

Fast Uncertainty Quantification for Deep Object Pose Estimation

Guanya Shi^{1,2}, Yifeng Zhu^{1,3}, Jonathan Tremblay¹,
Stan Birchfield¹, Fabio Ramos^{1,4}, Animashree Anandkumar^{1,2}, Yuke Zhu^{1,3}

Abstract—Deep learning-based object pose estimators are often unreliable and overconfident especially when the input image is outside the training domain, for instance, with sim2real transfer. Efficient and robust uncertainty quantification (UQ) in pose estimators is critically needed in many robotic tasks. In this work, we propose a simple, efficient, and plug-and-play UQ method for 6-DoF object pose estimation. We ensemble 2-3 pre-trained models with different neural network architectures and/or training data sources, and compute their average pairwise disagreement against one another to obtain the uncertainty quantification. We propose four disagreement metrics, including a learned metric, and show that the average distance (ADD) is the best learning-free metric and it is only slightly worse than the learned metric, which requires labeled target data. Our method has several advantages compared to the prior art: 1) our method does not require any modification of the training process or the model inputs; and 2) it needs only one forward pass for each model. We evaluate the proposed UQ method on three tasks where our uncertainty quantification yields much stronger correlations with pose estimation errors than the baselines. Moreover, in a real robot grasping task, our method increases the grasping success rate from 35% to 90%. Video and code are available at <https://sites.google.com/view/fastuq>.

I. INTRODUCTION

Robust and efficient estimation of object poses is a fundamental robot perception task that enables robots to make autonomous decisions in unstructured environments. Recent advances in deep learning-based machine perception have substantially improved the accuracies of pose estimation from real-world sensory data [2–6]. The estimated 6-DoF object pose, represented by the translation and rotation in $SE(3)$, serves as a compact and informative state representation for a variety of downstream tasks, such as robot grasping and manipulation [7], human-robot interactions [8], online camera calibration [9], and tele-presence robot control [10].

While these deep learning-based pose estimation models have attained remarkable performance, their applicability in mission-critical robotic tasks is hindered by their brittleness in dealing with inputs that are outside the training domain [11] or with perturbations in the visual observations [12]. These issues are especially relevant in pose estimation with sim2real transfer [2, 13, 14], where deep-learning models trained on synthetic data tend to be often wrong but overconfident on real data. To enhance the reliability of deep pose estimation models in risk-sensitive domains, it is critical to endow these models with the capability of accurately assessing the *uncertainty* of their own predictions and identifying failure cases automatically.

¹NVIDIA, ²California Institute of Technology, ³The University of Texas at Austin, ⁴The University of Sydney.



Fig. 1. (a) Illustration and examples of using an ensemble of heterogeneous models for uncertainty quantification. We calculate the average disagreement of K pose predictions from K different estimators as an estimation of uncertainty. The examples show the Milk object results of two images in the real-world HOPE dataset [1]. (b-c) Milk object examples in the synthetic ViSII (b) and NDDS (c) training datasets.

Uncertainty quantification (UQ) in robotics has been a long-standing research topic for many decades [15] due to the strict safety and reliability requirements of real-world systems. UQ can take on different meanings in robotic applications. 1) UQ can be in the form of a *probabilistic* distribution of possible outcomes on a prediction. 2) Alternatively, UQ can be a *deterministic* confidence measure of a specific prediction. In this work, we study the latter case, *i.e.*, we focus on a deterministic approximation of the quality of a certain pose prediction (see the formal problem statement in Section III-A). We refer to this type of UQ as *predictive uncertainty*. The predictive uncertainty is easier to estimate and use than the distribution of possible outcomes, since it is a scalar quantity and it directly reflects the risk of a certain pose estimation.

For the first type, Deng *et al.* estimated the distributions of possible poses for different objects, such as objects with symmetries, using particle filters [16]. Lee *et al.* (our baseline) also estimated the distributions of pose estimations by sampling the belief maps of the neural network output [17]. Such distributions can be embedded with other modules such as exploration in reinforcement learning [17]. For the

second type, one example is [2] (our baseline), where a trained deep object pose estimation model directly outputs a confidence score via keypoint belief maps. However, both the distribution estimation [17] and the confidence estimation [2] are not necessarily trustworthy and there is often a mismatch with the true uncertainty, as demonstrated in this paper.

Our goal is to develop a robust and computationally efficient method for quantifying the uncertainty of 6-DoF pose predictions from deep object pose estimation models. This is a challenging task from several directions. First, the predictive uncertainty is difficult to evaluate, since the ground truth for such measurements is non-existent. It thus requires us to develop a practical metric to quantify the uncertainty estimates. Second, deep pose estimation models are typically trained on synthetically rendered data [2, 13], due to the prohibitive costs of annotating 3D bounding boxes on a large set of real images. This sim2real gap requires the UQ method to robustly handle the discrepancies between the simulated data and real images and to make reliable estimates of predictive uncertainty across the two domains.

