

Deterministic Object Pose Confidence Region Estimation

Jinghao Wang^{1,2} Zhang Li^{1,2} Zi Wang^{1,2*} Banglei Guan^{1,2} Yang Shang^{1,2} Qifeng Yu^{1,2}

¹ National University of Defense Technology

² Hunan Provincial Key Laboratory of Image Measurement and Vision Navigation

Abstract

6D pose confidence region estimation has emerged as a critical direction, aiming to perform uncertainty quantification for assessing the reliability of estimated poses. However, current sampling-based approach suffers from critical limitations that severely impede their practical deployment: 1) the sampling speed significantly decreases as the number of samples increases. 2) the derived confidence regions are often excessively large. To address these challenges, we propose a deterministic and efficient method for estimating pose confidence regions. Our approach uses inductive conformal prediction to calibrate the deterministically regressed Gaussian keypoint distributions into 2D keypoint confidence regions. We then leverage the implicit function theorem to propagate these keypoint confidence regions directly into 6D pose confidence regions. This method avoids the inefficiency and inflated region sizes associated with sampling and ensembling. It provides compact confidence regions that cover the ground-truth poses with a user-defined confidence level. Experimental results on the LineMOD Occlusion and SPEED datasets show that our method achieves higher pose estimation accuracy with reduced computational time. For the same coverage rate, our method yields significantly smaller confidence region volumes, reducing them by up to 99.9% for rotations and 99.8% for translations. The code will be available soon.

1. Introduction

Determining the 6D pose of an object from an RGB image is a fundamental task in computer vision, with extensive applications in autonomous driving [1], robotic manipulation [2], augmented reality [3], and space robotics [4]. The majority of pose estimation research [5–11] primarily focuses on delivering a single optimal pose. It is often for applications that tolerate occasional inaccuracies. However, it underperforms where high reliability and precise uncertainty quantification (UQ) are essential. Because visual ambiguities

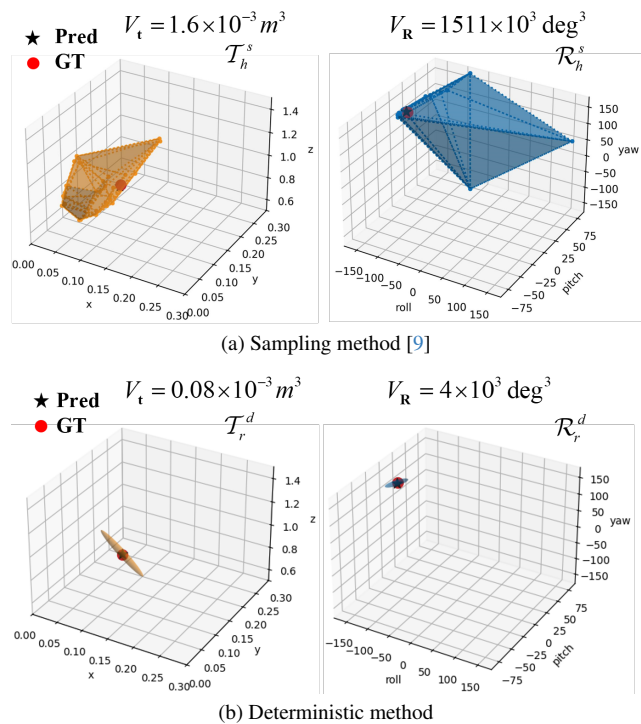


Figure 1. Narrower confidence regions of our deterministic approach compared to the sampling method [9]. (a) A heatmap-guided keypoints are established for 2D confidence regions \mathcal{K}^h via ICP, with 6D pose confidence regions \mathcal{T}_h^s and \mathcal{R}_h^s inferred through sampling. (b) We directly regress keypoint Gaussian distributions and generated \mathcal{K}^r by ICP. To avoid inefficiency and broadened regions due to sampling, we apply a Jacobian matrix from the IFT to propagate \mathcal{K}^r into the 6D pose.

undermine the accuracy of even the best algorithms, leading to disasters in safety-critical scenarios.

Driven by the need for safety and reliability in real-world applications, UQ has been a critical research focus in robotics for decades. UQ is vital for ensuring successful object grasping and prevents collisions, and it can be categorized as either a probabilistic distribution of outcomes or a scalar confidence [12]. However, both distribution estimation [13–16] and confidence estimation [17, 18] are not always reliable and often fail to provide statistically guaran-

*Corresponding author. wangzi16@nudt.edu.cn

teed confidence regions. Additionally, evaluating predictive uncertainty, especially in 6D pose estimation, is inherently challenging due to the lack of ground truth uncertainty.

A practical approach involves deriving 6D pose prediction confidence regions, which are designed to cover the ground truth poses with a user-specified probability. As depicted in Fig. 1, the method proposed by Yang et al. [9] uses inductive conformal prediction (ICP) [19] to conformalize heatmaps into keypoint confidence regions \mathcal{K}^h . ICP provides distribution-free, finite-sample guarantees by constructing confidence regions that adapt based on the non-conformity between predictions and observations. By applying ICP to keypoint detection, we can specify a threshold to generate a confidence region that covers the true position with a certain probability. Subsequently, a sampling-based method constructs 6D pose confidence regions $\mathcal{C}_h^s = \{\mathcal{R}_h^s, \mathcal{T}_h^s\}$, consisting of confidence regions for rotation (\mathcal{R}_h^s) and translation (\mathcal{T}_h^s). For the confidence region, the superscripts s denote the sampling approach, while the subscripts h represent the heatmap method. Although \mathcal{C}_h^s can cover the ground truth, the sampling process increases the confidence region volume and reduces computational efficiency. Furthermore, common metrics for pose estimation, such as reprojection, $5^\circ/5\text{cm}$, ADD(-S) [20], VSD, MSSD, and MSPD [21], assess correctness based on whether errors fall below specific thresholds. These metrics focus solely on measurement accuracy but do not account for the estimation regions' size.

To address these issues, we present a deterministic method for estimating pose confidence regions as shown in Fig. 1. We incorporate uncertainty into the loss function to efficiently regress keypoint Gaussian distributions. Utilizing ICP, these distributions are calibrated into keypoint confidence regions \mathcal{K}^r . To avoid expanding the confidence region due to low-quality sampled poses, we deterministically propagate \mathcal{K}^r into the pose confidence region \mathcal{C}_r^d using the Jacobian derived from the implicit function theorem (IFT). The superscripts d denote the deterministical approach, while the subscripts r represent the regress method. Additionally, we introduce metrics for evaluating measurement uncertainty, including coverage rates and radii for 2D keypoint confidence regions, as well as the coverage rates and volumes of 6D pose confidence region. These metrics provide a more comprehensive assessment of the estimated pose. On the SPEED [22] and LMO [23] datasets, our method reduces the pose confidence region volume by 63.8%, 99.9% for rotations and 92.2%, 99.8% for translations. In summary, our work makes the following contributions:

1. We present an ICP-based method for predicting keypoint confidence regions, which relies on deterministic regression of Gaussian keypoint distributions.
2. Using IFT, we propagate the keypoint confidence regions

directly into the 6D pose, while maintaining a pose coverage rate comparable to that of [9].

