

DiffusionNOCS: Managing Symmetry and Uncertainty in Sim2Real Multi-Modal Category-level Pose Estimation

Takuya Ikeda^{*1}, Sergey Zakharov^{*2}, Tianyi Ko¹, Muhammad Zubair Irshad², Robert Lee¹, Katherine Liu², Rares Ambrus², Koichi Nishiwaki¹

Abstract—This paper addresses the challenging problem of category-level pose estimation. Current state-of-the-art methods for this task face challenges when dealing with symmetric objects and when attempting to generalize to new environments solely through synthetic data training. In this work, we address these challenges by proposing a probabilistic model that relies on diffusion to estimate dense canonical maps crucial for recovering partial object shapes as well as establishing correspondences essential for pose estimation. Furthermore, we introduce critical components to enhance performance by leveraging the strength of the diffusion models with multi-modal input representations. We demonstrate the effectiveness of our method by testing it on a range of real datasets. Despite being trained solely on our generated synthetic data, our approach achieves state-of-the-art performance and unprecedented generalization qualities, outperforming baselines, even those specifically trained on the target domain. Our code and data for the generalization benchmark can be found at <https://woven-planet.github.io/DiffusionNOCS/>.

I. INTRODUCTION

In the fields of computer vision and robotics, object pose estimation is a fundamental problem. Recently, approaches based on neural networks [1]–[3] have shown great progress for instance-level object pose estimation, quickly saturating the existing benchmarks [4]. Consequently, a more challenging benchmark was proposed [5], relaxing instance-level object pose estimation to category-level and evaluating the pose of unseen objects within a given category. While a significant number of methods have made progress in category-level pose estimation [6]–[8], many challenges such as handling ambiguity and generalization persist.

The challenges inherent in modelling the ambiguity of object pose are particularly apparent in symmetric objects, where a set of poses appear to be equally feasible given a partial observation. Most current methods regress a single pose when presented with a partially observed input, and therefore must craft methods to handle symmetry issues such as extensive data labeling and heuristic operations to avoid training instabilities that can stem from the underlying multi-modality of the pose. In some cases, ambiguity can be resolved given additional information beyond geometry. For example, adding an image-space observation of a laptop can help to disambiguate the keyboard from the screen

¹ All authors are with the Woven by Toyota. 3 Chome-2-1 Nihonbashimurumachi, Chuo City, Tokyo 103-0022, Japan, [firstname.lastname]@woven.toyota

² All authors are with the Toyota Research Institute, 4440 El Camino Real, Los Altos, CA 94022, USA, [firstname.lastname]@tri.global

* Equal contribution.

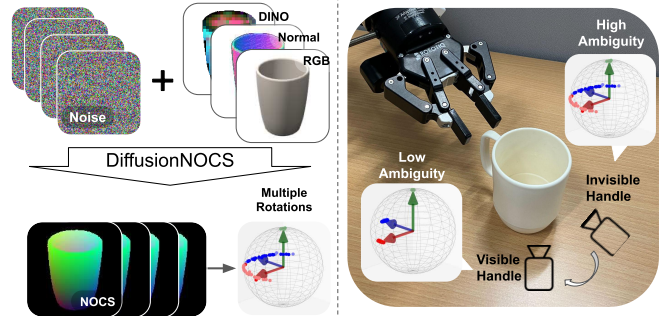


Fig. 1. **Overview:** Our method can estimate multiple possible poses via diffusion models with a single observation. The diffusion models can be conditioned on any image, and naturally handle the ambiguity present from symmetry, where the predictions are certain when the mug handle is visible in the input image and are uncertain when it is invisible. The projected Normalized Object Coordinate Space (NOCS) map is denoised by the diffusion model and used as dense correspondences to estimate the poses.

[9]. However, there are also common object instances for which visual texture cannot disambiguate the pose; consider the case of a uniformly colored cylindrical cup, where any rotation around the upright axis is a plausible pose candidate. Additionally, the existence and quality of different input modalities may change at inference time due to variations in sensor configuration and performance. We posit that utilizing a probabilistic method for pose estimation alongside multi-modal input representations can significantly improve the performance when ambiguities are present at deployment.

Generalization to the diversity within a class and across environments is another crucial challenge in building robust pose estimation. Collecting sufficient real-world training data for 3D pose estimation robust to visual texture and environment changes is a prohibitively consuming and costly process. Many existing works in the field of category-level pose estimation are trained on a specific set of real-world data and evaluated in a single environment [5]. Promising methods of mitigating the data scarcity issue include leveraging semantic features learned at scale for pose estimation [10] and utilizing readily available labels from synthetic data [9]. However, there is a notable lack of existing benchmarks suitable for evaluating the zero-shot performance of category-level pose estimation in novel, real-world settings.

In this work, we introduce a novel category-level dense pose estimation method based on diffusion models, with multi-modal input representations as shown in Fig. 1. The multi-modal inputs and the probabilistic nature of our method allow it to naturally handle uncertainty, improving robustness

and demonstrating state-of-the-art performance on various real datasets despite being trained solely on synthetic data. The contributions of this paper are as follows:

- 1) A novel pipeline for category-level object pose estimation via **diffusion models with selectable multi-modal inputs**, including different combinations of RGB, surface normals and learned semantic features. Our approach’s probabilistic nature enables it to seamlessly handle a range of symmetric objects without relying on hand-crafted heuristics.
- 2) A **zero-shot generalization benchmark** demonstrating how existing state-of-the-art methods perform in challenging real-world environments featuring pose distribution, lighting, and occlusion changes. Despite being trained only on synthetic data, our proposed approach outperforms all baselines using synthetic and real data on average.

