

Real-time Holistic Robot Pose Estimation with Unknown States

Shikun Ban¹, Juling Fan¹, Xiaoxuan Ma¹,
Wentao Zhu¹, Yu Qiao⁵, and Yizhou Wang^{1,2,3,4}

¹ Center on Frontiers of Computing Studies, School of Computer Science,
Peking University

² Inst. for Artificial Intelligence, Peking University

³ Nat'l Eng. Research Center of Visual Technology

⁴ Nat'l Key Lab of General Artificial Intelligence

⁵ Department of Automation, School of Electronics, Information and Electrical
Engineering, Shanghai Jiao Tong University

Abstract. Estimating robot pose from RGB images is a crucial problem in computer vision and robotics. While previous methods have achieved promising performance, most of them presume full knowledge of robot internal states, *e.g.* ground-truth robot joint angles. However, this assumption is not always valid in practical situations. In real-world applications such as multi-robot collaboration or human-robot interaction, the robot joint states might not be shared or could be unreliable. On the other hand, existing approaches that estimate robot pose without joint state priors suffer from heavy computation burdens and thus cannot support real-time applications. This work introduces an efficient framework for real-time robot pose estimation from RGB images without requiring known robot states. Our method estimates camera-to-robot rotation, robot state parameters, keypoint locations, and root depth, employing a neural network module for each task to facilitate learning and sim-to-real transfer. Notably, it achieves inference in a single feed-forward pass without iterative optimization. Our approach offers a 12× speed increase with state-of-the-art accuracy, enabling real-time holistic robot pose estimation for the first time. Code and models are available at <https://github.com/Oliverbansk/Holistic-Robot-Pose-Estimation>.

1 Introduction

Robot pose estimation is a fundamental task in computer vision and robotics as it not only lays the groundwork for multiple downstream tasks, *e.g.* visually-guided object grasping and manipulation [21, 40, 49], but also paves the way for a variety of applications, such as autonomous navigation [1, 13, 33, 34], human-robot interaction [3, 39, 43, 44, 46, 48] and multi-robot collaboration [20, 29, 35, 45].

✉Correspondence to: Wentao Zhu <wtzhu@pku.edu.cn> and Yu Qiao <qiaoyu@sjtu.edu.cn>.




Fig. 1: The majority of previous robot pose estimation methods assume known robot joint states and focus solely on estimating the camera-to-robot pose, *i.e.* camera-to-robot rotation and translation. In contrast, the holistic robot pose estimation problem requires estimating both joint states and the camera-to-robot pose, given only an RGB image without known joint states. For holistic robot pose estimation, RoboPose [15] uses costly test-time optimization (Render-and-Compare) to iteratively refine the predictions. In contrast, our feed-forward method achieves state-of-the-art accuracy with a $12\times$ speed boost.

Nonetheless, most previous works for monocular robot pose estimation [17, 18, 22, 40] rely on a known set of robot joint angles as an additional input to alleviate the difficulty of the estimation process. Yet, this prerequisite information of the robot state is not always available in practical scenarios:

- For instance, low-cost robotic arms without sophisticated sensors or internal monitors can not provide the visual calibration process with readily accessible robot state measurements [4, 10, 49].
- In the context of multi-robot collaboration, the real-time synchronization of joint state data may be hindered by environmental limitations [24], necessitating pose estimation without robot joint state information, and additional monitoring of the joint states.
- In domains such as human-robot interaction [36, 41] and competitive robot gaming [32, 42], the reliability of the state information from counterpart robots is not always guaranteed, as trust issues may arise in robotic systems [11, 12, 14]. In this sense, failing to obtain independent estimation of the counterpart robot joint state may lead to potential safety issues.

These scenarios further emphasize the urgent need for monocular robot pose estimation with unknown robot joint state. In this work, we address the challenge of *Holistic Robot Pose Estimation* with unknown internal states for articulated robots. Specifically, given a monocular RGB image, we aim to simultaneously solve for 6D robot pose (3D rotation and translation relative to the camera) and fine-grained robot joint state parameters (including joint angles for revolute joints and displacement for prismatic joints).

Nonetheless, learning holistic robot pose estimation without joint state information presents significant challenges. The primary challenge is the dual ambiguity of unknown joint states and 6D pose. Specifically, estimating the 6D pose

of an articulated robot is more complex than for a rigid object due to shape and scale variability from unknown joint states. Conversely, estimating joint states is complicated by the need for 2D-to-3D correspondences and the unknown 6D pose. This intertwined issue requires novel solutions to decouple joint states from 6D pose estimation. Furthermore, the sophisticated robot morphologies vary from model to model, requiring non-trivial model design. In addition, The robot parts are frequently invisible in the monocular image as a result of self-occlusion and truncation, posing more challenges for robust estimation. Finally, the domain gap between synthetic training data and real-world images introduces additional challenges for learning-based approaches.

It is worth noting that some previous works [15,49] have also addressed holistic estimation with unknown states, tackling some challenges. However, these methods all necessitate iterative optimization at inference, which is computation-intensive and unsuitable for real-time applications. Motivated by the need for efficient holistic robot pose estimation without knowledge of the robot state, we propose an end-to-end framework that directly estimates camera-to-robot pose and robot joint state parameters, requiring only a single feed-forward pass during inference.

In pursuit of this goal, we factorize the holistic robot pose estimation task into several sub-tasks, including the estimation of camera-to-robot rotation, robot joint states, root depth, and root-relative keypoint locations. We then design corresponding neural network modules for each sub-task. Specifically, the standalone root depth estimator disentangles the estimation of scale-variant and scale-invariant variables. The keypoint estimator bridges the gap between the pixel space and the robotic parameter space. Additionally, we utilize self-supervised consistency regularizations which could be performed on the real images without ground truth (GT).

We demonstrate that our designs of modular feed-forward networks improve model generalization while maintaining high computation efficiency. Consequently, our approach achieves state-of-the-art estimation accuracy with unknown robot states and a $12\times$ speedup compared to the iterative Render-and-Compare (RnC) [15] approach, as depicted in Fig. 1. The key insight behind our approach is that we could get rid of the cumbersome RnC [15] and error-prone Perspective-n-Points (PnP) [18,22,40] schemes to establish 2D-3D correspondences. Instead, we demonstrate that feed-forward networks could well handle the problem with proper problem decomposition and meticulous design.

In summary, our contributions are three-fold:

1. We present an efficient end-to-end learning framework for the holistic robot pose estimation problem with unknown robot states.
2. We propose to factorize and decompose the holistic robot pose estimation problem into several sub-tasks and tackle them with respective modules.
3. We achieve state-of-the-art performance on various robot models while delivering a significant $12\times$ speedup, paving the way for real-time applications.

2 Related Work

2.1 Hand-eye calibration

The goal of Hand-Eye Calibration (HEC) is to determine the transformation between the position of a robot’s end-effector (hand) and its vision system (eye). This task can be categorized into two types: *eye-in-hand* where the camera is fixed at the robot end-effector [27], and *eye-to-hand* which refers to the configurations where the camera is placed at a distance from the robot [5, 23, 30]. This work focuses on a similar setting with eye-to-hand. Traditional HEC methods involve attaching a fiducial marker to the end effector [6, 7, 28], placing it in several known positions, detecting the marker positions in images, and then solving for the camera-to-robot transformation. Despite its broad applicability, the classical HEC pipeline necessitates a complete offline calibration process every time the camera-robot transformation changes, which can be tedious in real-world applications. Consequently, recent works, such as DREAM [18], propose an online, markerless calibration problem, as discussed below.