3. We propose thorough metrics to evaluate the estimated pose confidence region.

2. Related Work

2.1. Single Point Object Pose Estimation

Current pose estimation methods can be categorized into regression-based and correspondence-based approaches. Regression-based methods directly recover object poses from images, either in a coupled [24–26] or decoupled manner [5–8]. They show better computational efficiency but still struggle with the nonlinear nature of rotation representation. Correspondence-based methods establish dense [27, 28] or sparse correspondences [29, 30] between the input image and a 3D model, then use algorithms like PnP [31] for pose estimation. Dense methods offer higher accuracy but at the cost of significant computational overhead, while sparse methods predict predefined keypoints via voting [11], heatmaps [9, 10], or direct regression [30, 32]. Recent approaches integrate differentiable PnP [33–35] into end-to-end pipelines to enhance performance.

While both regression-based and correspondence-based approaches perform well on standard benchmarks, a key limitation, especially in safety-critical applications, is their inability to provide reliable statistical guarantees and pose confidence regions. Inspired by [9], we propose a deterministic 6D pose confidence region estimation approach to address these issues.

Additionally, current 6D pose evaluation metrics primarily emphasize precision assessment through geometric alignment criteria. The reprojection metric assesses keypoint localization accuracy by measuring compliance with predefined pixel-error thresholds, while the $5^\circ/5\text{cm}$ benchmark requires simultaneous satisfaction of angular (5°) and translational (5cm) error bounds. For detailed geometric verification, the ADD(-S) metric computes mean keypoint displacement normalized by object diameter (10% threshold). However, these conventional metrics systematically neglect the critical dimension of measurement uncertainty in pose estimation. To bridge this critical gap, we propose novel quantification metrics for confidence region in Section 4.2, considering both of ground-truth coverage probability and volumetric compactness.

2.2. Uncertainty Quantification of 6D Pose

The increasing application of DNNs in safety-critical fields has heightened the demand for reliable uncertainty quantification (UQ) methods. These methods primarily address two types of uncertainties: aleatoric uncertainty, which arises from data noise, and epistemic uncertainty, associated with model parameters [36]. Both types of uncertainties have

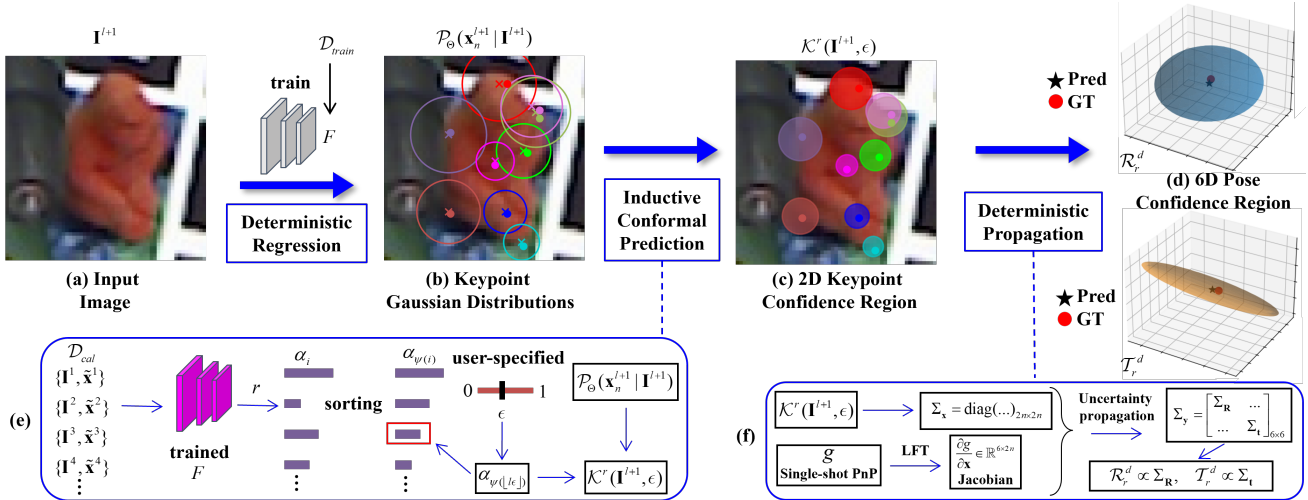


Figure 2. Given an input image, we deterministically regress semantic keypoints (crosses) along with their uncertainties (hollow circles) to generate 2D keypoint confidence regions (translucent circles). These regions are calibrated to ensure coverage of the ground truth keypoints (solid dots), achieving, e.g., 90% coverage using inductive conformal prediction. We then apply the PnP algorithm once to predict the 6D pose (pentagram). Through IFT, we propagate 2D into 6D confidence regions around the predicted pose, maintaining the same coverage rate for the ground truth (red dots).

been estimated across various visual applications, including semantic segmentation [37], optical flow [38], keypoint detection [39, 40], and object detection [41]. In this study, we focus specifically on aleatoric uncertainty.

There are three main approaches to UQ in deep learning. Deep Ensembles train multiple networks, aggregating their outputs through diverse initializations and training data [12]. MC-Dropout, grounded in variational inference [42], approximates Bayesian posterior distributions using multiple stochastic forward passes with active dropout layers. Both methods face scalability challenges due to high computational and memory requirements. Another approach is Direct Modeling, where a probability distribution over network outputs is assumed, and the network’s output layers directly estimate the parameters of this distribution.

By applying Direct Modeling approaches, recent advances in 6D pose estimation have systematically adopted the Bingham distribution for modeling rotational uncertainty. Gilitschenski et al. [13] introduced a differentiable Bingham loss function, enabling direct learning of rotation distributions. Building on this, Okorn et al. [14] demonstrated effective modeling of asymmetric object rotations through isotropic Bingham parametrization, while Deng et al. [15] proposed Deep Bingham Networks to account for a family of pose hypotheses. However, these methods do not provide confidence regions for 6D pose estimation.

The most related work to ours is by Yang et al. [9]. It uses ICP for keypoint confidence region prediction and a sampling-based method to propagate uncertainty to 6D pose. However, sampling methods are computationally in-

efficient and increase confidence region volume. The GPU-accelerated sampling by Gao et al. [43] merely re-samples the output of [9], failing to reduce time while still producing conservative confidence regions. So we propose a deterministic approach using the IFT to propagate keypoint confidence regions directly to pose confidence regions.

3. Method

As depicted in Fig. 2, we present a deterministic method for 6D pose confidence region estimation. Starting with an input image, we regress the Gaussian distributions of keypoints (Sec. 3.1) to create keypoint confidence regions (Sec. 3.2) that ensure a specified coverage of ground truth keypoints. We then apply the single-shot PnP to predict the 6D pose and use the IFT to propagate the keypoint confidence regions into 6D pose (Sec. 3.3).