II. RELATED WORK

A. Object Pose Estimation

Most existing pose estimation methods can be grouped into three categories: direct pose regression [11], [12] (learning to output the pose directly from data), template-matching [13]–[15] (comparing inputs with a predefined set of templates), and correspondence-based methods [1], [2] (learning to estimate correspondences between input data and corresponding 3D models of interest) - each having their own benefits and drawbacks. Overall, the current state-of-the-art instance and category-level methods [3], [5], [6] for object pose estimation are almost exclusively correspondence-based approaches. Correspondence-based methods have clear benefits when handling occlusions and generalizing to different cameras. However, they typically require carefully curated datasets and hand-crafted operations to handle symmetrical objects to solve the ambiguity of 6D object pose. For example, Wang et al. [5] designed loss functions for specific symmetry types. This is because the deterministic nature of their training cannot manage multiple feasible pose hypotheses. Since real-world objects often exhibit symmetry, this poses a challenge when scaling to larger training datasets and applying these methods in the real world.

B. Handling Symmetry and Uncertainty

While common pose estimation methods have shown great promise, they often lack the capability to provide information about uncertainty and are over-confident in their predictions [16]. In particular, they often rely on extensive prior knowledge of object geometry to address ambiguities coming from symmetry and occlusions. Such methods often require a dataset of symmetry-labeled objects [5], which is impractical in real-world applications. Another family of methods uses a probabilistic formulation that offers a means to capture uncertainty in the instance of interest directly. Kendall et al. [17] augmented PoseNet [11] with uncertainty by sampling the posterior to approximate probabilistic inference. Bishop et al. [18] learned to predict parameters of a Gaussian mixture distribution by using a neural network, which in turn

can be used to infer the uncertainty based on the variance of the predicted distribution. Gilitschenski et al. [19] presented orientation uncertainty learning via Bingham distribution which is suited to represent the rotational ambiguity. However, despite its effectiveness in representing axial symmetry, managing discrete symmetry is challenging as the multi-modal distribution needs to be handled [20]. On the other hand, diffusion models offer a promising alternative, as these models can directly handle the multi-modality that is present in the training data [21]. Recently, Zhang et al. [22] used a generative diffusion approach to estimate poses from partial point clouds, naturally handling all symmetry types. Our method also uses diffusion, but we instead estimate dense canonical maps from multi-modal image inputs that recover both pose and partial geometry, which not only enables us to handle symmetry and other types of ambiguity arising from object occlusions and truncation but also provides robustness against the potential degradation of noisy input data.

III. METHOD

Task Description. We seek to estimate the 6D pose of unknown instances in known categories from a single image. At inference time, we assume a depth image, 2D bounding box and category ID encompassing detected objects are given, but can flexibly accommodate more inputs, such as RGB, if available. To train the model, we use only synthetic data generated via rendering openly available 3D models, resulting in respective RGB, depth images, and category IDs.

Overview. We propose a dense canonical map estimator via diffusion models that does not require special operations to handle any type of symmetry. We name our method DiffusionNOCS as the normalized object coordinate space (NOCS) [5] is leveraged as the dense canonical map, which our diffusion model is trained to predict. In this section, we first explain the network architecture and its input representations, then propose a synthetic data generation method and an entire pipeline for category-level object pose estimation.

A. NOCS Preliminaries

We utilize the NOCS as proposed in previous work [5], where each object is canonicalized such that all points exist in some space $x, y, z \in [0, 1]$. The NOCS map is then the projection of the observed surface of the object into the camera frame. The resulting representation is an image-like quantity for which the channel dimension encodes the 3D coordinate of the surface point represented by the pixel. Given a predicted NOCS map, the pose of an object can be extracted via standard pose estimation techniques (i.e., PnP if only RGB is available and point cloud registration if depth information is available).

B. NOCS Estimator with Diffusion Models

One of the key insights of this work is to leverage Denoising Diffusion Probabilistic Models (DDPMs) [25] for NOCS map estimation, enabling the prediction of multiple possible maps via multiple noise samples. Although DDPMs [25] are commonly used for producing photo-realistic images

