

# CleanPose: Category-Level Object Pose Estimation via Causal Learning and Knowledge Distillation

Xiao Lin<sup>1</sup> Yun Peng<sup>1</sup> Liuyi Wang<sup>1</sup> Xianyou Zhong<sup>1</sup> Minghao Zhu<sup>1</sup>

Yi Feng<sup>1</sup> Jingwei Yang<sup>1</sup> Chengju Liu<sup>1,2\*</sup> Qijun Chen<sup>1,2\*</sup>

<sup>1</sup>College of Electronic and Information Engineering, Tongji University, Shanghai, China

<sup>2</sup>State Key Laboratory of Autonomous Intelligent Unmanned Systems

{linx\_xx, pengyun, wly, zhongxianyou, zmhh\_h, jw\_yang,  
fengyi0109, liuchengju, qjchen}@tongji.edu.cn

## Abstract

In the effort to achieve robust and generalizable category-level object pose estimation, recent methods primarily focus on learning fundamental representations from data. However, the inherent biases within the data are often overlooked: the repeated training samples and similar environments may mislead the models to over-rely on specific patterns, hindering models' performance on novel instances. In this paper, we present **CleanPose**, a novel method that mitigates the data biases to enhance category-level pose estimation by integrating causal learning and knowledge distillation. By incorporating key causal variables (structural information and hidden confounders) into causal modeling, we propose the causal inference module based on front-door adjustment, which promotes unbiased estimation by reducing potential spurious correlations. Additionally, to further confront the data bias at the feature level, we devise a residual-based knowledge distillation approach to transfer unbiased semantic knowledge from 3D foundation model, providing comprehensive causal supervision. Extensive experiments across multiple benchmarks (REAL275, CAMERA25 and HouseCat6D) highlight the superiority of proposed CleanPose over state-of-the-art methods. Code will be available at <https://github.com/chrislin0621/CleanPose>.

## 1. Introduction

Category-level object pose estimation (COPE) [42] aims to predict the 9Dof pose for arbitrary objects within predefined categories. This task, unlike instance-level pose estimation [8, 11, 21, 41, 58], does not require high-quality CAD models, making it feasible to perceive a broader range

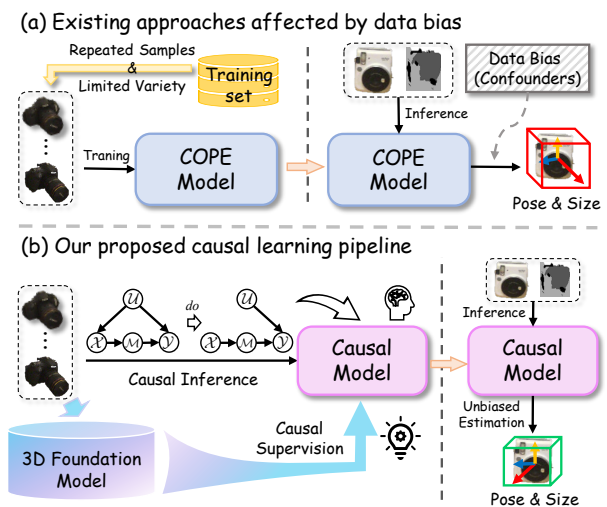


Figure 1. Comparison of (a) existing pose estimation approaches and (b) the proposed causal learning pipeline. The causal inference is employed to mitigate the negative effect raised by bias, and 3D semantic knowledge distilled from foundation model is leveraged to provide comprehensive causal guidance.

of novel objects instead of a single instance. The complex intra-category variations have long been regarded as main challenge of COPE, thus recent methods [1, 3, 4, 17, 19, 20, 22, 23, 56] mainly focus on model designing to capture robust category features from data. Despite remarkable development, the performance of COPE models remains limited [20, 23], suggesting that feature learning may not be the fundamental issue. This prompts us to identify the key hindrance blocking further progress. We observe that inherent biases are pervasive in current datasets [42], e.g., repeated training samples and limited pose variety, which may mislead COPE models to overfit to familiar object's appearance and poses during data fitting, as shown in Fig. 1(a). However, existing efforts often overlook these underlying data

\*Corresponding authors.

biases and the resulting negative impacts.

One way to tackle bias issue is to create broader and more diverse datasets [6, 15, 44, 54]. Though valuable, achieving a perfectly balanced dataset free of bias remains nearly impossible. Additionally, the dataset’s scale is still significantly constrained by the cost of 3D data annotation [55]. Therefore, the extension of datasets does not fundamentally solve the hindrance, developing a causal COPE models that can effectively confront and alleviate biases becomes a primary challenge.

To address this issue, we direct our attention to the human observation mechanism. In fact, the reason why humans can effectively handle variations among intra-category objects is that we can leverage the analogical association to infer the structural features and underlying causal relations. Moreover, the observations of numerous objects enable humans to maintain a stable perception of novel objects across different environments and viewpoints.

Motivated by this finding, for the first time, we propose to incorporate the causal inference [32] into the formulation of COPE, exploiting the concepts of *intervention* [33] to mitigate the negative effects caused by confounders. Here, confounders are variables that influence both input and outcomes, just as dataset biases in this task simultaneously affect input sampling and output pose. The causal learning can investigate the causal relation among variables and equip COPE models with similar cognitive abilities that humans have. However, it is non-trivial to directly incorporate the causal inference in COPE applications because of the following challenges: (i) The unique modality of 3D data makes it impractical to directly apply existing causal modeling techniques [14]. (ii) Subsequently, the confounders in pose estimation task are inherently unobserved and elusive, which further increases the challenge of identifying these confounders. (iii) Moreover, limited training samples raise the difficulties in learning comprehensive causal relations.

To overcome above challenges, we present **CleanPose**, a concise yet effective framework with **causal learning** and **knowledge distillation** to enhance the category-level **pose** estimation, as shown in Fig. 1(b). Faced with the challenge that hidden confounders are elusive or even unobserved in our task, we have two key designs: (i) The first design is a causal inference module based on *front-door adjustment*, which can effectively approximate the predicted expectations by potential confounders. Moreover, to suitably represent confounders, we devise and maintain a dynamic queue to efficiently update training samples, similar with MoCo [10]. To further mitigate biases, (ii) our second design is a residual-based feature knowledge distillation module, which can transfer unbiased point cloud category information of the 3D foundation model, ULIP-2 [47], into our model. Since the foundation model is pretrained on a diverse range of objects, it captures robust semantic fea-

tures that can implicitly provide comprehensive causal supervision at the feature level. As demonstrated by extensive experiments, our findings reveal the impact of integrating causal learning to reduce biases, and providing comprehensive causal supervision, enhancing the model’s robustness and generalization.

To summarize, our main contributions are as follows:

- We propose CleanPose, the first solution to mitigate the confounding effect in category-level pose estimation via causal learning. Taking inspiration from human observation mechanism, we propose to identify the causal effect to achieve unbiased estimation, recovering correct pose of novel instances within predefined categories.
- To provide comprehensive causal supervision and enhance intra-category generalization, we develop a residual category knowledge distillation approach to transfer rich and unbiased 3D semantic knowledge from 3D foundation models into COPE networks.
- Our proposed CleanPose achieves state-of-the-art performance on three mainstream challenge benchmarks (REAL275 [42], CAMERA25 [42] and House-Cat6D [15]). For instance, the accuracy attains **61.7%** on rigorous metric  $5^{\circ}2\text{ cm}$  of REAL275 dataset, surpassing the current best method with a large margin by 4.7%.

