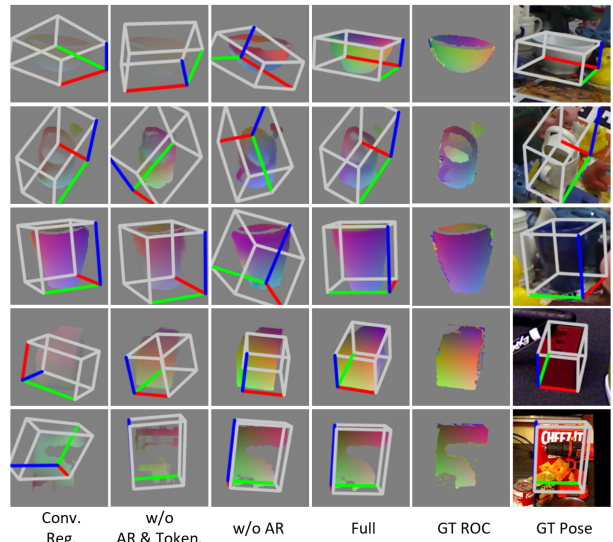


Figure 5: Qualitative Results of Ablation Study. We visualize the outputs of four decoder variants: (1) convolutional decoder, (2) full decoder without autoregressive tokenization, (3) full decoder without tokenization, and (4) full decoder.

| Methods | Modality | GPU | Time (s) | ADD(-S) |
|----------------|----------|----------|-------------|-------------|
| FoundationPose | 1-CAD | RTX 4090 | 2.70 | 48.3 |
| One2Any | RGBD | RTX4090 | 0.09 | 56.2 |
| Ours-64 steps | RGBD | RTX 4090 | 0.63 | 75.0 |
| Ours-16 steps | RGBD | RTX 4090 | 0.25 | 75.0 |
| Ours-4 steps | RGBD | RTX 4090 | 0.13 | 74.7 |
| Ours-1 step | RGBD | RTX 4090 | 0.10 | 74.3 |

Table 5: Inference time comparison. The runtimes of FoundationPose (Wen et al. 2024) and One2Any (Liu et al. 2025a) are taken from One2Any (Liu et al. 2025a).

Runtime Analysis

Our model supports trade-offs between accuracy and computational efficiency by adjusting the number of generation steps. Inference speed comparisons, along with their corresponding ADD(-S) on the LINEMOD dataset, are provided in Tab. 5. Impressively, our model can achieve near real-time speed (0.10 seconds per frame) with even a single step while maintaining comparable accuracy.

Conclusions

In this paper, we propose the first autoregressive framework for one-reference 6D pose estimation of novel objects. By formulating correspondence prediction as an autoregressive probabilistic token decoding task and introducing modality-decoupled encoding for visual understanding, CoordAR achieves superior performance on standard benchmarks. Extensive experiments demonstrate significant improvements in handling symmetry, occlusion, and novel objects. This work establishes autoregressive coordinate modeling as a promising direction for robust 6D pose estimation.

Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant 2022YFB3903801, in part by the National Science Foundation of China under Grant 62073214, and in part by the Corporate Research Center, State Key Laboratory of High-end Heavy-load Robots, Midea Group (3D Robot Vision Project). We thank Kun Wang, Li Shen, Yuhui Ni, and Yikun Zeng for providing the reproduced results of baseline methods for qualitative comparison.