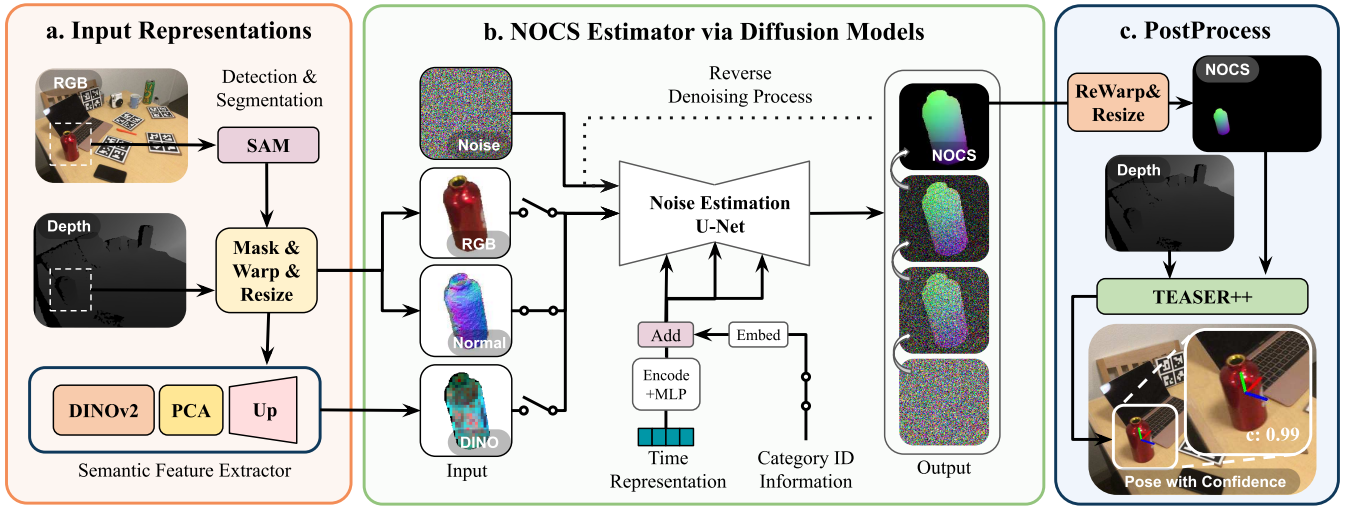


Fig. 2. **Pipeline Overview:** To prepare input representations from RGB, depth and given 2D bounding box, the operations of masking, warping, resizing are conducted sequentially. For the segmentation, SAM [23] with box prompts is used. To extract low-dimensional DINOv2 features with fixed shape, PCA and resizing are applied. Then, the diffusion model is conditioned by available inputs and estimates NOCS maps from noise images. Lastly, 6D pose and confidence are estimated via the point registration between denoised NOCS maps and partial point clouds from mask, and given intrinsics [24].

from noise, we make two modifications for our purpose. First, we change the output from photo-realistic images to NOCS maps. Second, we condition the denoising process on images and category IDs. As such, we use the following loss for training:

$$L = \text{PixelwiseMSE}(\epsilon^k, \epsilon_\theta(I, C, N^0 + \epsilon^k, k)), \quad (1)$$

where k represents the denoising iteration. ϵ^k is a random Gaussian noise map with proper variance for iteration k . N^0 represents the raw NOCS maps, which are the ground truth from the dataset. ϵ_θ is a noise estimation network with parameters θ . I and C represent images and category IDs, respectively, for the conditioning of the noise estimator. For the training details, we refer reader to DDPMs [25]. At the training phase, the discrete denoising scheduler [25] with 1000 iterations is used as a noise scheduler.

The network overview is shown in the middle of Fig. 2. Our main network is based on 2D U-Net [26] with ResNet blocks [27]. Images to be conditioned are concatenated to the noise before being passed to the U-Net. Additionally, two more components are utilized to manage the representation of time and category IDs. The first component is a sinusoidal positional encoding [28] and an MLP for time representation. The second is an embedding layer that serves as a lookup table for the category IDs. These outputs are then added and used as input for each layer of the U-Net model.

While the DDPM scheduler is utilized during training with 1000 iterations, repeating such a large number of iterations at inference time can be quite time-consuming. This is a significant concern for practical robotic applications where time efficiency is crucial. As such, during inference we utilize the DPM-solver [29], which is a fast ODE-solver for diffusion models. This allows us to significantly decrease the iterations from 1000 to 10, improving the time efficiency to enable more effective use for real-world robotics applications.

C. Representations for Conditioning

To estimate 6D poses for a large variety of objects with limited amount of real or synthetic data, the choice of input representation is critical. In our work, we choose to use four types of conditioning information: surface normals, RGB, DINOv2 [30] semantic features, and category IDs.

a) Surface normals: Depth information is essential for category-level pose estimation pipelines. Not only because it can help resolve scale-translation ambiguities, but also because it gives a good estimate of the partial object geometry. However, using unprocessed depth with learning-based methods poses certain challenges since it entangles scale, translation and camera intrinsic parameters, requiring large amounts of data to avoid overfitting. In our work, we decouple estimation of these parameters and condition our NOCS estimator network on scale/translation invariant surface normals computed directly from input depth data.

b) RGB images: Geometric information alone can be highly ambiguous due to depth sensor’s sensitivity to reflective objects, or when dealing with semantically rich but geometrically simple objects. For example, CPPF [9] discusses the difficulty in distinguishing the lid of a laptop from the keyboard base from potentially noisy geometric information alone, such as partial point clouds. Therefore, the authors propose utilizing an additional network that takes RGB inputs. To be able to handle such cases, our method also uses RGB as a conditioning input.

c) DINOv2 features: While RGB data can be highly useful, it is also a highly challenging modality. Even a single class of objects can contain a large variety of materials resulting in a diverse set of possible object appearances, while being very similar geometrically. Moreover, material appearance can further dramatically change under different lighting conditions. As a result, it requires large amounts of data to train methods that faithfully generalize to these



Fig. 3. **Semantic Features:** DINOv2 features via 3-dimensional PCA with input RGB to understand the consistency across the different instances.

changes, which can be very expensive. Thankfully, powerful feature extractors trained on vast amounts of RGB data already exist and are readily available. In our work, we leverage DINOv2 [30]—a foundation model trained in a self-supervised manner, which proved to be useful for a variety of image-level visual tasks [10].

To use DINOv2 features as a conditioning modality, we first preprocess them. The original feature dimension of the ViT-S-based DINOv2 is 384, which is much larger than all the previous modalities combined. This high-dimensional input not only makes inference and training speed slower, but also requires a larger network resulting in a larger memory footprint. To solve this problem, we follow Goodwin et al. [10] and utilize principal component analysis (PCA) to reduce the dimensionality from N to M ($N > M$). The DINOv2 features via PCA with $N = 384$, $M = 3$ are visualized in Fig. 3, and show consistent coloring for different object instances from both real and synthetic images.

d) Category IDs: Category IDs are also included as an input modality. They are readily available from 2D object detectors and serve as a valuable global semantic prior.