## 2. Related Works

**Category-level Object Pose Estimation.** The objective of this task encompasses predicting the 9DoF pose for unseen objects within predefined categories. To address this challenging task, pioneer method NOCS [42] suggests mapping input shape to a normalized canonical space and recovering the pose via Umeyama algorithm [39]. SPD [38] proposes a method for deriving and utilizing the shape prior for each category. This crucial insight inspires many subsequent prior-based works [5, 17, 18], which further improve the use of shape priors, continuously improving the pose estimation performance. More recently, prior-free methods [3, 4, 20, 22, 56] have achieved impressive performance. VI-Net [20] separates rotation into viewpoint and in-plane rotations, while SecondPose [3] propose to extract hierarchical panel-based geometric features for point cloud. AG-Pose [22] achieves current state-of-the-art performance by explicitly extract local and global geometric keypoint information of different instances. However, above methods have not fully considered the dataset biases behind the task, which may leads the models to learn spurious correlations, damaging the generalization ability to unseen instances within the categories.

**Causal Inference.** Causality [32] is an emerging technique referring to the modeling the relationships between factors in a task from a human perspective. Recent methods have incorporated causal inference to improve the performance of DNNs in computer vision domains, *e.g.*, object detec-

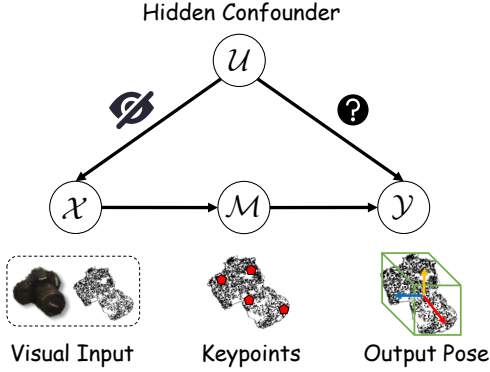


Figure 2. Illustration of the structural causal model of COPE.

tion [45, 51], image captioning [24, 48] and vision-language task [43, 49, 53]. Though valuable, these methods can flexibly identify confounders due to natural prompts in data, e.g., the keywords in instructions [43] and salient regions in an image [45], which are inapplicable for point cloud data.

In 3D domain, several works have focused on enhancing the robustness of point cloud classification [14] or 3D reconstruction [27] with causal inference. However, it is non-trivial to adapt these approaches to pose estimation due to the inherent differences in human modeling among these tasks. In this work, we propose a specific causal learning approach based on front-door adjustment for category-level pose estimation for the first time.

**Knowledge Distillation.** Knowledge distillation [12] is a technique that transfers knowledge from a teacher model to a student model. Transferring the knowledge of foundation models for downstream tasks has been proven to be effective [7, 13, 60]. ViLD [7] distill knowledge from CLIP [36] to achieve open-vocabulary object detection, while some methods [13, 60] utilize the generalizability of VLMs to address the video recognition task. In this work, we leverage the 3D foundation model ULIP-2 [47] as a teacher model to guide the student model in correcting biases at the feature level, enhancing the robustness and generalization capabilities of COPE model from multiple perspectives.

### 3. Preliminaries

#### 3.1. Task Formulation

Given an RGB-D image containing objects from a predefined set of categories, off-the-shelf segmentation models such as MaskRCNN [9] are employed to obtain masks and category labels for each object in a image. Then, the segmentation masks can be utilized to get the cropped RGB image  $\mathcal{I}_{obj} \in \mathbb{R}^{H \times W \times 3}$  and the point cloud  $\mathcal{P}_{obj} \in \mathbb{R}^{N \times 3}$ , where  $N$  is the number of points and  $\mathcal{P}_{obj}$  is acquired by back-projecting the cropped depth image with camera intrinsics followed by a downsampling process. With the input  $\mathcal{I}_{obj}$  and  $\mathcal{P}_{obj}$ , the objective of COPE [42] is to recover

the 9DoF poses, including the 3D rotation  $\mathcal{R} \in SO(3)$ , the 3D translation  $t \in \mathbb{R}^3$ , and 3D metric size  $s \in \mathbb{R}^3$ .

#### 3.2. The Causal Modeling of CleanPose

To quantify the underlying logic behind human observation, we construct a structural causal model [32, 33] capturing the relationship among the key variables in COPE: visual input  $\mathcal{X}$ , output pose  $\mathcal{Y}$ , mediator  $\mathcal{M}$  and hidden confounders  $\mathcal{U}$ . We illustrate the causal model in Fig. 2, where each direct link denotes a causal relationship between two nodes.

- $\mathcal{X} \rightarrow \mathcal{M} \rightarrow \mathcal{Y}$  (*Front-door path*): Typically, humans first recognize the structural information of an object, i.e., keypoints [22, 56], and then determine the object’s pose based on the similar poses of other objects within this category. This process involves identifying keypoints of the object and their relative positions, and leveraging this information to perform pose estimation. We use mediator  $\mathcal{M}$  to represent such structural information and describe such process via the causal path  $\mathcal{X} \rightarrow \mathcal{M} \rightarrow \mathcal{Y}$ , which is also referred to as the *front-door path*.
- $\mathcal{X} \leftarrow \mathcal{U} \rightarrow \mathcal{Y}$  (*Hidden confounders*): The confounders are extraneous variables that influence both inputs and outputs, e.g., dataset biases [15, 54] or category specific attributes [38].  $\mathcal{U} \rightarrow \mathcal{X}$  exists because the input data is inevitably affected by the limited resources in real world and sampling noises from sensors when collection and simulation. Moreover,  $\mathcal{U} \rightarrow \mathcal{Y}$  emerges because collected scenes, annotation bias, or the variety of pose also affect the probability of pose distributions. Existing COPE methods tend to model the statistical correlations  $P(\mathcal{Y}|\mathcal{X})$ , given the optimization goal of maximizing the pose accuracy [57]. With the lack of modeling of hidden confounders, no matter how large the amount of training data is, the model can not identify the true causal effect from  $\mathcal{X}$  to  $\mathcal{Y}$ .

## 4. Methodology

### 4.1. Feature Extractor

Following [22], we utilize the PointNet++ [35] to extract point feature  $\mathcal{F}_P \in \mathbb{R}^{N \times C_1}$  of input point cloud  $\mathcal{P}_{obj}$ . As for the RGB image  $\mathcal{I}_{obj}$ , we adopt DINOv2 (ViT-S/14) [31] as our image feature extractor, which has been proven to extract abundant semantic-aware information from RGB images [34]. We select those pixel features corresponding to  $\mathcal{P}_{obj}$  and utilize linear interpolation to propagate the original DINOv2 features into the final RGB features  $\mathcal{F}_I \in \mathbb{R}^{N \times C_2}$ . Eventually, we concatenate  $\mathcal{F}_P$  and  $\mathcal{F}_I$  to form  $\mathcal{F}_{obj} \in \mathbb{R}^{N \times C}$  as the input for the subsequent networks.

