

---

# Probabilistic Orientation Estimation with Matrix Fisher Distributions

---

**David Mohlin**  
KTH/Tobii  
Stockholm, Sweden  
davmo@kth.se

**Gérald Bianchi**  
Tobii  
Danderyd, Sweden  
gerald.bianchi@tobii.com

**Josephine Sullivan**  
KTH  
Stockholm, Sweden  
sullivan@kth.se

## Abstract

This paper focuses on estimating probability distributions over the set of 3D rotations ( $SO(3)$ ) using deep neural networks. Learning to regress models to the set of rotations is inherently difficult due to differences in topology between  $\mathbb{R}^N$  and  $SO(3)$ . We overcome this issue by using a neural network to output the parameters for a matrix Fisher distribution since these parameters are homeomorphic to  $\mathbb{R}^9$ . By using a negative log likelihood loss for this distribution we get a loss which is convex with respect to the network outputs. By optimizing this loss we improve state-of-the-art on several challenging applicable datasets, namely Pascal3D+, ModelNet10- $SO(3)$ . Our code is available at [https://github.com/Davmo049/Public\\_prob\\_orientation\\_estimation\\_with\\_matrix\\_fisher\\_distributions](https://github.com/Davmo049/Public_prob_orientation_estimation_with_matrix_fisher_distributions)

## 1 Introduction

Estimating the 3D rotation of an object from 2D images is one of the fundamental problems in computer vision. Several applications relying on 3D rotation estimation have been developed such as a robot grasping an object [23], a self driving vehicle constantly sensing its surrounding environment [18], an augmented reality system combining computer-generated information onto the real world [16], or a system detecting the face orientation to enhance human-computer interactions [25].

Advances of deep learning techniques have resulted in improvements in estimation of 3D orientation. However, precise orientation estimation remains an open problem. The main problem is that the space of all 3D rotations lies on a nonlinear and closed manifold, referred to as the special orthogonal group  $SO(3)$ . This manifold has a different topology than unconstrained values in  $\mathbb{R}^N$ , where neural network outputs exist. As a result it is hard to design a loss function which is continuous without disconnected local minima. For example using euler angles as an intermediate step causes problems due to the so-called gimbal lock. Quaternions have a double embedding giving rise to the existence of two disconnected local minimas. Some more complicated methods use Gram-Schmidt [30] which has a continuous inverse, but the function is not continuous with a discontinuity when the input vectors do not span  $\mathbb{R}^3$ .

Despite these issues various deep learning based solutions have been suggested. One approach is to use one of the rotation representations and model the constraint in the loss function or in the network architecture [14]. An alternative is to construct a mapping function, which directly converts the network output to a rotation matrix [30].

Quantifying the 3D orientation uncertainty when dealing with noisy or otherwise difficult inputs is also an important task. Uncertainty estimation provides valuable information about the quality of the prediction during the process of decision making. Only recent efforts have been made on modeling the uncertainty of 3D rotation estimation [20, 5]. However, those methods still rely on complex solutions to fulfil the constraints required by their parameterization.

In this paper, we instead propose a deep learning approach to estimate the 3D rotation uncertainty by using the matrix Fisher probability density function developed in the field of directional statistics [17]. This unimodal distribution has been selected because of its relevant properties in regards to the problem of 3D orientation estimation: i) The parameterization is unconstrained so there is no need for complex functions to enforce constraints. ii) It is possible to create a loss for this distribution which has desirable properties such as convexity iii) The mode of the distribution can subsequently be estimated along with the uncertainty around that mode for further analysis.

Our method offers a simple solution to the problem of 3D orientation estimation, where a neural network learns to regress the parameters of the matrix Fisher distribution. While several other losses used for rotation estimation are discontinuous with non path connected sublevel sets, with respect to the network output. In this paper we instead propose a loss which is convex with bounded gradient magnitudes, resulting in a stable training. In addition, we also implement a method for computing the the non-trivial normalizing constant of the distribution. Finally, the proposed method is evaluated on multiple problem domains and compared with the latest published approaches. The results show that our method outperforms all previous methods for several problems.