D. Selectable Input Representations

For the learning-based category-level pose estimation, depth only or RGB-depth is a common input representation. To make our method more flexible, we propose an approach that can freely take the available inputs at inference time with a single network without re-training. To do this, we drop each condition at a certain rate during training as follows:

$$\tilde{\epsilon}_\theta = \epsilon_\theta(I_{normal} \vee \emptyset, I_{RGB} \vee \emptyset, I_{dino} \vee \emptyset, C \vee \emptyset, N^0 + \epsilon^k, k) \quad (2)$$

I in Eq. 1 is separated to I_{normal} , I_{RGB} , and I_{dino} . When a condition is dropped, it is replaced with a *null* value \emptyset , which in practice means a tensor of the same dimension as the original condition, but filled with zeros. For example, dropping an image conditioning variable means that all pixels are filled by zeroes, and we let the zeroth ID represent an unknown category. Even if all conditioning modalities are dropped, the model attempts to generate NOCS maps from only noise, i.e., it performs unconditional diffusion. On the other hand, if an RGB image is available, the model generates NOCS by conditioning on the RGB image as shown in Fig. 4. At inference time, available inputs can be used without re-training. We observe this conditioning method is both powerful and flexible, enabling new useful representations to be incorporated as they are developed.

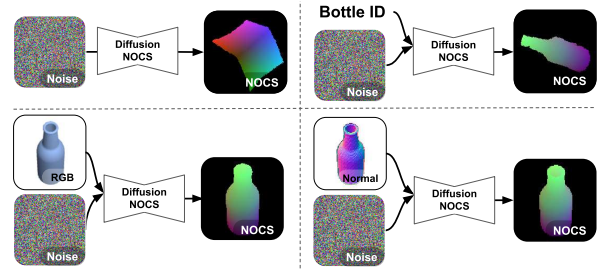


Fig. 4. **Estimated NOCS maps from Different Conditions:** As a generative model, our method is capable of recovering plausible NOCS maps even when provided with only noise image input, i.e. the top left example shows a NOCS map highly resembling the “laptop” object. Adding a category ID to the input limits the output distribution to the objects of a specified category (top right). When RGB images or surface normals are provided, the silhouettes of the resulting NOCS maps are aligned with the inputs. Moreover, the output NOCS faithfully preserves geometric information provided by the surface normals as can be seen from the bottom right example where a hollow bottle is reconstructed. In contrast, the filled bottle is recovered when RGB information alone is provided (bottom left).

E. Synthetic Data Generation

To make DiffusionNOCS generalize across various real environments using just a single dataset, we introduce a domain-agnostic synthetic data generation method. We assume 3D models with aligned coordination system are given. RGB and depth images from 162 camera poses using the 3D models are rendered on white background as shown in Fig. 5. The camera poses are created via icosahedron sampling with 2 subdivision levels. Then, we create normal and groundtruth NOCS maps given rendered depth as well as the known poses and sizes of objects. We estimate the DINOv2 features for all RGB images in the training dataset and apply principle component analysis (PCA) to reduce the dimension. As the resolution of DINOv2 features becomes lower than input RGB, we resize with nearest neighbor interpolation to ensure consistent dimensions. To aid in generalization, we augment the data in three ways. The first is to apply in-plane rotation for the images at the training time instead of making camera poses more various at rendering time. The second is to apply lighting augmentation based on Phong reflection model [31] with the normal maps at training time. Because of this, the lighting variation in the rendered RGB images can be less since it can be added after the rendering. The third is conducting Cutout [32] by random shapes that are rectangle, circle, triangle, and ellipse to improve the robustness against occlusion. At inference time, we also apply a warping operation to move the object to the center of image to ensure consistency with the training images. We show how well the generalization of DiffusionNOCS trained by the synthetic data works across various real environments in Sec. IV-B.




Fig. 5. **Synthetic Data for DiffusionNOCS:** A visualization of camera poses and generated synthetic images used to train DiffusionNOCS.

F. Pipeline for Category-Level Pose Estimation

The overview of entire inference pipeline for the 6D pose estimation is described in Fig. 2. For the condition of DiffusionNOCS, we aim to prepare fixed square shape normal maps, RGB images, and DINOv2 features with white background from given RGB, depth images and intrinsic parameters with 2D bounding box from detection. To create the normal maps and RGB images with fixed square shape, the cropping, warping and resizing operations are conducted sequentially on given RGB and depth images. For the cropping and warping, detected square bounding boxes and their center are used. To generate segmentations, a pretrained SAM [23] with 2D bounding box prompts is used. For the creation of low-dimensional DINOv2 features, the processed RGB image are fed to DINOv2 network, then PCA and resizing to the same size of the RGB image are applied. After the NOCS map estimation via DiffusionNOCS, the square NOCS maps are reverted to the original image shape by re-warping and resizing, then a point cloud registration is conducted between points from masked depth image with intrinsic and points from NOCS maps to estimate 6D pose and scale. TEASER++, a fast point cloud registration algorithm [24], is utilized for the registration. After the registration, we treat the inlier rate from TEASER++ as the confidence of the estimated poses. This confidence is an important value for the next section.

G. Handling Uncertainty for Robust Estimation

We propose an effective approach to boost the performance of 6D pose estimation at inference time via multiple noise sampling. DiffusionNOCS can estimate multiple possible NOCS maps from the different input noises, even if with the same conditioning. In other words, possible multiple poses can be estimated via multiple noise inputs as shown in Fig. 6. At inference time, we can input multiple noises as a batch, then select one estimated pose with the highest confidence as a final result. In this work, we take the inliers rate from the TEASER++ rotation estimate as the confidence value. This approach significantly improves the robustness of pose estimation, as shown in Table III. Moreover, multiple estimated poses can be leveraged by various downstream tasks, such as grasping and placing. For example, if multiple high-confidence poses exist, the best grasping or placing pose can be selected by considering collision or kinematics constraints. As another example, to grasp the mug handle properly, the shape of rotation distribution can be used to measure the pose ambiguity as shown in Fig. 1 and Fig. 8.

IV. EXPERIMENTS

A. Experimental Setup

Benchmarks. We evaluate our method on two benchmarks comprising a total of four real-world datasets, namely NOCS Real 275 [5], TYO-L [33], YCB-V [34], and HOPE [35].