### 4.2. Causal Inference of CleanPose

**The Adjustment Formulation.** To avoid making the pose estimation process overfit to specific hidden confounders,

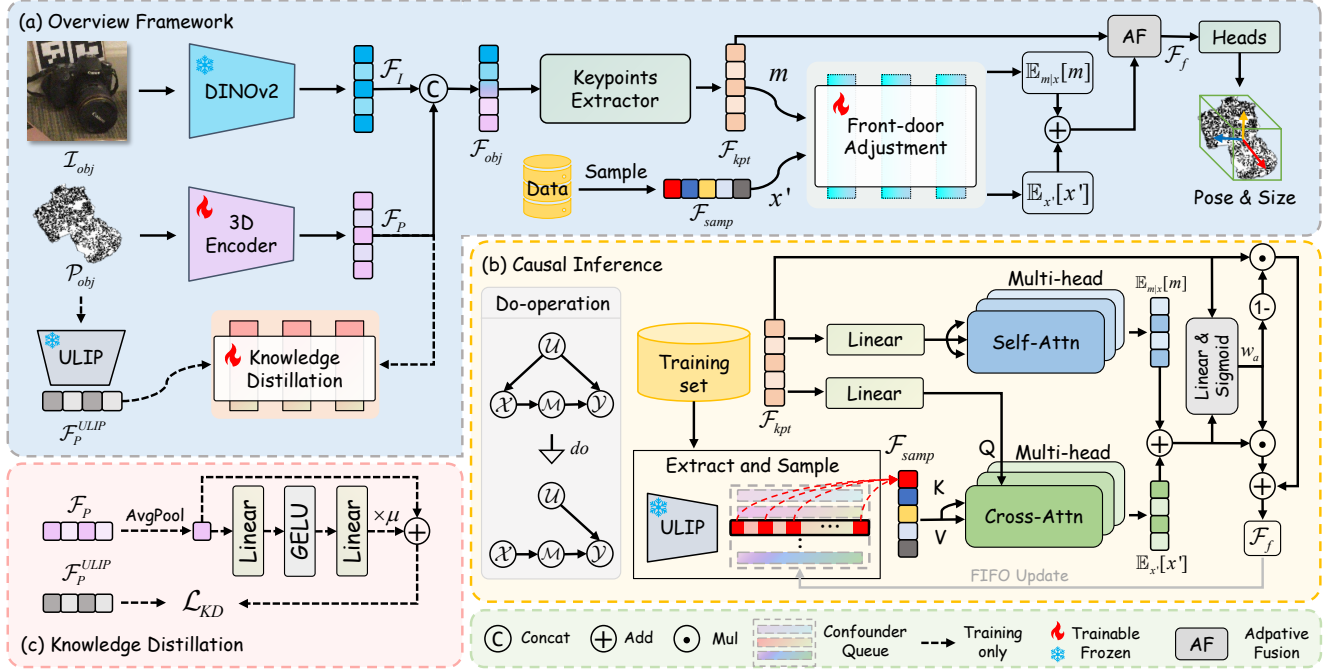


Figure 3. (a) Framework of CleanPose. (b) Causal inference based on front-door adjustment is employed for mitigating potential spurious correlations, promoting unbiased pose estimation. Moreover, the proposed (c) residual-based knowledge distillation module can efficiently provide comprehensive category-specific guidance.

we propose to perform causal modeling according to the causal graph and the concepts of *intervention* [32]. Specifically, the process of *intervention* can be reflected as *do-operation* [33], which provides scientifically sound methods for determining causal effects, as shown in Fig. 3(b). The specific casual modeling via *do-operation* is as follow. Based on Bayes’s theorem, the typical observational likelihood is as:

$$P(\mathcal{Y}|\mathcal{X}) = \sum_u P(\mathcal{Y}|\mathcal{X}, u)P(u|\mathcal{X}), \quad (1)$$

where  $P(u|\mathcal{X})$  would bring biased weights. Then we perform the *do-operation* on inputs  $\mathcal{X}$ , which amounts to removing all edges directed into that variable in the graph model, *i.e.*, from hidden confounders  $\mathcal{U}$ . According to the invariance and independence rules [33], we have:

$$P(\mathcal{Y}|do(\mathcal{X})) = \sum_u P(\mathcal{Y}|do(\mathcal{X}), u)P(u|do(\mathcal{X})) \quad (2)$$

$$= \sum_u P(\mathcal{Y}|\mathcal{X}, u)P(u). \quad (3)$$

In this way, the intervention is realized by blocking the causal path  $\mathcal{U} \rightarrow \mathcal{X}$ . The Eq. (3) is also known as the *adjustment formulation* [32].

**Front-door Adjustment Causal Learning.** In the previous section, we have clarified meaning of the *front-door path*  $\mathcal{X} \rightarrow \mathcal{M} \rightarrow \mathcal{Y}$ . In COPE task, the DNNs-based models  $P(\mathcal{Y}|\mathcal{X}) = \sum_m P(\mathcal{Y}|m)P(m|\mathcal{X})$  will choose the suitable knowledge  $\mathcal{M}$  from input  $\mathcal{X}$  for pose estimation results  $\mathcal{Y}$ .

Due to the presence of the mediators  $\mathcal{M}$ , applying adjustment Eq. (3) solely to the inputs  $\mathcal{X}$  is insufficient to fully block the confounding effects [43]. Therefore, we implement continuous intervention in both stages of front-door path [33], *i.e.*, by simultaneously applying *do-operation* to  $\mathcal{X}$  and  $\mathcal{M}$ . First, the effect of  $\mathcal{X}$  on  $\mathcal{M}$  is identifiable:

$$P(\mathcal{M}|do(\mathcal{X})) = P(\mathcal{M}|\mathcal{X}). \quad (4)$$

Then, for the second stage  $\mathcal{M} \rightarrow \mathcal{Y}$ , we apply the adjustment formulation and have:

$$P(\mathcal{Y}|do(\mathcal{M})) = \sum_{x'} P(\mathcal{Y}|\mathcal{M}, x')P(x'). \quad (5)$$

Here,  $x'$  denotes potential inputs of the whole representation space, different from current inputs  $\mathcal{X} = x$ . To chain together these two partial effects to obtain the overall effect of  $\mathcal{X}$  on  $\mathcal{Y}$ , we sum over all states  $m$  of  $\mathcal{M}$  to form overall front-door adjustment as follow:

$$P(\mathcal{Y}|do(\mathcal{X})) = \sum_m P(\mathcal{Y}|do(m))P(m|do(\mathcal{X})) \quad (6)$$

$$= \sum_{x'} P(x') \sum_m P(\mathcal{Y}|m, x')P(m|\mathcal{X}). \quad (7)$$

From the above derivation, we can observe that the front-door adjustment in causal inference involves the intermediate knowledge of current input  $\mathcal{X} = x$  (*i.e.*, the current features) as well as the cross-sampling features of other samples from the entire training set. We utilize the bold symbol  $m$  and  $x'$  to represent these two components respectively.