3.1. Keypoint Deterministic Regression

Heatmap-based methods, while accurate, require significant computational and storage resources, limiting their flexibility [44]. To address this, we adopt an efficient and deterministic regression-based keypoint detection approach [45], which directly maps the input image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ to a set of N Gaussian distributions, $\mathcal{P}_\Theta(\mathbf{x}_n | \mathbf{I})$, where $n = 1 \dots N$ and N is the total number of predefined semantic keypoints. Each distribution $\mathcal{P}_\Theta(\mathbf{x}_n | \mathbf{I})$ represents the probability of the true keypoint $\tilde{\mathbf{x}}_n \in \mathbb{R}^2$ being at position $\mathbf{x}_n \in \mathbb{R}^2$. To handle the aleatoric uncertainty, the variance term is integrated into the training loss, enabling the unsupervised generation of these distributions. The objective is to optimize model

parameters Θ by maximizing the likelihood of annotated keypoints. This approach helps account for ambiguities in labeled data, due to perspective variations and occlusion. We utilize the Negative Log-Likelihood (NLL) loss function to quantify the divergence between the predicted distributions and the ground-truth Dirichlet distribution [46].

$$\mathcal{L}_{NLL} = - \sum_{n=1}^N \log \mathcal{P}_{\Theta}(\mathbf{x}_n | \mathbf{I}) \Big|_{\mathbf{x}_n = \tilde{\mathbf{x}}_n} \quad (1)$$

The trained model on the train dataset \mathcal{D}_{train} serves as the prediction function F for conformal prediction, and is applied for keypoint confidence region estimation in the next section. The network architecture follows the design in [45], where the model outputs $F(\mathbf{I})_n$ represent a Gaussian distribution $\mathcal{N}(\mathbf{x}_n, \sigma_n^2)$. Its probability density function (PDF) is parameterized as $\mathcal{P}_{\Theta}(\mathbf{x}_n | \mathbf{I})$. Where \mathbf{x}_n is the mean vector $\mathbf{x}_n \in \mathbb{R}^{2 \times 2}$ is the covariance matrix, both are predicted through multi-layer perceptrons (MLPs).

3.2. Conformal Keypoint Confidence Region

We adopt an ICP framework to predict the conformal keypoint confidence region, denoted by \mathcal{K}^r as shown in Fig. 2 (e). The ICP process for the trained prediction function F involves two key steps: conformal calibration on the calibration dataset $\mathcal{D}_{cal} = \{\mathbf{I}^i, \tilde{\mathbf{x}}^i\}_{i=1}^l$ and conformal prediction on the test dataset \mathcal{D}_{test} . l denotes the total number of samples in \mathcal{D}_{cal} . During calibration, a sequence of nonconformity scores is generated. These scores are applied in the prediction stage to compute \mathcal{K}^r . Given a new sample $(\mathbf{I}^{l+1}, \tilde{\mathbf{x}}^{l+1}) \in \mathcal{D}_{test}$ that satisfies the exchangeability condition [9], ICP predicts a confidence region $\mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon)$, parameterized by a user-specified error rate ϵ , such that:

$$\mathbb{P}[\tilde{\mathbf{x}}^{l+1} \in \mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon)] \geq 1 - \epsilon, \quad (2)$$

This implies that the prediction region $\mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon)$ for the new sample covers the true keypoint location $\tilde{\mathbf{x}}^{l+1}$ with a probability of at least $1 - \epsilon$.

The ICP framework addresses a bi-objective problem: while ensuring the coverage condition stated Eq. (2), the goal is to minimize the size of the prediction region $\mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon)$. By balancing the above two objectives, ICP provides a principled approach to constructing confidence regions that are both reliable and as tight as possible.

The most critical aspect of ICP conformal calibration is designing the nonconformity function, which is denoted by r . It is constructed to quantify the discrepancy between the new sample and the training dataset \mathcal{D}_{train} . The neural network trained using \mathcal{D}_{train} is referred to as the function F . Hence, the nonconformity function can be denoted as:

$$r(\tilde{\mathbf{x}}, F(\mathbf{I})) = \max \{\phi(\tilde{\mathbf{x}}_n, F(\mathbf{I})_n)\}_{n=1}^N \quad (3)$$

where the function ϕ computes the nonconformity score for each keypoint, with the highest score representing the overall score. For the n -th keypoint, ϕ is defined as follows:

$$\phi(\tilde{\mathbf{x}}_n, F(\mathbf{I})_n) = \frac{\|\tilde{\mathbf{x}}_n - \mathbf{x}_n\|}{\det(\sigma_n)} \quad (4)$$

For the i -th sample in the calibration set \mathcal{D}_{cal} , the nonconformity score, calculated using Eq. (3), is denoted as $\alpha_i = r(\tilde{\mathbf{x}}^i, F(\mathbf{I}^i))$ for $i = 1, \dots, l$. These scores are then ranked in descending order as $\alpha_{\psi(1)} \geq \dots \geq \alpha_{\psi(l)}$, where $\psi(\cdot)$ denotes the ranking index. Next, given $\epsilon \in (0, 1)$, the $\lfloor l\epsilon \rfloor$ -th largest nonconformity score is denoted as $\alpha_{\psi(\lfloor l\epsilon \rfloor)}$. Using this score and considering Eq. (2) and Eq. (3), the keypoint confidence regions $\mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon)$ of the newly introduced test image \mathbf{I}^{l+1} from \mathcal{D}_{test} is derived as:

$$\begin{aligned} \mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon) &= \{\mathbf{x}' \mid \max\{\phi(\mathbf{x}'_n, F(\mathbf{I}^{l+1})_n)\}_{n=1}^N \leq \alpha_{\psi(\lfloor l\epsilon \rfloor)}\} \\ &= \{\mathbf{x}' \mid \phi(\mathbf{x}'_n, F(\mathbf{I}^{l+1})_n) \leq \alpha_{\psi(\lfloor l\epsilon \rfloor)}, \forall n\}, \end{aligned} \quad (5)$$

Note that $\max\{\phi_1, \dots, \phi_n, \dots, \phi_N\} \leq \alpha$ holds if and only if $\phi_n \leq \alpha$ for every n . We then substitute Eq. (4) into Eq. (5) and arrive at the desired result.

$$\mathcal{K}^r(\mathbf{I}^{l+1}, \epsilon) = \{\mathbf{x}' \mid \|\mathbf{x}'_n - \mathbf{x}_n\| \leq \det(\sigma_n) \alpha_{\psi(\lfloor l\epsilon \rfloor)}, \forall n\} \quad (6)$$

Eq. (6) describes an ball region for the k -th keypoint centered at \mathbf{x}_n . The area of the ball is proportional to the determinant of the covariance matrix $\det(\sigma_n)$ and a scaling factor $\alpha_{\psi(\lfloor l\epsilon \rfloor)}$. Note that the size of the confidence region increases when the keypoint's Gaussian distributions are uncertain. It can be indicated by a large determinant in the covariance matrix, and the keypoint regression perform poorly on \mathcal{D}_{cal} , leading to a large $\alpha_{\psi(\lfloor l\epsilon \rfloor)}$.