2.2 Image-based robot pose estimation

Previous works can be divided into two categories based on whether the robot states are known or not. Most previous works estimate the 6D camera-to-robot pose with known robot states. For example, DREAM [18] first trains a 2D keypoint estimator and then applies Perspective- n -Point (PnP) [19] to obtain camera-to-robot pose, showing comparable results with the classical methods. Some work further utilizes segmentation masks and edges to aid pose estimation [8, 16]. Chen *et al.* [2] apply rendering-based camera pose optimization and consistency-based joint space exploration. SGTAPose [40] leverages robot structure prior and temporal information. Lu *et al.* [22] propose to train on real-world data using foreground segmentation and differentiable rendering in a self-supervised manner.

Recently, several works have shifted to estimating robot pose with unknown joint states. Zuo *et al.* [49] estimate the robot pose and joint angles of a simple 4-DoF robotic arm from a single RGB image by solving a nonlinear nonconvex optimization problem. However, this method is not well-suited for robot arms with higher complexity. Labbé *et al.* [15] propose a render-and-compare approach to iteratively refine the estimation for both known and unknown states, and apply the method to different robot models up to 15-DoF. However, the iterative optimization process limits its applications in real-time scenarios. In contrast, our approach can be deployed with a single feed-forward at inference and achieves real-time efficiency.

3 Method

3.1 Overview

Given an RGB image \mathbf{I} , we employ a suite of neural network modules, including DepthNet, JointNet, RotationNet, and KeypointNet. These modules are used

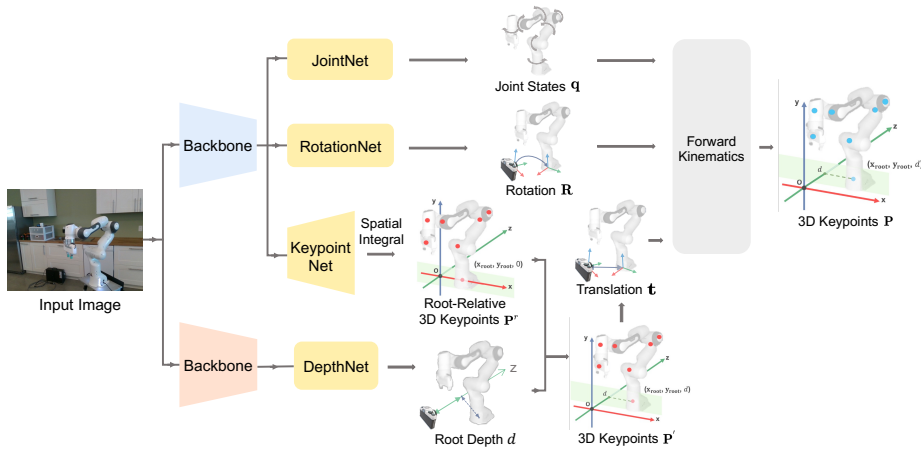


Fig. 2: Framework overview. The JointNet and the RotationNet regress joint state parameters \mathbf{q} and camera-to-robot rotation \mathbf{R} , respectively. The KeypointNet estimates root-relative 3D keypoint locations \mathbf{P}^r . The DepthNet’s estimation of root depth d is combined with \mathbf{P}^r to acquire absolute 3D keypoint locations \mathbf{P}' and camera-to-robot translation \mathbf{t} . Joint state parameters \mathbf{q} , rotation \mathbf{R} and translation \mathbf{t} are used to compute 3D keypoint locations \mathbf{P} via forward-kinematics.

to respectively estimate the absolute root depth d , joint state parameters \mathbf{q} , camera-to-robot rotation $\mathbf{R} \in SO(3)$, and root-relative 3D keypoint locations $\mathbf{P}^r \in \mathbb{R}^{N \times 3}$, where N represents the total number of keypoints. The keypoints are defined at the joint articulations, in line with previous works [15, 18, 22], and one keypoint is designated as the robot root. The root joint is critical for computing the camera-to-robot rotation \mathbf{R} and translation \mathbf{t} , as well as serving as the origin for the forward kinematics (FK) process.

Our modularized design offers several key insights. Firstly, DepthNet is tailored to learn the scale-variant camera-robot spatial relationship, while the other modules concentrate on the scale-invariant robot features. Secondly, while the joint state parameters \mathbf{q} provide a compact representation of robot morphologies subject to robotic structure constraints, they are highly abstract and thus challenging to directly regress from image features. To address this, we utilize the 3D keypoint locations as an intermediary. On the one hand, they can be effectively estimated with pixel-aligned image features; on the other hand, they are inherently linked to \mathbf{q} through an FK process.

An overview of the proposed framework is illustrated in Fig. 2. Specifically, DepthNet estimates the absolute root depth d , which, combined with KeypointNet’s output $\mathbf{P}^r \in \mathbb{R}^{N \times 3}$, provides the absolute 3D keypoint locations $\mathbf{P}' \in \mathbb{R}^{N \times 3}$. The root keypoint location in \mathbf{P}' is the camera-to-robot translation \mathbf{t} . The estimated joint state parameters \mathbf{q} , along with the camera-to-robot rotation \mathbf{R} and translation \mathbf{t} , are then used to perform FK, yielding the final

3D keypoint estimations $\mathbf{P} \in \mathbb{R}^{N \times 3}$. Please refer to the appendix for a detailed illustration of the relations of these variables.

In the following, we will first introduce the architecture design of the network modules in Sec. 3.2, and then explain the training loss design consisting of ground-truth supervision and self-supervision in Sec. 3.3.

3.2 Architecture

DepthNet Monocular object depth estimation suffers from significant ambiguities, as the absolute depth value and camera intrinsics are deeply intertwined [25]. For instance, a robot in a fixed pose relative to the camera may appear larger in photos taken with cameras having larger focal lengths. Scaling both the absolute depth and focal lengths equally results in identical object size and appearance in different images. To address this challenge, we implement DepthNet following the approach in [26] to disentangle depth from camera intrinsics. Initially, we use camera intrinsics to calculate a normalized, coarse depth value d_c :

$$d_c = \sqrt{\frac{f_x \cdot f_y \cdot A_{real}}{A_{bbox}}}, \quad (1)$$

where f_x and f_y represent the focal lengths along the x- and y-axes, A_{real} is the area of the robot in the x-y plane of the camera frame (in mm^2), and A_{bbox} is the area of the robot’s bounding box (in $pixel^2$). This method of calculating d_c normalizes it in 3D space and provides an approximate depth estimation for subsequent refinement. The primary function of DepthNet is to predict a correction factor λ based on the input image. We multiply the estimated λ by d_c to obtain the final absolute root depth d :

$$d = \lambda \cdot d_c. \quad (2)$$

With a fixed camera, the relationship between the normalized coarse depth d_c and the robot’s size in the image is linear. However, robots can assume poses with the same depth but different sizes in the image, making d_c inadequate for accurate root depth representation. The correction factor λ introduces non-linear flexibility to the depth estimation, addressing this issue by adjusting d_c based on image features.

JointNet The goal of JointNet is to estimate joint state parameters $\mathbf{q} = (q_1, q_2, \dots, q_J)$, where J denotes the number of robot joints. As this work deals with 1-DoF joints (revolute or prismatic), J equals the robot’s DoF. JointNet shares the same feature extractor with RotationNet and KeypointNet. We extract feature map $\mathbf{F} \in \mathbb{R}^{C \times H \times W}$ (where C denotes the number of channels, and $H \times W$ denotes the spatial resolution) and use global average pooling to obtain $\mathbf{f} \in \mathbb{R}^C$. Finally, we employ several MLP blocks to regress the joint state parameters \mathbf{q} from \mathbf{f} .

RotationNet The goal of RotationNet is to estimate the camera-to-robot rotation $\mathbf{R} \in SO(3)$. We employ an MLP to estimate \mathbf{R} from the image feature \mathbf{f} . We utilize the 6D continuous rotation representation [47] to parameterize \mathbf{R} .