Here, the intermediate knowledge is extracted from current input  $\mathbf{m} \sim x$ . Detailed methods for sampling  $\mathbf{x}'$  are introduced in subsequent sections.

For convenience, we define the network module as a linear function  $f(x, \mathbf{x}')$  to model the computation of  $x$  and  $\mathbf{x}'$  to obtain  $\mathcal{Y}$ . Meanwhile, according to the Bayes' rule and the definition of Expected values, the front-door adjustment  $\mathcal{FD}(\cdot)$  of Eq. (7) can be expressed as:

$$P(\mathcal{Y}|do(\mathcal{X})) = \mathcal{FD}(x, \mathbf{x}') \quad (8)$$

$$= \mathbb{E}_{\mathbf{x}'} \mathbb{E}_{\mathbf{m}|x} [f(x, \mathbf{x}')], \quad (9)$$

where the linear function is used as  $f(x, \mathbf{x}') = f_x(x) + f_{\mathbf{x}'}(\mathbf{x}')$ . Hence, based on the linear mapping model, the Eq. (9) can be formulated as  $\mathbb{E}_{\mathbf{x}'}[\mathbf{x}'] + \mathbb{E}_{\mathbf{m}|x}[\mathbf{m}]$ . Since the closed-form solution of expected values are difficult to obtain in the complex representation space [43], we employ the query mechanism to achieve the estimation:

$$\mathbb{E}_{\mathbf{x}'}[\mathbf{x}'] \approx \sum_{\mathbf{x}'} P(\mathbf{x}'|\mathbf{g}_1)\mathbf{x}' = \sum_i \frac{\exp(\mathbf{g}_1 \mathbf{x}'_i^T)}{\sum_j \exp(\mathbf{g}_1 \mathbf{x}'_j^T)} \mathbf{x}'_i \quad (10)$$

$$\mathbb{E}_{\mathbf{m}|x}[\mathbf{m}] \approx \sum_{\mathbf{m}} P(\mathbf{m}|\mathbf{g}_2)\mathbf{m} = \sum_i \frac{\exp(\mathbf{g}_2 \mathbf{m}_i^T)}{\sum_j \exp(\mathbf{g}_2 \mathbf{m}_j^T)} \mathbf{m}_i, \quad (11)$$

where  $\mathbf{g}_1 = q_1(x)$  and  $\mathbf{g}_2 = q_2(x)$  are two embedding functions [29, 48] that transmit input  $x$  into two query sets. Then, the front-door adjustment  $\mathcal{FD}(\cdot)$  is approximated as follows:

$$\mathcal{FD}(x, \mathbf{x}') = \mathbb{E}_{\mathbf{x}'}[\mathbf{x}'] + \mathbb{E}_{\mathbf{m}|x}[\mathbf{m}]. \quad (12)$$

As we realize an approximate estimation through a query-mechanism, multi-head attention [40] can be efficiently employed to handle the aforementioned process.

**Specific Network Design.** Although the derivation is clear, how to implement the causal inference in networks presents a significant technical challenge. Considering the characteristics of COPE, we propose leveraging the keypoint features of object as the intermediate knowledge  $\mathbf{m}$ , since the keypoints that are evenly distributed on the surface can effectively capture structure and pose information of the object. Specifically, with the fusion input feature  $\mathcal{F}_{obj}$ , we detect  $N_{kpt}$  local keypoints and perform global information aggregation to extract features  $\mathcal{F}_{kpt} \in \mathbb{R}^{N_{kpt} \times C}$ , following previous keypoint-based methods [22, 25]. As for the cross-sampling features  $\mathbf{x}'$ , one straightforward way is to construct a memory bank of all features and randomly sample a certain number of features from it as the cross-sampling features. Afterward, the memory bank is updated at the corresponding positions using the features from the current mini-batch [46]. However, this intuitive approach is not applicable to our task, since the memory bank should cover all samples, it may not capture the dynamic variations of features during training. The lack of feature consistency

could introduce additional confounders, which will negatively impact the models to focus on correct causal relations.

To this end, we draw inspiration from MoCo [10] and devise a features sampling approach based on a dynamic queue. Specifically, we first construct a 2D queue with a shape of  $N_c \times N_q$ , where  $N_c$  and  $N_q$  represent the number of categories and the queue length of each category, respectively. We then select training instances at random equal to  $N_q$  for each category, extract their features using 3D encoders of ULIP-2 [47], and store them into the queue to complete the initialization. Subsequently, we randomly sample  $N_s$  features from the queue to form  $\mathbf{x}'$ , which is denoted as  $\mathcal{F}_{samp} \in \mathbb{R}^{N_s \times C}$  in the networks. To maintain feature consistency within the queue, we update it using the output features  $\mathcal{F}_f$  from the current epoch using a first-in-first-out (FIFO) strategy. Following Eq. (10) to Eq. (12), the causality-enhanced features  $\mathcal{F}_f$  are obtained as follow:

$$\mathcal{F}_s = SA(\mathcal{F}_{kpt}), \mathcal{F}_c = CA(\mathcal{F}_{kpt}, \mathcal{F}_{samp}) \quad (13)$$

$$\mathcal{F}_f = LN(\mathcal{F}_s + \mathcal{F}_c), \quad (14)$$

where  $SA(\cdot)$  and  $CA(\cdot)$  represent multi-head self-attention and cross-attention, respectively.  $LN(\cdot)$  denotes the layer normalization. Furthermore, to strengthen the stability of learning by fusing causality-enhanced features with original keypoint features of objects, we introduce an adaptive weight fusion method:

$$w_a = \sigma(\mathcal{F}_f W_f + \mathcal{F}_{kpt} W_k) \quad (15)$$

$$\mathcal{F}_f \leftarrow w_a \odot \mathcal{F}_f + (1 - w_a) \odot \mathcal{F}_{kpt}, \quad (16)$$

where  $\sigma$  and  $\odot$  mean Sigmoid function and element-wise multiplication. Suppose  $W_{f/k} \in \mathbb{R}^{C \times 1}$  is learnable weight parameter.