The **NOCS Real Benchmark** [5] is a de facto standard benchmark for category-level pose estimation. It contains objects of 6 target categories (bottle, can, bowl, laptop, camera, mug) placed in 6 different environments spanning




Fig. 6. **Final Pose Selection:** By different noise images, multiple NOCS maps can be estimated. TEASER++ predicts poses and their confidence by the NOCS maps. The final pose is selected based on the confidence value.

a total of 2.75K test images. For a fair comparison to previous works [6], [22], we use 2D bounding boxes and category IDs from precomputed MaskRCNN results as input during inference. The fine masks are estimated by pretrained SAM [23] with the 2D bounding box prompts.

To demonstrate how existing state-of-the-art (SOTA) methods perform in various challenging real-world environments, we introduce a zero-shot **Generalization Benchmark** consisting of three datasets commonly used for instance-level pose estimation, TYO-L [33], YCB-V [34], and HOPE [35]. Each dataset poses different challenges to the baselines. TYO-L dataset features 21 single-object scenes with five different lighting conditions and four different backgrounds. HOPE dataset contains 10 highly occluded multi-object scenes with a very diverse pose distribution. The YCB-V dataset contains 12 multi-object scenes with light occlusions and is most similar to NOCS in terms of pose distribution. Figure 7 shows sample images from the used datasets. To ensure fair evaluation, we standardize object models to be consistently canonicalized with the NOCS dataset and only use the object categories present in NOCS, resulting in 4 categories for TYO-L and YCB-V (bottle, bowl, mug, can), and 2 categories for HOPE (bottle, can). We utilize the GT masks for all baselines and ours during inference.

Data Generation. For training, we exclusively use synthetic data rendered using the Pyrender [36] library. To



Fig. 7. **Visualization of estimated poses on NOCS Real 275, TYO-L, YCB-V, and HOPE datasets:** Despite training only on background-free synthetic data, we show remarkable zero-shot generalization on datasets featuring diverse pose, lighting, and occlusion distributions. HOPE dataset masks are colored as translucent white to highlight the occluded parts.

ensure a non-biased class distribution, we render an equal number (20) of 3D models per target category from the ShapeNet dataset [37], which is the subset of 3D models used in NOCS CAMERA’s training dataset. As described in III-E, cameras are sampled using the icosahedron sampling resulting in 162 images per instance. The all images have a size of 160 x 160 pixels.

Metrics. Following original NOCS [5] protocol, we use the mean Average Precision (mAP) in n° and $m\text{ cm}$ to evaluate the accuracy of estimated poses. For axial-symmetric objects, the rotation error around the axis of symmetry is ignored. In the generalization benchmark, we additionally adjust the metric for the “can” category, the only object in the dataset that exhibits both axial and reflective symmetry, to be invariant to the direction of the axis of symmetry since our benchmark contains objects in arbitrary poses, whereas NOCS Real 275 contains only objects standing upright.

$$\theta_{error} = \arccos\left(\frac{|Y_{GT} \cdot Y_{pred}|}{\|Y_{GT}\|_2 \|Y_{pred}\|_2}\right), \quad Y_{GT}, Y_{pred} \in \mathbb{R}^3 \quad (3)$$

Y_{GT} and Y_{pred} represent the unit vector of the Y-axis which is aligned with the symmetry axis.

Implementation Details. At training time, we randomly drop the input representation at a 25% rate via the procedure described by Eq. 2, and the time steps of DDPM are set to 1000. We apply PCA to DINOv2 features, reducing their dimension from 382 to 6 for both training and inference. At inference time, we use 6 different noises to generate 6 possible NOCS maps. Then, a final pose is selected based on the highest confidence. Our choices of PCA dimension and the number of noises are informed by the ablation shown in Table III. We use 10-time steps of a DPM-solver [29], which takes 0.855s on the Tesla V100 GPU (corresponding to the 5th row in Table III). As noted in Table III, other faster configurations may be chosen as a trade-off with accuracy.

B. Comparisons to State of the Art

NOCS Real 275 Benchmark. Table I reports a comparison between our method and SOTA approaches. Our method

shows significantly better results when compared to all other methods trained on synthetic data, achieving mAPs of 35.0 ($5^\circ 5\text{cm}$), 66.7 ($10^\circ 5\text{cm}$), and 77.1 ($15^\circ 5\text{cm}$), whereas the former best methods score 16.9 ($5^\circ 5\text{cm}$), 44.9 ($10^\circ 5\text{cm}$), and 52.8 ($15^\circ 5\text{cm}$) respectively. Even though only synthetic data is used for our training, our method also demonstrates strong results when compared to the methods trained on a combination of both real and synthetic datasets outperforming NOCS [5], ShapePrior [6], and CenterSnap [8].

We also show the importance of multi-modal input representations. The combination of surface normals and DINO features shows over 10 percent improvement compared with only normals on both $10^\circ 5\text{cm}$ and $15^\circ 5\text{cm}$ metrics. Also, Table I shows the DINO features are more effective than RGB. For clarity, a single network was used to generate all score sets (N, RGB-N, DINO-N, RGB-DINO-N) without re-training since our method supports selectable inputs.

Generalization Benchmark. To analyze zero-shot generalization capabilities of existing SOTA baselines, we computed their mAP scores on our newly introduced generalization benchmark as shown in Table II. For all baselines, we used respective official implementations and provided pre-trained weights. Similarly to the NOCS benchmark, we significantly outperform the SOTA synthetic baseline ShapePrior [6] scoring average mAPs of 29.2 ($5^\circ 5\text{cm}$), 51.7 ($10^\circ 5\text{cm}$), and 57.1 ($15^\circ 5\text{cm}$) against ShapePrior’s 14.4 ($5^\circ 5\text{cm}$), 31.2 ($10^\circ 5\text{cm}$), and 38.4 ($15^\circ 5\text{cm}$) in rightmost block in Table II. Moreover, we outperform all real methods when averaging over all three datasets, demonstrating our robustness to occlusions, changing lighting conditions, and a wide variety of object poses. Table II also shows real methods dominate the YCB-V scores. We explain this by this dataset’s similarity with NOCS: it only contains objects standing upright placed in well-lit scenes placed close to the camera. Overall, as the mean and standard deviation scores demonstrate, our method achieves the best overall performance and stability even when compared to the methods trained on synthetic and real data.