References

- Cai, J.; He, Y.; Yuan, W.; Zhu, S.; Dong, Z.; Bo, L.; and Chen, Q. 2024. Ov9d: Open-vocabulary category-level 9d object pose and size estimation. *arXiv preprint arXiv:2403.12396*.
- Caraffa, A.; Boscaini, D.; Hamza, A.; and Poiesi, F. 2024. FreeZe: Training-free zero-shot 6D pose estimation with geometric and vision foundation models. In *European Conference on Computer Vision (ECCV)*.
- Chen, D.; Li, J.; Wang, Z.; and Xu, K. 2020. 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*, 11973–11982.
- Chen, K.; and Dou, Q. 2021. Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2773–2782.
- Corsetti, J.; Boscaini, D.; Oh, C.; Cavallaro, A.; and Poiesi, F. 2024. Open-vocabulary object 6D pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18071–18080.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Fan, Z.; Pan, P.; Wang, P.; Jiang, Y.; Xu, D.; and Wang, Z. 2024. Pope: 6-dof promptable pose estimation of any object in any scene with one reference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7771–7781.
- Gümeli, C.; Dai, A.; and Nießner, M. 2023. Objectmatch: Robust registration using canonical object correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13082–13091.
- He, X.; Sun, J.; Wang, Y.; Huang, D.; Bao, H.; and Zhou, X. 2022a. Onepose++: Keypoint-free one-shot object pose estimation without CAD models. *Advances in Neural Information Processing Systems*, 35: 35103–35115.
- He, Y.; Wang, Y.; Fan, H.; Sun, J.; and Chen, Q. 2022b. Fs6d: Few-shot 6d pose estimation of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6814–6824.
- Hinterstoisser, S.; Holzer, S.; Cagniart, C.; Ilic, S.; Konolige, K.; Navab, N.; and Lepetit, V. 2011. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *2011 international conference on computer vision*, 858–865. IEEE.
- Hodan, T.; Barath, D.; and Matas, J. 2020. Epos: Estimating 6d pose of objects with symmetries. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11703–11712.
- Hodaň, T.; Michel, F.; Brachmann, E.; Kehl, W.; Glent Buch, A.; Kraft, D.; Drost, B.; Vidal, J.; Ihrke, S.; Zabulis, X.; Sahin, C.; Manhardt, F.; Tombari, F.; Kim, T.-K.; Matas, J.; and Rother, C. 2018. BOP: Benchmark for 6D Object Pose Estimation. *European Conference on Computer Vision (ECCV)*.
- Huang, S.; Gojcic, Z.; Usvyatsov, M.; Wieser, A.; and Schindler, K. 2021. Predator: Registration of 3d point clouds with low overlap. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 4267–4276.
- Jiang, K.; and Huang, J. 2024. A Survey on Vision Autoregressive Model. *arXiv preprint arXiv:2411.08666*.
- Labbé, Y.; Manuelli, L.; Mousavian, A.; Tyree, S.; Birchfield, S.; Tremblay, J.; Carpentier, J.; Aubry, M.; Fox, D.; and Sivic, J. 2022. MegaPose: 6D Pose Estimation of Novel Objects via Render & Compare. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*.
- Lee, T.; Wen, B.; Kang, M.; Kang, G.; Kweon, I. S.; and Yoon, K.-J. 2025. Any6D: Model-free 6D Pose Estimation of Novel Objects. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 11633–11643.
- Li, T.; Tian, Y.; Li, H.; Deng, M.; and He, K. 2024. Autoregressive image generation without vector quantization. *Advances in Neural Information Processing Systems*, 37: 56424–56445.
- Li, Z.; Wang, G.; and Ji, X. 2019. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7678–7687.
- Liu, M.; Li, S.; Chhatkuli, A.; Truong, P.; Van Gool, L.; and Tombari, F. 2025a. One2Any: One-Reference 6D Pose Estimation for Any Object. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 6457–6467.
- Liu, X.; Zhang, R.; Zhang, C.; Wang, G.; Tang, J.; Li, Z.; and Ji, X. 2025b. GDRNPP: A Geometry-guided and Fully Learning-based Object Pose Estimator. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Liu, Y.; Wen, Y.; Peng, S.; Lin, C.; Long, X.; Komura, T.; and Wang, W. 2022. Gen6D: Generalizable Model-Free 6-DoF Object Pose Estimation from RGB Images. In *ECCV*.
- Micheli, V.; Alonso, E.; and Fleuret, F. 2022. Transformers are sample-efficient world models. *arXiv preprint arXiv:2209.00588*.
- Nguyen, V. N.; Groueix, T.; Ponimatkin, G.; Hu, Y.; Marlet, R.; Salzmann, M.; and Lepetit, V. 2024a. Nope: Novel object pose estimation from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17923–17932.

- Nguyen, V. N.; Groueix, T.; Salzmann, M.; and Lepetit, V. 2024b. Gigapose: Fast and robust novel object pose estimation via one correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9903–9913.
- Park, K.; Mousavian, A.; Xiang, Y.; and Fox, D. 2020. Latentfusion: End-to-end differentiable reconstruction and rendering for unseen object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10710–10719.
- Su, Y.; Saleh, M.; Fetzer, T.; Rambach, J.; Navab, N.; Busam, B.; Stricker, D.; and Tombari, F. 2022. Zebrapose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6738–6748.
- Sun, J.; Wang, Z.; Zhang, S.; He, X.; Zhao, H.; Zhang, G.; and Zhou, X. 2022. OnePose: One-Shot Object Pose Estimation without CAD Models. *CVPR*.
- Umeyama, S. 1991. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13(04): 376–380.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, G.; Manhardt, F.; Tombari, F.; and Ji, X. 2021. GDR-Net: Geometry-Guided Direct Regression Network for Monocular 6D Object Pose Estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16611–16621.
- Wang, H.; Sridhar, S.; Huang, J.; Valentin, J.; Song, S.; and Guibas, L. J. 2019a. 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*, 2642–2651.
- Wang, H.; Sridhar, S.; Huang, J.; Valentin, J.; Song, S.; and Guibas, L. J. 2019b. Normalized Object Coordinate Space for Category-Level 6D Object Pose and Size Estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wen, B.; Yang, W.; Kautz, J.; and Birchfield, S. 2024. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17868–17879.
- Xiang, Y.; Schmidt, T.; Narayanan, V.; and Fox, D. 2018. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes.
- Xiong, J.; Liu, G.; Huang, L.; Wu, C.; Wu, T.; Mu, Y.; Yao, Y.; Shen, H.; Wan, Z.; Huang, J.; et al. 2024. Autoregressive models in vision: A survey. *arXiv preprint arXiv:2411.05902*.
- Yu, L.; Cheng, Y.; Sohn, K.; Lezama, J.; Zhang, H.; Chang, H.; Hauptmann, A. G.; Yang, M.-H.; Hao, Y.; Essa, I.; et al. 2023. Magvit: Masked generative video transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10459–10469.
- Zhang, J. Y.; Ramanan, D.; and Tulsiani, S. 2022. Relpose: Predicting probabilistic relative rotation for single objects in the wild. In *European Conference on Computer Vision*, 592–611. Springer.