KeypointNet The goal of KeypointNet is to estimate root-relative 3D keypoint locations $\mathbf{P}^r \in \mathbb{R}^{N \times 3}$ for N robot keypoints. We transform the feature map \mathbf{F} to a 3D heatmap $\mathbf{H} \in \mathbb{R}^{N \times D \times H' \times W'}$ through convolutional layers, where $D \times H' \times W'$ is the spatial resolution discretizing the 3D space. The 3D heatmap encodes per-voxel likelihood of containing each robot keypoint. We derive the root-relative 3D positions of each keypoint \mathbf{P}_n^r by performing spatial integral techniques [38] over the corresponding heatmap \mathbf{H}_n in a differentiable manner:

$$\mathbf{P}_n^r = \sum_{k=1}^D \sum_{i=1}^{H'} \sum_{j=1}^{W'} (k, i, j) \cdot \mathbf{H}_n(k, i, j). \quad (3)$$

By combining \mathbf{P}^r with the estimated root depth d , we obtain the regressed 3D keypoint locations \mathbf{P}' and root translation \mathbf{t} . We further calculate the FK-based 3D keypoint locations \mathbf{P} with differentiable FK parameterized by \mathbf{q} , \mathbf{R} , and \mathbf{t} .

3.3 Training loss

Ground-truth supervision For training on synthetic datasets where we have access to all the ground-truth information, we apply separate supervision for all the sub-tasks to train the corresponding network modules:

$$\mathcal{L}_{\text{depth}} = \|d - \hat{d}\|_1, \quad (4)$$

$$\mathcal{L}_{\text{joint}} = \|\mathbf{q} - \hat{\mathbf{q}}\|_2, \quad (5)$$

$$\mathcal{L}_{\text{rot}} = \|\mathbf{R} - \hat{\mathbf{R}}\|_2, \quad (6)$$

$$\mathcal{L}_{\text{trans}} = \|\mathbf{t} - \hat{\mathbf{t}}\|_2, \quad (7)$$

$$\mathcal{L}_{\text{kpts}} = \|\mathbf{P} - \hat{\mathbf{P}}\|_2 + \|\mathbf{p} - \hat{\mathbf{p}}\|_2, \quad (8)$$

$$\mathcal{L}'_{\text{kpts}} = \|\mathbf{P}' - \hat{\mathbf{P}}'\|_2 + \|\mathbf{p}' - \hat{\mathbf{p}}'\|_2, \quad (9)$$

where d , \mathbf{q} , \mathbf{R} , \mathbf{t} , \mathbf{P} , and \mathbf{P}' are GT root depth, joint states, camera-to-robot rotation and translation, FK-based keypoint locations and regressed keypoint locations, respectively; \mathbf{p} and \mathbf{p}' are 2D projections of \mathbf{P} and \mathbf{P}' using the camera intrinsics; $\hat{\cdot}$ denotes model estimations.

We then train all the network modules end-to-end using

$$\mathcal{L}_{\text{GT}} = \mathcal{L}_{\text{joint}} + \mathcal{L}_{\text{rot}} + \mathcal{L}_{\text{trans}} + \lambda(\mathcal{L}_{\text{kpts}} + \mathcal{L}'_{\text{kpts}}), \quad (10)$$

where we set $\lambda = 10.0$ to balance the loss terms.

Self-supervision Apart from GT supervision on synthetic data, our framework also enables scaling the training pipeline to unlabeled real-world data in a self-supervised manner [22]. We employ two consistency-based training regularizations to improve model generalization.

Despite the absence of GT joint states and keypoint locations, we regularize the consistency of the FK-based keypoints $\hat{\mathbf{P}}$ and regressed keypoints $\hat{\mathbf{P}}'$ by:

$$\mathcal{L}_{\text{kc}} = \|\hat{\mathbf{P}} - \hat{\mathbf{P}}'\|_2. \quad (11)$$

We introduce this keypoint consistency regularization to effectively reconcile the predictions subject to robot structure constraints $\hat{\mathbf{P}}$ and predictions based on the pixel-aligned feature $\hat{\mathbf{P}}'$. In addition, similar to [22], we also regularize the consistency between robot rendering mask $\mathbf{M}_{\text{render}}$ and foreground segmentation mask \mathbf{M}_{seg} using their Intersection-over-Union (IoU):

$$\mathcal{L}_{\text{mc}} = 1 - \frac{\mathcal{S}_{\text{intersection}}}{\mathcal{S}_{\text{union}}}, \quad (12)$$

where $\mathcal{S}_{\text{intersection}}$ and $\mathcal{S}_{\text{union}}$ is the area of the intersection and the union of $\mathbf{M}_{\text{render}}$ and \mathbf{M}_{seg} , respectively. $\mathbf{M}_{\text{render}}$ is obtained with differentiable rendering using \mathbf{q} , \mathbf{R} , and \mathbf{t} [15]; \mathbf{M}_{seg} is obtained with an image segmentation model [22].

The overall self-supervision training objective is as follows:

$$\mathcal{L}_{\text{self}} = \mathcal{L}_{\text{kc}} + \lambda_{\text{mc}}\mathcal{L}_{\text{mc}} \quad (13)$$

where we set $\lambda_{\text{mc}} = 1.0$. We apply end-to-end self-supervised training for the sim-to-real domain adaptation stage.

4 Experiments

4.1 Dataset and metrics

We utilize the benchmark dataset DREAM [18], for both training and testing in our study. This dataset comprises images from three widely used robotic arms, namely, Franka Emika Panda (Panda), Kuka iiwa7 (Kuka), and Rethink Robotics Baxter (Baxter). The training dataset for each robot consists of approximately 100k synthetic images generated using the domain randomization (DR) technique. These images encompass scenarios with varying camera-to-robot poses and joint states. The testing datasets include both DR-generate synthetic images for each robot and photorealistic synthetic images (denoted as "Photo") for Panda and Kuka. Additionally, real-world image datasets of Panda are provided. Panda-3Cam comprises a total of 17k image frames captured by three different cameras, namely Azure Kinect (AK), XBOX 360 Kinect (XK), and Realsense (RS). This dataset is presented as three sub-datasets, each corresponding to one camera. Another real-world dataset, Panda-ORB, contains 32,315 images captured from 27 viewpoints using a Realsense camera.

We evaluate the accuracy of our estimations using the Average Distance (ADD) metric, which represents the average Euclidean distances (in millimeters) between predicted 3D keypoint locations and their corresponding ground truth locations. Our reported metrics include both the Area Under the Curve (AUC) value and the mean value of ADD. The AUC integrates ADD over different thresholds (with the maximum threshold set at 100mm), while the mean value effectively represents the general prediction accuracy.

4.2 Implementation details

We first pre-train DepthNet on the synthetic training dataset for 100 epochs, employing a learning rate of $1e-4$ and relying solely on ground truth supervision for depth. Subsequently, we conduct training for our entire model based on the pretrained DepthNet for an additional 100 epochs with a learning rate of $1e-4$ and a decay rate of 0.95. The Adam optimizer is employed in each stage to optimize the network parameters, with the momentum set to 0.9. As depth influences the global offset of all keypoints, a pre-trained DepthNet serves as a reasonable starting point for training other networks. On the Panda real datasets, we further apply self-supervised training on real-world images to overcome the sim-to-real domain gap with a learning rate of $1e-6$. This fine-tuning process uses only image data and no ground-truth labels.

4.3 Evaluation and Comparison

Table 1: Comparison of AUC \uparrow of the ADD distribution curve and mean \downarrow of the ADD, on each dataset in the DREAM datasets. Panda 3CAM datasets and Panda ORB are real-world datasets and the rest are synthetic datasets. \dagger denotes using ground-truth joint state parameters.