The distinction between pose confidence region \mathcal{K}^r proposed in this paper and those by Yang et al. \mathcal{K}^h [9] lies in our introduction of a novel nonconformity function ϕ related to $\mathcal{P}_{\Theta}(\mathbf{x}_n | \mathbf{I})$ in Sec. 3.1. As shown in Tab. 1, \mathcal{K}^r is generated faster compared to \mathcal{K}^h . Additionally, \mathcal{K}^r is more stable, avoiding the generation of excessively large confidence regions compared with \mathcal{K}^h (e.g., keypoint 4 of object 2 and keypoints 1 and 9 of object 8 in Fig. 6).

3.3. Deterministic Pose Confidence Region

The sampling-based approach [9] for propagating keypoint confidence regions (\mathcal{K}^h) to a pose confidence region (\mathcal{C}_h^s) is inefficient and presents several challenges. Randomly sampled keypoints are not only computationally expensive but also struggles to find sufficient valid samples in difficult cases. Moreover, the low accuracy of the P3P algorithm leads to pose samples with large deviations from the ground truth, unnecessarily inflating the volume of the final confidence region \mathcal{C}_h^s .

To address the above issues, as shown in Fig. 2 (f), we calculate the Jacobians of 6D pose with respect to the 2D keypoints based on implicit function theorem (IFT) [47]. Then we deterministically propagate \mathcal{K}^r to \mathcal{C}_r^d using the uncertainty propagation theorem. The Jacobian is difficult to compute due to the non-linear relationship in perspective geometry. The IFT implicitly computes Jacobians by leveraging geometric constraints, avoiding the need for an explicit solution. Following [34], we define a single-shot PnP solver, which outputs the pose $\mathbf{y} = [y_1, \dots, y_m]^T$, denoted as g :

$$\mathbf{y} = g(\mathbf{x}, \mathbf{z}, \mathbf{K}) \quad (7)$$

where $\mathbf{x} \in \mathbb{R}^{2 \times N}$ and $\mathbf{z} \in \mathbb{R}^{3 \times N}$ represent N 2D-3D correspondences, and $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ is the camera intrinsic matrix. The objective function of g is defined as follow:

$$O(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{K}) = \sum_{n=1}^N \|\mathbf{r}_n\|_2^2 \quad (8)$$

where $\mathbf{r}_n = \mathbf{x}_n - \pi_n$ is the reprojection error of the n -th correspondence. $\pi_n = \Pi(\mathbf{z}_n | \mathbf{y}, \mathbf{K})$ is the reprojected points, calculated by the projective function Π .

A stationary condition is formulated by computing the first-order derivative of O with respect to \mathbf{y} .

$$\frac{\partial O(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{K})}{\partial \mathbf{y}} = 0 \quad (9)$$

The IFT constraint function f is constructed based on Eq. (9).

$$f(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{K}) = [f_1, \dots, f_m]^T \quad (10)$$

The dimensionality of the pose, i.e., m , depends on the parameterization of the rotation $SO(3)$. Chen et al. [34] use the less intuitive axis-angle representation $m = 6$. For better user comprehension, we adopt the Euler angle in this work and $m = 6$. For each parameter of the pose representation, the constraint function f_j is expressed as:

$$f_j = \frac{\partial O(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{K})}{\partial y_j} = 2 \sum_{n=1}^N \langle \mathbf{r}_n, -2 \frac{\partial \pi_n}{\partial y_j} \rangle \quad (11)$$

Furthermore, based on the IFT, we can apply the constraint function f to compute $\frac{\partial g}{\partial \mathbf{x}}$.

$$\frac{\partial g}{\partial \mathbf{x}} = - \left[\frac{\partial f}{\partial \mathbf{y}} \right]^{-1} \left[\frac{\partial f}{\partial \mathbf{x}} \right] \quad (12)$$

Subsequently, the covariance matrix $\Sigma_{\mathbf{x}} \in \mathbb{R}^{2N \times 2N}$ representing \mathcal{K}^r is directly propagated to $\Sigma_{\mathbf{y}} \in \mathbb{R}^{6 \times 6}$ standing for \mathcal{C}_r^d using the uncertainty propagation.

$$\Sigma_{\mathbf{y}} = \frac{\partial g}{\partial \mathbf{x}} \Sigma_{\mathbf{x}} \frac{\partial g}{\partial \mathbf{x}}^T \quad (13)$$

The covariance matrices for the Euler angles and the translation vector are denoted as $\Sigma_{\mathbf{R}} = \Sigma_{\mathbf{y}}(1 : 3, 1 : 3)$ and $\Sigma_{\mathbf{t}} = \Sigma_{\mathbf{y}}(4 : 6, 4 : 6)$, respectively.

In summary, as shown in Fig. 2, our method eliminates the need for inefficient sampling and integration processes by developing a deterministic method for tighter 6D pose confidence regions $\mathcal{C}_r^d = \{\mathcal{R}_r^d, \mathcal{T}_r^d\}$.

$$\mathcal{R}_r^d = \{\mathbf{x} \in \mathbb{R}^3 : (\mathbf{x} - \mathbf{y}_{1:3})^\top \Sigma_{\mathbf{R}}^{-1} (\mathbf{x} - \mathbf{y}_{1:3}) \leq 1\} \quad (14)$$

$$\mathcal{T}_r^d = \{\mathbf{x} \in \mathbb{R}^3 : (\mathbf{x} - \mathbf{y}_{4:6})^\top \Sigma_{\mathbf{t}}^{-1} (\mathbf{x} - \mathbf{y}_{4:6}) \leq 1\} \quad (15)$$

In contrast to the sample-based approach \mathcal{C}_h^s , we apply the IFT to derive parameterized, ellipsoidal confidence regions for both rotation and translation.

4. Experiment

4.1. Datasets and implementation details

We conduct experiments on the LMO [23] and the SPEED [22] datasets. Both are designed for 6D pose estimation task. LMO includes photorealistic rendered training images of randomly cluttered scenes. We split the dataset following [9], with 200 images allocated to \mathcal{D}_{cal} and 1,214 images to \mathcal{D}_{test} . LMO includes 8 labeled objects with significant occlusion. SPEED is used to test the safety-critical applicability of the proposed pose confidence region. To ensure a rigorous analysis, we apply a sixfold cross-validation (CV) method, dividing the 12,000 simulated images into six subsets. Five-sixths are designated as \mathcal{D}_{train} , while one-sixth is split equally into \mathcal{D}_{cal} and \mathcal{D}_{test} .