Our contributions include: 1) a method for estimating a probability distributions over the set of rotations with neural networks by using the matrix Fisher distribution, 2) a loss associated with this distribution and show it is convex with bounded gradients, and 3) an extensive analysis encompassing several datasets and recent orientation estimation works, where we demonstrate the superiority of our method over the state-of-the-art.

## 2 Related Work

3D rotation estimation has been studied over the last two decades. A common method estimates 3D rotation by aligning two sets of 3D feature points where each data set is matched and defined in a different coordinate system [4]. Another well-known approach matches 2D keypoints extracted from images with features of a known 3D model and recovers the 3D pose given the 2D-3D correspondences [21, 10]. With the advances of deep learning, the detection of 2D keypoints has significantly been improved [2]. The keypoints can be associated either with physical points [19] or with virtual points such as the corners of an object’s bounding box [6].

Recent methods using deep networks often predict 3D rotation directly from images without the knowledge of a 3D model of the object. Those methods can be grouped in two categories. The first one divides the set of rotations into regions and subsequently solves the 3D orientation estimation as a classification task. Subsequently, the classification network output is refined by a regression network [15, 12].

The second category transforms the network output to a 3D rotation representation and learns to directly regress the 3D rotation given an image input. Commonly, quaternions [28] or Euler angles [14] representation are used. However, the paper [30] shows any rotation representation of dimensions four or lower is discontinuous, which makes it difficult for the neural network to generalize over the set of rotations. They propose two continuous 5D and 6D rotation representations and construct a function that maps those representations to a rotation matrix.

Recently some studies have investigated the prediction of rotation uncertainty using probability distributions over rotations. In [20] the parameters of a mixture of von Mises distribution using a biternion network are estimated. In [5], the Bingham distribution over quaternions is used to jointly estimate a probability distribution over all rotation axes. However, their parameters have to be positive semidefinite due to their choice of probability distribution.

In this paper, we propose a solution which learns to regress the probability distribution with unconstrained parameters leading to a simple formulation of the problem of 3D rotation estimation.

## 3 Method

We train a neural network to estimate the 3D orientation of objects in an input image. Specifically, the network outputs the parameters of the matrix Fisher distribution, which is a distribution over  $SO(3)$ . From the predicted parameters we can obtain the maximum likelihood estimate of the input’s orientation. In the rest of this section we review the matrix Fisher distribution and provide




Figure 1: **Visualizing the matrix Fisher distribution on  $SO(3)$ .** We follow the convention of [9] and recreate their figures to explain the approach, similarly for figure 2. For the above plots the parameter matrix is  $F = \text{diag}(5, 5, 5)$ . Let  $e_1, e_2$  and  $e_3$  correspond to the standard basis of  $\mathbb{R}^3$  and is shown by the black axes. (a) This plot shows the probability distribution of  $Re_1$  when  $R \sim \mathcal{M}(F)$ . Thus the pdf shown on the sphere corresponds to the probability of where the  $x$ -axis will be transformed to after applying  $R \sim \mathcal{M}(F)$ . (b) and (c) Same comment as (a) except consider  $e_2$  and  $e_3$  instead of  $e_1$ . (d) A compact visualization of the plots in (a), (b) and (c) is obtained by summing the three marginal distributions and displaying them on the 3D sphere. All four plots are plotted within the same scale and a *jet* colormap is used.

some visualizations to help the reader’s interpretation of its parameters. Then we derive the loss, based on maximizing the likelihood of the labelled data, and finally explain how we deal with the distribution’s complex normalizing constant when we calculate our loss and calculate the gradient during back-propagation.