Method	Metric	Baxter DR	Kuka DR	Kuka Photo	Panda DR	Panda Photo	Panda 3CAM-AK	Panda 3CAM-XK	Panda 3CAM-RS	Panda ORB
\dagger DREAM-F [18]	AUC \uparrow	-	-	-	81.3	79.5	68.9	24.4	76.1	61.9
\dagger DREAM-Q [18]	AUC \uparrow	75.5	-	-	77.8	74.3	52.4	37.5	78.0	57.1
\dagger DREAM-H [18]	AUC \uparrow	-	73.3	72.1	82.9	81.1	60.5	64.0	78.8	69.1
RoboPose [15]	AUC \uparrow	32.7	80.2	73.2	82.9	79.7	70.4	77.6	74.3	70.4
Ours	AUC \uparrow	58.8	75.1	73.9	82.7	82.0	82.2	76.0	75.2	75.2
\dagger DREAM-F [18]	Mean \downarrow	-	-	-	552.0	785.8	11413.1	491911.4	2077.4	95319.1
\dagger DREAM-Q [18]	Mean \downarrow	59.0	-	-	359.0	569.0	78089.3	54178.2	27.2	64247.6
\dagger DREAM-H [18]	Mean \downarrow	-	5613.2	11099.4	325.6	518.8	56.5	7381.6	23.6	25685.3
RoboPose [15]	Mean \downarrow	87.9	34.2	89.0	21.4	32.2	34.3	22.3	26.0	30.1
Ours	Mean \downarrow	45.0	28.0	29.0	18.0	18.4	18.9	24.0	24.8	24.8

Accuracy We first quantitatively evaluate the estimation accuracy of our method and compare it with the state-of-the-art (SOTA) approaches. To the best of our knowledge, among prior works, only RoboPose [15] addresses a scenario similar to ours, where both the pose and joint state parameters for an industry-level robot are estimated from a single RGB image. We also present the results of DREAM [18], a representative feed-forward method using ground-truth joint state parameters, as a reference. Table 1 presents the comparisons of AUC and the mean value of ADD. Our method outperforms SOTA on most datasets, significantly reducing joint localization errors. Specifically, DREAM [18] detects the 2D keypoint locations first, then employs a PnP process to solve the 2D-to-3D transformation. However, once the 2D keypoint detector fails due to truncations and occlusions, the PnP process produces unreliable numerical solutions. RoboPose [15] utilizes iterative render-and-compare, which is prone to falling into local optima on complex robot models with symmetric structures (*e.g.*, Baxter). In contrast, our approach directly regresses both the robot pose and joint state parameters, enabling end-to-end optimization from multi-facet supervision signals. We also visualize the complete ADD distribution. As shown in Fig. 3, our approach achieves higher accuracy across most ADD thresholds. Notably, our approach attains near 100% accuracy with a 60mm threshold, indicating more reliable estimation results on challenging cases with fewer outliers.




Fig. 3: Comparison of ADD distributions on the real-world datasets between our approach and RoboPose [15]. The y-axis represents the accuracy of our estimation at different ADD thresholds.

Finally, we visualize the predictions of our approach and compare them with SOTA to better understand the performance improvement. Figure 4 illustrates our approach’s ability to produce high-quality estimates across different robot morphologies. Particularly, our approach demonstrates superior performance in highly challenging cases, including self-occlusions, truncations, and extreme lighting conditions. Additionally, we compare the model performance qualitatively on in-the-wild lab images and quantitatively under various levels of truncation severity, illustrating the superior generalization ability and robustness of our approach. Please refer to the appendix for more detailed results.

Efficiency Furthermore, we benchmark the computation efficiency of our method and compare it with previous works. We measure the average inference time,

Table 2: Comparison of computation efficiency on the Panda 3CAM-AK dataset. † denotes using ground-truth joint state parameters.

Method	AUC ↑	Time (ms) ↓	FPS ↑	FLOPS (G) ↓
† DREAM-F 18	68.9	66.0	15.1	90.5
RoboPose 15	70.4	570.8	1.8	57.3
Ours	82.2	44.3	22.6	25.2

frame per second (FPS), and Floating Point Operations (FLOPs). For a fair comparison, we benchmark all methods on a Linux machine with CPU Intel(R) Xeon(R) Gold 6240 CPU @ 2.60GHz and GPU NVIDIA Tesla V100S. Batch size is set to be 1 for all methods. Table [2](#) shows that our method exhibits a remarkable advantage in terms of inference speed and computation loads, enabling real-time inference while achieving state-of-the-art accuracy. This advantage stems from the fact that our method involves only a single feed-forward pass without costly iterative optimization during inference. We demonstrate that our approach serves as an efficient and effective solution for real-time holistic robot pose estimation, which holds great potential in real-world applications.

4.4 Comparison under known joints states

Although the primary focus of this work is holistic robot pose estimation with unknown states, our approach can also be easily adapted to the known joint states setting. We conduct additional experiments on four real datasets for robot pose estimation with ground-truth joint states available. We simply replace the predicted joint state with the ground-truth joint state during the self-supervised training process. As shown in Tab. [3](#), our approach also achieves comparable performance with previous SOTA under this well-explored setting.

4.5 Ablation study

Ablation on network modules We first conduct an ablation study on the proposed network modules to investigate their effectiveness. Since JointNet and RotationNet are the major regression targets of our task, we focus on the ablation of two assistive network modules, DepthNet and KeypointNet. We conduct the experiments on the synthetic datasets, utilizing only ground-truth supervision to simplify the process.

Firstly, we delve into the implementation of DepthNet. In Tab. [4](#), we evaluate the design choice of whether to share the backbone of DepthNet with other modules or not. The results indicate that separating the DepthNet backbone consistently improves performance. This improvement could be explained by the fact that DepthNet involves learning scale-variant features, while the other network modules all focus on learning scale-invariant features. We also investigate whether to pre-train the DepthNet or jointly train it with other modules

Table 3: Comparison of AUC \uparrow of the ADD distribution curve and Mean \downarrow of the ADD, on the real-world datasets of DREAM, with all methods using **known ground-truth joint state parameters**.

Method	Metric	Panda 3CAM-AK	Panda 3CAM-XK	Panda 3CAM-RS	Panda ORB
DREAM-F [18]	AUC \uparrow	68.9	24.4	76.1	61.9
DREAM-Q [18]	AUC \uparrow	52.4	37.5	78.0	57.1
DREAM-H [18]	AUC \uparrow	60.5	64.0	78.8	69.1
RoboPose [15]	AUC \uparrow	76.5	86.0	76.9	80.5
CtRNet [22]	AUC \uparrow	89.9	79.5	90.8	85.3
SGTAPose [40]	AUC \uparrow	67.8	2.1	87.6	72.3
Ours	AUC \uparrow	90.2	81.2	91.9	87.6
DREAM-F [18]	Mean \downarrow	11413.1	491911.4	2077.4	95319.1
DREAM-Q [18]	Mean \downarrow	78089.3	54178.2	27.2	64247.6
DREAM-H [18]	Mean \downarrow	56.5	7381.6	23.6	25685.3
RoboPose [15]	Mean \downarrow	24.2	14.0	23.1	19.4
CtRNet [22]	Mean \downarrow	13.0	32.0	10.0	21.0
SGTAPose [40]	Mean \downarrow	35.7	157.7	12.8	35.0
Ours	Mean \downarrow	12.9	19.9	11.0	15.9

Table 4: Ablation studies on the effect of separating DepthNet image feature extractor.

Dataset	Separate	AUC \uparrow	Mean (mm) \downarrow
Panda DR (a)	–	79.98	21.9
Panda DR (b)	+	82.68	18.0
Panda Photo (a)	–	79.66	22.0
Panda Photo (b)	+	81.98	18.4

Table 5: Ablation studies on the effect of pre-training DepthNet in advance of jointly training.

Dataset	Pretrain	AUC \uparrow	Mean (mm) \downarrow
Panda DR (a)	–	74.29	26.8
Panda DR (b)	+	82.68	18.0
Panda Photo (a)	–	73.34	27.9
Panda Photo (b)	+	81.98	18.4

from scratch. Table [5](#) suggests that a pre-training stage leads to considerable performance improvement. This improvement likely occurs because a well-trained DepthNet reliably aids other modules in better training, as it determines the global translation of all joints. Additionally, we select the keypoint closest to the robot center as the root keypoint requiring depth prediction by DepthNet. Please refer to the S2.3 section of the appendix for rationale behind this selection.