The runtimes are evaluated on a workstation equipped with Nvidia A6000 GPUs. The network architecture is based on [45]. The regression model undergoes 96 epochs of training, utilizing a pre-trained model based on the COCO dataset [48]. The training employs the AdamW optimizer [49]. The fixed learning rate is set as 3×10^{-5} . The upper bound for the keypoint confidence region radius is set to the image diagonal length, as this value sufficiently covers the entire image, making larger radii unnecessary.

4.2. Evaluation metrics

We adopt keypoint reprojection accuracy [11], denoted by Acc, to evaluate the accuracy. However, the existing metrics for pose estimation mainly focus on accuracy, neglecting the evaluation of confidence regions. Hence, we propose 4 metrics beyond accuracy, including the size of the confidence regions and coverage rate.

2D keypoint coverage rate: The validity of the region is assessed by the probability of all ground truth keypoints $\tilde{\mathbf{x}}$ falling within \mathcal{K}^h or \mathcal{K}^r aiming to align with $1 - \epsilon$, as stated in (2). For example, the coverage rate η^{kpt} of \mathcal{K}^h for

\mathcal{D}_{test} containing K images is:

$$\eta^{kpt} = \frac{1}{K} \sum_{k=1}^K \mathbb{I}(\tilde{\mathbf{x}}^k \in \mathcal{K}^h(\mathbf{I}^k, \epsilon)) \quad (16)$$

2D keypoint confidence region radius: We denote the region size by the keypoint confidence region radius, $\det(\sigma_n) \alpha_{\psi(\lfloor l\epsilon \rfloor)}$ in Eq. (6). The greater keypoints' uncertainty, the larger the confidence region radius.

6D pose coverage rate: Yang et al. [9] employ 2D points sampled within \mathcal{K}^h and use the PnP method to obtain pose samples, the convex hull of these samples formed a 6D pose confidence region $\mathcal{C}_h^s = \{\mathcal{R}_h^s, \mathcal{T}_h^s\}$, as shown in Fig. 1. In contrast, we deterministically propagate \mathcal{K}^r to ellipsoidal pose confidence region $\mathcal{C}_r^d = \{\mathcal{R}_r^d, \mathcal{T}_r^d\}$ based on IFT. For both Euler angles and translation vectors, we calculate the coverage rate of the ground truth. Taking \mathcal{C}_h^s as an example, the coverage rate for the Euler angles $\eta^{\mathbf{R}}$ is as follows:

$$\eta^{\mathbf{R}} = \frac{1}{K} \sum_{k=1}^K \mathbb{I}(\mathbf{R} \in \mathcal{R}_h^s(\mathbf{I}^k, \epsilon)) \quad (17)$$

For translation coverage rate $\eta^{\mathbf{t}}$ follows a similar form.

6D pose confidence region volume: However, if achieving a high coverage rate requires an excessively large confidence region, its practical utility may be greatly diminished. Taking the Euler angles of rotation as an example, if the region spans $360^\circ \times 180^\circ \times 360^\circ$, it will inevitably contain any true Euler angle. Therefore, we also compute the volume of the pose confidence region as a metric. The volume of \mathcal{C}_r^d is:

$$V_{\mathbf{R}} = \frac{4}{3}\pi\sqrt{\det(\Sigma_{\mathbf{R}})}, V_{\mathbf{t}} = \frac{4}{3}\pi\sqrt{\det(\Sigma_{\mathbf{t}})} \quad (18)$$

We set the volume threshold $\tau_{\mathbf{R}} = 90^3 \text{ deg}^3$ for $V_{\mathbf{R}}$, and $\tau_{\mathbf{t}} = 1^3 \text{ m}^3$ for $V_{\mathbf{t}}$. Any sample-based confidence region \mathcal{C}_h^s whose volume—defined as the convex hull of its pose samples—exceeds a predefined threshold is excluded from the mean volume calculation and simultaneously considered a pose coverage failure.

	Ours		[9]	
	\mathcal{K}^r	\mathcal{C}_r^d	\mathcal{K}^h	\mathcal{C}_h^s
LMO [23]	0.0038	0.0361	0.0076	0.0550
SPEED [22]	0.0032	0.0358	0.0064	0.0521

Table 1. Time consumption for keypoint and pose confidence region estimation (unit: s).

4.3. Experimental Results

Accuracy and speed: We provide a concise comparison of Acc with baseline methods in Tab. 2 (a). While prior methods such as [9] rely on heatmap calibration and occlusion-

Acc	Ours \mathcal{K}^r	[9] \mathcal{K}^h		[11]
		$\epsilon = 0.1$	$\epsilon = 0.4$	
LMO mean	74.45	67.33	70.71	61.06
SPEED [22]	97.09	57.80	57.40	57.46

(a) Acc of baseline methods and our approach

η^{kpt}	Ours \mathcal{K}^r		[9] \mathcal{K}^h	
	$\epsilon = 0.1$	$\epsilon = 0.4$	$\epsilon = 0.1$	$\epsilon = 0.4$
LMO mean	90.65	60.38	91.16	63.38
SPEED [22]	89.66	61.25	88.88	62.64

(b) η^{kpt} of baseline methods and our approach

Table 2. 2D keypoint confidence region results

LMO Objects	Ours		[9] + Samp.		[9] + Det.	
	\mathcal{T}_r^d	\mathcal{R}_r^d	\mathcal{T}_h^s	\mathcal{R}_h^s	\mathcal{T}_h^d	\mathcal{R}_h^d
1	70.52	91.26	97.26	N/A	76.88	93.38
2	88.73	89.98	99.25	N/A	99.59	98.18
3	77.28	87.55	98.29	N/A	88.02	90.59
4	85.09	96.62	97.12	N/A	90.28	98.85
5	97.18	79.70	99.81	N/A	90.13	76.41
6	98.36	1.46	77.81	N/A	98.63	1.37
7	79.73	87.76	98.64	N/A	89.99	91.84
8	69.01	86.94	99.92	N/A	98.51	98.02
mean	83.24	77.66	96.61	N/A	91.50	81.08
SPEED	86.69	88.81	6.40	N/A	87.10	90.92

Table 3. $\eta^{\mathbf{R}}$ and $\eta^{\mathbf{t}}$ of baselines and ours with $\epsilon = 0.1$. 'N/A': The volume in all regions exceeds the threshold.

aware sampling to improve accuracy in challenging scenarios, their improvements diminish in occlusion-free environments (e.g., on SPEED [22]). Our method universally surpasses [9, 11] across both cases with various ϵ . The comprehensive results for the 8 LMO objects are presented in supplementary Tab. 5.

For heatmap-based keypoint detection and sampling-based confidence region prediction, increasing the size of heatmaps and the number of samples significantly raises inference and sampling time. As shown in Tab. 1, our regression-based keypoint detection and deterministic confidence region prediction reduce time consumption by 50% and 31.3% on SPEED, and by 50% and 34.4% on LMO, compared to [9]. In [9], Faster R-CNN and the Stacked Hourglass are used for bounding box detection and keypoint detection, respectively. For a fair comparison, we use the same bounding box results as [9]. The Stacked Hourglass has 26.4 million parameters and 26.77 GFlops, while our method has only 11.1 million parameters and 7.19 GFlops.