### 3.1 The matrix Fisher distribution on $SO(3)$

We model 3D rotation matrices probabilistically with the matrix Fisher distribution [3, 8]. This distribution has probability density function

$$p(R | F) = \frac{1}{a(F)} \exp(\text{tr}(F^T R)) \quad (1)$$

where  $F$  is an unconstrained matrix in  $\mathbb{R}^{3 \times 3}$  parametrising the distribution,  $R \in SO(3)$ , and  $a(F)$  is the distribution’s normalizing constant. We will denote that  $R$  is distributed according to a matrix Fisher distribution with  $R \sim \mathcal{M}(F)$ . The distribution is unimodal but visualizing the distribution in equation (1) is hard as it has a 3D domain. Fortunately, [9] describes a helpful visualization scheme, see figure 1 for details, which we use throughout the paper.

Also not immediately apparent is how the shape of the distribution varies as  $F$  varies. From [4] we know the mode of the distribution can be computed from the singular value decomposition of  $F = USV^T$ , where the singular values are sorted in descending order, and setting

$$\hat{R} = U \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \det(UV) \end{bmatrix} V^T \quad (2)$$

This is done to ensure that the orientation  $\hat{R}$  has determinant 1 and is orthonormal. Similar results are available in [3]. Figure 2 displays examples of the distribution for simple  $F$  matrices. These figures show that larger singular values correspond to more peaked distributions. To further help understanding of how the shape of the distribution relates to  $F$  please consult section 3 of the supplementary material and [9].

Finally, the normalizing function  $a(F)$  in equation (1) can be defined as in equation (3). This function, as well as its gradients, can be computed by doing an integral over Bessel functions[9]. For implementation details regarding this see supplementary material section 5.

$$a(F) = \int_{R \in SO(3)} \exp(\text{tr}(F^T R)) dR \quad (3)$$

z




Figure 2: **Visualization of the matrix Fisher distribution for simple  $F$  matrices.** (a) For a spherical  $F$  the mode of the distribution is the identity. The distribution for each axis is circular and identical. (b) Here the axis distributions are more peaked than in (a) as the singular values are larger. (c) The distribution for the  $y$ - and  $z$ -axes are more elongated than for the  $x$ -axis as the first singular value dominates. (d)  $A$  is the rotation matrix obtained by rotating around the  $z$ -axis by  $-\pi/6$  degrees and thus the mode rotation is  $A$  shown by the red axes. The shape of the axes distributions though remains as in (c).

### 3.2 A negative log-likelihood loss function

Assume we have a labelled training example  $(x, R_x)$  where  $x$  is the input and  $R_x \in SO(3)$  its ground truth 3D rotation matrix. To train a neural network that estimates  $F_x$  for input  $x$ , it is necessary to define a loss function measuring the compatibility between  $F_x$  and  $R_x$ . As the pdf in equation (1) has support in all of  $SO(3)$ , we use the negative log-likelihood of  $R_x$  given  $F_x$  as the loss:

$$\mathcal{L}(F_x, R_x) = -\log(p(R_x | F_x)) = \log(a(F_x)) - \text{tr}(F_x^T R_x) \quad (4)$$

This loss has several interesting properties such as it is Lipschitz continuous, convex and has Lipschitz continuous gradients which makes it suitable for optimization. See supplementary material section 4 for proofs.

In practice the loss in equation 4 has an equilibrium far from the origin, in some experiments we believe this led to instability. To alleviate this problem we used a regularizing term which was 2.5% larger than what is analytically correct to move the equilibrium closer to the origin.

## 4 Experimental Details

We test our proposed approach on three separate datasets Pascal3D+, ModelNet10- $SO(3)$  and UPNA head pose. We briefly describe these datasets, the pre-processing we applied to the images from each dataset before training and then the evaluation metrics.

### 4.1 Datasets & Pre-processing

**Pascal3D+** [27] has 12 rigid object classes and contains images from Pascal VOC and ImageNet of these classes. Each image is annotated with the object’s class, bounding box and 3D pose. The latter is found by having an annotator align a 3D CAD model to the object in the image.

We pre-process each image by applying a homography so that the transformed image appears to come from a camera with known intrinsics and a principal axis pointing towards the object. This approach is similar to [13] and [29]. More details are given in section 6 of the supplementary material. We perform data augmentation similar to the data augmentations introduced in [14], but adapted for our preprocessing. At test time we apply the same type of homography transformation as applied during training, but no data augmentation.