### 4.3. Knowledge Distillation

As mentioned in above section, the foundation model is trained on a broader dataset covering a wide range of categories and scenarios, enabling it to be exposed to diverse data distributions and to implicitly learn more robust debiased feature representations, which can supervise the COPE model to confront biases at the feature level. Inspired by the feature alignment design [13, 59, 60] to transfer visual-language knowledge to different tasks, we exploit 3D foundation model ULIP-2 [47] as the teacher model to provide causal guidance for COPE model. Specifically, we introduce a residual-based approach to distill debiasing knowledge from 3D foundation models into COPE network, as shown in Fig. 3(c). Unlike CLIPose [23] exploits contrastive learning to extract semantic knowledge from text and image modalities, we leverage the pre-trained 3D encoder of ULIP-2, which possesses strong awareness of point cloud structures and category information. Specifically, given the input point cloud of an object  $\mathcal{P}_{obj}$ , we can directly feed it into the frozen 3D Encoder  $\Phi_{ULIP}$  of ULIP-2, which can be written as:

$$\mathcal{F}_P^{ULIP} = \Phi_{ULIP}(\mathcal{P}_{obj}), \quad (17)$$

where  $\mathcal{F}_P^{ULIP} \in \mathbb{R}^{1 \times C_3}$  denotes the [CLS] token in Point Transformer. The feature of [CLS] token embeds representation of the whole point cloud. To obtain the representation of our model, we employ a simple average pooling on point features and get  $\mathcal{F}_P^{avg} \in \mathbb{R}^{1 \times C_1}$ , we have:

$$\mathcal{F}_P^{avg} = \text{AvgPool}(\mathcal{F}_P). \quad (18)$$

Then, we apply a modified residual distillation network on  $\mathcal{F}_P^{avg}$  to transform with two MLP projectors  $K_1$  and  $K_2$ :

$$\widehat{\mathcal{F}}_P^{avg} = \mathcal{F}_P^{avg} + \mu \times K_2(\delta(K_1(\mathcal{F}_P^{avg}))), \quad (19)$$

where  $K_{1/2} \in \mathbb{R}^{C_1 \times C_1}$ ,  $\delta$  represents GeLU function and  $\mu$  is a balancing parameter. Moreover, we initialize the parameters of the second layer  $K_2$  as zeros. Therefore,  $\widehat{\mathcal{F}}_P^{avg}$  initially contains only  $\mathcal{F}_P^{avg}$  and is gradually updated, which avoids introducing additional confounders. Lastly, we can utilize L2 loss to supervise the category knowledge distillation:

$$\mathcal{L}_{KD} = \frac{1}{B} \sum_i^B \left\| \mathcal{F}_P^{ULIP} - \psi(\widehat{\mathcal{F}}_P^{avg}) \right\|_2, \quad (20)$$

where  $\psi \in \mathbb{R}^{C_1 \times C_3}$  is a learnable embedding layer and  $B$  denotes the batch size. In this way, we effectively transfer the debiasing knowledge from the foundation model to the COPE model, further mitigating the interference of data biases. Moreover, the tuned point cloud features of our model can effectively receive supervision from generalized ones, facilitating comprehensive category knowledge learning and enhancing the model’s intra-category generalization.

#### 4.4. Pose Estimation and Overall Loss Function

With the obtained causality-enhanced features, following previous works [19, 22], we recover the final pose and size  $\mathcal{R}, t, s$  via set of keypoint-level correspondences containing global features and points. A simply L1 Loss is used to supervise the predicted pose, in formula:

$$\mathcal{L}_{pose} = \|\mathcal{R}_{gt} - \mathcal{R}\|_2 + \|t_{gt} - t\|_2 + \|s_{gt} - s\|_2, \quad (21)$$

where  $\mathcal{R}_{gt}, t_{gt}, s_{gt}$  means the ground truth rotation, translation and size. For more details please refer to supplementary or [22]. Hence, the overall loss function is as follow:

$$\mathcal{L}_{all} = \alpha_1 \mathcal{L}_{pose} + \alpha_2 \mathcal{L}_{KD}, \quad (22)$$

where  $\alpha_1, \alpha_2$  are hyper-parameters to balanced the contribution of each term. We omit some loss terms for brevity. Please refer to appendix for more details.

## 5. Experiments

**Datasets.** Following previous works [3, 20, 22, 23, 57], we conduct experiments not only on two mainstream NOCS benchmarks, REAL275 [42] and CAMERA25 [42] as well as HouseCat6D [15] datasets. REAL275 is a challenge real-world dataset that contains objects from 6 categories.

The training data consists of 4.3k images from 7 scenes, while testing data includes 2.75k from 6 scenes and 3 objects from each category. CAMERA25 is a synthetic dataset that contains the same categories as REAL275. It provides 300k synthetic RGB-D images of objects rendered on virtual background, with 25k images are withheld for testing. HouseCat6D is a comprehensive multi-modal real-world dataset and encompasses 10 household categories. The training set consists of 20k frames from 34 scenes and testing set consists of 3k frames across 5 scenes. With a total of 194 objects, each category contains 19 objects on average.

**Evaluation Metrics.** Following [3, 22], we evaluate the model performance with two metrics. (i) **3D IoU.** As for NOCS dataset, we report mean average precision (mAP) of CATRE [28] Intersection over Union (IoU) with the thresholds of 75%. For the HouseCat6D dataset, we report the mAP of 3D IoU under thresholds of 25% and 50%. (ii) **n°m cm.** We also utilize the combination of rotation and translation metrics of 5°2 cm, 5°5 cm, 10°2 cm and 10°5 cm, which means the estimation is considered correct when the error is below a threshold.

**Implementation Details.** For a fair comparison, we utilize the same segmentation masks as AG-Pose [22] and DPDN [19] from MaskRCNN [9] and resize them to  $224 \times 224$ . For model parameters, the feature dimensions are set as  $C_1 = C_2 = 128$ ,  $C_3 = 768$  and  $C = 256$ , respectively. The number of point  $N$  in point cloud is 1024 and the number of keypoints  $N_{kpt}$  is set as 96. In dynamic queue, we set the size of queue  $N_q$  to 80 and randomly select  $N_s = 12$  features. For the use of pre-trained 3D models of ULIP-2, we apply PointBert [50] for knowledge distillation and employ PointNet++ [35] for queue construction. For the hyper-parameters setting, the balancing parameter  $\mu$  in residual network is set as 0.1 following [13], and  $\alpha_1, \alpha_2$  in overall loss function are 1 and 0.01, respectively. For model optimizing, we employ the same data augmentation approach as previous works [19, 22], which leverage random rotation degree sampled from  $U(0, 20)$  and rotation  $\Delta t \sim U(-0.02, 0.02)$  and scaling  $\Delta s \sim U(-0.18, 1.2)$ . The network is training on a single NVIDIA L40 GPU for a total of 120k iterations by the Adam [16] optimizer, with a mini training batch is 24 and a learning rate range from  $2e-5$  to  $5e-4$  based on triangular2 cyclical schedule [37].

### 5.1. Comparison with State-of-the-Art Methods

As shown in Fig. 4, our method outperforms prior-free SOTA methods across key metrics in different datasets.