2D Keypoint Confidence Region: After applying ICP, the confidence region coverage rates of \mathcal{K}^h and \mathcal{K}^r closely align with $1 - \epsilon$ in both SPEED and LMO datasets, as shown in Tab. 2 (b). The comprehensive results for the each object in the LMO dataset are presented in supplementary Tab. 6. However, the 2D keypoint coverage can not fully reflect the 6D pose coverage. Therefore, we proceed to evaluate the confidence region coverage and volume for the 6D pose.

6D Pose Confidence Region Coverage Rates: To compare the sampling-based method with our deterministic approach, Tab. 3 presents the 6D pose coverage rates for three pose confidence regions: regions obtained by direct propa-

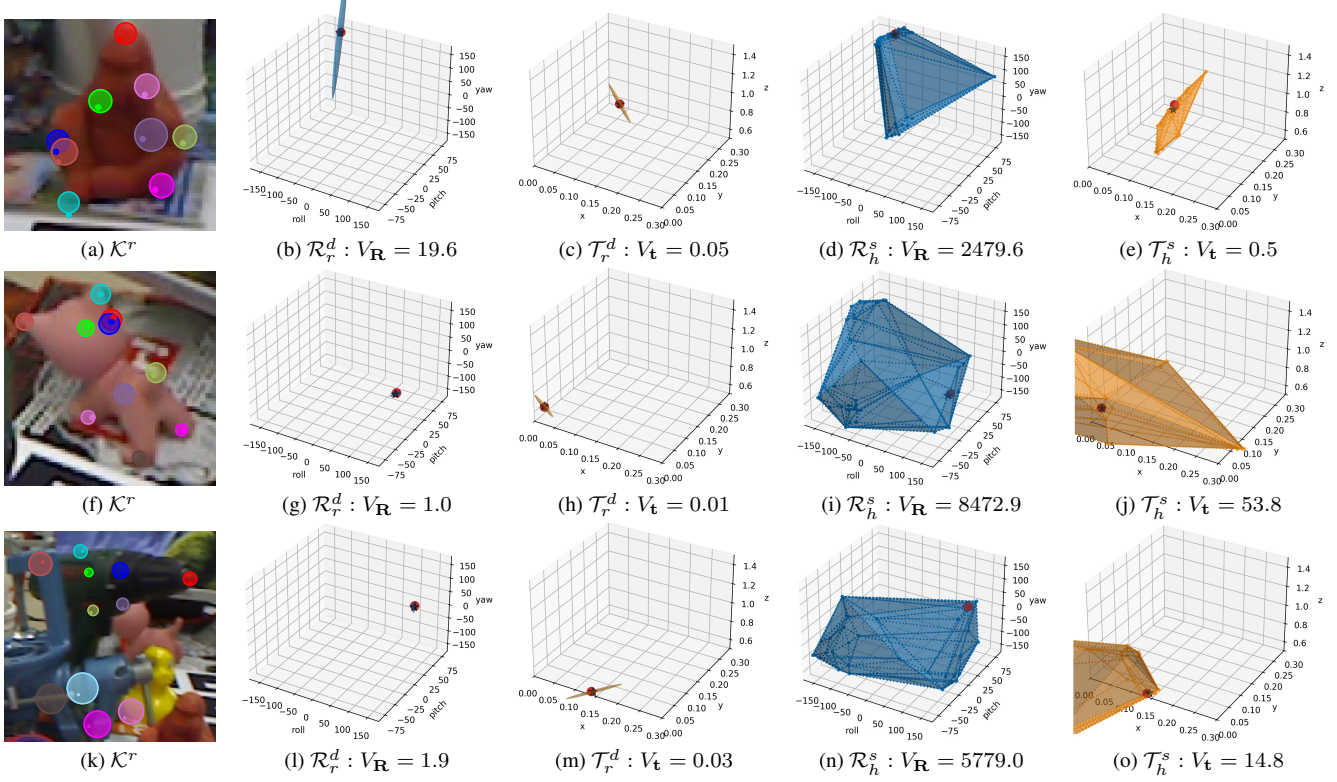


Figure 3. Visualization of the 6D confidence region for objects in LMO. Our \mathcal{K}^r can cover the true keypoint within a user-defined ϵ , while the volumes of \mathcal{R}_r^d and \mathcal{T}_r^d are significantly smaller than that of \mathcal{R}_h^s and \mathcal{T}_h^s .

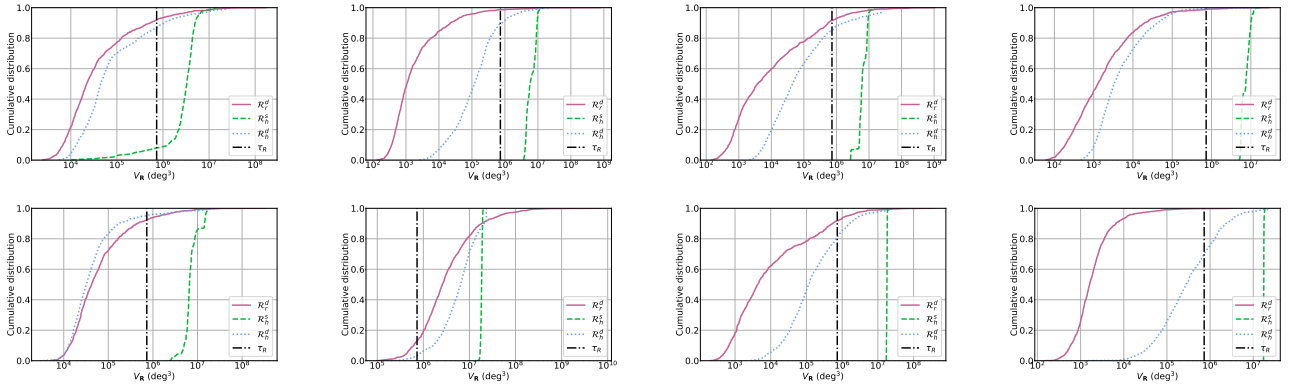


Figure 4. Cumulative distribution function of the V_R with $\epsilon = 0.1$ for the 8 objects in the LMO, the unit for V_R is 10^3 deg^3 . A CDF curve that is closer to the upper-left corner indicates a smaller V_R .

gation (\mathcal{C}_h^d) and sampling (\mathcal{C}_h^s) from \mathcal{K}^h generated by [9], and our deterministically propagated \mathcal{C}_r^d from \mathcal{K}^h . Our single-shot PnP algorithm for pose estimation (Sec. 6.3) does not utilize sampling or the RANSAC method. For translation in the LMO dataset, although \mathcal{T}_h^s achieves a higher coverage rate compared to \mathcal{T}_r^d (96.61% vs. 83.24%), its volume is significantly larger, as shown in Tab. 4 (109.9 dm^3 vs. 8.6 dm^3). On the SPEED dataset, the larger range

of translation causes more \mathcal{T}_h^s volumes to exceed the threshold, thereby lowering its pose coverage rate.