Contributions. We develop a simple, efficient, and plug-and-play technique for estimating predictive uncertainty using an ensemble of 2-3 pre-trained models with different architectures and/or training data sources, see Figure 1(a). Unlike prior approaches [12, 18, 19], our method does not require any modification of the training process or the model input, and our method only takes one forward pass for each model at inference time. Thus, we can readily integrate our UQ method with off-the-shelf deep pose estimation methods [2]. Our method uses only a few (2-3) deep pose estimators, computes their average pairwise disagreement to estimate the predictive uncertainty. We propose four metrics for computing the disagreement: translation, rotation, average distance (ADD) and a learned metric. We use the Spearman’s rank correlation coefficient to choose the best disagreement metric for predictive uncertainty. This is the first time, to the best of our knowledge, that such a correlation measure is used to quantify the extent to which the predictive uncertainty is consistent with the true uncertainty.

In our experiments, we evaluate the proposed UQ method on three tasks. In the first task, we evaluate the correlation between the estimated predictive uncertainties and pose estimation errors on the HOPE dataset [1]. Our UQ method has much stronger Spearman’s correlations than the baselines. Using the Spearman’s correlation, we find that ADD is the best data-free disagreement metric and that it is only slightly worse than the learned metric, which requires labeled target data. In the second task, we apply the proposed UQ to a camera perspective selection task. We evaluate the performance on a synthetic dataset generated by a photorealistic renderer ViSII [20]. The most confident camera perspectives selected by our method lead to a 30–40% reduction of pose estimation errors over the baselines. Finally, we demonstrate that our UQ method trained only on simulated data can guide a real robot to select the best point of view to robustly grasp objects. Our method dramatically improves the grasping success rate from 25%

(by baseline [2]) or 35% (by [17]) to 90%.

II. RELATED WORK

A. Object Pose Estimation

The problem of object pose estimation is vibrant within the robotics and computer vision communities [2, 3, 21–27]. Recent leading methods rely on an approach similar to the one used in our work: A network is trained to predict object keypoints in the 2D image, followed by PnP [28] to estimate the pose of the object in the camera coordinate frame [2, 5, 24, 25, 27, 29]. Other methods have regressed directly to the pose [3, 26], but these methods bake the camera intrinsics into the learned weights, although geometric post-processing can address this limitation [26]. In a related strand, researchers have used keypoint detection for human pose estimation [30–34]. Nevertheless, in robotics applications, it is not uncommon for objects to be detected via fiducial markers [35–37]. Closely related to our work, Deng *et al.* [16] used a particle filter-based approach to predicting the distribution of possible object poses. In contrast, our work studies pose estimators with deterministic predictions, and we quantify and validate their uncertainty.

B. Uncertainty Quantification (UQ)

Several methods have been proposed to quantify uncertainty in deep learning models. A relatively simple method is to directly use the output confidence of the trained model [2]. However, this has been shown to have limited robustness [38], and often yielding overconfident estimates in our experiments. A more principled approach is to adopt Bayesian principles [39] and posterior inference approximations such as the dropout approximation [18], variational inference [40], and Markov chain Monte Carlo [19]. In general, Bayesian approaches require the specification of prior distributions over model parameters and are more computationally expensive. In contrast, alternative sampling approaches [12] can estimate uncertainty without changing the network training process by adopting a frequentist strategy. However, it requires multiple forward passes and significantly increases the computational cost during inference. Ensemble methods [41], on the other hand, are much easier to run at inference time with modern parallel computational resources. However, they still require modifications in the training procedure and the members in the ensemble need to have the same architecture. To the best of our knowledge, ensemble-based UQ has not been previously explored in context of object pose estimation. We leverage this idea but consider a much simpler class of ensembles through disagreement metrics.

III. PROBLEM STATEMENT AND METHOD

We first define the 6-DoF pose estimation problem and the uncertainty quantification (UQ) problem. We also introduce the metrics used for UQ. Then we introduce our ensemble-based method in Sections III-B and III-C, as depicted in Figure 1(a). Section III-B formulates our ensemble method with a general disagreement metric, and Section III-C focused on a special case where the metric is learned.

A. Problem Statement and Metrics

1) *6-DoF pose estimation and ADD*: In the 6-DoF pose estimation task, the pose estimator g takes an image x as the input, and outputs a deterministic pose prediction $p = g(x)$ in $SE(3)$ – rotation and translation. We focus on model-based pose estimation of known objects, where each object has a separate trained pose estimator. Suppose the ground truth pose is given by \bar{p} . We measure the error of the pose estimation p by the average distance (ADD) [21] between poses p and \bar{p} , which is denoted by $\text{ADD}(p, \bar{p}, \mathcal{X})$:

$$\text{ADD}(p, \bar{p}, \mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{t \in \mathcal{X}} \|(p_R t + p_T) - (\bar{p}_R t + \bar{p}_T)\|_2$$

where \mathcal{X} is the object point cloud (for simplicity we will omit this from notation), p_R and p_T refer to the pose rotation and translation respectively. In other words, for a specific object, $\text{ADD}(p, \bar{p})$ is defined by the average 3D Euclidean distance between this object's two point clouds, corresponding to poses p and \bar{p} .