C. Analysis on Symmetric Objects

In this ablation, we qualitatively demonstrate the ability of our method to handle symmetrical objects. As shown in Fig. 8, we visualize 20 poses on the RGB image and 20 rotations on the unit sphere that are estimated from 20 different noise images. For the bowl object, the rotations are widely distributed in one-axis since it’s an axial-symmetric object. On the other hand, the distribution of the mug with a visible handle has well defined peaks, confirming that the handle resolves the pose ambiguity and makes it possible to establish a unique pose. For the mug with an invisible handle, the distribution is highly correlated with the possible poses of the mug correctly identifying a possible direction of a handle (x-axis). These results qualitatively demonstrate the capability of our method to handle arbitrary types of symmetric objects without the need for elaborate dataset symmetry labeling and heuristic operations at training time.

TABLE I

EVALUATION RESULTS ON NOCS REAL275. DATA MEANS THE TYPE OF TRAINING DATA. S DENOTES SYNTHETIC SHAPENET OBJECTS ONLY, WHILE S(B) DENOTES SHAPENET MODELS RENDERED WITH REAL BACKGROUNDS. R DENOTES REAL IMAGES IN NOCS REAL275. N AND DINO REPRESENT NORMAL AND DINOv2 FEATURES RESPECTIVELY.

	Input	Data	$5^\circ 5\text{cm} \uparrow$	$10^\circ 5\text{cm} \uparrow$	$15^\circ 5\text{cm} \uparrow$
NOCS [5]	RGB-D	S(B)+R	9.8	24.1	34.9
ShapePrior [6]	RGB-D	S(B)+R	21.4	43.6	66.7
CenterSnap [8]	RGB-D	S(B)+R	27.2	58.8	-
DualPoseNet [7]	RGB-D	S+R	35.9	66.8	76.3
ShaPO [38]	RGB-D	S(B)+R	48.8	66.8	-
GenPose [22]	D	S(B)+R	50.6	84.0	89.7
Chen et al. [39]	RGB-D	S	0.7	3.6	9.1
Gao et al. [40]	D	S	7.8	17.1	26.5
CPPF [9]	D	S	16.9	44.9	50.8
ShapePrior [6]	RGB-D	S	12.0	37.9	52.8
Ours (N)	D	S	28.7	56.3	65.5
Ours (RGB-N)	RGB-D	S	33.5	62.8	72.2
Ours (DINO-N)	RGB-D	S	34.5	66.7	77.0
Ours (RGB-DINO-N)	RGB-D	S	35.0	66.6	77.1

TABLE II

EVALUATION RESULTS ON GENERALIZATION BENCHMARK. THE RESULTS SHOW THE ZERO-SHOT POSE ESTIMATION PERFORMANCE ON TYO-L, YCB-V, AND HOPE REAL-WORLD DATASETS. OUR METHOD OUTPERFORMS ALL METHODS USING **SYNTHETIC DATA** AND **REAL DATA** ON AVERAGE.

	Data	TYO-L			YCB-V			HOPE			Mean \uparrow \pm Std \downarrow		
		5°5cm \uparrow	10°5cm \uparrow	15°5cm \uparrow	5°5cm \uparrow	10°5cm \uparrow	15°5cm \uparrow	5°5cm \uparrow	10°5cm \uparrow	15°5cm \uparrow	5°5cm	10°5cm	15°5cm
ShapePrior [6]	S+R	14.9	27.5	32.1	27.4	54.3	66.9	14.1	41.0	57.0	18.8 \pm 6.1	40.9\pm10.9	52.0\pm14.6
DualPoseNet [7]	S+R	22.5	28.8	32.2	36.6	78.8	87.5	3.9	9.6	14.1	21.0 \pm 13.4	39.1 \pm 29.2	44.6 \pm 31.2
GenPose [22]	S+R	20.1	23.7	24.7	48.1	68.8	73.8	18.8	23.2	25.2	29.0\pm13.5	38.6 \pm 21.4	41.2 \pm 23.0
CPPF [9]	S	13.2	18.8	20.8	8.8	35.2	46.9	3.2	15.3	20.5	8.4 \pm 4.1	23.1 \pm 8.7	29.4 \pm 12.4
ShapePrior [6]	S	17.3	26.5	33.4	11.0	40.3	49.1	15.0	26.8	32.7	14.4 \pm 2.6	31.2 \pm 6.4	38.4 \pm 7.6
Ours (N)	S	33.8	43.9	49.4	20.9	46.2	53.0	17.9	37.3	45.7	24.2 \pm 6.9	42.5 \pm 3.8	49.4 \pm 3.0
Ours (RGB-N)	S	40.3	51.0	56.6	20.0	47.3	51.0	19.1	40.0	46.8	26.5 \pm 9.8	46.1 \pm 4.6	51.5 \pm 4.0
Ours (DINO-N)	S	43.5	60.2	66.0	23.4	54.8	58.2	20.1	40.2	47.1	29.0 \pm 10.3	51.7\pm8.4	57.1\pm7.8
Ours (RGB-DINO-N)	S	43.7	59.0	64.7	23.6	53.7	56.6	20.4	38.5	46.5	29.2\pm10.3	50.4 \pm 8.7	55.9 \pm 7.4

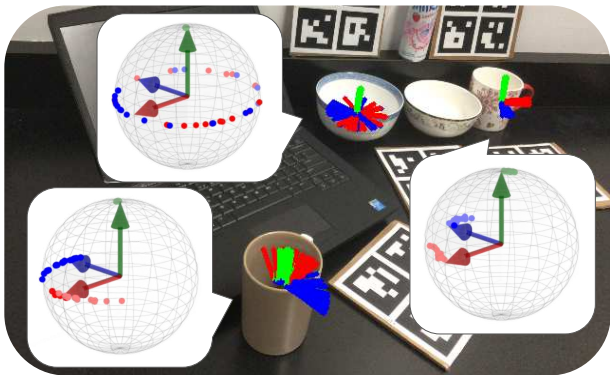


Fig. 8. **Results of Multiple Estimated Poses:** A visualization of 20 estimated poses from 20 different noises. The corresponding rotations are plotted on a unit sphere to show the distribution shape of symmetry objects.