6D Pose Confidence Region Volume: As shown in Tab. 4, replacing the sampling method of [9] with our deterministic propagation reduces the volume of \mathcal{R}_h^d below the required threshold while maintaining coverage. While the heatmap-based approach of [9] performs comparably on the LMO, it yields significantly larger rotation (V_R) and trans-




Figure 5. A boxplot analysis of baseline and proposed method for V_t on the LMO dataset. The boxplot for T_r^d exhibits consistently lower values across all key statistical measures compared to T_h^s . After applying T_h^d , V_t also decreases.

LMO Objects	Ours				[9] + Samp.				[9] + Det.			
	T_r^d	Out	\mathcal{R}_r^d	Out	T_h^s	Out	\mathcal{R}_h^s	Out	T_h^d	Out	\mathcal{R}_h^d	Out
1	0.9	0	28.96	69	12.6	0	252.2	1041	0.5	0	50.3	64
2	0.7	0	9.9	9	28.3	0	N/A	1207	0.9	0	45.0	12
3	2.8	3	50.7	113	66.8	0	N/A	1052	1.4	0	42.9	73
4	1.5	0	6.5	12	52.0	0	N/A	1214	0.1	0	3.5	2
5	0.4	0	52.8	43	63.1	0	N/A	1064	0.2	0	24.7	24
6	57.6	11	514.4	1055	82.2	1	N/A	1095	59.9	10	515.7	1057
7	4.8	0	62.9	98	252.4	0	N/A	809	0.8	0	67.5	56
8	0.03	0	4.6	3	322.2	0	N/A	1210	0.3	0	51.6	9
mean	8.6	2	91.3	175	109.9	0	252.2	1070	8.0	1	100.1	162
SPEED	0.7	0	0.2	0	287.8	920	210.6	898	73.4	50	4.4	18

Table 4. V_t and V_R with $\epsilon = 0.1$: The unit for the rotational and translational confidence region are 10^3 deg^3 and $10^{-3} m^3$. 'N/A' means all confidence regions are over the threshold, while 'Out' is the count of images that are.

lation (V_t) volumes on the occlusion-free SPEED. Consequently, our method achieves a significant reduction in both V_R and V_t compared to [9] (Fig. 3). Other visualizations can be found in the supplementary material. The largest pose confidence region volume is observed with \mathcal{C}_h^s , where most \mathcal{R}_h^s volumes exceed the threshold. The mean volume of T_h^s remains within the threshold but is still the highest, indicating that sampling-based confidence regions fail to cover the ground truth pose within a narrow range.

Transitioning from \mathcal{C}_h^s to \mathcal{C}_r^d reduces V_R by 63.8% and 99.9%, and V_t by 92.2% and 99.8% across the two datasets. In the occlusion-free SPEED, our proposed method reduces the volumes of T_r^d and \mathcal{R}_r^d to as low as $0.7 \times 10^{-3} m^3$ and $0.2 \times 10^3 \text{ deg}^3$. To elucidate the statistical disparities in region volumes across diverse methods, we visualize the CDF curves of V_R (Fig. 4) and the boxplot of V_t (Fig. 5) within the LMO. In the CDF curve, \mathcal{R}_r^d is closer to the top-left corner, while most \mathcal{R}_h^s exceeds τ_R , validating the efficacy of our deterministic propagation method in reducing region size. Furthermore, as shown in Fig. 5, V_t of T_r^d is reduced by several orders of magnitude compared to T_h^s .

5. Conclusions

Our framework introduces an efficient pipeline for estimating compact 6D pose confidence regions that achieve real-time operational capability. Crucially, we develop novel

evaluation metrics that consider both ground-truth coverage probability and volumetric compactness of the confidence regions. Benchmark evaluations on the LineMOD Occlusion and SPEED datasets demonstrate significant improvements: 1) 33% faster inference speeds compared to the baseline, 2) up to 99% decrease in confidence region volume, while maintaining comparable coverage rates across test scenarios. But our Gaussian pose assumption is a key limitation, as the non-linear 2D-to-6D transformation challenges its coverage guarantees, which we will address by exploring more suitable distributions.

Acknowledgements

This paper is supported in part by the National Natural Science Foundation of China under Nos. 12472189 and 12302252, in part by the Research Program of National University of Defense Technology, China (No. ZK24-31), in part by the Science and Technology Innovation Program of Hunan Province under No. 2022RC1196, in part by the Natural Science Foundation of Hunan Province under Nos. 2025JJ60412 and 2025JJ60054, and in part by the Postgraduate Scientific Research Innovation Project of Hunan Province under Grant XJQY2025054.