Moreover, we conduct ablation experiments to investigate the impact of Key-pointNet on other network modules, namely JointNet and RotationNet. We utilize Mean Joint Error, which represents the mean error of the joint state parameter predictions, and the Mean Rotation Angle Error [\[15\]](#), which represents the mean error of the Euler angle predictions $(\theta_x^{pred}, \theta_y^{pred}, \theta_z^{pred})$ for the camera-to-robot rotation, to quantify the accuracy of the JointNet and RotationNet predictions, respectively. It is important to note that a Panda robot consists of seven revolute joints with angle parameters (representing one-DoF rotations) and one prismatic joint with a length parameter (representing one-DoF transla-

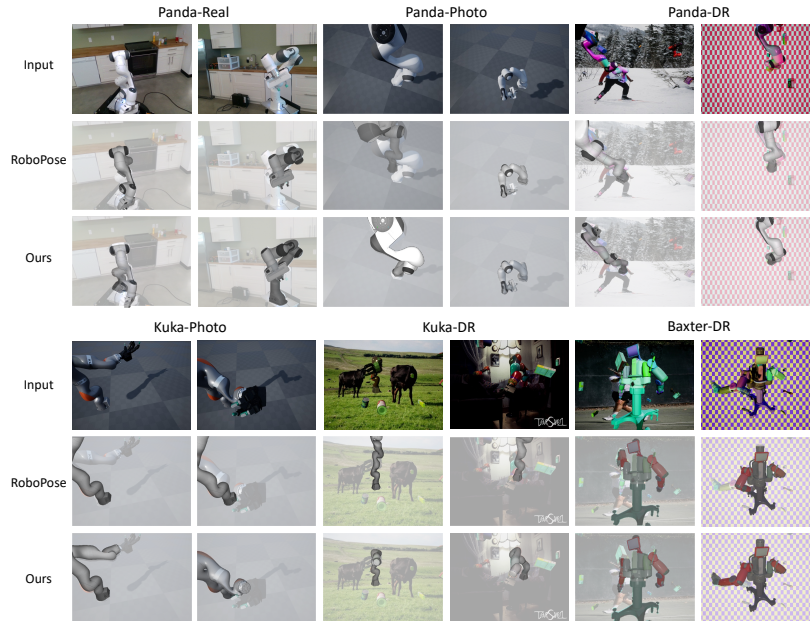



Fig. 4: Qualitative comparison between our method and RoboPose [15] on both real and synthetic datasets.

tion). Therefore, we calculate and present the Mean Joint Error separately for the revolute and prismatic joints of Panda. We present the detailed ablation results in the S2.1 section of the appendix, in which Tables 6 and 7 show that the presence of KeypointNet significantly improves the prediction accuracy of JointNet and RotationNet. By sharing a feature extractor with JointNet and RotationNet, KeypointNet guides the network towards learning pixel-aligned robot visual features, which further assists in estimating robot pose and joint states.

Finally, we examine the overall effect of DepthNet and KeypointNet. To replace DepthNet, we implement a straightforward alternative, directly regressing the camera-to-robot translation \mathbf{t} from image features \mathbf{f} using an MLP, similar to RotationNet. Figure 5a illustrates that both modules play an important role in prediction accuracy. The DepthNet notably enhances accuracy by improving root depth estimation, a global offset that affects the precision of all joints. Meanwhile, KeypointNet further benefits the estimation results.

The contribution of these modules to the overall performance validates our design. DepthNet disentangles depth from camera intrinsics, facilitating a learning-based approach to depth estimation. This method exhibits greater expressive capability compared to the analytical, non-parameterized PnP methods employed in previous work [18, 22, 40]. On the other hand, KeypointNet enhances both JointNet and RotationNet by steering the shared backbone towards learning more pixel-aligned features, thus complementing the learning objectives defined in the robot kinematic space.

Ablation on self-supervision We also perform an ablation study of our self-supervised learning objectives. We conduct the experiments on the Panda real datasets. Figure 5b demonstrates that both the keypoint consistency regularization and the mask consistency regularization contribute to the overall performance. Specifically, the mask consistency regularization provides pixel-wise supervision signal with the help of differentiable rendering, while the keypoint consistency regularization unifies the predictions in the pixel space and that in the robot kinematic space. We suppose that pixel-aligned 2D tasks, *e.g.* keypoint detection and mask segmentation, generally generalize better, and could therefore improve the 3D estimation quality through consistency regularization, further improving the overall performance. As a result, our multi-stage training scheme yields favorable convergence ability. Please refer to the S4 section of the appendix for the error analysis of each module across each training stage.



(a) Ablation study on network modules. ‘Baseline’ represents direct regression, ‘DN’ and ‘KN’ refer to DepthNet and KeypointNet, respectively.

(b) Ablation study on self-supervision. ‘Baseline’ represents no self-supervision, ‘MC’ and ‘KC’ refer to mask and keypoint consistency regularization, respectively.

Fig. 5: Ablation studies on network modules and self-supervision.

5 Conclusion

This work presents a novel end-to-end framework for real-time holistic robot pose estimation without prior knowledge of internal robot states, providing substantial assistance in real-world scenarios where the joint state of the estimation target is unavailable or unreliable. By decomposing this complex task into specific sub-tasks and tailoring neural network modules for each, the method effectively addresses the challenges while retaining high computation efficiency. Our approach, which circumvents iterative optimization, significantly accelerates the estimation process and achieves state-of-the-art accuracy. This advancement holds promise for enhancing real-world robotics applications, offering a practical solution for real-time robot pose estimation and joint state monitoring.

Acknowledgments

This work is supported by National Science and Technology Major Project (2022ZD0114904). We thank Hai Ci for the valuable discussions during the early stages of this work. We thank Haoran Lu and Ruihai Wu for their assistance with real-world image captures.

References

1. Bultmann, S., Memmesheimer, R., Behnke, S.: External camera-based mobile robot pose estimation for collaborative perception with smart edge sensors. In: ICRA (2023)
2. Chen, L., Qin, Y., Zhou, X., Su, H.: EasyHeC: Accurate and Automatic Hand-eye Calibration via Differentiable Rendering and Space Exploration. In: IEEE Robotics and Automation Letters (2023)
3. Christen, S., Yang, W., Pérez-D’Arpino, C., Hilliges, O., Fox, D., Chao, Y.W.: Learning human-to-robot handovers from point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9654–9664 (2023)
4. Deisenroth, M.P., Rasmussen, C.E., Fox, D.: Learning to control a low-cost manipulator using data-efficient reinforcement learning. *Robotics: Science and Systems VII* **7**, 57–64 (2011)
5. Feniello, A., Dang, H.N., Birchfield, S.: Program synthesis by examples for object repositioning tasks. 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems pp. 4428–4435 (2014)
6. Fiala, M.: Artag, a fiducial marker system using digital techniques. In: CVPR (2005)
7. Garrido-Jurado, S., Muñoz-Salinas, R., Madrid-Cuevas, F.J., Marín-Jiménez, M.J.: Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* (2014)
8. Hao, R., Özgüner, O., Çavuşoğlu, M.C.: Vision-based surgical tool pose estimation for the da vinci® robotic surgical system. In: 2018 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 1298–1305. IEEE (2018)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
10. Icli, C., Stepanenko, O., Bonev, I.: New method and portable measurement device for the calibration of industrial robots. *Sensors* **20**(20) (2020)
11. Kim, W., Kim, N., Lyons, J.B., Nam, C.S.: Factors affecting trust in high-vulnerability human-robot interaction contexts: A structural equation modelling approach. *Applied Ergonomics* **85**, 103056 (2020)
12. Kok, B.C., Soh, H.: Trust in robots: Challenges and opportunities. *Current Robotics Reports* **1**, 297–309 (2020)
13. Kothari, N., Gupta, M., Vachhani, L., Arya, H.: Pose estimation for an autonomous vehicle using monocular vision. In: 2017 Indian Control Conference (ICC). pp. 424–431 (2017)
14. Kuipers, B.: How can we trust a robot? *Communications of the ACM* **61**(3), 86–95 (2018)
15. Labbé, Y., Carpentier, J., Aubry, M., Sivic, J.: Single-view robot pose and joint angle estimation via render & compare. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR) (2021)