**ModelNet10- $SO(3)$** [12] is a synthetic dataset. It is created by rendering rotated 3D models from ModelNet10 [26] with uniformly sampled rotations. The task is to estimate the applied rotation matrix.

We do not use any preprocessing for these images since the object is already centered and of a reasonable size. We do not use any data augmentation as the original paper did not use any and we want a fair comparison between the losses.

**UPNA head pose** [1] consists of videos with synchronized annotations of keypoints for the face in the image as well as its 3D rotation and position. The dataset has 10 people each with 11 recordings.

From the keypoint annotations we create a face bounding box for each image. After this we perform a small random perturbation of this bounding box to degrade the quality of the bounding box to be similar to what one would expect to get from a face detector. Using this artificial bounding box enables us to use the same data augmentation and preprocessing as we used for Pascal3D+.

There is no official training/test split for this dataset. We use a test split with the same people held out as in prior work [5]. We did not use a validation set since we did not do a new hyperparameter search for the dataset.

## 4.2 Details of network & training

We run experiments with ResNet-101 as our backbone network. The ResNet-101 parameters are initialized from pre-trained ImageNet weights. The object’s class is encoded by an embedding layer that produces a 32-dimensional vector and which is appended to the ResNet’s activations obtained from the final average pooling layer. We apply 3 fully connected layers to this vector with [512, 512, 9] nodes output at each layer. We use pytorch’s implementation of SVD for forward and backward propagation.

We fine-tune the embedding and fully connected layer weights for 2 epochs. We use SGD and start with a learning rate of 0.01. We use a batch size of 32 and train for 120 epochs. For Pascal3D+ we reduce this learning rate by a factor 10 at epochs 30, 60 and 90. For ModelNet10-SO(3) we train for 50 epochs and reduce the learning rate by a factor of 10 at epochs 30, 40 and 45. For UPNA head pose we use the same hyperparameters as for Pascal3D+, except we do not use a class embedding since there are only faces in this dataset.

## 4.3 Evaluation metrics

The evaluation metrics used are based on the geodesic distance:

$$d(R, \hat{R}) = \arccos\left(\frac{1}{2}\left(\text{tr}(R^T \hat{R}) - 1\right)\right) \quad (5)$$

where  $R$  and  $\hat{R}$  are the ground truth and estimated rotation respectively. This metric returns an angle error and we measure it in degrees. For a test set  $\mathcal{X}$ , containing tuples  $(x, R_x)$  of input  $x$  and its ground truth rotation  $R_x$ , we summarize performance on  $\mathcal{X}$  with the median angle error and  $\text{Acc}@Y$ :

$$\text{Acc}@Y = \frac{1}{|\mathcal{X}|} \sum_{(x, R_x) \in \mathcal{X}} \mathbb{1}(d(R_x, \hat{R}_x) < Y) \quad (6)$$

where  $\mathbb{1}(\cdot)$  is the indicator function,  $\hat{R}_x$  is the estimated rotation for input  $x$  and  $|\cdot|$  is the cardinality of a set. To compute the overall performance on a dataset the median angle error and  $\text{Acc}@Y$  are first computed per class and then averaged across all classes. For the UPNA dataset we use the mean geodesic error angle instead of the median to allow more direct comparison with the results in [5].

We have done all development and hyper-parameter optimization where the full training set was partitioned into a training and validation set. After hyper-parameter optimization, we have used the full training set for training and evaluated on the test set to get the numbers presented in the tables. For the Pascal3D+ dataset, we use the ImageNet validation split for the test set. Some samples of Pascal3D+ are labeled as “truncated”, “difficult” or “occluded”. We exclude these samples from our evaluations similar to other reported results [12]. This implementation detail had only a very slight effect on performance.