2) *Uncertainty quantification and correlation analysis*: For a pose estimator g , performance may vary with respect to different input images x . Intuitively, an ideal uncertainty measurement would yield a high confidence (low uncertainty) for accurate pose predictions, and a low confidence (high uncertainty) for inaccurate predictions, such that we can trust g in the scenarios when it performs well. For this reason, we aim to design a measurement d that has a high correlation with the performance metric $\text{ADD}(p, \bar{p})$. The value of d is treated as a UQ. This task is inherently challenging because the ground truth pose \bar{p} is unknown and x is in the *target domain*, which may have a large gap to the training data (see Figure 1 for examples).

To understand how well a measurement d characterizes the true pose estimation error $\text{ADD}(p, \bar{p})$, we analyze the correlation between d and the underlying pose estimation errors. The goal is to have a *strong positive correlation* between them. In particular, suppose that there are N images $\{x^{(1)}, \dots, x^{(N)}\}$ with pose predictions $\{p^{(1)} = g(x^{(1)}), \dots, p^{(N)} = g(x^{(N)})\}$ and their ground truth poses $\{\bar{p}^{(1)}, \dots, \bar{p}^{(N)}\}$. We also compute the ADD errors of these predictions $\{\text{ADD}(p^{(1)}, \bar{p}^{(1)}), \dots, \text{ADD}(p^{(N)}, \bar{p}^{(N)})\}$ and the uncertainty quantification $\{d^{(1)}, \dots, d^{(N)}\}$. We use the Spearman's rank correlation coefficient between $\{\text{ADD}(p^{(1)}, \bar{p}^{(1)}), \dots, \text{ADD}(p^{(N)}, \bar{p}^{(N)})\}$ and $\{d^{(1)}, \dots, d^{(N)}\}$ as our metric to evaluate d . The Spearman's rank correlation coefficient is a nonparametric measure of rank correlation, *i.e.*, statistical dependence between the rankings of two variables. It assesses how well the relationship between two variables can be described using a monotonic function. Compared to the classic Pearson correlation, the Spearman's rank correlation supports any nonlinear dependence and it is also more robust [42].

In summary, the Spearman's rank metric is a scalar from -1 to 1 , and reflects the dependence of the pose estimation error $\text{ADD}(p, \bar{p})$ on the UQ d . In particular, a larger correlation coefficient suggests d is a better uncertainty

measurement, in the sense that a larger d corresponds to a larger pose estimation error with higher chance.

B. Heterogeneous Ensemble for Uncertainty Quantification

Our ensemble method simultaneously infers the same image x using K different pose estimation models g_1, \dots, g_K , as shown in Figure 1(a). Our method makes no assumption to these K pose estimators. They are potentially heterogeneous in terms of model architecture and training data (see Section IV). These K models have K different pose estimations: $\{p_1 = g_1(x), \dots, p_K = g_K(x)\}$. With these K poses, we propose the following average disagreement based on the metric f :

$$d_f = \frac{1}{\binom{K}{2}} \sum_{1 \leq i < j \leq K} f(p_i, p_j) \quad (1)$$

where $f(p_i, p_j)$ is some distance metric between two poses p_i and p_j , *e.g.*, we can define $f(p_i, p_j) = \text{ADD}(p_i, p_j)$.

In other words, d_f represents the average disagreement of these K different models against each other on a certain image input x . We use d_f as an uncertainty quantification of these K models on this image. The higher d_f is, the more uncertain these K models are about this image. For example, as shown in Figure 1(a), we use three models (NDDS DOPE, NDDS DOPE full, and ViSII DOPE) to detect the Milk object, and get three different poses represented by three bounding boxes with different colors. In the first example (left) in Figure 1(a), these three models disagree with each other much more than the second example (right) in Figure 1(a), which implies these three models are much more uncertain about the first example.