D. Analysis on Multimodal Input Representations

It is common knowledge that structured light depth sensors often struggle to output consistent depth when it comes to reflective surfaces. For example, the depth information of a black laptop in the NOCS Real 275 is partially lacking as shown in Fig. 9. Estimating poses of such sparse point clouds is difficult or sometimes impossible, especially for methods not using RGB. This is also a problem for other tasks, such as detecting object-grasping locations from the point clouds. On the other hand, our method can robustly recover the missing information thanks to its multi-modality and using both depth and RGB inputs as shown in Fig. 9. It’s useful for various applications, such as not only point registration but also grasp detectors or collision avoidance.

E. Ablation Study About the Parameters of PCA and Noise

To better understand the effect of our method’s parameters on the final score, we conduct an ablation study and report the results based on RGB-DINO-N inputs in Table III. In particular, we show how the number of (1) PCA dimensions for DINOv2 outputs and (2) noise samples affect the pose quality (mAP) and inference time (s). Keeping the number of noise samples at 1, we observe that the performance of PCA peaks at 6 dimensions showing the best scores while maintaining minimal inference overhead. Changing the noise number from 1 to 6 with 6-dimensional PCA boosts the performance significantly, resulting in more than a 10 percent

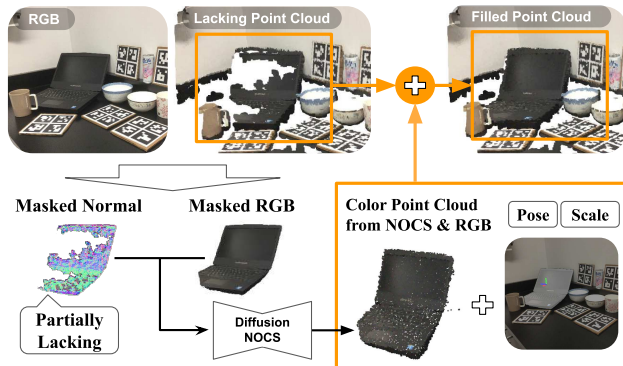


Fig. 9. **Point Clouds Completion by DiffusionNOCS:** DiffusionNOCS with RGB-D information can estimate the NOCS map correctly even with a lack of depth information since RGB image can cover the information. The lacking point cloud in target objects can be reconstructed by an estimated NOCS map, pose, and scale. It’s useful for various robotics applications.

TABLE III
ABLATION RESULTS ON NOCS REAL 275 BENCHMARK WITH RGB-DINO-N INPUTS.

Noise	PCA	5°5cm \uparrow	10°5cm \uparrow	15°5cm \uparrow	Time (s) \downarrow
1	3	27.8	48.7	56.5	0.372
1	6	30.0	54.1	63.4	0.386
1	32	27.9	51.0	58.9	0.392
3	6	34.5	65.0	75.2	0.522
6	6	35.0	66.6	77.1	0.855
9	6	35.3	66.6	77.1	1.147

improvement on 10°5cm and 15°5cm metric. The 6 noises are chosen as a default for our pipeline from both mAP and inference time perspectives. Thus, the 5th row in Table III corresponds the last row in Table I.

V. CONCLUSION

We propose a novel category-level dense pose estimation pipeline via diffusion models with multi-modal selectable inputs. Thanks to its probabilistic nature, our approach is able to handle symmetrical objects without a need for special data annotations and heuristics at the training time typical for many SOTA methods. Moreover, we demonstrate that the use of multiple pose hypotheses helps to boost the performance at inference time. Outputting dense correspondences from multi-modal inputs allow us to resolve ambiguities better

and faithfully recover partial object geometry even in the presence of lacking depth. Finally, the evaluation results demonstrate that our method significantly outperforms all SOTA baselines trained on synthetic data on the typical NOCS benchmark, and outperforms all SOTA baselines on average including the ones trained on synthetic and real data on our novel zero-shot generalization benchmark. As future work, we will focus on further improving the inference speed and enhancing the variety and quality of the synthetic dataset.

VI. ACKNOWLEDGEMENT

We would like to thank Prof. Greg Shakhnarovich, Thomas Stewart, and Dr. Vitor Guizilini for their valuable feedback.