**Performance on REAL275 dataset.** The comparisons between proposed CleanPose and previous methods on challenging REAL275 dataset are shown in Tab. 1. As can be easily observed, our CleanPose achieves the state-of-the-art performance in all metrics and outperforms all previous

Methods	Venue/Source	Shape Prior	$IoU_{75}^*$ ↑	$5^\circ 2cm$ ↑	$5^\circ 5cm$ ↑	$10^\circ 2cm$ ↑	$10^\circ 5cm$ ↑
DPDN[19]	ECCV'22	✓	54.0	46.0	50.7	70.4	78.4
MH6D[26]	TNNLS'24	✓	54.2	53.0	61.1	72.0	82.0
GCE-Pose[17]	CVPR'25	✓	-	<u>57.0</u>	<u>65.1</u>	<u>75.6</u>	<b>86.3</b>
HS-Pose[56]	CVPR'23	✗	39.1	45.3	54.9	68.6	83.6
VI-Net[20]	ICCV'23	✗	48.3	50.0	57.6	70.8	82.1
CLIPose[23]	TCSVT'24	✗	-	48.5	58.2	70.3	85.1
GenPose[52]	NeurIPS'23	✗	-	52.1	60.9	72.4	84.0
SecondPose[3]	CVPR'24	✗	49.7	56.2	63.6	74.7	86.0
AG-Pose[22]	CVPR'24	✗	<u>61.3</u>	<u>57.0</u>	64.6	75.1	84.7
<b>CleanPose (ours)</b>		✗	<b>62.7</b>	<b>61.7</b>	<b>67.6</b>	<b>78.3</b>	<b>86.3</b>

Table 1. **Comparisons with state-of-the-art methods on REAL275 dataset.** ↑: a higher value indicating better performance. ‘\*’ denotes CATRE [28] IoU metrics and ‘-’ means unavailable statistics. Overall best results are in **bold** and the second best results are underlined.

Methods	$IoU_{75}^*$	$5^\circ 2cm$	$5^\circ 5cm$	$10^\circ 2cm$	$10^\circ 5cm$
HS-Pose[56]	-	73.3	80.5	80.4	89.4
CLIPose[23]	-	74.8	82.2	82.0	91.2
GeoReF[57]	79.2	<u>77.9</u>	<u>84.0</u>	83.8	90.5
AG-Pose[22]	<b>81.2</b>	<u>79.5</u>	83.7	<u>87.1</u>	<u>92.6</u>
<b>CleanPose (ours)</b>	<u>80.7</u>	<b>80.3</b>	<b>84.2</b>	<b>87.7</b>	<b>92.7</b>

Table 2. **Comparisons with state-of-the-art methods on CAMERA25 dataset.** A higher value indicating better performance. ‘\*’ denotes CATRE [28] IoU metrics and ‘-’ means unavailable statistics. Overall best results are in **bold** and the second best results are underlined.

Methods	$IoU_{25}$	$IoU_{50}$	$5^\circ 2cm$	$5^\circ 5cm$	$10^\circ 2cm$	$10^\circ 5cm$
FS-Net[2]	74.9	48.0	3.3	4.2	17.1	21.6
GPV-Pose[4]	74.9	50.7	3.5	4.6	17.8	22.7
VI-Net[20]	80.7	56.4	8.4	10.3	20.5	29.1
SecondPose[3]	83.7	66.1	11.0	13.4	25.3	35.7
AG-Pose[22]	<u>88.1</u>	<u>76.9</u>	<u>21.3</u>	<u>22.1</u>	<u>51.3</u>	<u>54.3</u>
<b>CleanPose (ours)</b>	<b>89.2</b>	<b>79.8</b>	<b>22.4</b>	<b>24.1</b>	<b>51.6</b>	<b>56.5</b>

Table 3. **Comparisons with prior-free state-of-the-art methods on HouseCat6D dataset.** A higher value indicating better performance. Overall best results are in **bold** and the second best results are underlined.

methods on REAL275 dataset. Significantly, we achieve the precision of **61.7%**, **67.6%** and **78.3%** in the rigorous metric of  $5^\circ 2cm$ ,  $5^\circ 5cm$  and  $10^\circ 2cm$ , surpassing the current prior-free state-of-the-art method AG-Pose [22] with a large margin by 4.7%, 3.0% and 3.2%, respectively. Our method even outperforms prior-based sota method GCE-Pose [17] by 4.7%, 2.5% and 2.7% in above three rigorous metric. Moreover, the qualitative results of AG-Pose and proposed CleanPose are shown in Fig. 5. It can be observed that our method achieves significantly higher precision. These exceptional outcomes further support the efficacy of our approach. Please refer to appendix for detailed comparison with SOTA method.




Figure 4. Comparison of ours with prior-free SOTA methods.




Figure 5. Qualitative comparison on REAL275 [42]. We compare the predictions of CleanPose and the baseline AG-Pose [22]. The ground truth is marked by white borders.

**Performance on CAMERA25 dataset.** The comparison results are presented in Tab. 2. From the observation of results, it can be deduced that CleanPose ranks top performance across all the metrics except  $IoU_{75}$ , in which our method also achieves comparable performance with the best (80.7% vs. 81.2%). In detail, CleanPose achieves 80.3% on  $5^\circ 2cm$ , 84.2% on  $5^\circ 5cm$  and 87.7% on  $10^\circ 2cm$ , respectively. This superior performance on synthetic dataset further proves the effectiveness of CleanPose.

**Performance on HouseCat6D dataset.** In Tab. 3, we evaluate our method on HouseCat6D [15] dataset. The proposed

ID	Causal	Distillation	5°2cm	5°5cm	10°2cm	10°5cm
1	✗	✗	57.0	64.6	75.1	84.7
2	✓	✗	59.7	66.5	77.5	86.0
3	✗	✓	57.9	65.5	76.1	85.2
4	✓	✓	<b>61.7</b>	<b>67.6</b>	<b>78.3</b>	<b>86.3</b>

Table 4. Effect of causal learning and knowledge distillation.

CleanPose again achieves state-of-the-art performance in all metrics. In detail, our method outperforms the current best method AG-Pose [22] by 2.9% on  $IoU_{50}$ , 1.1% on 5°2 cm, 2.0% on 5°5 cm and 2.1% on 10°5 cm, respectively. The overall and category-wise evaluation of 3D IoU on House-Cat6D dataset is provided in appendix.

## 5.2. Ablation Studies

We conduct ablation experiments to demonstrate the effectiveness of the proposed method on REAL275 dataset [42].

### Effect of Causal Learning and Knowledge Distillation.

In Tab. 4, we perform ablations of the proposed two main components. We adopt AG-Pose [22] as our baseline, which serves as the current sota framework in COPE. The results indicate that the integration of causal inference or knowledge distillation leads to significant enhancements in the model’s performance. This strongly demonstrates causal learning’s considerable potential and the effectiveness of comprehensive category information guidance in COPE.

**Effect of different feature storage methods and update strategies.** In Tab. 5a, we investigate the impact of distinct combinations of feature storage and update strategies. ‘Similarity’ represents updating the closest features via similarity computation. The results indicate that employing the dynamic queue and FIFO update mechanism yields the best performance. We owe the advantage to the dynamic queue’s superior ability to capture feature variations compared to a memory bank. Moreover, the FIFO ensures the removal of outdated features, which is beneficial for causal inference.

**Effect of adaptive fusion method.** As shown in Tab. 5b, our model achieves sota performance even without feature fusion. Introduction of the adaptive feature fusion method further improves the results in rigorous metrics.

### Effect of distinct category knowledge learning policies.

In Tab. 5c, we assess the impact of category knowledge learning policies. ‘Concat’ means directly concatenating  $\mathcal{F}_P^{ULIP}$  and  $\mathcal{F}_P$ , while ‘Contrastive’ indicates aligning above two features via contrastive learning. Our proposed residual distillation policy achieves the best performance since the residual layers effectively balance feature learning and category knowledge transfer.