The disagreement metric f plays an important role in d_f . We examine two properties of different choices of f : *object-variance* and *model-variance*. *Object-variance* allows flexibility across different objects, *e.g.*, Milk and Ketchup may use different disagreement metrics. *Model-variance* allows flexibility across different pose estimators, *e.g.*, NDDS DOPE and ViSII DOPE may use different disagreement metrics. We consider four types of disagreement metrics f with different properties (see a summary in Table I):

- **Translational disagreement** computes the translation error between p_i and p_j and ignores the rotation. It is an object-invariant and model-invariant disagreement, as the translation between p_i and p_j does not depend on object and model types.
- **Rotational disagreement** computes the rotation error between p_i and p_j and ignores the translation. It is also an object-invariant and model-invariant disagreement.
- **ADD disagreement** directly uses the $\text{ADD}(p_i, p_j)$ metric as $f(p_i, p_j)$. It is object-variant but model-invariant, as computing the ADD metric requires the object 3D geometry, but does not depend on the pose estimator. Empirically, we find that ADD disagreement is much better than translational/rotational disagreement.
- **Learned disagreement** is a data-driven disagreement. In this case, we need a dataset of labeled images in the target domain to train $f(p_i, p_j)$, which is represented

by a neural network (see Section III-C for details). Note that we can train different f for different pose estimators and objects, so this disagreement allows full flexibility across objects and models, *i.e.*, it is object-variant and model-variant.

TABLE I
FOUR TYPES OF DISAGREEMENT. SEE SECTION III-B FOR THE
DEFINITIONS OF OBJECT-VARIANT AND MODEL-VARIANT.

	Translational	Rotational	ADD	Learned
Object-Variant	No	No	Yes	Yes
Model-Variant	No	No	No	Yes
Require Target Data	No	No	No	Yes

C. Learned Disagreement Metric for Ensemble

Differently from the translational/rotational/ADD disagreement metrics, the learned disagreement metric is data driven and needs labeled data in the target domain. The idea of learning to combine different models is inspired by model stacking in statistics [43], where the goal is to boost the performance of an ensemble of heterogeneous models by learning a combination of their outputs in the target domain. In particular, for a certain object in the target domain we are given images $\{x^{(1)}, \dots, x^{(M)}\}$ with pose supervision $\{\bar{p}^{(1)}, \dots, \bar{p}^{(M)}\}$. With K models/pose estimators g_1, \dots, g_K in the ensemble, we have their pose predictions on each image in $\{x^{(1)}, \dots, x^{(M)}\}$:

$$\left\{ p_1^{(m)} = g_1(p^{(m)}), \dots, p_K^{(m)} = g_K(p^{(m)}) \right\}_{m=1}^M$$

For model g_k we also compute its ADD errors:

$$\left\{ \text{ADD}(p_k^{(m)}, \bar{p}^{(m)}) \right\}_{m=1}^M$$

With these K models' predictions and the ADD errors of the model g_k , our goal is to minimize the following loss function with respect to the disagreement metric f_k :

$$\min_{f_k \in \mathcal{F}} \sum_{m=1}^M \left| \frac{\sum_{1 \leq i < j \leq K} f_k(p_i^{(m)}, p_j^{(m)})}{\binom{K}{2}} - \text{ADD}(p_k^{(m)}, \bar{p}^{(m)}) \right|^2 \quad (2)$$

where \mathcal{F} is the function class. Equation (2) finds the optimal disagreement metric f_k^* for the pose estimator g_k , given M labeled data in the target domain. Note that different pose estimators may have different optimal disagreement metrics, *i.e.*, f_k^* depends on k and the learned disagreement is model-variant (see the definition in Section III-B). In practice, we use a simple feed-forward neural network to represent f_k and stochastic gradient descend is used to optimize the loss function in Equation (2). Empirically, we find that training with $\{x^{(1)}, \dots, x^{(M)}\}$ learns a disagreement metric that can generalize to other images in the target domain, and that it outperforms all other model-invariant and learning-free disagreements (translational/rotational/ADD).

IV. IMPLEMENTATION DETAILS

For UQ, our ensemble method uses multiple pose estimators from possibly different biases, *e.g.*, different training datasets. We use the newly available real-world HOPE

dataset [1] (see Figure 1(a) for examples), which consists of 28 grocery toy items with associated 3D models. It has 238 unique images with 914 unique object poses for evaluation, and the dataset does not provide any training images. For each object presented in the HOPE dataset, we generated different sets of 60,000 images, which uses visual domain randomization [44] to bridge the reality gap. We generated two different datasets from two different synthetic data generators available online, NDDS [45] (see Figure 1(c) for one example) and ViSII [20] (see Figure 1(b)). The former one uses rasterization whereas the latter one uses ray-tracing. We trained the open-source pose estimator DOPE [2] on both datasets as well as a modified DOPE architecture on NDDS only.

The DOPE pose estimator first finds the projected cuboid keypoints in the image frame, and using the known size of the object and the camera intrinsics it regresses to the object pose using PnP [46]. The keypoints are represented as heatmaps, more precisely, nine heatmaps for the cuboid vertex and its centroid. A greedy algorithm finds the keypoint positions in the original image. For each model we report the area under the curve (AUC) for ADD threshold curves at 10 cm for all objects in the HOPE dataset. In this work we use the following models:

- **NDDS DOPE model.** We directly use the model architecture in [2] and train all 28 models on NDDS synthetic data with domain randomization. This model achieves an AUC of 0.37.
- **NDDS DOPE full model.** Inspired by previous work [9], we use the auto-encoder architecture that has more capacity than the original DOPE model architecture. This model achieves an AUC of 0.45.
- **ViSII DOPE model.** This model has the same architecture to the NDDS DOPE model, but is trained on ViSII synthetic data. This model achieves an AUC of 0.37.