References

- [1] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8:58443–58469, 2020. [1](#)
- [2] Lucas Manuelli, Wei Gao, Peter Florence, and Russ Tedrake. kpm: Keypoint affordances for category-level robotic manipulation. In *The International Symposium of Robotics Research*, volume 20, pages 132–157. Springer, 2019. [1](#)
- [3] Yongzhi Su, Jason Rambach, Nareg Minaskan, Paul Lesur, Alain Pagani, and Didier Stricker. Deep multi-state object pose estimation for augmented reality assembly. In *IEEE International Symposium on Mixed and Augmented Reality Adjunct*, pages 222–227, 2019. [1](#)
- [4] Bo Chen, Jiewei Cao, Alvaro Parra, and Tat-Jun Chin. Satellite pose estimation with deep landmark regression and non-linear pose refinement. In *Int. Conf. Comput. Vis. Worksh.*, pages 2816–2824, 2019. [1](#)
- [5] Jimmy Wu, Bolei Zhou, Rebecca Russell, Vincent Kee, Syler Wagner, Mitchell Hebert, Antonio Torralba, and David MS Johnson. Real-time object pose estimation with pose interpreter networks. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 6798–6805. IEEE, 2018. [1](#), [2](#)
- [6] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *Int. Conf. Comput. Vis.*, pages 1521–1529, 2017.
- [7] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Robotics: Science and Systems*, 2018.
- [8] Tsu-Ching Hsiao, Hao-Wei Chen, Hsuan-Kung Yang, and Chun-Yi Lee. Confronting ambiguity in 6d object pose estimation via score-based diffusion on se (3). In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 352–362, 2024. [2](#)
- [9] Heng Yang and Marco Pavone. Object pose estimation with statistical guarantees: Conformal keypoint detection and geometric uncertainty propagation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8947–8958, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [11](#)
- [10] Karl Schmeckpeper, Philip R Osteen, Yufu Wang, Georgios Pavlakos, Kenneth Chaney, Wyatt Jordan, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Semantic keypoint-based pose estimation from single rgb frames. *Field Robotics*, 2(1):147–171, 2022. [2](#)
- [11] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4561–4570, 2019. [1](#), [2](#), [5](#), [6](#), [11](#)
- [12] Guanya Shi, Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, Fabio Ramos, Animashree Anandkumar, and Yuke Zhu. Fast uncertainty quantification for deep object pose estimation. In *IEEE Int. Conf. Robot. Autom.*, pages 5200–5207, 2021. [1](#), [3](#)
- [13] Igor Gilitschenski, Roshni Sahoo, Wilko Schwarting, Alexander Amini, Sertac Karaman, and Daniela Rus. Deep orientation uncertainty learning based on a bingham loss. In *International Conference on Learning Representations*, 2020. [1](#), [3](#)
- [14] Brian Okorn, Mengyun Xu, Martial Hebert, and David Held. Learning orientation distributions for object pose estimation. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 10580–10587, 2020. [3](#)
- [15] Haowen Deng, Mai Bui, Nassir Navab, Leonidas Guibas, Slobodan Ilic, and Tolga Birdal. Deep bingham networks: Dealing with uncertainty and ambiguity in pose estimation. *IJCV*, 130(7):1627–1654, 2022. [3](#)
- [16] Hiroya Sato, Takuya Ikeda, and Koichi Nishiwaki. A probabilistic rotation representation for symmetric shapes with an efficiently computable bingham loss function. In *IEEE Int. Conf. Robot. Autom.*, pages 6923–6929, 2023. [1](#)
- [17] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3338–3347, 2019. [1](#)
- [18] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *2nd Annual Conference on Robot Learning*, volume 87, pages 306–316, 2018. [1](#)
- [19] Harris Papadopoulos. Inductive conformal prediction: Theory and application to neural networks. In *Tools in Artificial Intelligence*, 2008. [2](#)
- [20] Sida Peng, Xiaowei Zhou, Yuan Liu, Haotong Lin, Qixing Huang, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof object pose estimation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(6):3212–3223, 2022. [2](#)
- [21] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, et al. Bop: Benchmark for 6d object pose estimation. In *Eur. Conf. Comput. Vis.*, pages 19–34, 2018. [2](#)
- [22] Mate Kisantal, Sumant Sharma, Tae Ha Park, Dario Izzo, Marcus Mörtens, and Simone D’Amico. Satellite pose estimation challenge: Dataset, competition design, and results. *IEEE Transactions on Aerospace and Electronic Systems*, 56(5):4083–4098, 2020. [2](#), [5](#), [6](#), [11](#)
- [23] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *Eur. Conf. Comput. Vis.*, volume 8690, pages 536–551, 2014. [2](#), [5](#), [6](#), [11](#)
- [24] Fabian Manhardt, Diego Martín Arroyo, Christian Rupprecht, Benjamin Busam, Tolga Birdal, Nassir Navab, and Federico Tombari. Explaining the ambiguity of object detection and 6d pose from visual data. In *Int. Conf. Comput. Vis.*, pages 6840–6849, 2019. [2](#)
- [25] Kilian Kleeberger and Marco F Huber. Single shot 6d object pose estimation. In *IEEE Int. Conf. Robot. Autom.*, pages 6239–6245, 2020.
- [26] Stefan Hein Bengtson, Hampus Åström, Thomas B. Moeslund, Elin A. Topp, and Volker Krueger. Pose estimation from rgb images of highly symmetric objects using a novel

- multi-pose loss and differential rendering. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 4618–4624, 2021. 2
- [27] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *Int. Conf. Comput. Vis.*, pages 7678–7687, 2019. 2
- [28] Tomas Hodan, Daniel Barath, and Jiri Matas. Epos: Estimating 6d pose of objects with symmetries. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11703–11712, 2020. 2
- [29] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-dof object pose from semantic keypoints. In *IEEE Int. Conf. Robot. Autom.*, pages 2011–2018. IEEE, 2017. 2
- [30] Mahdi Rad and Vincent Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *Int. Conf. Comput. Vis.*, pages 3848–3856, 2017. 2
- [31] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 2
- [32] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 292–301, 2018. 2
- [33] Eric Brachmann and Carsten Rother. Learning less is more-6d camera localization via 3d surface regression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4654–4662, 2018. 2
- [34] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8100–8109, 2020. 5
- [35] Shun Iwase, Xingyu Liu, Rawal Khiredkar, Rio Yokota, and Kris M Kitani. Repose: Fast 6d object pose refinement via deep texture rendering. In *Int. Conf. Comput. Vis.*, pages 3303–3312, 2021. 2
- [36] Di Feng, Ali Harakeh, Steven L Waslander, and Klaus Dietmayer. A review and comparative study on probabilistic object detection in autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):9961–9980, 2021. 2
- [37] Janis Postels, Francesco Ferroni, Huseyin Coskun, Nassir Navab, and Federico Tombari. Sampling-free epistemic uncertainty estimation using approximated variance propagation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2931–2940, 2019. 3
- [38] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 652–667, 2018. 3
- [39] Zhenglin Zhou, Huaxia Li, Hong Liu, Nanyang Wang, Gang Yu, and Rongrong Ji. Star loss: Reducing semantic ambiguity in facial landmark detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15475–15484, 2023. 3
- [40] Abhinav Kumar, Tim K Marks, Wenxuan Mou, Ye Wang, Michael Jones, Anoop Cherian, Toshiaki Koike-Akino, Xiaoming Liu, and Chen Feng. Luvli face alignment: Estimating landmarks’ location, uncertainty, and visibility likelihood. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8236–8246, 2020. 3
- [41] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2888–2897, 2019. 3
- [42] Yarín Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016. 3
- [43] Yihui Gao, Yukai Tang, Han Qi, and Heng Yang. CLOSURE: Fast quantification of pose uncertainty sets. In *Robotics: Science and Systems (RSS)*, Delft, Netherlands, jul 2024. 3
- [44] Jiefeng Li, Siyuan Bian, Ailing Zeng, Can Wang, Bo Pang, Wentao Liu, and Cewu Lu. Human pose regression with residual log-likelihood estimation. In *Int. Conf. Comput. Vis.*, pages 11025–11034, 2021. 3
- [45] Jinghao Wang, Yang Li, Zhang Li, Zi Wang, and Qifeng Yu. Monocular satellite pose estimation based on uncertainty estimation and self-assessment. *IEEE Transactions on Aerospace and Electronic Systems*, 2024. 3, 4, 5, 11
- [46] Yee Whye Teh. Dirichlet process. In *Encyclopedia of machine learning and data mining*, pages 361–370. Springer, 2017. 4
- [47] Steven G. Krantz and Harold R. Parks. *The Implicit Function Theorem, History, Theory, and Applications*. Birkhäuser, New York, NY, 2012. 5
- [48] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Eur. Conf. Comput. Vis.*, pages 740–755. Springer, 2014. 5
- [49] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Int. Conf. Learn. Represent.*, 2019. 5