**Effect of different 3D encoders.** Tab. 5d illustrates the impact of using different 3D encoders in the distillation. Objectively, the pre-trained PointBert [50] realizes the sota point cloud zero-shot classification results [47]. Our model

Data	Update	5°2cm	5°5cm	10°2cm	10°5cm
Queue	FIFO	<b>61.7</b>	<b>67.6</b>	<b>78.3</b>	<b>86.3</b>
Queue	w/o update	57.1	66.1	77.5	85.7
Queue	Similarity	61.0	67.1	77.6	86.0
Memory bank	w/o update	57.6	66.4	75.6	85.5
Memory bank	Similarity	59.1	66.9	77.5	86.1

(a) Effect of different feature storage methods and update strategies.

AF	5°2cm	5°5cm	10°2cm	10°5cm
✗	60.5	66.5	<b>78.8</b>	<b>86.5</b>
✓	<b>61.7</b>	<b>67.6</b>	78.3	86.3

(b) Effect of adaptive fusion in front-door adjustment.

Policies	5°2cm	5°5cm	10°2cm	10°5cm
Concat	60.0	66.7	77.9	85.9
Contrastive	59.8	66.5	77.6	85.8
Distillation	<b>61.7</b>	<b>67.6</b>	<b>78.3</b>	<b>86.3</b>

(c) Effect of distinct category knowledge learning policies.

3D Encoder	5°2cm	5°5cm	10°2cm	10°5cm
PointNet++[35]	59.6	66.8	77.7	85.8
PointMLP[30]	60.8	67.2	77.4	86.0
PointBert[50]	<b>61.7</b>	<b>67.6</b>	<b>78.3</b>	<b>86.3</b>

(d) Effect of different 3D encoders of ULIP-2 [47] for distillation.

Table 5. Ablation studies on key details. Default settings are colored in gray.

is also more performant using PointBert. Moreover, since our method selects PointNet++ as 3D encoder, we hypothesize that when the teacher and student models share similar architectures (*e.g.*, both using PointNet++), the distillation may mislead the student model to focus on feature structure similarity rather than transferring category knowledge. Please refer to appendix for more ablations.

## 6. Conclusion

In this paper, we present CleanPose, the first solution that addresses the dataset biases in category-level pose estimation from the perspective of causal learning. Motivated by pivotal observation that humans can learn inherent causality beyond biases, we formulate the modeling of crucial causal variables and develop a causal inference framework in COPE task. Furthermore, we devise a residual knowledge distillation network to transfer unbiased semantics knowledge from 3D foundation model, providing comprehensive causal guidance to achieve unbiased estimation. Extensive experiments on challenging benchmarks demonstrate that CleanPose can significantly improve performance, showing the effectiveness of our method.

## Acknowledgments

This paper is supported by the National Natural Science Foundation of China under Grants (62233013, 62173248, 62333017, 624B2105).

## References

- [1] Kai Chen and Qi Dou. Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 2773–2782, 2021. 1
- [2] Wei Chen, Xi Jia, Hyung Jin Chang, Jinming Duan, Linlin Shen, and Ales Leonardis. Fs-net: Fast shape-based network for category-level 6d object pose estimation with decoupled rotation mechanism. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1581–1590, 2021. 7
- [3] Yamei Chen, Yan Di, Guangyao Zhai, Fabian Manhardt, Chenyangguang Zhang, Ruida Zhang, Federico Tombari, Nassir Navab, and Benjamin Busam. Secondpose: Se (3)-consistent dual-stream feature fusion for category-level pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9959–9969, 2024. 1, 2, 6, 7
- [4] Yan Di, Ruida Zhang, Zhiqiang Lou, Fabian Manhardt, Xiangyang Ji, Nassir Navab, and Federico Tombari. Gpv-pose: Category-level object pose estimation via geometry-guided point-wise voting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6781–6791, 2022. 1, 2, 7
- [5] Zhaoxin Fan, Zhengbo Song, Jian Xu, Zhicheng Wang, Kejian Wu, Hongyan Liu, and Jun He. Acr-pose: Adversarial canonical representation reconstruction network for category level 6d object pose estimation. *arXiv preprint arXiv:2111.10524*, 2021. 2
- [6] Yang Fu and Xiaolong Wang. Category-level 6d object pose estimation in the wild: A semi-supervised learning approach and a new dataset. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:27469–27483, 2022. 2
- [7] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *International Conference on Learning Representations (ICLR)*, 2024. 3
- [8] Yang Hai, Rui Song, Jiaojiao Li, and Yinlin Hu. Shape-constraint recurrent flow for 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4831–4840, 2023. 1
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (CVPR)*, pages 2961–2969, 2017. 3, 6
- [10] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9729–9738, 2020. 2, 5
- [11] Yisheng He, Haibin Huang, Haoqiang Fan, Qifeng Chen, and Jian Sun. Ffb6d: A full flow bidirectional fusion network for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3003–3013, 2021. 1
- [12] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [13] Xiaohu Huang, Hao Zhou, Kun Yao, and Kai Han. Froster: Frozen clip is a strong teacher for open-vocabulary action recognition. In *International Conference on Learning Representations (ICLR)*, 2024. 3, 5, 6
- [14] Yuanmin Huang, Mi Zhang, Daizong Ding, Erling Jiang, Zhaoxiang Wang, and Min Yang. Causalpc: Improving the robustness of point cloud classification by causal effect identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19779–19789, 2024. 2, 3
- [15] HyunJun Jung, Shun-Cheng Wu, Patrick Ruhkamp, Guangyao Zhai, Hannah Schieber, Giulia Rizzoli, Pengyuan Wang, Hongcheng Zhao, Lorenzo Garattoni, Sven Meier, et al. Housecat6d-a large-scale multi-modal category level 6d object perception dataset with household objects in realistic scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22498–22508, 2024. 2, 3, 6, 7
- [16] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [17] Weihang Li, Hongli Xu, Junwen Huang, Hyunjun Jung, Peter KT Yu, Nassir Navab, and Benjamin Busam. Gce-pose: Global context enhancement for category-level object pose estimation. *arXiv preprint arXiv:2502.04293*, 2025. 1, 2, 7
- [18] Haitao Lin, Zichang Liu, Chilam Cheang, Yanwei Fu, Guodong Guo, and Xiangyang Xue. Sar-net: Shape alignment and recovery network for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6707–6717, 2022. 2
- [19] 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 *European Conference on Computer Vision (ECCV)*, pages 19–34. Springer, 2022. 1, 6, 7
- [20] Jiehong Lin, Zewei Wei, Yabin Zhang, and Kui Jia. Vi-net: Boosting category-level 6d object pose estimation via learning decoupled rotations on the spherical representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 14001–14011, 2023. 1, 2, 6, 7
- [21] Xiao Lin, Deming Wang, Guangliang Zhou, Chengju Liu, and Qijun Chen. Transpose: 6d object pose estimation with geometry-aware transformer. *Neurocomputing*, 589:127652, 2024. 1
- [22] Xiao Lin, Wenfei Yang, Yuan Gao, and Tianzhu Zhang. Instance-adaptive and geometric-aware keypoint learning for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21040–21049, 2024. 1, 2, 3, 5, 6, 7, 8

- [23] Xiao Lin, Minghao Zhu, Ronghao Dang, Guangliang Zhou, Shaolong Shu, Feng Lin, Chengju Liu, and Qijun Chen. Clipose: Category-level object pose estimation with pre-trained vision-language knowledge. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 2024. 1, 5, 6, 7
- [24] Bing Liu, Dong Wang, Xu Yang, Yong Zhou, Rui Yao, Zhiwen Shao, and Jiaqi Zhao. Show, deconfound and tell: Image captioning with causal inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18041–18050, 2022. 3
- [25] Jianhui Liu, Yukang Chen, Xiaoqing Ye, and Xiaojuan Qi. Ist-net: Prior-free category-level pose estimation with implicit space transformation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13978–13988, 2023. 5
- [26] Jian Liu, Wei Sun, Chongpei Liu, Hui Yang, Xing Zhang, and Ajmal Mian. Mh6d: Multi-hypothesis consistency learning for category-level 6-d object pose estimation. *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*, 2024. 7
- [27] Weiyang Liu, Zhen Liu, Liam Paull, Adrian Weller, and Bernhard Schölkopf. Structural causal 3d reconstruction. In *European Conference on Computer Vision (ECCV)*, pages 140–159. Springer, 2022. 3
- [28] Kingyu Liu, Gu Wang, Yi Li, and Xiangyang Ji. Catre: Iterative point clouds alignment for category-level object pose refinement. In *European Conference on Computer Vision (ECCV)*, pages 499–516. Springer, 2022. 6, 7
- [29] Yang Liu, Guanbin Li, and Liang Lin. Cross-modal causal relational reasoning for event-level visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 45(10):11624–11641, 2023. 5
- [30] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. In *International Conference on Learning Representations (ICLR)*, 2022. 8
- [31] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research Journal (TMLR)*, pages 1–31, 2024. 3
- [32] J Pearl. *Causality*. Cambridge university press, 2009. 2, 3, 4
- [33] Judea Pearl. *Causal inference in statistics: a primer*. John Wiley & Sons, 2016. 2, 3, 4
- [34] Yun Peng, Xiao Lin, Nachuan Ma, Jiayuan Du, Chuangwei Liu, Chengju Liu, and Qijun Chen. Sam-lad: Segment anything model meets zero-shot logic anomaly detection. *arXiv preprint arXiv:2406.00625*, 2024. 3
- [35] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017. 3, 6, 8
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763. PMLR, 2021. 3
- [37] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 464–472. IEEE, 2017. 6
- [38] Meng Tian, Marcelo H Ang, and Gim Hee Lee. Shape prior deformation for categorical 6d object pose and size estimation. In *European Conference on Computer Vision (ECCV)*, pages 530–546. Springer, 2020. 2, 3
- [39] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 13(04):376–380, 1991. 2
- [40] 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 (NeurIPS)*, 30, 2017. 5
- [41] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3343–3352, 2019. 1
- [42] 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 (CVPR)*, pages 2642–2651, 2019. 1, 2, 3, 6, 7, 8
- [43] Liuyi Wang, Zongtao He, Ronghao Dang, Mengjiao Shen, Chengju Liu, and Qijun Chen. Vision-and-language navigation via causal learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13139–13150, 2024. 3, 4, 5
- [44] Pengyuan Wang, HyunJun Jung, Yitong Li, Siyuan Shen, Rahul Parthasarathy Srikanth, Lorenzo Garattoni, Sven Meier, Nassir Navab, and Benjamin Busam. Phocal: A multi-modal dataset for category-level object pose estimation with photometrically challenging objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21222–21231, 2022. 2
- [45] Tan Wang, Chang Zhou, Qianru Sun, and Hanwang Zhang. Causal attention for unbiased visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3091–3100, 2021. 3
- [46] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3733–3742, 2018. 5
- [47] Le Xue, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, et al. Ulip-2: Towards scalable multimodal pre-training for 3d understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27091–27101, 2024. 2, 3, 5, 8
- [48] Xu Yang, Hanwang Zhang, and Jianfei Cai. Deconfounded image captioning: A causal retrospect. *IEEE Transactions*

- on *Pattern Analysis and Machine Intelligence (TPAMI)*, 45 (11):12996–13010, 2021. 3, 5
- [49] Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. Causal attention for vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9847–9857, 2021. 3
- [50] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19313–19322, 2022. 6, 8
- [51] Hua Zhang, Liqiang Xiao, Xiaochun Cao, and Hassan Foroosh. Multiple adverse weather conditions adaptation for object detection via causal intervention. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 46 (3):1742–1756, 2022. 3
- [52] Jiyao Zhang, Mingdong Wu, and Hao Dong. Genpose: generative category-level object pose estimation via diffusion models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NeurIPS)*, pages 54627–54644, 2023. 7
- [53] Letian Zhang, Xiaotong Zhai, Zhongkai Zhao, Yongshuo Zong, Xin Wen, and Bingchen Zhao. What if the tv was off? examining counterfactual reasoning abilities of multi-modal language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21853–21862, 2024. 3
- [54] Mengchen Zhang, Tong Wu, Tai Wang, Tengfei Wang, Ziwei Liu, and Dahua Lin. Omni6d: Large-vocabulary 3d object dataset for category-level 6d object pose estimation. *arXiv preprint arXiv:2409.18261*, 2024. 2, 3
- [55] Xiangdong Zhang, Shaofeng Zhang, and Junchi Yan. Pcp-mae: Learning to predict centers for point masked autoencoders. *arXiv preprint arXiv:2408.08753*, 2024. 2
- [56] Linfang Zheng, Chen Wang, Yinghan Sun, Esha Dasgupta, Hua Chen, Ales Leonardis, Wei Zhang, and Hyung Jin Chang. Hs-pose: Hybrid scope feature extraction for category-level object pose estimation. *arXiv preprint arXiv:2303.15743*, 2023. 1, 2, 3, 7
- [57] Linfang Zheng, Tze Ho Elden Tse, Chen Wang, Yinghan Sun, Hua Chen, Ales Leonardis, Wei Zhang, and Hyung Jin Chang. Georef: Geometric alignment across shape variation for category-level object pose refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10693–10703, 2024. 3, 6, 7
- [58] Guangliang Zhou, Deming Wang, Yi Yan, Huiyi Chen, and Qijun Chen. Semi-supervised 6d object pose estimation without using real annotations. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 32(8): 5163–5174, 2021. 1
- [59] Minghao Zhu, Xiao Lin, Ronghao Dang, Chengju Liu, and Qijun Chen. Fine-grained spatiotemporal motion alignment for contrastive video representation learning. In *Proceedings of the 31st ACM International Conference on Multimedia (ACMMM)*, pages 4725–4736, 2023. 5
- [60] Minghao Zhu, Zhengpu Wang, Mengxian Hu, Ronghao Dang, Xiao Lin, Xun Zhou, Chengju Liu, and Qijun Chen.

Mote: Reconciling generalization with specialization for visual-language to video knowledge transfer. *arXiv preprint arXiv:2410.10589*, 2024. 3, 5