## 5 Results

**Quantitative results** Table 1 compares the performance of our method and recent high performing approaches. Table 2 compares per class performance for some classes on Pascal3D+ with the previous state-of-the-art method [15]. Our method significantly outperforms all the prior approaches. When

Table 1: **Performance on Pascal3D+**. Results are reported for the median angle error,  $\text{Acc}@ \pi/6$  and  $\text{Acc}@ \pi/12$ . The last column indicates if the training set was augmented with the synthetic dataset from [22].

| Method                  | MedErr      | Acc@ $\pi/6$ (%) | Acc@ $\pi/12$ (%) | Use synth. |
|-------------------------|-------------|------------------|-------------------|------------|
| Mahendran et al. [14]   | 15.38       | –                | –                 | ×          |
| Pavlakos et al. [19]    | 14.16       | –                | –                 | ×          |
| Tulsiani and Malik [24] | 13.60       | 81.0             | –                 | ×          |
| Su et al. [22]          | 11.70       | 82.0             | –                 | ✓          |
| Grabner et al. [6]      | 10.90       | 83.9             | –                 | ×          |
| Prokudin et al. [20]    | 10.40       | 83.9             | –                 | ×          |
| Mahendran et al. [15]   | 10.10       | 85.9             | –                 | ✓          |
| Liao et al. [12]        | 9.20        | 88.7             | –                 | ✓          |
| Ours                    | 9.11        | 90.9             | 73.4              | ×          |
| Ours                    | <b>8.17</b> | <b>92.8</b>      | <b>77.8</b>       | ✓          |

the training set is augmented with the synthetic dataset from [22], we further reduce the mean over medians angle error by approximately 1 degree.

The results reported in table 3 show that our method also achieves state-of-the-art performance on ModelNet10-SO(3).

On the UPNA head pose dataset our algorithm gives a mean angle error of 6.5 degrees. This is on par with current state of the art who quote a performance of 6.3 degrees [5]. We do not think the differences between these two methods are significant since there are only 4 persons in the test set.

**Qualitative results** Figure 3 displays and discusses interesting qualitative results on Pascal3D+ which highlight the probabilistic performance of our method.

**Behaviour for classes with rotational symmetries** Several classes in the datasets used have rotational symmetries or effectively have rotational symmetries due to very similar appearance at several distinct viewpoints. Some examples of these classes are canoes, bathtubs, tables, desks, and bottles. Our modelling though is based on a unimodal distribution and here we describe how the model copes with the inherent ambiguity of rotational symmetric objects.

For Pascal3D+ our method performs well, somewhat surprisingly, for the *dining table* class, see table 2. However, when the synthetic data is used to augment the training data the performance on this class drops. We suspect the manual labeling process introduces biases for this class with one of the ambiguous poses being labelled much more frequently. But the synthetic data added does not have these biases. This discrepancy between the distribution of training and test set label results in the drop in performance. For ModelNet10-SO(3) the *table* and *bathtub* classes have rotational symmetries and thus these two classes have much higher median errors than the other classes, see table 4. In figure

Table 2: **Pascal3D+ per-class performance** of our method, with or without using extra synthetic training data, compared to the competitive method Mahendran et al. [15]. The top three rows report the median angle error per class measured in degrees. The bottom three rows report  $\text{Acc}@ \pi/6$  measured as a percentage.

| Method    | aero        | bike        | boat        | bottle      | bus         | chair       | dtable      | sofa        | train       | mean        |
|-----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| [15]      | 8.5         | 14.8        | 20.5        | <b>7.0</b>  | <b>3.1</b>  | 9.3         | 11.3        | 10.2        | 5.6         | 10.1        |
| Ours w/o  | 10.1        | 14.6        | 13.2        | 8.0         | 3.3         | 7.4         | <b>8.2</b>  | 8.2         | 5.8         | 9.1         |
| Ours with | <b>6.6</b>  | <b>12.5</b> | <b>11.6</b> | 7.7         | 3.5         | <b>6.6</b>  | 11.2        | <b>7.4</b>  | <b>5.3</b>  | <b>8.2</b>  |
| [15]      | 87.0        | 81.0        | 64.0        | <b>96.0</b> | 97.0        | 92.0        | 67.0        | 97.0        | 82.0        | 85.9        |
| Ours w/o  | 87.7        | 83.2        | 75.6        | 94.9        | 98.6        | 93.9        | <b>82.3</b> | 97.4        | 97.9        | 87.7        |
| Ours with | <b>92.9</b> | <b>88.5</b> | <b>80.7</b> | 95.1        | <b>99.0</b> | <b>98.7</b> | 76.5        | <b>99.0</b> | <b>98.0</b> | <b>92.8</b> |

Table 3: **Performance on ModelNet10-SO(3)**. \* indicates the numbers reported in the original paper, but † denotes the revised numbers [11] where the evaluation metric uses the distance defined in equation (5). Thus we compare the performance of our method to the latter numbers.

| Method            | MedErr (deg) | Acc@ $\pi/6$ (%) | Acc@ $\pi/12$ (%) | Acc@ $\pi/24$ (%) |
|-------------------|--------------|------------------|-------------------|-------------------|
| Liao et al. [12]* | 20.3         | 70.9             | 58.9              | 38.4              |
| Liao [11]†        | 28.7         | 65.8             | 49.6              | 35.2              |
| Ours              | <b>18.0</b>  | <b>75.2</b>      | <b>68.5</b>       | <b>53.9</b>       |

Table 4: **Per class performance on ModelNet10-SO(3)** of our method.

| Metric        | bathtub | bed  | chair | desk | dress. | t.v. | n. stand | sofa | table | toilet |
|---------------|---------|------|-------|------|--------|------|----------|------|-------|--------|
| MedErr        | 86.5    | 4.4  | 5.2   | 13.7 | 6.9    | 6.1  | 15.4     | 4.1  | 34.3  | 3.9    |
| Acc@ $\pi/6$  | 41.1    | 90.0 | 93.7  | 67.7 | 72.8   | 85.8 | 59.1     | 94.8 | 49.2  | 98.2   |
| Acc@ $\pi/12$ | 32.2    | 87.0 | 88.7  | 52.9 | 65.9   | 78.4 | 49.5     | 91.7 | 43.4  | 95.8   |

4(f) the histogram of angle errors for the *table* class has a “U” shape and the median is in the middle of this “U”. This histogram indicates that at test time the network predicts one of the relevant poses.

To further illuminate this point, plots (b)-(g) of figure 4 show the evolution of the distribution predicted during training for one specific *table* class test image. The axis which has no associated ambiguity is identified correctly and confidently early on in training. The other two directions are predicted to have an almost uniform distribution on the plane spanned by the ambiguous axes. This is arguably the best way for the unimodal distribution to describe the situation. In the latter stages of training the network correctly identifies the object’s full pose on the training set and uncertainty becomes small. Such behaviour should be considered as a deterioration of the network’s probabilistic modelling as it effectively randomly chose one pose from the set of plausible poses and report it is very confident about this decision. Continuing to improve the accuracy on the test set while overfitting the loss often occurs with cross-entropy training of classification networks as well [7]. The dataset’s accuracy and loss plots, in the supplementary material section 2, show our loss is susceptible to this trend too.

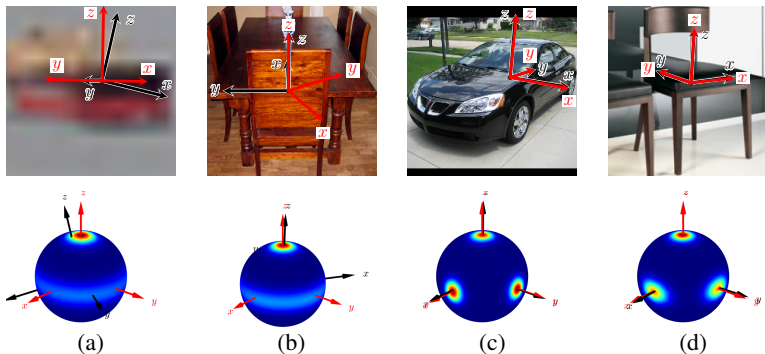


Figure 3: **Interesting qualitative results on Pascal3D+**. The top row displays example input images with the projected axes displaying the predicted pose (red) and labelled pose (black) of the object. The bottom row shows a visualization of the pdf estimated by the network. The red axis show the maximum likelihood estimate of the rotation matrix estimated from the predicted  $F$  matrix, while the axis in black corresponds to the ground truth rotation/pose. For clarity we have aligned the predicted pose with the standard axis. Each probability plot has been scaled independently. The examples shown have been specifically chosen to highlight our algorithm’s performance for certain cases: (a)-(b) Examples where model has high uncertainty for the azimuth either due to low resolution or rotation symmetry (c)-(d) Examples where model predicts rotations with high certainty and reasonably low errors.




Figure 4: **Evolution of the estimated pdf during training for a rotational symmetric object.** The leftmost figure is a test image from the table class. (b)-(e) Each plot displays the predicted distributions for the object’s pose after  $e$  epochs of training. The mode of the distribution is shown in red and the ground truth rotation in black. (f) Histogram of the test error angle for the table class after 50 epochs. See the main text for comments.

**Ablation experiments** We run ablation experiments on Pascal3D+ to identify the importance of the individual components of our approach. The factors considered are data-augmentation, the class embedding and pre-processing the image via a homography. The results are shown in table 5. Row one in table 5 shows the performance for a very simple method which uses a standard network architecture and plain cropping without data augmentation as preprocessing. This method gets higher performance than several more complicated methods, See table 1. This indicates that the loss we have introduced is a significant improvement by itself.

By comparing Row one and two in table 5 we see that the class embedding does not seem to give any improvement for this dataset. This can also be seen by comparing row four and five. Rows one and three in table 5 show that our warping does not provide a significant improvement by itself on Pascal3D+ compared to cropping. We believe though this warping could be advantageous in many situations and therefore should be used irrespective of these results. In theory this pre-processing should allow our method to generalize across all pinhole cameras with known intrinsic parameters and negligible radial distortion rather than for cameras with the same intrinsics as Pascal3D+.

Comparing row three with row four or comparing row one with row five show that data augmentation gives a significant improvement. This is consistent with prior work.

Table 5: **Results of ablation experiments on Pascal3D+ for our method.**

| Data aug. | Class embed | Crop | Warp | MedErr (deg) | Acc@ $\pi/6$ (%) | Acc@ $\pi/12$ (%) |
|-----------|-------------|------|------|--------------|------------------|-------------------|
| ×         | ×           | ✓    | ×    | 10.5         | 87.7             | 68.6              |
| ×         | ✓           | ✓    | ×    | 10.5         | 87.1             | 67.9              |
| ×         | ✓           | ×    | ✓    | 10.5         | 87.0             | 68.8              |
| ✓         | ✓           | ×    | ✓    | 9.2          | <b>90.9</b>      | 73.4              |
| ✓         | ×           | ×    | ✓    | <b>9.0</b>   | 90.5             | <b>74.1</b>       |

## 6 Conclusion & Future work

In this paper we have introduced a way to use neural networks to output probability distributions over the set of rotations with the matrix Fisher distribution. We show when optimizing the negative log likelihood of this distribution we end up with a convex loss. When applying this method on several datasets we get state-of-the-art performance. Our ablation studies show the relative robustness of the approach.

Since the matrix Fisher distribution is unimodal it poorly models classes which have rotational symmetries. It could be interesting to try to create a loss supporting multimodal distributions while keeping the desirable optimization properties of our loss. It could be possible to use these estimated probabilities for time tracking filters such as the one described in [9].

We have not done any quantification of how well the estimated variances correspond to the actual errors. Doing this as well as calibrating the uncertainties similar to [7] is potential future work.



## Broader Impact

The methods described in this paper has obvious applications in fields which some consider ethically questionable such as for surveillance and military systems. One example could be determining heading for ships or airplanes for tactical planning. That being said, the orientation of objects is a fundamental property of objects in the real world and being able to accurately estimate this property should be helpful for many applications of either an ethically desirable or undesirable nature. In our opinion improving the techniques used for orientation estimation has a similar societal impact as improving the techniques used for classification or object detection.

The persons in the UPNA dataset are unlikely to be sampled from a uniform distribution of people across the world, for this reason one can not expect the reported performance to be accurate for the world population in general, that being said due to the small test size this reported performance might not reflect the average performance for any population. We do not believe this is an issue since models for predicting head pose which are deployed on a wider scale are very unlikely to use this dataset due to its small size and non-commercial licence. The method itself is not reliant on any population specific feature.

## Acknowledgments and Disclosure of Funding

DM is an industrial PhD student at Tobii AB and GB is a senior algorithm engineer at Tobii AB. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. Computational resources was provided by KTH.

## References

- [1] Mikel Ariz, José J. Bengoechea, Arantxa Villanueva, and Rafael Cabeza. A novel 2d/3d database with automatic face annotation for head tracking and pose estimation. *Computer Vision and Image Understanding*, 148:201–210, 2016.
- [2] Andreas Doumanoglou, Vassileios Balntas, Rigas Kouskouridas, and Tae-Kyun Kim. Siamese regression networks with efficient mid-level feature extraction for 3d object pose estimation. *arXiv preprint arXiv:1607.02257*, 2016.
- [3] Thomas D Downs. Orientation statistics. *Biometrika*, 59(3):665–676, 1972.
- [4] David W Eggert, Adele Lorusso, and Robert B Fisher. Estimating 3-d rigid body transformations: a comparison of four major algorithms. *Machine vision and applications*, 9(5-6):272–290, 1997.
- [5] Igor Gilitschenski, Roshni Sahoo, Wilko Schwarting, Alexander Amini, Sertac Karaman, and Daniela Rus. Deep orientation uncertainty learning based on a bingham loss. In *International Conference on Learning Representations*, 2019.
- [6] Alexander Grabner, Peter M Roth, and Vincent Lepetit. 3d pose estimation and 3d model retrieval for objects in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3022–3031, 2018.
- [7] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1321–1330. JMLR. org, 2017.
- [8] C. Khatri and K. Mardia. The von mises-fisher matrix distribution in orientation statistics. *Journal of the Royal Statistical Society: Series B*, 39(1):95–106, 1977.
- [9] Taeyoung Lee. Bayesian attitude estimation with the matrix fisher distribution on so (3). *IEEE Transactions on Automatic Control*, 63(10):3377–3392, 2018.
- [10] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. EPnP: An Accurate O(n) Solution to the PnP Problem. *International Journal of Computer Vision*, 2009.
- [11] Shuai Liao. spherical regression, 2020. URL [https://github.com/leoshine/Spherical\\_Regression](https://github.com/leoshine/Spherical_Regression).

- [12] Shuai Liao, Efstratios Gavves, and Cees GM Snoek. Spherical regression: Learning viewpoints, surface normals and 3d rotations on n-spheres. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9759–9767, 2019.
- [13] Erik Lindén, Jonas Sjostrand, and Alexandre Proutiere. Learning to personalize in appearance-based gaze tracking. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019.
- [14] Siddharth Mahendran, Haider Ali, and René Vidal. 3d pose regression using convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2174–2182, 2017.
- [15] Siddharth Mahendran, Haider Ali, and Rene Vidal. A mixed classification-regression framework for 3d pose estimation from 2d images. *arXiv preprint arXiv:1805.03225*, 2018.
- [16] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose Estimation for Augmented Reality: A Hands-On Survey. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2633 – 2651, December 2016.
- [17] Kanti V Mardia and Peter E Jupp. *Directional Statistics*. John Wiley & Sons, 2009.
- [18] Xiaoli Meng, Heng Wang, and Bingbing Liu. A robust vehicle localization approach based on gnss/imu/dmi/lidar sensor fusion for autonomous vehicles. *Sensors*, 17:2140, Sep 2017.
- [19] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-dof object pose from semantic keypoints. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2011–2018. IEEE, 2017.
- [20] Sergey Prokudin, Peter Gehler, and Sebastian Nowozin. