# Alignist: CAD-Informed Orientation Distribution Estimation by Fusing Shape and Correspondences

Shishir Reddy Vutukur<sup>1</sup>, Rasmus Laurvig Haugaard<sup>2</sup>, Junwen Huang<sup>1</sup>, Benjamin Busam<sup>1,3</sup>, and Tolga Birdal<sup>4</sup>

<sup>1</sup>Technical University of Munich <sup>2</sup>University of Southern Denmark <sup>3</sup>Munich Center for Machine Learning (MCML) <sup>4</sup>Imperial College London

Abstract. Object pose distribution estimation is crucial in robotics for better path planning and handling of symmetric objects. Recent distribution estimation approaches employ contrastive learning-based approaches by maximizing the likelihood of a single pose estimate in the absence of a CAD model. We propose a pose distribution estimation method leveraging symmetry respecting correspondence distributions and shape information obtained using a CAD model. Contrastive learning-based approaches require an exhaustive amount of training images from different viewpoints to learn the distribution properly, which is not possible in realistic scenarios. Instead, we propose a pipeline that can leverage correspondence distributions and shape information from the CAD model, which are later used to learn pose distributions. Besides, having access to pose distribution based on correspondences before learning pose distributions conditioned on images, can help formulate the loss between distributions. The prior knowledge of distribution also helps the network to focus on getting sharper modes instead. With the CAD prior, our approach converges much faster and learns distribution better by focusing on learning sharper distribution near all the valid modes, unlike contrastive approaches, which focus on a single mode at a time. We achieve benchmark results on SYMSOL-I and T-Less datasets.

**Keywords:** Object Pose Estimation · Pose Distribution · 6D Pose Estimation · Uncertainty · Ambiguity

## 1 Introduction

Real world is quintessentially VUCA, *i.e.*, volatile, uncertain, complex and ambiguous [26], significantly complicating the task of any methodology aimed at perceiving and interpreting our natural environments. In fact, this inherent nature renders the quest for a universally applicable, unique solution to computer vision, elusive. Still, modern AR/VR systems and robotics applications require to navigate in our VUCA world with a high degree of accuracy, robustness, safety and reliability [44]. In such cases, it becomes more pragmatic to reason about uncertainty and ambiguities rather than seeking a one-size-fits-all solution.

In the realm of computer vision, uncertainty often stems from incomplete, insufficient, or noisy data, while ambiguity arises when visual information can



Fig. 1: Training and Inference Pipeline: We employ a training mechanism where supervision is generated from pre-trained SurfEmb $(f_{SE})$  and SDF $(f_{SDF})$  blocks. The CAD model,  $\mathbf{X}_0$ , undergoes a projection to render an image aligned point cloud,  $\mathbf{X}$ . The image aligned point cloud, **X**, is rotated with ground truth rotation,  $\mathbf{R}_{at}$ , and passed through  $f_{\rm SE}$  block to estimate canonical features. Similarly, X, is rotated with a random rotations,  $\mathbf{R}_k$ , and passed through  $f_{SE}$  block to generate features that are compared with canonical features to estimate the score  $\mu_{\rm SE}$  for the rotation. Similarly, the rotated point cloud with a random rotation is passed through  $f_{\text{SDF}}$  to estimate the SDF values of the point cloud. An  $L_0$  norm is applied to the SDF values to compute  $\mu_{\text{SDF}}$  score for the rotation. These scores are used to supervise the Dual-branch MLP. The Dual-branch MLP network takes an image and the same rotation matrix,  $\mathbf{R}_k$ , as input and infers two scores  $\mu_{\theta}$  and  $\mu_{\phi}$ . This process is carried out for K rotations for a given image and a Generalized KL divergence loss (GKL) is formulated between inferred scores from the right block and estimated scores from the left block to train the Dual-branch MLP network. The Dual-branch MLP is part of both training and inference. During inference, an image and rotations sampled from a grid are passed through the network to estimate the full distribution on the grid.

be interpreted in multiple ways. For example, a shadow might be mistaken for an object, or a partially occluded object might be hard to identify. Both of these notions can be captured by a *multi-modal probability distribution*, whose modes correspond to ambiguous yet probable solutions, and the variation captures a notion of uncertainty. This motivated a plethora of pose estimation approaches [8, 16, 40] to infer densities over the pose space rather than predicting point estimates as done classically [17, 56].

The rotational components of the 6D pose are particularly impacted by ambiguities, which mostly arise due to symmetry or self-occlusions [38]. Consequently, state-of-the-art approaches aim to characterize the density over SO(3) (rotations), either directly in the form of mixtures of matrix distributions [8, 60, 62] or through implicit neural fields [40] and invertible neural networks [37]. These methods generally assume that the training data consists of posed images, and, in one way or another, learn to infer densities conditioned on an input image. Such an approach necessitates a large amount of training data, as the networks need to encounter the same appearance under various rotations.

In this paper, we introduce, **Alignist**, a new way of learning distributions over SO(3), leveraging the CAD model that is readily available for a majority of 3D pose tasks and provided with typical datasets together with object poses either in instance or category form. To make use of this additional piece of shape information, we first establish the proportionality of image-conditional densities over rotations to densities over the visible parts of the CAD model. With this change of variables, we can sample the space of rotations during training and obtain a sharp and full density over the entire SO(3) in the form of an unnormalized, empirical measure (scores). In particular, we model the desired density over the model as a products of experts [18], reflecting the spatial and feature-space alignment to the original CAD model. We choose two expert probabilities to follow Bolzmann densities induced by (i) the distance to the original CAD model measured by the of norm of signed distances, (ii) the similarity of symmetry-respecting features obtained from SurfEMB [15]. The first expert specifies to 3D geometric alignment whereas the second one is additionally informed by the appearance, as SurfEMB is trained also using the image cues. The final product-probability supervises a dual-branch MLP, which, given an input image, infers both of the distributions in the form of empirical distributions or measures by sampling over SO(3). As a loss between the computed and inferred distributions, we utilize the recently proposed generalized KL divergence [39], capable of comparing unnormalized probability measures. This training scheme, illustrated in Fig. 1, is agnostic to the network architecture and not only facilitates stronger cues for learning but also aids in generalization in a low-data regime, as the available CAD model is exploited to generate densely sampled distributions. Lastly, as classical positional encoding methods do not generalize to the manifold of SO(3) [11], we introduce a new rotation encoding method by transforming the vertices of a unit cube under the given rotation.

In summary, our contributions are:

- To the best of our knowledge, we present the first method to use a CAD model to train an implicit network in order to infer densities over SO(3).
- We propose to utilize geometry-aware (SDF) and symmetry-aware (SurfEMB) experts to obtain rich and informative supervision cues during training.
- Backed by a novel positional encoding, we propose a dual-branch MLP to infer two distributions, one on the SDF and one on the surface features.
- Experiments demonstrate the advantages of our approach in capturing ambiguities and uncertainties on textureless objects, especially in low-data regimes, thanks to powerful 3D object priors our method utilizes.

We make our implementation publicly available here.

## 2 Related Work

While a plethora of 6D object pose estimation exist [17,24,35,52–54,56], we will specifically review works which deal with rotational symmetries in the context of deep learning.

Rotation representations for deep networks. Deep neural networks typically produce feature vectors in a Euclidean space, making it hard to reason on nonlinear manifolds, such as that of SO(3). Early deep learning models opted for rather classical parameterizations for regressing rotations, like Eulerangles [30, 48], direction cosine matrices (DCM) [23, 59], axis-angle [10, 51], and quaternions [28, 56, 63]. Some scholars [5, 27, 49] instead discretize the space and classify, however, pose space represents a continuous Riemannian manifold [4]. Recently, [64] showed that the typical representations are discontinuous, *i.e.*, do not admit homeomorphic maps to SO(3). Instead, 6D [64], 9D [34] and even 10D [42] representations were proposed to resolve the discontinuity issue and improve the regression accuracy. Brégier [2] has thoroughly examined different manifold mappings, finding out that SVD orthogonalization [34] performs the best when regressing arbitrary rotations.

Besides considering rotations as manifold-valued labels, there are multiple ways they are used in deep networks. Regressing rotations lead to geometric gradients, which require rethinking of backpropagation [6, 50]. When used as input data, the sinusoidal basis functions that are used to encode Euclidean coordinates become ill-suited for SO(3) [11]. In this work, we also contribute a symmetry-respecting *positional encoding* for rotations, inspired by the metrics in [1] and also follow a regression approach, however, we investigate pose distribution prediction rather than single pose estimation to incorporate multimodal solutions induced by object or projection ambiguities.

**Representing belief over rotations in deep networks**. Extending Bui *et al.* [3] to object pose estimation, Deng *et al.* [8], along with [12], were the first to address pose ambiguity and uncertainty prediction via deep networks. In these works, as well as extensions such as [42], a Bingham distribution was used to represent the belief over SO(3). The suitability of Bingham distribution was later questioned by follow-up works, proposing Matrix-Fisher distributions [60] or heavy-tailed variants such as Laplace distributions [61, 62] for increased robustness. Implicit-PDF (iPDF) [40] used a rotation-conditioned neural network to implicitly learn the orientation distribution and SpyroPose [16] proposed a coarse-to-fine version of iPDF using a hierarchical grid. HyperPosePDF [22] learned the weights of an iPDF via a hyper-network and spherical convolutions are used in I2S [29] to map a deep feature-map onto a distribution over SO(3). Recently, Liu *et al.* [37] proposed a normalizing flow based approach to map an initial distribution to a target one. This approach has later been generalized to manifolds other than SO(3) via free-form flows [47].

We follow the ideas of implicit distribution learning by incorporating a signed distance function [41] alongside a method for ambiguity aware pose description [15] into our approach.

Uncertainty and ambiguity aware object pose estimation. Object shape, symmetry, or occlusions can cause ambiguities in object poses given only the perspective projection. Corona *et al.* [7] and Pitteri *et al.* [43] assumed the availability of object symmetry information during training. As obtaining this information is challenging, others tried to explicitly learn it from data [33]. Stateof-the-art approaches bypass the requirement of prior knowledge about object symmetries. Manhardt et al. [38] modeled the ambiguity in object pose estimation by predicting multiple hypotheses, for a given object's visual appearance, and Shi et al. [46] utilize an ensemble of object pose estimators to derive uncertainties. SurfEmb [15], EPOS [19] and NeRF-Feat [52] learned a dense distribution of 2D-3D correspondences. EPOS [19] handled symmetries implicitly by discretizing the surface and predicting a probability distribution over fragments per pixel, whereas SurfEmb, NeRF-Feat estimated a continuous distribution over the object surface. Similarly, Ki-pode [25] used object keypoints as an intermediary to derive the probability density, explicitly and [57] makes use of the conformal prediction framework to propagate uncertainties from keypoints to object poses.

In our pipeline, we incorporate a learned continuous distribution [15] of 2D-3D correspondences between object model and image.

# 3 Method

**Problem setting.** We consider an input image crop  $\mathbf{I} \in \mathbb{R}^{M \times N \times 3}$  parameterized by coordinates  $\mathbf{p} \in \mathbb{R}^2$  on the image lattice, containing the projection of a 3D object  $\mathcal{M} = (\mathbf{X}, \mathbf{T})$  represented in normalized object coordinate space (NOCS) [55]  $\mathbf{X} = \{\mathbf{x}_i \in [-1, 1]^3 \mid i = 1, \dots, N_X\}$  together with triangle faces  $\mathbf{T} = \{\mathbf{t}_i \in \mathbb{Z}^3 \mid i = 1, \dots, N_T\}$ . We formulate the rotation estimation problem as a probabilistic inference, where we are interested in the following two quantities:

- 1. Single-best orientation obtained by the maximum a-posteriori (MAP) estimate:  $\mathbf{R}^{\star} = \arg \max_{\mathbf{R}} p(\mathbf{R} \mid \mathbf{I})$ , where  $p(\mathbf{R} \mid \mathbf{I})$  models the conditional probability distribution for the orientation of  $\mathbf{I}$ .
- 2. Multi-modal orientation estimate in the form of a full posterior distribution:  $p(\mathbf{R} | \mathbf{I}) \propto p(\mathbf{I} | \mathbf{R}) \times p(\mathbf{R}) = p(\mathbf{I}, \mathbf{R}).$

A plethora of the 6D object pose estimation methods solve (1) (excluding translations) to get a point estimate [21, 27, 49]. While solving (2) can provide important additional information, such as *uncertainty* and *ambiguity*, perhaps not surprisingly it is a much harder task. Note that, in addition to the usual difficulties associated with posterior characterization, in our context we are facing extra challenges due to the Riemannian nature of SO(3) [4]. This is arguably the reason why the task of *multiple simultaneous rotation estimations* remains unsolved, to date. In this work, we focus on estimating posterior densities over rotations, which are the main sources of ambiguity.

**Model.** Inspired by the energy based models [31,32], we write down the posterior density of interest as  $p(\mathbf{R} | \mathbf{I}) = \exp(-U(\mathbf{R}))/p(\mathbf{R})$  where U is the potential energy with the form  $U(\mathbf{R}) = -(\log p(\mathbf{R} | \mathbf{I}) + \log p(\mathbf{R}))$ . Estimating  $p(\mathbf{R})$  by

6 Vutukur et al.

discretizing the SO(3) by  $N_V$  partitions of volume V and marginalizing, leads to the following approximate posterior density [40]:

$$p(\mathbf{R} \mid \mathbf{I}) \approx \frac{1}{V} \exp\left(-U(\mathbf{R})\right) / \sum_{i=1}^{N_{V}} \exp\left(-U(\mathbf{R}_{i})\right) = \frac{1}{V} \operatorname{softmax}\left(-U(\mathbf{R})\right). \quad (1)$$

To design  $U(\mathbf{R})$ , we assume the availability of a  $\mathbf{X}_0$ , s.t.  $\mathbf{X}' = \mathbf{R}\Pi(\mathbf{R}_{gt}^{\top}\mathbf{X}_0; \mathbf{I}, \mathcal{M})$  is obtained by rotating the *image-aligned* model  $\mathbf{X}$  back to the CAD space using  $\mathbf{R}$ .  $\mathbf{R}_{gt}$  is the ground truth pose of the image under consideration and the rendering operator  $\mathbf{X} = \Pi(\mathbf{R}_{gt}^{\top}\mathbf{X}_0; \mathbf{I}, \mathcal{M})$  extracts the points of the transformed mesh, visible in the image  $\mathbf{I}$  by rendering the model mesh  $\mathcal{M}$ . In order to maintain the invariance under symmetries, *i.e.*,  $\mathbf{R}\mathbf{X} \equiv \mathbf{X}$  for any *stabilizer*  $\mathbf{R}$  of  $\mathbf{X}$ , we relate orientation distributions to distributions over object coordinates:

**Proposition 1.** The probability  $p(\mathbf{R}|\mathbf{I}) \propto p(\mathbf{X}'|\mathbf{I})$  where  $\mathbf{X}' = \mathbf{R}\Pi(\mathbf{R}_{gt}^{\top}\mathbf{X}_0; \mathbf{I}, \mathcal{M})$ .

*Proof.* We will consider the vectorized versions of our variables  $\mathbf{r} := \operatorname{vec}(\mathbf{R})$ and  $\mathbf{x}' := \operatorname{vec}(\mathbf{X}')$  and assume  $p(\mathbf{r} \mid \mathbf{I}) = p(\mathbf{R} \mid \mathbf{I})$  and  $p(\mathbf{x}' \mid \mathbf{I}) = p(\mathbf{X}' \mid \mathbf{I})$ . Let  $\mathbf{C} = \Pi(\mathbf{R}_{gt}^{\top}\mathbf{X}_0; \mathbf{I}, \mathcal{M})$  denote the visible part of the model, w.l.o.g. assumed to be a constant. Then, the two distributions can be related by the *change of variables*:

$$p(\mathbf{r} \mid \mathbf{I}) = p(\mathbf{x}' \mid \mathbf{I}) \cdot |\mathbf{J}|, \qquad (2)$$

where

$$\mathbf{J} := \frac{\mathrm{d}\mathbf{x}'}{\mathrm{d}\mathbf{r}} = \frac{\mathrm{dvec}(\mathbf{R}\mathbf{C})}{\mathrm{d}\mathbf{r}} = \frac{\mathrm{d}\left(\left(\mathbf{C}^{\top} \otimes \mathbf{I}_{N_{X}}\right)\mathbf{r}\right)}{\mathrm{d}\mathbf{r}} = \mathbf{C}^{\top} \otimes \mathbf{I}_{N_{X}}, \quad (3)$$

where  $\mathbf{I}_{N_X}$  is the  $N_X \times N_X$  identity matrix. Hence the determinant of the Jacobian,  $|\mathbf{J}| := |\det \mathbf{J}|$ , is independent of  $\mathbf{R}$ :

$$|\mathbf{J}| = |\mathbf{C}^{\top} \otimes \mathbf{I}_{N_X}| = \sqrt{|(\mathbf{C}^{\top} \otimes \mathbf{I}_{N_X})^{\top} (\mathbf{C}^{\top} \otimes \mathbf{I}_{N_X})|} = \sqrt{|\mathbf{C}\mathbf{C}^{\top}|}.$$
 (4)

As the canonical model  $\mathcal{M}$  (hence  $\mathbf{X}_0$ ) and  $\mathbf{R}_{gt}$  are known,  $|\mathbf{J}|$  remains constant leading to:

$$p(\mathbf{R} \mid \mathbf{I}) = p(\mathbf{X}' \mid \mathbf{I}) \cdot \sqrt{|\mathbf{C}\mathbf{C}^{\top}|} \propto p(\mathbf{X}' \mid \mathbf{I}).$$

With this proportionality, we can instead model  $p(\mathbf{X}'|\mathbf{I})$  in a symmetry aware fashion via the *product of experts* (PoEs) [18]:

$$p(\mathbf{X}' \mid \mathbf{I}) = \frac{1}{\int \prod_{i} \hat{p}_{i}(\mathbf{X}' \mid \mathbf{I}) \,\mathrm{d}\mathbf{X}'} \prod_{i} \hat{p}_{i}(\mathbf{X}' \mid \mathbf{I}) = \frac{1}{Z} \hat{p}_{\mathrm{SDF}}(\mathbf{X}' \mid \mathbf{I}) \hat{p}_{\mathrm{SE}}(\mathbf{X}' \mid \mathbf{I}), \quad (5)$$

where  $Z = \int \hat{p}_{\text{SDF}}(\mathbf{X}' | \mathbf{I}) \hat{p}_{\text{SE}}(\mathbf{X}' | \mathbf{I}) d\mathbf{X}'$ . Note that the product factorizes the conditional probability p into a set of simpler distributions, *experts*,  $\hat{p}_i$  which are normalized after multiplication. The first term,  $\hat{p}_{\text{SDF}}$ , will measure the deviation

from the model in the space of 3D coordinates. The second term,  $\hat{p}_{\rm SE}$ , will ensure that the symmetry-aware features of the transformed model align with the canonical ones under a given **R** (thus **X**). The PoE framework can work with unnormalized densities,  $\hat{p}(\cdot)$ , allowing us to model  $\hat{p}_{\rm SDF}$  and  $\hat{p}_{\rm SE}$  by a Bolzmann distribution:

$$\hat{p}_{\text{SDF}} \propto \exp\left(-\|f_{\text{SDF}}\left(\mathbf{X}'\right)\|_{0}\right) \text{ and } \hat{p}_{\text{SE}} \propto \exp\left(-\|f_{\text{SE}}\left(\mathbf{X}'\right) - f_{\text{SE}}\left(\mathbf{X}_{0}\right)\|_{\text{F}}\right), (6)$$

where  $f_{\text{SDF}} : \mathbb{R}^3 \to \mathbb{R}$  and  $f_{\text{SE}} : \mathbf{X} \to \mathbb{R}^{N_F \times N_X}$  are chosen to be the pretrained Deep-SDFs [41] and SurfEmb [15], respectively. Deep-SDF parameterizes the signed distance value to the closest surface point via a neural network. Using  $L_0$ -norm amplifies the penalty on more distant points. SurfEmb models continuous 2D-3D correspondence ( $\mathbf{X} \ni \mathbf{x} \leftrightarrow \mathbf{p} \in \mathbb{R}^2$ ) distributions  $p(\mathbf{x}, \mathbf{p} \mid \mathbf{I}) = p(\mathbf{p} \mid \mathbf{I})p(\mathbf{x} \mid \mathbf{I}, \mathbf{p})$  over the surface of objects, where  $p(\mathbf{p} \mid \mathbf{I})$  denotes the discrete distribution over image coordinates. SurfEmb consists of a network supervised to maximize the probability of the GT coordinates. The specific forms of  $f_{\text{SDF}}$  and  $f_{\text{SE}}$  are precised later in Sec. 3.1.

In our framework, we model these distributions by two MLPs  $f_{\theta}$  and  $f_{\phi}$  with learnable parameters  $\theta$  and  $\phi$ , which can be multiplied and passed through the softmax in Eq. (1) to yield  $p(\mathbf{R} | \mathbf{I})$ .

**Training.** During training, we have access to  $\{\mathbf{R}_{gt}^n\}_n$  for all images and thus corresponding  $\{\mathbf{X}'_n\}_n$ . While maximizing  $p(\mathbf{R} \mid \mathbf{I})$  by minimizing the negative log-likelihood for a single ground truth pose is viable, it requires more data samples to cover all symmetries and train the network. Another alternative is contrastive training (CT) as done in [16]. Yet, besides the challenge of sampling *hard* negatives over SO(3) during training, CT assumes a single positive label per sample. This resembles learning a multi-modal distribution by matching several unimodal distributions. Instead, we take a different approach and align the entire distribution. To learn the parameters of  $(f_{\theta}, f_{\phi})$  this way, we first obtain two unnormalized empirical measures  $(\boldsymbol{\mu}_{\text{SDF}}, \boldsymbol{\mu}_{\text{SE}})$  supported on SO(3), corresponding to  $\hat{p}_{\text{SDF}}$  and  $\hat{p}_{\text{SE}}$ , by randomly sampling SO(3) to get the locations and computing the individual terms in Eq. (6) as weights of these measures. Similarly, we let  $(\boldsymbol{\mu}_{\theta}, \boldsymbol{\mu}_{\phi})$  denote the inferred unnormalized measures (see "Inference") and use the generalized KL divergence [39] to align unbalanced distributions:

$$\operatorname{GKL}(\boldsymbol{\mu} \parallel \boldsymbol{\nu}) := \sum_{i=1}^{m} \left( -\log\left(\frac{a_i^{\boldsymbol{\nu}}}{a_i^{\boldsymbol{\mu}}}\right) + \left(\frac{a_i^{\boldsymbol{\nu}}}{a_i^{\boldsymbol{\mu}}}\right) - 1 \right) a_i^{\boldsymbol{\mu}}, \tag{7}$$

where  $\mathbf{a} = \{a_i^{\boldsymbol{\mu}}\}\)$  and  $\{\mathbf{R}_i^{\boldsymbol{\mu}}\}\)$ ,  $i = 1, \dots, m$  denote the weights and locations of a discrete measure  $\boldsymbol{\mu} = \sum_{i=1}^m a_i \delta_{\mathbf{R}_i^{\boldsymbol{\mu}}}, a_i \geq 0$ , where  $\delta_{\mathbf{R}^{\boldsymbol{\mu}}}$  is a Dirac delta at  $\mathbf{R}^{\boldsymbol{\mu}}$ . Finally, we train our networks by optimizing the following objectives:

$$\boldsymbol{\theta}^{\star} = \underset{\boldsymbol{\theta}}{\arg\min} \mathbb{E}_{p(\mathbf{I})} \mathrm{GKL}(\boldsymbol{\mu}_{\mathrm{SDF}} \parallel \boldsymbol{\mu}_{\boldsymbol{\theta}}), \quad \boldsymbol{\phi}^{\star} = \underset{\boldsymbol{\phi}}{\arg\min} \mathbb{E}_{p(\mathbf{I})} \mathrm{GKL}(\boldsymbol{\mu}_{\mathrm{SE}} \parallel \boldsymbol{\mu}_{\boldsymbol{\phi}}), \quad (8)$$

where  $p(\mathbf{I})$  denotes the data distribution. In summarry (see Fig. 1), to train our networks, we sample the SO(3) randomly or over a grid and use these samples

8 Vutukur et al.

in (i) explicitly computing a product-distribution by leveraging the spatial and feature-space distances and (ii) inferring these conditioned on the image. The divergence between the resulting measures becomes our supervision signal.

**Inference**. During test time, given an input image **I**, we infer the potentially multimodal distribution  $p(\mathbf{R} \mid \mathbf{I})$  by densely sampling SO(3) and evaluating  $-\log p(\mathbf{R} \mid \mathbf{I}) = f_{\theta}(\mathbf{I}, \mathbf{R}) + f_{\phi}(\mathbf{I}, \mathbf{R})$ . The individual predictions can be thought of as the empirical measures  $\mu_{\theta}$  and  $\mu_{\phi}$ , which are combined to get the final measure. We then pass the result through softmax to arrive at the full posterior.

#### 3.1 Network Design, Positional Encoding and Impl. Details

Network architecture. Our architecture is similar to IPDF with a ResNet50 image encoder and an MLP each for inferring unnormalized SDF and SE measures. The image encoder encodes the image as a global latent vector which is concatenated with the positional encoded rotation matrix and passed through two MLPs to get unnormalized SDF and SE measures. For a given image, we sample K rotations, to learn the distribution at K discrete points in SO(3) space. To train the network, for a given image, we generate the supervision signal using pre-trained SE and SDF MLPs to gather scores for all rotations for supervising the network. The pre-trained SE and SDF MLPs are two layer Siren MLPs that predict per-point feature and SDF value respectively from a 3D point input.

Sampling. Using our CAD-based priors, SurfEmb and SDF MLPs, we can generate the distribution for viewpoints before we start the training. We sample rotations on an equivolumetric partitioning of SO3 grid proposed by Yershova [58]. We subdivide the grid up to 5 levels and estimate distribution for certain viewpoints. Using this precomputed distribution, we can sample more rotations near the modes which helps us in learning distribution better. In our approach, instead of estimating distribution for the entire dataset, we estimate distributions for a few viewpoints. It is a tradeoff between memory and accuracy. We save distribution for a few viewpoints at a higher resolution grid than saving distribution for every viewpoint at a lower level grid. This is because the distribution can be transferred to another viewpoint by rotating the grid accordingly using the relative rotation between the given viewpoint and the viewpoint for which the distribution is estimated. We randomly sample a viewpoint from the precomputed distributions and convert it to the current viewpoint and sample the rotations from this distribution. Even though the viewpoints don't exactly have the same distribution in some cases, this won't affect the learning as this will be reflected in the estimated probability and still a bad rotation will be learned as such. Even if we sample rotations away from modes sometimes, the supervision signal will provide a low probability, but it is important to sample more near the modes to learn a sharper distribution. We precompute both shape and feature based distribution for a few viewpoints and use them during training for better sampling. Shape-based precomputed distribution will be valid even in the case of conditionally symmetric objects with ambiguities that break symmetry.

**Positional encoding(PE)**. IPDF employs a PE that generates high frequency components of individual rotation matrix elements separately. PE is not applied to the rotation manifold leading to errors in the training where multiple rotations that are far apart in SO3 space with similar element-wise absolute values are closer in positional encoding space in some cases. Specifically, when some negative values are lower in magnitude in the rotation matrix, they get a similar PE to when the sign is flipped in another rotation matrix since sines are getting suppressed because of low value and cosine is getting amplified at the same time and also suppressing the sign. The PE is not exactly the same, but the network treats them as similar because the change is minute in cartesian space. To alleviate this issue, we take the 3D corner coordinates of a unit cube and apply positional encoding on the rotated 3D corner coordinates of the cube. The positionally encoded corner coordinates are concatenated to generate a vector that serves as a PE representation of the rotation matrix. This reduces the noise as shown in Figure 4 and also brings a boost in accuracy.

## 4 Experiments

We evaluate on Symsol [40], T-Less [20], and ModelNet10-SO3 [36]. Symsol has two subsets, Symsol-I with untextured symmetric objects and Symsol-II with conditionally symmetric objects with some texture that breaks symmetry from certain viewpoints. We adopt the metrics from IPDF, log-likelihood(LL), and Average Recall(AR) for evaluation. LL measures how strongly the learned network captures the distribution. It is calculated by taking the log mean of the probabilities at the ground truth rotations in our captured distribution. As we can't estimate distribution in continuous rotation space, we compute our distribution on a discrete equivolumetric grid of rotations similar to IPDF. We use 5 levels of subdivision to obtain a rotation grid with around 2M rotations for evaluation. We compute the distribution on the grid and evaluate the LL for rotations at the discrete rotation samples on the grid closest to ground truth rotations. LL evaluation can be carried out with a single ground truth annotation or with multiple ground truth annotation for a single image. To quantify the rotation error on ModelNet, we employ AR with  $30^{\circ}$  rotation error threshold. We employ distribution visualization proposed in IPDF where each point on sphere indicates the axis of rotation and the color indicates the tilt about that axis.

### 4.1 SYMSOL-I

Symsol-I comprises textureless geometric shapes, Cube, Tetrahedron, Cone, Icosahedron, and Cylinder that exhibit different types of symmetries and pose a challenging task to capture all the symmetric configurations with high probability. The dataset provides 45k images for each object and they do not provide camera intrinsics or a CAD model. However, these generic CAD models are easily available from any 3D library and are already canonically aligned with the dataset. We employ PBR data [9] and SDF samples from CAD models to

**Table 1:** SYMSOL I results: We present results on Symsol-I evaluated using the loglikelihood metric. Models refers to the number of models trained per dataset. Iterations refer to training iterations and Images refer to the amount of training images employed per object. Note that we follow the convention employed in the Normalizing Flow(NF) paper by adding 2.29 to the scores to make a uniform distribution to have a zero score instead of having -2.29. We adjusted the scores for us and Spyropose(SP) and other approaches were already adjusted in NF paper. We present results for Deng [8], Gil [13], Prok [45], IPDF [40], SP [16] and NF [37].

	Deng	Gil.	Prok.	IPDF	$\mathbf{SP}$	NF	NF	NF	NF	Ours	Ours
Models	1	1	1	1	5	1	1	1	1	1	1
Iterations	100k	100k	100k	100k	100k	100k	900k	100k	900k	100k	100k
Images	45k	45k	45k	45k	45k	10k	10k	45k	45k	10k	45k
cone	2.45	6.13	-1.05	6.74	9.91	8.45	8.94	8.42	10.05	9.66	10.10
cube	-2.15	0.00	1.79	7.10	10.92	5.02	9.01	7.13	11.64	11.29	12.24
cyl	1.34	3.17	1.01	6.55	8.75	8.04	6.41	7.83	9.54	9.32	9.40
icosa	-0.16	0.00	-0.10	3.57	7.52	-2.14	-6.03	2.03	8.26	7.99	9.54
tet	2.56	0.00	0.43	7.99	10.98	5.91	10.79	8.98	12.43	11.39	11.96
avg	0.81	1.86	0.42	6.39	9.62	5.06	5.82	6.88	10.38	9.69	10.64

train  $f_{SE}$ ,  $f_{SDF}$  respectively. We use the frozen  $f_{SE}$ ,  $f_{SDF}$  MLPs for training our pipeline. The test set provides ground truth annotations for all the symmetric configurations and is considered in the evaluation for log-likelihood and spread. Prior approaches focus solely on learning distribution from images and show no way to incorporate a CAD model into the training pipeline. Spyropose tries to incorporate a CAD model, but only uses it to strengthen the encoder instead of using a Resnet based encoder. However, their approach requires camera intrinsics. So, they render their version of symsol to use the encoder. Moreover, their cad-based encoder requires the object instance to be known to encode the image.

To have a common baseline, we compare all the approaches using Resnet image encoder. Using a basic architecture such as IPDF, we achieve benchmark results with log-likelihood scores of 10.64 and 9.69 in full-data and low-data regimes using our training pipeline by distilling knowledge from a CAD model to learn distribution as shown in Table 1. Our approach converges faster and performs better than NF, SpyroPose with just 100k iterations. In a low-data regime with



**Fig. 2:** LL vs. Training Data for NF, Ours on SYMSOL-I.

10k images, the accuracy drop(0.95) is not significant depicting the benefits of a CAD prior. Besides, NF has weaker performance with low data and negatively affects some objects with increasing iterations. This demonstrates the advantage of employing a CAD model which is crucial as real data is not easily available in large amounts. Our approach can transfer learned features from the CAD model and help learn the distribution better in the real domain even with fewer images.

**Table 2:** SYMSOL-II results: We report results using log-likelihood metric(LL). We adjust the scores for us and SpyroPose using the convention of normalizing flows(NF) by adding 2.29 to make a uniform distribution to have an average LL of score zero.-10k refers to experiments in low-data regime using only 10k images instead of 45k.

Obj	Deng	Gil.	Prok.	IPDF	$\mathbf{SP}$	$\operatorname{SP-10k}$	NF-10k	$\mathbf{NF}$	Ours-10k	Ours
SphX	3.41	5.61	-1.90	9.59	11.36	7.67	7.62	12.37	6.32	10.93
cylO	5.28	7.17	6.45	9.20	11.61	9.11	6.99	12.92	11.57	12.18
$\mathrm{tet}\mathbf{X}$	5.90	5.19	3.77	10.78	11.70	6.48	3.52	13.53	11.53	12.38
LL	4.86	5.99	2.77	9.86	11.56	7.76	6.04	12.94	9.80	11.83

As shown in Fig. 2, the gap between the LLs attained by our method and those by NF increases rapidly as the amount of training data decreases. We present some qualitative results in Fig. 3 as proposed in IPDF.

### 4.2 SYMSOL-II

Symsol-II comprises conditionally symmetric objects, tetrahedron, cylinder, and sphere, with markers. So, the modes are reduced when the markers are visible. For conditional symmetric objects, the pose distribution not only shifts but also changes completely based on the visibility of markers. We fit the meshes with silhouette images and estimate intrinsics which is later used to fit the texture of the mesh using color images through differentiable rendering. We train  $f_{\rm SE}$ ,  $f_{\rm SDF}$ networks with the obtained mesh. Our approach has slightly weaker performance compared to NF as shown in Table 2. On the Sphere object, our SDF expert doesn't contribute anything as it learns a uniform distribution and it is a limitation of our approach. Hence, we have to rely only on the SurfEmb expert to explain the distribution which decreases the performance. However, we perform better than NF and Spyropose in the low-data regime with 10k images. Besides, LL metric inherently favors NF as it provides exact LL estimation while IPDF and our approach evaluate LL on a discrete rotation grid, making our methods limited by the finite resolution. The performance drop is less significant on textured cylinder and tetrahedron indicating that the shape component can bring a significant boost to the distribution. LL scores for tetrahedron and cylinder are higher than the Symsol-I dataset since Surfemb features provide information about the texture in these objects and perform better in conjunction with shape experts to capture the textured viewpoints with much finer distribution.

#### 4.3 T-Less

T-Less contains 30 texture-less symmetric objects. We follow [13] to split data into 75% for training and 25% for testing. We train SurfEmb,  $f_{\rm SE}$ , using PBR and 75% of the real data.  $f_{\rm SDF}$  network is trained using the CAD model. We perform better than the benchmark approach, SpyroPose, as shown in Table 3. The dataset has minute symmetry breaking features that can turn a multimodal



Fig. 3: Pose distribution visualization for objects in Symsol-I and Symsol-II. a) Untextured Cylinder has continuous symmetry b) Textured Cylinder with marker has a unimodal distribution when the marker is visible, c) Broken continuous symmetry on the textured cylinder when the marker is not visible d) Untextured Tetrahedron in Symsol-I has 12 modes which are captured appropriately e) Textured Tetrahedron has three modes when orange face is visible. f) Textured terahedron has 6 modes when the orange face is not visible. Note that only one ground truth annotation is provided in Symsol-II and hence only one mode is circled.

distribution into an unimodal distribution. Although they provide symmetry annotations, the CAD models are not perfectly symmetric like the objects from Symsol-I. So, we follow the other approaches, and the evaluation is performed only on a single ground truth pose similar to Symsol-II. We perform an ablation with training data used for SurfEmb. SurfEmb provides better features with PBR+Real data compared to using just PBR data and is reflected in an increase in LL of 0.2. The results indicate the ability to notice the symmetry breaking features to capture the distribution more sharply when they are visible.

**Table 3:** T-Less Results: We present T-Less results using the log-likelihood measure. P refers to using only PBR data and P+R refers to using both PBR and Real data for training the Surfemb network. Please note that we still use real data to train our distribution network in all cases. L1 and GKL refer to the loss used for training.

Method	Prokudin [45]	Gilitschenski	[13] IPDF [40]	Spyropose [16]		Our	s
SE data	-	-	-	-	Р	Р	+R
Loss	-	-	-	-	L1	L1	GKL
LL	11.0	9.1	12.0	14.1	13.6	13.8	14.53

### 4.4 ModelNet10-SO3

ModelNet10-SO3 is an unimodal category-level dataset comprising images of CAD models of 10 categories. We choose a single CAD model per category to train  $f_{\rm SE}$ ,  $f_{\rm SDF}$ . We employ images from all CADs for training our pipeline, but the distribution supervision comes from the chosen CAD model. The distribution supervision should come ideally from the CAD model of the specific instance in the image, but it is hard to train  $f_{\rm SE}$  for all CADs in the category. We compute the distribution on grid and select the rotation with the highest score for evaluation. We achieve an AR of 70.5% at 30° rotation error threshold compared

to NF(77.4%) despite using distribution supervision from a single CAD model per category. Besides, the dataset contains some symmetric models, but only one GT rotation is considered for evaluation leading to lower accuracy.

#### 4.5 Ablations

**Positional encoding(PE)**. We perform an ablation on different PE of rotation matrices. The PE of elements of rotation matrix leads to noisy distribution as shown in Figure 4. This happens when rotation matrices with similar absolute values sometimes get a similar position encoding leading to noise. This could be avoided by parameterizing the encoding of the rotation matrix on manifold by employing Wigner matrices [11]. We propose a PE where corners of a cube are rotated with rotation matrices and then PE is applied on the transformed 3d coordinates. Both Wigner matrices and our cube position encoding(cube PE) were able to remove noise from the distribution. However, our cube PE performs better compared to Wigner and IPDF PE as shown in Table 4.



**Fig. 4:** Pose distribution visualization of our approach with different positional encodings. Our cubePE and wigner matrices based encoding removes noise in the distribution compared to encoding elements of rotation matrices from IPDF. Our encoding removes noise and also has a better performance compared to Wigner matrices. a) Input Image, b)IPDF positional encoding, c) Rotation encoded as Wigner matrices d) Our positional encoding with cube vertices. e) color wheel indicating the tilt about the axis.

**Sampling**. During training, we can precompute the distribution for each image through shape and feature based distribution estimation. IPDF employs random sampling that leads to learning a smoother distribution making it difficult to capture modes sharply. As we can convert a single ground truth rotation into the entire distribution, we can sample more rotations near modes and focus on learning the distribution sharply. In Table 4, random sampling performs worse compared to grid-5 and grid-6 sampling. The grid level refers to HEALPix grid [14] level at which we estimate our distribution. Sampling on a finer grid leads to improved accuracy. Mode focused sampling is crucial to get sharper modes.

Geometry and SurfEmb based features. An ablation to see how Shape and SurfEmb feature based probability affect LL in the Symsol-I. LL from the shape and feature modules achieve decent performance separately but perform better when combined to estimate the final distribution. Feature-based probability is appropriate for capturing shape and texture information. Shape based probability doesn't consider texture and it cannot perform alone on Symsol-II. Feature based probability is essential to capture the distribution better in the presence of texture. On the other hand, learning could be affected by the quality of the features from another learned network(SurfEmb) which is not true for shape

#### 14 Vutukur et al.

**Table 4:** Ablation table: Ablations on SYMSOL-I dataset with positional encoding, sampling strategy, SurfEmb Feature and Shape components, loss terms used for training. Encoding refers to the rotation encoding employed in the approach. Sampling refers to the strategy employed for sampling rotations during training. Random refers to random sampling. Module refer to module employed to estimate the final distribution. S refers to SDF based module and F refers to the SurfEmb Feature based module. The combined result of SDF and Feat modules provides the best accuracy.

Encoding	IPDF	Wigner	Cube							
Sampling	grid-5	grid-5	random	grid-5	grid-6					
Module	F+S	F+S	F+S	F+S	$\mathbf{S}$	F	F+S	F+S		
Loss	L1	L1	L1	L1	GKL	GKL	L1	GKL		
LL	10.09	8.59	6.94	10.2	10.32	10.18	10.48	10.648		

based probability as it is based on ground truth shape and more reliable for capturing the distribution sharply. In essence, both of the scoring mechanisms together contribute to the betterment of the approach as shown in Table 4.

Loss function. Generalized KL divergence (GKL) performs better than the L1 loss as shown in Table 4. This is reflected in T-Less dataset as well from Table 3 with an increase in LL of 0.6. This indicates that the choice of GKL is well suited for formulating loss between unnormalized probability distributions compared to L1 loss, which fails to be a natural fit for comparing distributions.

## 5 Conclusion

We proposed Alignist, a novel approach to learn pose distributions over images by mapping the conditional estimation from a given image to one over a prototypical CAD model, acquired either as ground truth or via a 3D reconstruction. Our probabilistic framework uses a product of experts corresponding to losses over SDF and SurfEmb features, and anchors on full distributional distances computed by a generalized KL divergence, rather than considering individual samples as in normalizing flows. Precomputable distribution with the help of the CAD helps in better sampling leading to sharper distribution. Our novel positional encoding further reduces noise around modes and allows for learning cleaner distributions. We achieve benchmark accuracy in the Symsol-I and T-Less datasets, especially outperforming the state-of-the-art in low data regimes. Limitations and future work. Our approach does not explicitly utilize the texture cues either coming from the image or the CAD model. Instead, it exploits the features from SurfEmb which implicitly depend upon both texture and geometry, simultaneously. Employing another expert that models the texture cues explicitly will enable our model to better handle objects like SphereX on Symsol-II where pure geometry cannot contribute towards the distribution. Future work also involves extending our work to diffusion models.

Acknowledgements. T. Birdal acknowledges support from the Engineering and Physical Sciences Research Council [grant EP/X011364/1].

## References

- Brégier, R., Devernay, F., Leyrit, L., Crowley, J.L.: Defining the pose of any 3d rigid object and an associated distance. International Journal of Computer Vision 126(6), 571–596 (2018)
- Brégier, R.: Deep regression on manifolds: A 3d rotation case study. In: 2021 International Conference on 3D Vision (3DV). pp. 166–174 (2021)
- Bui, M., Birdal, T., Deng, H., Albarqouni, S., Guibas, L., Ilic, S., Navab, N.: 6d camera relocalization in ambiguous scenes via continuous multimodal inference. In: European Conference on Computer Vision (ECCV) (2020)
- 4. Busam, B., Birdal, T., Navab, N.: Camera pose filtering with local regression geodesics on the riemannian manifold of dual quaternions. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 2436–2445 (2017)
- Busam, B., Jung, H.J., Navab, N.: I like to move it: 6d pose estimation as an action decision process. arXiv preprint arXiv:2009.12678 (2020)
- Chen, J., Yin, Y., Birdal, T., Chen, B., Guibas, L.J., Wang, H.: Projective manifold gradient layer for deep rotation regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6646–6655 (2022)
- Corona, E., Kundu, K., Fidler, S.: Pose estimation for objects with rotational symmetry. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7215–7222. IEEE (2018)
- Deng, H., Bui, M., Navab, N., Guibas, L., Ilic, S., Birdal, T.: Deep bingham networks: Dealing with uncertainty and ambiguity in pose estimation. International Journal of Computer Vision pp. 1–28 (2022)
- Denninger, M., Winkelbauer, D., Sundermeyer, M., Boerdijk, W., Knauer, M., Strobl, K.H., Humt, M., Triebel, R.: Blenderproc2: A procedural pipeline for photorealistic rendering. Journal of Open Source Software 8(82), 4901 (2023)
- 10. Do, T., Cai, M., Pham, T., Reid, I.D.: Deep-6dpose: Recovering 6d object pose from a single RGB image. CoRR (2018)
- Esteves, C., Lu, T., Suhail, M., Chen, Y.f., Makadia, A.: Generalized fourier features for coordinate-based learning of functions on manifolds. ICLR 2022 Submission (2021)
- Gilitschenski, I., Sahoo, R., Schwarting, W., Amini, A., Karaman, S., Rus, D.: Deep orientation uncertainty learning based on a bingham loss. In: International Conference on Learning Representations (2020)
- Gilitschenski, I., Sahoo, R., Schwarting, W., Amini, A., Karaman, S., Rus, D.: Deep orientation uncertainty learning based on a bingham loss. In: International Conference on Learning Representations (2020)
- Gorski, K.M., Hivon, E., Banday, A.J., Wandelt, B.D., Hansen, F.K., Reinecke, M., Bartelmann, M.: Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. The Astrophysical Journal 622(2), 759 (2005)
- Haugaard, R.L., Buch, A.G.: Surfemb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6749–6758 (2022)
- Haugaard, R.L., Hagelskjær, F., Iversen, T.M.: Spyropose: Se (3) pyramids for object pose distribution estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2082–2091 (2023)

- 16 Vutukur et al.
- Hinterstoisser, S., Lepetit, V., Ilic, S., Holzer, S., Bradski, G., Konolige, K., Navab, N.: Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In: Computer Vision–ACCV 2012: 11th Asian Conference on Computer Vision, Daejeon, Korea, November 5-9, 2012, Revised Selected Papers, Part I 11. pp. 548–562. Springer (2013)
- Hinton, G.E.: Products of experts. In: Artificial Neural Networks, 1999. ICANN 99. Ninth International Conference on (Conf. Publ. No. 470). vol. 1. IET (1999)
- Hodan, T., Barath, D., Matas, J.: Epos: Estimating 6d pose of objects with symmetries. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11703–11712 (2020)
- Hodan, T., Haluza, P., Obdrzalek, S., Matas, J., Lourakis, M., Zabulis, X.: T-LESS: An RGB-D dataset for 6D pose estimation of texture-less objects. IEEE Winter Conference on Applications of Computer Vision (WACV) (2017)
- Hodan, T., Michel, F., Brachmann, E., Kehl, W., GlentBuch, A., Kraft, D., Drost, B., Vidal, J., Ihrke, S., Zabulis, X., et al.: Bop: Benchmark for 6d object pose estimation. In: Proceedings of the European conference on computer vision (ECCV). pp. 19–34 (2018)
- Höfer, T., Kiefer, B., Messmer, M., Zell, A.: Hyperposepdf-hypernetworks predicting the probability distribution on so (3). In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2369–2379 (2023)
- Huang, J., Wang, H., Birdal, T., Sung, M., Arrigoni, F., Hu, S.M., Guibas, L.J.: Multibodysync: Multi-body segmentation and motion estimation via 3d scan synchronization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7108–7118 (2021)
- 24. Huang, J., Yu, H., Yu, K.T., Navab, N., Ilic, S., Busam, B.: Matchu: Matching unseen objects for 6d pose estimation from rgb-d images (2024)
- Iversen, T.M., Haugaard, R.L., Buch, A.G., Robotics, S.: Ki-pode: Keypoint-based implicit pose distribution estimation of rigid objects. In: BMVC (2022)
- Johansen, B., Euchner, J.: Navigating the vuca world. Research-Technology Management 56(1), 10–15 (2013)
- 27. Kehl, W., Manhardt, F., Tombari, F., Ilic, S., Navab, N.: Ssd-6d: Making rgbbased 3d detection and 6d pose estimation great again. In: Proceedings of the IEEE international conference on computer vision. pp. 1521–1529 (2017)
- Kendall, A., Grimes, M., Cipolla, R.: Posenet: A convolutional network for realtime 6-dof camera relocalization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (December 2015)
- Klee, D., Biza, O., Platt, R., Walters, R.: Image to sphere: Learning equivariant features for efficient pose prediction. In: The Eleventh International Conference on Learning Representations (2022)
- Kundu, A., Li, Y., Rehg, J.M.: 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
- 31. Landau, L.D., Lifshitz, E.M.: Course of theoretical physics. Elsevier (2013)
- 32. LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., Huang, F.: A tutorial on energybased learning. Predicting structured data 1(0) (2006)
- Lee, T., Jang, Y., Kim, H.J.: Object-based slam utilizing unambiguous pose parameters considering general symmetry types. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2120–2126 (2023)
- Levinson, J., Esteves, C., Chen, K., Snavely, N., Kanazawa, A., Rostamizadeh, A., Makadia, A.: An analysis of svd for deep rotation estimation. arXiv preprint arXiv:2006.14616 (2020)

- 35. Li, F., Vutukur, S.R., Yu, H., Shugurov, I., Busam, B., Yang, S., Ilic, S.: Nerfpose: A first-reconstruct-then-regress approach for weakly-supervised 6d object pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2123–2133 (2023)
- 36. Liao, S., Gavves, E., Snoek, C.G.M.: Spherical regression: Learning viewpoints, surface normals and 3d rotations on n-spheres. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Long Beach, USA (2019)
- 37. Liu, Y., Liu, H., Yin, Y., Wang, Y., Chen, B., Wang, H.: Delving into discrete normalizing flows on so (3) manifold for probabilistic rotation modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21264–21273 (2023)
- Manhardt, F., Arroyo, D.M., Rupprecht, C., Busam, B., Birdal, T., Navab, N., Tombari, F.: Explaining the ambiguity of object detection and 6d pose from visual data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6841–6850 (2019)
- Miller, B.K., Federici, M., Weniger, C., Forré, P.: Simulation-based inference with the generalized kullback-leibler divergence. In: 1st Workshop on the Synergy of Scientific and Machine Learning Modeling@ ICML2023 (2023)
- Murphy, K.A., Esteves, C., Jampani, V., Ramalingam, S., Makadia, A.: Implicitpdf: Non-parametric representation of probability distributions on the rotation manifold. In: Proceedings of the 38th International Conference on Machine Learning. pp. 7882–7893 (2021)
- 41. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2019)
- 42. Peretroukhin, V., Giamou, M., Rosen, D.M., Greene, W.N., Roy, N., Kelly, J.: A Smooth Representation of SO(3) for Deep Rotation Learning with Uncertainty. In: Proceedings of Robotics: Science and Systems (RSS'20) (Jul 12–16 2020)
- Pitteri, G., Ramamonjisoa, M., Ilic, S., Lepetit, V.: On object symmetries and 6d pose estimation from images. In: 2019 International conference on 3D vision (3DV). pp. 614–622. IEEE (2019)
- 44. Placed, J.A., Strader, J., Carrillo, H., Atanasov, N., Indelman, V., Carlone, L., Castellanos, J.A.: A survey on active simultaneous localization and mapping: State of the art and new frontiers. IEEE Transactions on Robotics (2023)
- Prokudin, S., Gehler, P., Nowozin, S.: Deep directional statistics: Pose estimation with uncertainty quantification. In: European Conference on Computer Vision (ECCV). vol. 9, pp. 542–559 (Sep 2018)
- 46. Shi, G., Zhu, Y., Tremblay, J., Birchfield, S., Ramos, F., Anandkumar, A., Zhu, Y.: Fast uncertainty quantification for deep object pose estimation. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 5200–5207. IEEE (2021)
- Sorrenson, P., Draxler, F., Rousselot, A., Hummerich, S., Köthe, U.: Learning distributions on manifolds with free-form flows. arXiv preprint arXiv:2312.09852 (2023)
- 48. Su, H., Qi, C.R., Li, Y., Guibas, L.J.: Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (December 2015)
- Su, Y., Saleh, M., Fetzer, T., Rambach, J., Navab, N., Busam, B., Stricker, D., Tombari, F.: Zebrapose: Coarse to fine surface encoding for 6dof object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6738–6748 (2022)

- 18 Vutukur et al.
- Teed, Z., Deng, J.: Tangent space backpropagation for 3d transformation groups. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10333–10342. IEEE Computer Society (2021)
- Ummenhofer, B., Zhou, H., Uhrig, J., Mayer, N., Ilg, E., Dosovitskiy, A., Brox, T.: Demon: Depth and motion network for learning monocular stereo. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
- Vutukur, S.R., Brock, H., Busam, B., Birdal, T., Hutter, A., Ilic, S.: Nerf-feat: 6d object pose estimation using feature rendering. In: 2024 International Conference on 3D Vision (3DV). pp. 1146–1155 (2024)
- 53. Vutukur, S.R., Shugurov, I., Busam, B., Hutter, A., Ilic, S.: Welsa: Learning to predict 6d pose from weakly labeled data using shape alignment. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII. pp. 645–661. Springer (2022)
- Wang, G., Manhardt, F., Tombari, F., Ji, X.: Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16611– 16621 (2021)
- 55. Wang, H., Sridhar, S., Huang, J., Valentin, J., Song, S., Guibas, L.J.: Normalized object coordinate space for category-level 6d object pose and size estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2642–2651 (2019)
- Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In: Robotics: Science and Systems (RSS) (2018)
- 57. Yang, H., Pavone, M.: Object pose estimation with statistical guarantees: Conformal keypoint detection and geometric uncertainty propagation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8947–8958 (2023)
- Yershova, A., LaValle, S.M., Mitchell, J.C.: Generating Uniform Incremental Grids on SO(3) Using the Hopf Fibration, pp. 385–399. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
- 59. Yi, L., Huang, H., Liu, D., Kalogerakis, E., Su, H., Guibas, L.: Deep part induction from articulated object pairs. ACM Transactions on Graphics **37**(6) (2019)
- Yin, Y., Cai, Y., Wang, H., Chen, B.: Fishermatch: Semi-supervised rotation regression via entropy-based filtering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11164–11173 (2022)
- Yin, Y., Lyu, J., Wang, Y., Wang, H., Chen, B.: Towards robust probabilistic modeling on so (3) via rotation laplace distribution. arXiv preprint arXiv:2305.10465 (2023)
- 62. Yin, Y., Wang, Y., Wang, H., Chen, B.: A laplace-inspired distribution on SO(3) for probabilistic rotation estimation. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023. OpenReview.net (2023)
- Zhao, Y., Birdal, T., Lenssen, J.E., Menegatti, E., Guibas, L., Tombari, F.: Quaternion equivariant capsule networks for 3d point clouds. In: European Conference on Computer Vision. pp. 1–19. Springer (2020)
- 64. 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)