REFERENCES

- [1] T. Hodan, D. Barath, and J. Matas, "Epos: estimating 6d pose of objects with symmetries," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2020, pp. 11 703–11 712.
- [2] S. Zakharov, I. Shugurov, and S. Ilic, "Dpod: Dense 6d pose object detector in rgb images," in *Int. Conf. on Computer Vision*, 2019.
- [3] Y. Su, M. Saleh, T. Fetzter, J. Rambach, N. Navab, B. Busam, D. Stricker, and F. Tombari, "ZebraPose: Coarse to fine surface encoding for 6dof object pose estimation," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2022, pp. 6738–6748.
- [4] T. Hodan and A. Melenovsky. (2019) Bop: Benchmark for 6d object pose estimation: <https://bop.felk.cvut.cz/home/>.
- [5] H. Wang, S. Sridhar, J. Huang, J. P. C. Valentin, S. Song, and L. J. Guibas, "Normalized object coordinate space for category-level 6d object pose and size estimation," *Conf. on Computer Vision and Pattern Recognition*, pp. 2637–2646, 2019.
- [6] M. Tian, M. H. Ang, and G. H. Lee, "Shape prior deformation for categorical 6d object pose and size estimation," in *European Conf. on Computer Vision*. Springer, 2020, pp. 530–546.
- [7] J. Lin, Z. Wei, Z. Li, S. Xu, K. Jia, and Y. Li, "Dualposenet: Category-level 6d object pose and size estimation using dual pose network with refined learning of pose consistency," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2021, pp. 3560–3569.
- [8] M. Z. Irshad, T. Kollar, M. Laskey, K. Stone, and Z. Kira, "Centersnap: Single-shot multi-object 3d shape reconstruction and categorical 6d pose and size estimation," in *Int. Conf. on Robotics and Automation*. IEEE, 2022, pp. 10 632–10 640.
- [9] Y. You, R. Shi, W. Wang, and C. Lu, "Cpof: Towards robust category-level 9d pose estimation in the wild," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2022, pp. 6856–6865.
- [10] W. Goodwin, I. Havoutis, and I. Posner, "You only look at one: Category-level object representations for pose estimation from a single example," *Conf. on Robot Learning*, 2023.
- [11] A. Kendall, M. Grimes, and R. Cipolla, "Posenet: A convolutional network for real-time 6-dof camera relocalization," in *Int. Conf. on Computer Vision*. IEEE, 2015, pp. 2938–2946.
- [12] F. Engelmann, K. Rematas, B. Leibe, and V. Ferrari, "From Points to Multi-Object 3D Reconstruction," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2021.
- [13] S. Hinterstoisser, S. Holzer, C. Cagniart, S. Ilic, K. Konolige, N. Navab, and V. Lepetit, "Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes," in *Int. Conf. on Computer Vision*. IEEE, 2011, pp. 858–865.
- [14] P. Wohlhart and V. Lepetit, "Learning descriptors for object recognition and 3d pose estimation," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2015.
- [15] S. Zakharov, W. Kehl, B. Planche, A. Hutter, and S. Ilic, "3d object instance recognition & pose estimation using triplet loss with dynamic margin," in *Int. Conf. on Intelligent Robots and Systems*. IEEE/RSJ, 2017, pp. 552–559.
- [16] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Int. Conf. on Machine Learning*. PMLR, 2017, pp. 1321–1330.
- [17] A. Kendall and R. Cipolla, "Geometric loss functions for camera pose regression with deep learning," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2017, pp. 5974–5983.
- [18] C. Bishop, "Mixture density networks," Aston University, Working-Paper, 1994.
- [19] I. Gilitschenski, R. Sahoo, W. Swarting, A. Amini, S. Karaman, and D. Rus, "Deep orientation uncertainty learning based on a bingham loss," in *Int. Conf. on Learning Representations*, 2019.
- [20] H. Deng, M. Bui, N. Navab, L. Guibas, S. Ilic, and T. Birdal, "Deep bingham networks: Dealing with uncertainty and ambiguity in pose estimation," *Int. Journal of Computer Vision*, vol. 130, no. 7, pp. 1627–1654, 2022.
- [21] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Robotics: Science and Systems*, 2023.
- [22] J. Zhang, M. Wu, and H. Dong, "GenPose: Generative category-level object pose estimation via diffusion models," in *Conf. on Neural Information Processing Systems*, 2023.
- [23] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al., "Segment anything," *Int. Conf. on Computer Vision*, pp. 3992–4003, 2023.
- [24] H. Yang, J. Shi, and L. Carlone, "TEASER: Fast and Certifiable Point Cloud Registration," *IEEE Trans. Robotics*, 2020.
- [25] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [26] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th Int. Conf., Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Conf. on Computer Vision and Pattern Recognition*. IEEE/CVF, 2016, pp. 770–778.
- [28] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [29] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, "Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5775–5787, 2022.
- [30] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, et al., "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.
- [31] B. T. Phong, "Illumination for computer generated pictures," in *Seminal graphics: pioneering efforts that shaped the field*, 1998, pp. 95–101.
- [32] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," *arXiv preprint arXiv:1708.04552*, 2017.
- [33] T. Hodan, F. Michel, E. Brachmann, W. Kehl, A. GlentBuch, D. Kraft, B. Drost, J. Vidal, S. Ihrke, X. Zabulis, et al., "Bop: Benchmark for 6d object pose estimation," in *European Conf. on Computer Vision*, 2018, pp. 19–34.
- [34] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," in *Robotics: Science and Systems*, 2018.
- [35] Y. Lin, J. Tremblay, S. Tyree, P. A. Vela, and S. Birchfield, "Multi-view fusion for multi-level robotic scene understanding," in *Int. Conf. on Intelligent Robots and Systems*. IEEE/RSJ, 2021, pp. 6817–6824.
- [36] M. Matl, "Pyrender," <https://github.com/mmatl/pyrender>, 2019.
- [37] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Sava, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.
- [38] M. Z. Irshad, S. Zakharov, R. Ambrus, T. Kollar, Z. Kira, and A. Gaidon, "Shapo: Implicit representations for multi-object shape appearance and pose optimization," in *European Conf. on Computer Vision*, 2022.
- [39] X. Chen, Z. Dong, J. Song, A. Geiger, and O. Hilliges, "Category level object pose estimation via neural analysis-by-synthesis," in *European Conf. on Computer Vision*, 2020.
- [40] G. Gao, M. Lauri, Y. Wang, X. Hu, J. Zhang, and S. Frintrop, "6d object pose regression via supervised learning on point clouds," in *Int. Conf. on Robotics and Automation*. IEEE, 2020, pp. 3643–3649.