We compare the ensemble method against two baselines. Both baselines are based on the internal UQ of a single neural network model, *i.e.*, these two baselines do not leverage information from multiple models. The baselines include:

- **Neural network confidence score (Confidence).** Along with the predicted 6-DoF pose, the DOPE model [2] also gives a confidence score to each pose estimation. In particular, this confidence score reflects the confidence level of the centroid of the detected object, *e.g.*, the height of the peak on the centroid's heatmap.
- **Randomly sampling belief maps (GUAPo).** This baseline [17] augments the peak estimation algorithm in DOPE by fitting a 2D Gaussian around each found peak. GUAPo then runs PnP algorithm on T sets of keypoints, where each set of keypoint is constructed by sampling from all the 2D Gaussians. This provides T possible poses of the object consistent with the detection algorithm, and the variance of these T possible poses (standard deviation) represents the uncertainty.




Fig. 2. Two test sequences generated by ViSII [20]. Each sequence contains 45 images with camera angle from 0° to 180° . We only present 4 of 45.

V. EXPERIMENTS

We evaluate the capabilities of our uncertainty quantification (UQ) approach with respect to the problem of 6-DoF pose estimation of known objects. In other words, we want to evaluate how likely a 3D position and orientation prediction is correct, and to that end, we seek to answer the following questions throughout the experiments: 1) How does our method compare to different ways to quantify uncertainty? 2) Can our method be used to ensure higher success-rate for grasping objects with a real-world manipulator?

A. Correlation Analysis

We evaluate our proposed method using the correlation between UQ values and the predicted pose errors against ground truth (see the Spearman’s rank correlation coefficient in Section III-A). This evaluation is conducted on the real-world HOPE dataset [1]. Please note that all our pose estimation models are only trained on synthetic data (either NDDS or ViSII). In particular, for each object in the 28 HOPE set, we collected all images in the HOPE dataset [1] that do not contain multiple instances of the same object. We first explore which disagreement metric f in eq. (1) offers the best correlation, and we explore the impact of different architectures/training data when using ensemble for UQ.

1) *Disagreement metric selection:* We evaluate six UQ methods: our ensemble-based method with four types of disagreement (ADD, rotation, translation, learned) as described in Section III-B, and the two baselines presented in Section IV. For the ensemble methods we use two heterogeneous models with different architectures (NDDS DOPE and NDDS DOPE full) in the ensemble.

For the learned disagreement, we train two neural networks on 1/3 of the HOPE dataset for each object to learn the disagreement metrics f for both NDDS DOPE and NDDS DOPE full (see Section III-C). We use four-layer multilayer perceptrons with ReLU activations to parameterize the disagreement metric, where the input dimension is 14 (two poses represented by quaternions and 3D vectors) and the output is a scalar (disagreement between two poses). The other five methods do not need labeled data.

Table II shows the correlation coefficients for the six UQ methods. The first/second column shows the result for NDDS DOPE/NDDS DOPE full. From this experiment we

TABLE II
CORRELATION ANALYSIS FOR DIFFERENT DISAGREEMENTS

	NDDS DOPE	NDDS DOPE full
Confidence [2]	0.26 ± 0.14	0.40 ± 0.16
GUAPO [17]	0.45 ± 0.18	0.44 ± 0.15
Translational Disagreement	0.52 ± 0.11	0.46 ± 0.12
Rotational Disagreement	0.39 ± 0.15	0.44 ± 0.13
ADD Disagreement	0.55 ± 0.11	0.50 ± 0.12
Learned Disagreement*	0.58 ± 0.09	0.52 ± 0.13

*requires labeled data in the target domain

see that 1) ensemble methods with all four disagreements perform better than baselines, except for an outlier from the rotational disagreement; 2) the learned disagreement achieves the best correlations because it is trained on the target domain and it is both object-variant and model-variant, *i.e.*, different objects and models may learn diverse disagreements accordingly; and 3) ensemble with ADD disagreement yields the best results in all data-free methods, and is only slightly worse than the learned disagreement ensemble method. The small performance gap between the *data-free* ADD disagreement and learned disagreement suggests that ADD is a suitable disagreement for UQ in an ensemble.

2) *Architecture and training data:* The second set of experiments considers the influence of heterogeneity from training data/architectures, and Table III shows the results. Note that the ViSII DOPE model is trained on the ViSII synthetic data, but the NDDS DOPE and NDDS DOPE full models are trained on the NDDS synthetic data. Here we focus on the performance of the ADD disagreement ensemble with three different ensemble combinations: (ViSII DOPE, NDDS DOPE), (ViSII DOPE, NDDS DOPE full), and (ViSII DOPE, NDDS DOPE, NDDS DOPE full). More precisely, we evaluate two baselines and ADD disagreement ensemble with these three combinations, as their correlation with respect to the ViSII DOPE model.

Table III shows the results on this experiment. We find that 1) all the ensemble combinations have much stronger correlations than the baselines; 2) the three-model ensemble is slightly worse than two-model ensembles, which implies adding more models in the ensemble does not necessarily improves UQ; and 3) compared with Table II, ViSII DOPE achieves better correlations than NDDS DOPE and NDDS DOPE full. In particular, the baseline method GUAPO for ViSII DOPE already outperforms ensemble methods for NDDS DOPE or NDDS DOPE full in Table II, which implies it is easier to quantify uncertainty for the ViSII [20] (based on ray-tracing) trained model than the NDDS [44] (based on rasterization) trained models, possibly due to realistic light modelling from ray-tracing [20].

B. Application I: Camera Perspective Selection

The estimate uncertainty can inform an active perception agent to explore an environment and select the most confident observation for decision making. In the context of pose estimation, visual observations from different camera perspectives may result in object occlusions and other environmental



Fig. 3. (a) Experimental setup. (b-c) The view selection and grasping results of the Ketchup object from our method and GUAPO [17]. We use the NDDS DOPE model (green box) as the pose estimator and we quantify its uncertainty. For our method we use the average ADD disagreement between NDDS DOPE and NDDS DOPE full as the uncertainty. The numbers in (b-c) show the uncertainty quantification values of GUAPO and our method.

TABLE III

CORRELATION ANALYSIS FOR ARCHITECTURES AND TRAINING DATA

	ViSII DOPE
Confidence [2]	0.38 ± 0.16
GUAPO [17]	0.55 ± 0.15
ADD (ViSII DOPE, NDDS DOPE)	0.64 ± 0.14
ADD (ViSII DOPE, NDDS DOPE full)	0.68 ± 0.14
ADD (All three models)	0.62 ± 0.14

TABLE IV

AVERAGE ADD ERROR OF THE MOST CONFIDENT FRAME (UNIT: CM)

	NDDS DOPE	NDDS DOPE full
Confidence [2]	4.9 ± 6.5	4.7 ± 6.1
GUAPO [17]	5.6 ± 6.9	5.1 ± 5.9
ADD (Ours)	3.5 ± 4.2	3.1 ± 3.6
Oracle	1.0 ± 1.2	1.2 ± 1.4

noises. Here we examine whether our UQ can select an optimal perspective for making the best pose predictions.

To quantitatively study the performances of different UQ methods in this problem, we generate 125 image sequences using ViSII [20] (see Figure 2 for examples). In particular, for each sequence, we randomly pick 3–6 objects from 28 HOPE objects, and randomize their collision-free placements. We also randomize the background, table texture, and illumination for each sequence. Finally, we move around the virtual camera in 180 degree and sample 45 images.

We use a simple greedy approach to select the optimal frame in the image sequence. For each object in each sequence, we directly select the frame of the least UQ value. For instance, in the case of the ensemble method, we will choose the object with the smallest d_f in eq. (1). Other than the ensemble method and two baselines, we also compute the performance of the oracle approach, which has access to the ground truth and will directly select the image with lowest ADD error. This oracle helps us understand the upper bound of the performance of the pose estimators. We use ADD as the disagreement metric in ensemble and we excluded ViSII DOPE model from this experiment since it is trained on images from ViSII renderer.

The results are presented in Table IV, where we average all 28 objects and all 125 sequences. Our method decreases the ADD error of the greedy choice by 30–40% compared to baselines, and the variance is also significantly reduced.

TABLE V

UNCERTAINTY-GUIDED REAL-WORLD ROBOT GRASPING SUCCESS RATE

	Confidence [2]	GUAPO [17]	ADD (Ours)
Stable Grasp	35%	25%	90%
Unstable Grasp	10%	25%	10%
Failure	55%	50%	0%

C. Application II: Real-World Robotic Grasping

To examine the utility of UQ of our proposed ensemble method for downstream tasks in the real world, we use an uncertainty-guided robotic grasping task to demonstrate its improvement over the baselines. We use a 7-DoF robotic arm, Franka Emika Panda, for our grasping task, see Figure 3(a). Similar to the experiments on ViSII synthetic data, we first control the robot to observe six images of the workspace with a clutter of objects from different perspectives, and we use the UQ value to select the most confident point of view. We command the robot to grasp the target object using the pose estimation from this point of view, see Figure 3(b-c). After robot execution, we evaluate if the grasp is a stable grasp, unstable grasp, or a complete failure. Unstable grasp corresponds to the cases where the gripper successfully makes contact with the object, but it slipped away. For failure cases, the gripper fails to touch the target object at all. We use the NDDS DOPE model as our pose estimator, and we use the average ADD disagreement between NDDS DOPE and NDDS DOPE full for UQ.