16. Lambrecht, J., Grosenick, P., Meusel, M.: Optimizing keypoint-based single-shot camera-to-robot pose estimation through shape segmentation. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 13843–13849. IEEE (2021)
17. Lambrecht, J., Kästner, L.: Towards the usage of synthetic data for marker-less pose estimation of articulated robots in rgb images. In: ICAR (2019)
18. Lee, T.E., Tremblay, J., To, T., Cheng, J., Mosier, T., Kroemer, O., Fox, D., Birchfield, S.: Camera-to-robot pose estimation from a single image. In: International Conference on Robotics and Automation (ICRA) (2020)
19. Lepetit, V., Moreno-Noguer, F., Fua, P.: Epnnp: An accurate o (n) solution to the pnp problem. *IJCV* (2009)
20. Li, S., De Wagter, C., De Croon, G.C.: Self-supervised monocular multi-robot relative localization with efficient deep neural networks. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 9689–9695. IEEE (2022)
21. Li, X., Zhang, X., Zhou, X., Chen, I.M.: Upg: 3d vision-based prediction framework for robotic grasping in multi-object scenes. *Knowledge-Based Systems* **270**, 110491 (2023)
22. Lu, J., Richter, F., Yip, M.C.: Markerless camera-to-robot pose estimation via self-supervised sim-to-real transfer. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
23. Mahler, J., Matl, M., Satish, V., Danielczuk, M., DeRose, B., McKinley, S., Goldberg, K.: Learning ambidextrous robot grasping policies. *Science Robotics* **4**(26), eaau4984 (2019)
24. Masehian, E., Jannati, M., Hekmatfar, T.: Cooperative mapping of unknown environments by multiple heterogeneous mobile robots with limited sensing. *Robotics and Autonomous Systems* **87**, 188–218 (2017)
25. Ming, Y., Meng, X., Fan, C., Yu, H.: Deep learning for monocular depth estimation: A review. *Neurocomputing* **438**, 14–33 (2021)
26. Moon, G., Chang, J., Lee, K.M.: Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In: The IEEE Conference on International Conference on Computer Vision (ICCV) (2019)
27. Morrison, D., Corke, P., Leitner, J.: Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach. In: Proc. of Robotics: Science and Systems (RSS) (2018)
28. Olson, E.: Apriltag: A robust and flexible visual fiducial system. In: ICRA (2011)
29. Papadimitriou, A., Mansouri, S.S., Nikolakopoulos, G.: Range-aided ego-centric collaborative pose estimation for multiple robots. *Expert Systems with Applications* **202**, 117052 (2022)
30. Park, D., Seo, Y., Chun, S.Y.: Real-time, highly accurate robotic grasp detection using fully convolutional neural networks with high-resolution images (2019)
31. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
32. Pinto, L., Davidson, J., Gupta, A.: Supervision via competition: Robot adversaries for learning tasks. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 1601–1608 (2017)
33. Prusak, A., Melnychuk, O., Roth, H., Schiller, I., Koch, R.: Pose estimation and map building with a time-of-flight-camera for robot navigation. *International Journal of Intelligent Systems Technologies and Applications* **5**(3-4), 355–364 (2008)
34. Rana, A., Vulpi, F., Galati, R., Milella, A., Petitti, A.: A pose estimation algorithm for agricultural mobile robots using an rgb-d camera. In: 2022 International Con-

- ference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME). pp. 1–5 (2022)
35. Rizk, Y., Awad, M., Tunstel, E.W.: Cooperative heterogeneous multi-robot systems. *ACM Computing Surveys (CSUR)* **52**, 1 – 31 (2019)
 36. Salvini, P., Paez-Granados, D.F., Billard, A.: Safety concerns emerging from robots navigating in crowded pedestrian areas. *International Journal of Social Robotics* **14**, 441 – 462 (2021)
 37. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: *CVPR* (2019)
 38. Sun, X., Xiao, B., Wei, F., Liang, S., Wei, Y.: Integral human pose regression. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2018)
 39. Svenstrup, M., Tranberg, S., Andersen, H.J., Bak, T.: Pose estimation and adaptive robot behaviour for human-robot interaction. In: *2009 IEEE International Conference on Robotics and Automation*. pp. 3571–3576 (2009)
 40. Tian, Y., Zhang, J., Yin, Z., Dong, H.: Robot structure prior guided temporal attention for camera-to-robot pose estimation from image sequence. In: *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)* (2023)
 41. Tzafestas, S.G.: 15 - mobile robots at work. In: Tzafestas, S.G. (ed.) *Introduction to Mobile Robot Control*, pp. 635–663. Elsevier, Oxford (2014)
 42. Uchibe, E., Asada, M., Noda, S., Takahashi, Y., Hosoda, K.: Vision-based reinforcement learning for robocup: Towards real robot competition. In: *Proc. of IROS*. vol. 96 (1996)
 43. Xia, Y., Tang, Y., Hu, Y., Hoffman, G.: Privacy-preserving pose estimation for human-robot interaction. *CoRR* **abs/2011.07387** (2020)
 44. Xu, C., Yu, X., Wang, Z., Ou, L.: Multi-view human pose estimation in human-robot interaction. In: *IECON 2020 The 46th Annual Conference of the IEEE Industrial Electronics Society*. pp. 4769–4775 (2020)
 45. Xun, Z., Huang, J., Li, Z., Xu, C., Gao, F., Cao, Y.: Crepes: Cooperative relative pose estimation towards real-world multi-robot systems. *arXiv preprint arXiv:2302.01036* (2023)
 46. Yang, W., Paxton, C., Mousavian, A., Chao, Y.W., Cakmak, M., Fox, D.: Reactive human-to-robot handovers of arbitrary objects. In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE (2021)
 47. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5745–5753 (2019)
 48. Zhu, W., Ma, X., Ro, D., Ci, H., Zhang, J., Shi, J., Gao, F., Tian, Q., Wang, Y.: Human motion generation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
 49. Zuo, Y., Qiu, W., Xie, L., Zhong, F., Wang, Y., Yuille, A.L.: Craves: Controlling robotic arm with a vision-based economic system. In: *CVPR* (2019)

Appendix

The content of our supplementary material is organized as follows.

1. Implementation details of our method.
2. Additional quantitative results.
3. Additional qualitative results.
4. Error analysis of each sub-module across training stages.
5. Multi-view fusion results.
6. Failure cases of our method.

S1. Implementation Details

S1.1 Model architecture

We implement the proposed pipeline using PyTorch [31]. We utilize HRnet-w32 [37] as the feature extractor for DepthNet. JointNet, RotationNet and KeypointNet share a ResNet-50 [9] network as the feature extractor.

S1.2 Method Details

We provide a detailed illustration of the relations of the variables mentioned in Sec. 3.1

$$\mathbf{P}'_i = \mathbf{P}^r_i + (0, 0, d), i = 0, 1, 2, \dots \quad (14)$$

$$\mathbf{t} = \mathbf{P}'_{RootId} \quad (15)$$

$$\mathbf{P} = FK(\mathbf{q}, \mathbf{R}, \mathbf{t}) \quad (16)$$

where d , $RootId$, \mathbf{q} , \mathbf{R} , \mathbf{t} , \mathbf{P}^r , \mathbf{P}' , and \mathbf{P} are the depth (z-axis value) of the root keypoint, index of the root keypoint, joint states, camera-to-robot rotation, camera-to-robot translation, root-relative keypoint locations, KeypointNet-derived keypoint locations and FK-based keypoint locations, respectively. FK denotes the forward-kinematics process.

Here we clarify the coordinate frames of the three sets of keypoints: \mathbf{P}^r is a set of root-relative keypoints in the root-relative frame, with absolute xy value and root-relative z-axis offsets. Both \mathbf{P}' and \mathbf{P} are absolute keypoints in the camera frame. \mathbf{P}' is derived from d and \mathbf{P}^r using Eq. (14), and \mathbf{P} is derived from forward kinematics as in Eq. (16).

S1.3 Robot arms

We leverage the datasets provided by DREAM [18], which contains images for three types of robot arms, namely Franka Emika Panda (Panda), Kuka iiwa7 (Kuka), and Rethink Robotics Baxter (Baxter).

Panda features 8 one-DoF joints and 7 keypoints, manually defined at joint articulations. The first 7 joints in Panda are revolute, facilitating pure rotational motion around a shared axis, and their internal state measures in degrees. The last joint is prismatic, allowing linear sliding along a shared axis without rotational motion, with its internal state measurable in millimeters.

Kuka comprises 7 one-DoF revolute joints and includes 8 defined keypoints. Baxter incorporates 15 one-DoF revolute joints and 17 defined keypoints. Both Kuka and Baxter have no prismatic joints.

Panda and Kuka are both single-arm robotic manipulators, while Baxter has multiple arms connected with a torso, resulting in a human-like appearance.

S1.4 Rendering method

We utilize the provided Unified Robot Description Format (URDF) file following previous works [15, 22]. With URDF parsed and joint configuration \mathbf{q} given, we determine the robot’s morphology via forward kinematics, using the URDF-Pytorch module provided by [15]. \mathbf{R} and \mathbf{t} determine the cam-to-robot-root transformation. As the robot’s own morphology and global pose are acquired, we render its mask with a given camera using a mesh renderer derived from [22].

S2. Additional Quantitative Results

Table 6: Ablation studies on the influence of KeypointNet on JointNet.

Dataset	KeypointNet	Mean Joint Error ↓
Panda DR	–	14.49° 3.45 mm
Panda DR	+	9.29° 2.63 mm
Panda Photo	–	15.14° 3.83 mm
Panda Photo	+	10.45° 2.69 mm
Kuka DR	–	9.59°
Kuka DR	+	8.25°
Kuka Photo	–	9.51°
Kuka Photo	+	8.22°
Baxter DR	–	13.73°
Baxter DR	+	12.95°

Table 7: Ablation studies on the influence of KeypointNet on RotationNet.

Dataset	KeypointNet	Mean Rot. Error (°) ↓
Panda DR	–	3.78
Panda DR	+	3.25
Panda Photo	–	4.03
Panda Photo	+	3.36
Kuka DR	–	8.14
Kuka DR	+	7.12
Kuka Photo	–	8.28
Kuka Photo	+	7.38
Baxter DR	–	4.36
Baxter DR	+	3.69

S2.1 Ablation results on the influence of KeypointNet on JointNet and RotationNet

From the results in Tabs. 6 and 7, we can infer that the presence of KeypointNet significantly improves the prediction accuracy of JointNet and RotationNet.

This strongly validates our model design. By sharing a feature extractor with JointNet and RotationNet, KeypointNet guides the multi-module network towards learning pixel-aligned visual features. This bridges the gap between the pixel space and the robotic parameter space, providing strong support for the holistic robot pose estimation task.

Table 8: Performance comparison with different numbers of in-frame keypoints (different levels of truncation severity).

Dataset	In-frame KPs	Images	RoboPose [15]		Ours	
			AUC \uparrow	Mean \downarrow	AUC \uparrow	Mean \downarrow
Panda 3CAM-AK	4	96	52.28	47.71	62.54	37.45
Panda 3CAM-AK	5	92	73.88	26.11	79.31	20.68
Panda 3CAM-AK	6	105	6.05	126.87	71.98	28.01
Panda 3CAM-AK	7	6041	72.32	31.58	82.88	18.29
Panda 3CAM-XK	7	4966	77.61	22.27	76.01	23.99
Panda 3CAM-RS	5	100	38.60	61.39	76.57	23.42
Panda 3CAM-RS	6	11	35.83	64.16	74.56	25.42
Panda 3CAM-RS	7	5833	74.90	25.42	75.20	24.79
Panda ORB	4	98	16.96	88.73	14.40	85.66
Panda ORB	5	476	47.75	52.42	52.02	47.97
Panda ORB	6	650	56.22	44.50	67.78	32.21
Panda ORB	7	31091	71.20	29.28	75.94	24.05
Panda Photo	2	4	36.05	186.04	38.57	69.22
Panda Photo	3	17	53.44	112.83	38.13	75.10
Panda Photo	4	491	66.65	49.85	66.52	34.80
Panda Photo	5	207	74.06	36.89	76.74	23.57
Panda Photo	6	202	73.80	43.72	78.43	21.82
Panda Photo	7	5076	84.12	24.76	84.01	16.21
Panda DR	2	4	0	17347.83	23.00	77.40
Panda DR	3	10	0	15646.74	61.24	38.75
Panda DR	4	391	52.19	58.09	68.04	33.56
Panda DR	5	231	71.76	33.37	78.21	21.83
Panda DR	6	206	78.16	25.66	79.59	22.30
Panda DR	7	5156	86.02	18.40	84.20	16.44

S2.2 Ablation results with varying levels of truncation severity

We observe that image truncation, *i.e.*, the partial capture of robot arms in images, considerably influences the estimation results. To evaluate our method’s robustness, we measure our model performance with varying numbers of in-frame keypoints (representing different levels of truncation severity) and compare the results with RoboPose [\[15\]](#) in the context of holistic estimation with

unknown joint states, as detailed in Tab. 8. Our approach generally surpasses RoboPose 15 across various datasets and truncation degrees, particularly in terms of mean estimation errors. This demonstrates the improved robustness of our approach.

S2.3 Ablation results with different selection of the reference root keypoint

In our experiments for all three robots, we consistently designate the keypoint located in close proximity to the geometric center of the robot as the reference root keypoint. This choice is motivated by the observation that the keypoints near the geometric center are less likely to be truncated.

Table 9: Ablation studies on the selection of reference root keypoint on the Panda DR dataset.

Reference Root Keypoint	Mean Depth Error (<i>mm</i>) ↓
P_0	18.07
P_1	14.58
P_2	13.00
P_3	9.45
P_4	15.57
P_5	18.13
P_6	20.38




Fig. 6: Visualization of the keypoints of robot Panda. We choose P_3 , the keypoint closest to the geometric center, as the reference root keypoint.

In addition, we quantitatively evaluate the selection of the reference root keypoint on the Panda DR dataset. We measure Mean Depth Error, which represents the mean error of the root depth predictions of DepthNet, under different reference root keypoints, denoted as \mathbf{P}_i ($i \in 0, 1, 2, 3, 4, 5, 6$). Table 9 and Figure 6 demonstrate that the central keypoint indeed leads to more precise depth estimations, validating our choices.

S2.4 Ablation results with different model backbones

In our SOTA model, we use HRNet-32 for DepthNet (scale-variant) and ResNet-50 for other modules (scale-invariant). Here we present a comparison of other options in Tab. 10, which validates our selection. Lighter backbones generally yield slightly inferior results, though they are still comparable and may be suitable in scenarios where speed is prioritized.

Table 10: Ablation studies on the choices of model backbone.

DepthNet	Other modules	Panda DR		Panda Photo	
		AUC \uparrow	Mean (mm) \downarrow	AUC \uparrow	Mean (mm) \downarrow
ResNet-34	ResNet-50	80.7	20.0	80.5	19.9
ResNet-50	ResNet-50	79.5	21.2	78.8	21.7
HRNet-32	ResNet-34	78.3	22.5	78.0	22.5
HRNet-32	ResNet-50	82.7	18.0	82.0	18.4
HRNet-32	HRNet-32	79.9	20.9	79.0	21.5

S3. Additional Qualitative Results

S3.1 Results with unknown joint states

Additionally, in the context of holistic estimation with unknown joint states, we present qualitative comparison results on freshly-collected in-the-wild lab images, demonstrating our method’s real-world generalization capabilities. The images are captured using an iPhone, and camera calibration is not performed. Solely based on RGB images of the robot Panda and predefined camera intrinsic parameters, we estimate both the robot’s pose (6DoF) and joint states (8DoF) using our approach. The estimated robot pose and states are rendered and visualized in Fig. 7. The marker mounted at the end-effector is not utilized.

For example, in the top row of images, our method provides a more accurate estimation of the pose of the robot’s base while RoboPose [15] fails to properly position the base. In the bottom row of images, our method provides a more precise estimation of the joint angles of the robot’s upper part, resulting a more accurate alignment between the rendered robot image and the input robot image.

The key component that contributes to our method’s generalization capabilities is the self-supervision of our framework. It enables our model to be fine-tuned

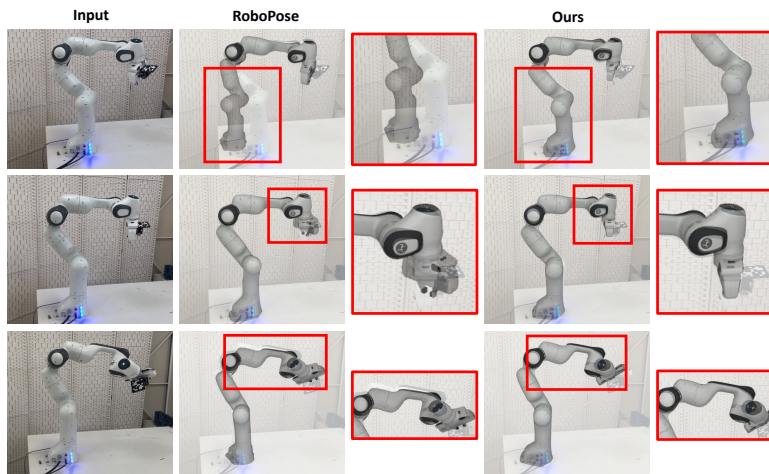


Fig. 7: Additional qualitative comparison results on in-the-wild images, with unknown joint states.

on unlabeled data, making our method well-suited for adaptation to various real-world use cases. To further investigate this, we also conducted self-supervised training with 50% real data from the correspondent DREAM [18] subset, resulting in only a slight decrease in AUC: 81.9 vs. 82.2 on Panda 3CAM-AK and 72.6 vs. 75.2 on Panda 3CAM-RS. This demonstrates that self-supervised domain adaptation can be effectively applied even with a small amount of data, and sets the foundation for more flexible real-world applications of our method.

S3.2 Results with known joint states

In the context of estimation with known joint states, we also present qualitative comparison results on several real world images of robot Panda from the DREAM [18] datasets, as shown in Fig. 8. Our approach can be easily adapted for known-joint estimation by simply substituting predicted joint states with ground-truth joint states.

Moreover, by eliminating the inefficient Render-and-Compare (RnC) optimization used RoboPose [15] and unstable Perspective-n-Points (PnP) methods used in CtrNet [22], our approach generally outperforms previous state-of-the-art methods in this well-explored setting. Free of the cumbersome RnC, our approach achieves a $12\times$ speed boost compared with RoboPose [15]. Without unstable and error-prone PnP methods, our approach demonstrates higher estimation accuracy and robustness compared with CtrNet [22]. Please kindly refer to the video on the project page for more video comparison results.

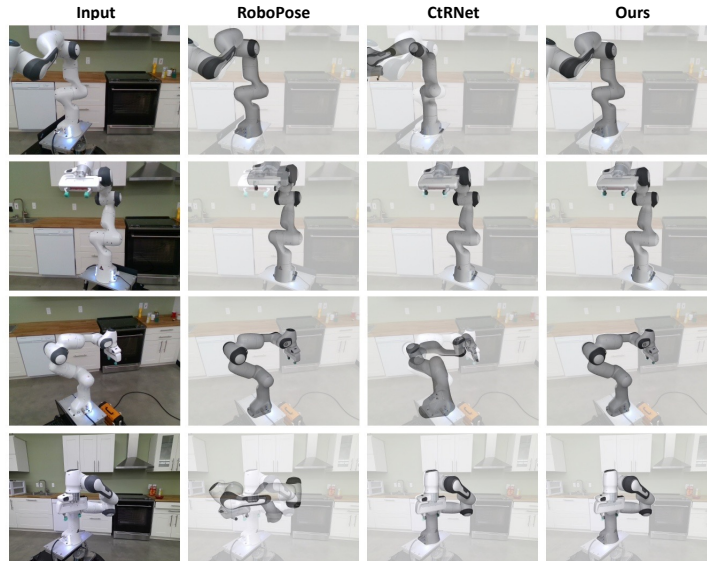


Fig. 8: Additional qualitative comparison results on real-world images, with known joint states.

Table 11: Error analysis of each module across training stages on Panda 3CAM-AK dataset.

Module	Metric	Synthetic	Self-supervised	Err. Change
JointNet	\mathcal{L}_1 Joint Angle Err. ($^\circ$)	9.68	8.83	\downarrow 8.8%
RotationNet	\mathcal{L}_1 Euler Rotation Angle Err. ($^\circ$)	2.94	2.01	\downarrow 31.6%
KeypointNet	Keypoint Euclidean Dis. (pixel)	2.38	1.97	\downarrow 17.2%
DepthNet	\mathcal{L}_1 Root Depth Err. (mm)	14.1	11.8	\downarrow 16.3%

S4. Error Analysis

We further provide error analysis of each module across training stages (training on synthetic data and self-supervised training on real data) on Panda 3CAM-AK in Tab. 11. Errors of the 4 modules (JointNet, RotationNet, KeypointNet and DepthNet) after each training stage are presented. The results show that our multi-stage training effectively reduces the errors, rather than accumulating them. Moreover, our end-to-end training could mitigate error propagation by jointly training all modules, allowing supervision signals to guide inter-module learning and prediction refinement. This contrasts with traditional RnC and non-differentiable PnP systems [15, 18], where error recovery is limited by the lack of integrated training. These empirical results clearly demonstrate the convergence capability of our approach.

S5. Multi-view Fusion Results

Although our primary focus is on monocular estimation to ensure a fair comparison with previous studies, we also extend our method to a multi-view setting through quantitative prediction fusion on the Panda-Orb subset, which contains RGB sequences of the movement of a Panda robot arm from 27 viewpoints. We ensemble the predictions from the first view with each of the remaining 26 views. Specifically, for each of the remaining 26 views whose predictions are being fused, we average the joint state predictions, transform the camera-to-robot translation predictions of the first view into the current viewpoint using relative camera poses, and average the pose-aligned translation predictions. For simplicity, we do not fuse the camera-to-robot rotation predictions, because the process of averaging two $SO(3)$ rotations has no analytical solution.

This fusion approach results in a 2.8% average increase in AUC \uparrow and a 9.1% decrease in mean ADD \downarrow for the remaining 26 views, demonstrating the usefulness of multi-view information in promoting robustness. Future research could further investigate the integration of multi-view information into the framework and model design.

S6. Failure Cases

We illustrate the typical failure cases of our method for each robot arm in Fig. 9. In the case of Panda, estimation becomes unstable under extremely severe self-occlusion. For the example of Kuka and Baxter, estimation fails due to an extremely severe level of truncation. When only a very small number of keypoints are visible, the estimations are generally more error-prone, as also evidenced in Tab. 8.




Fig. 9: Typical failure cases.