We evaluate on four objects (Ketchup, BBQ Sauce, Mayo, and Butter) of diverse shapes/colors/textures, and attempt 5 trials for each object and each method. For a fair comparison, we evaluate all three methods in the same environment setting for each trial. A summary of results is shown in Table V, suggesting that our method significantly outperforms the two baselines. Figure 3(b) shows a Ketchup example of view selection and grasping from our method, which leads to a successful grasping. Figure 3(c) shows the result from GUAPO under the same environment, where GUAPO is overconfident about a non-Ketchup pose estimation.

VI. CONCLUSIONS

We develop a simple, efficient, and plug-and-play method for quantifying the predictive uncertainty of 6-DoF pose estimators. Our method uses an ensemble of two or more

heterogeneous models and computes their average disagreement against one another as the uncertainty quantification. Our experimental results demonstrate that average distance (ADD) is a suitable learning-free disagreement metric, and the proposed method significantly outperforms baselines in different tasks, including real-robot grasping. Future directions include the theoretical analysis of ensemble-based UQ and the integration of our method with other algorithms, such as uncertainty-aware RL and pose estimation in videos.

ACKNOWLEDGEMENT

We would like to thank members of the NVIDIA AI Algorithms research team for their constructive feedback and Nathan V. Morrical for his help with the ViSII renderer.

REFERENCES

- [1] S. Tyree, J. Tremblay, T. To, J. Cheng, T. Mossier, and S. Birchfield, “6-DoF pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark,” in *ICCV Workshop on Recovering 6D Object Pose*, 2019.
- [2] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, “Deep object pose estimation for semantic robotic grasping of household objects,” in *CoRL*, 2018.
- [3] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, “PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes,” in *RSS*, 2018.
- [4] C. Wang, D. Xu, Y. Zhu, R. Martin-Martin, C. Lu, L. Fei-Fei, and S. Savarese, “DenseFusion: 6d object pose estimation by iterative dense fusion,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [5] D. J. Tan, N. Navab, and F. Tombari, “6d object pose estimation with depth images: A seamless approach for robotic interaction and augmented reality,” *arXiv preprint arXiv:1709.01459*, 2017.
- [6] A. Zeng, K.-T. Yu, S. Song, D. Suo, E. Walker, A. Rodriguez, and J. Xiao, “Multi-view self-supervised deep learning for 6d pose estimation in the Amazon Picking Challenge,” in *IEEE International Conference on Robotics and Automation*, 2017.
- [7] X. Chen, R. Chen, Z. Sui, Z. Ye, Y. Liu, R. Bahar, and O. C. Jenkins, “Grip: Generative robust inference and perception for semantic robot manipulation in adversarial environments,” *IROS*, 2019.
- [8] W. Yang, C. Paxton, M. Cakmak, and D. Fox, “Human grasp classification for reactive human-to-robot handovers,” *arXiv preprint arXiv:2003.06000*, 2020.
- [9] T. E. Lee, J. Tremblay, T. To, J. Cheng, T. Mosier, O. Kroemer, D. Fox, and S. Birchfield, “Camera-to-robot pose estimation from a single image,” in *International Conference on Robotics and Automation*, 2020.
- [10] A. Handa, K. Van Wyk, W. Yang, J. Liang, Y.-W. Chao, Q. Wan, S. Birchfield, N. Ratliff, and D. Fox, “DexPilot: Vision-based teleoperation of dexterous robotic hand-arm system,” in *IEEE International Conference on Robotics and Automation*, 2020.
- [11] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, “Concrete problems in AI safety,” *arXiv preprint arXiv:1606.06565*, 2016.
- [12] A. Loquercio, M. Segu, and D. Scaramuzza, “A general framework for uncertainty estimation in deep learning,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3153–3160, 2020.
- [13] C. Doersch and A. Zisserman, “Sim2real transfer learning for 3d human pose estimation: Motion to the rescue,” in *Advances in Neural Information Processing Systems*, 2019.
- [14] J. Lu, F. Richter, and M. Yip, “Robust keypoint detection and pose estimation of robot manipulators with self-occlusions via sim-to-real transfer,” *arXiv preprint arXiv:2010.08054*, 2020.
- [15] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, pp. 35–45, 1960.
- [16] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, “PoseRBPF: A Rao-Blackwellized particle filter for 6d object pose tracking,” *arXiv preprint arXiv:1905.09304*, 2019.
- [17] M. A. Lee, C. Florensa, J. Tremblay, N. Ratliff, A. Garg, F. Ramos, and D. Fox, “Guided uncertainty-aware policy optimization: Combining learning and model-based strategies for sample-efficient policy learning,” *arXiv preprint arXiv:2005.10872*, 2020.
- [18] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *International Conference on Machine Learning*, 2016.
- [19] C. Andrieu, N. de Freitas, A. Doucet, and M. I. Jordan, “An introduction to MCMC for machine learning,” *Machine Learning*, vol. 50, no. 1-2, pp. 5–43, 2003.
- [20] N. Morrical, J. Tremblay, S. Birchfield, and I. Wald, *ViSII: Virtual scene imaging interface*, <https://github.com/owl-project/ViSII/>, 2020.
- [21] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, “Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes,” in *ACCV*, 2012.
- [22] T. Hodaň, P. Haluza, Š. Obdržálek, J. Matas, M. Lourakis, and X. Zabulis, “T-LESS: An RGB-D dataset for 6D pose estimation of texture-less objects,” in *WACV*, 2017.
- [23] S. Zakharov, I. Shugurov, and S. Ilic, “DPOD: Dense 6D pose object detector in RGB images,” *arXiv preprint arXiv:1902.11020*, 2019.
- [24] Y. Hu, J. Hugonot, P. Fua, and M. Salzmann, “Segmentation-driven 6D object pose estimation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3385–3394.
- [25] S. Peng, Y. Liu, Q. Huang, X. Zhou, and H. Bao, “Pvnet: Pixel-wise voting network for 6dof pose estimation,” in *CVPR*, 2019.
- [26] M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, “Implicit 3d orientation learning for 6d object detection from rgb images,” in *ECCV*, 2018.
- [27] B. Tekin, S. N. Sinha, and P. Fua, “Real-time seamless single shot 6D object pose prediction,” in *CVPR*, 2018.
- [28] V. Lepetit, F. Moreno-Noguer, and P. Fua, “EPnP: An accurate $O(n)$ solution to the PnP problem,” *International Journal Computer Vision*, vol. 81, no. 2, 2009.
- [29] A. Dhall, D. Dai, and L. V. Gool, “Real-time 3D traffic cone detection for autonomous driving,” in *arXiv:1902.02394*, 2019.
- [30] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh, “Convolutional pose machines,” in *CVPR*, 2016.
- [31] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, “Realtime multi-person 2D pose estimation using part affinity fields,” in *CVPR*, 2017.
- [32] B. Xiao, H. Wu, and Y. Wei, “Simple baselines for human pose estimation and tracking,” in *European Conference on Computer Vision (ECCV)*, 2018, pp. 466–481.
- [33] W. Li, Z. Wang, B. Yin, Q. Peng, Y. Du, T. Xiao, G. Yu, H. Lu, Y. Wei, and J. Sun, “Rethinking on multi-stage networks for human pose estimation,” *arXiv preprint arXiv:1901.00148*, 2019.
- [34] K. Sun, B. Xiao, D. Liu, and J. Wang, “Deep high-resolution representation learning for human pose estimation,” in *CVPR*, 2019.

- [35] C. Liu and M. Tomizuka, "Robot safe interaction system for intelligent industrial co-robots," in *arXiv:1808.03983*, 2018.
- [36] C. H. Kim and J. Seo, "Shallow-depth insertion: Peg in shallow hole through robotic in-hand manipulation," in *ICRA*, 2019.
- [37] N. Tian, A. K. Tanwani, J. Chen, M. Ma, R. Zhang, B. H. K. Goldberg, and S. Sojoudi, "A fog robotic system for dynamic visual servoing," in *ICRA*, 2019.
- [38] H. Blum, P.-E. Sarlin, J. Nieto, R. Siegwart, and C. Cadena, "The fishyscapes benchmark: Measuring blind spots in semantic segmentation," *arXiv preprint arXiv:1904.03215*, 2019.
- [39] J. M. Bernardo and A. F. Smith, *Bayesian theory*. John Wiley & Sons, 2009, vol. 405.
- [40] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine Learning*, vol. 37, 1999.
- [41] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., 2017.
- [42] J. Hauke and T. Kossowski, "Comparison of values of pearson's and spearman's correlation coefficients on the same sets of data," *Quaestiones geographicae*, vol. 30, no. 2, pp. 87–93, 2011.
- [43] D. H. Wolpert, "Stacked generalization," *Neural networks*, vol. 5, no. 2, pp. 241–259, 1992.
- [44] J. Tremblay, T. To, A. Molchanov, S. Tyree, J. Kautz, and S. Birchfield, "Synthetically trained neural networks for learning human-readable plans from real-world demonstrations," in *ICRA*, 2018.
- [45] T. To, J. Tremblay, D. McKay, Y. Yamaguchi, K. Leung, A. Balanon, J. Cheng, and S. Birchfield, *NDDS: NVIDIA deep learning dataset synthesizer*, https://github.com/NVIDIA/Dataset_Synthesizer, 2018.
- [46] X.-S. Gao, X.-R. Hou, J. Tang, and H.-F. Cheng, "Complete solution classification for the perspective-three-point problem," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 8, pp. 930–943, Aug. 2003.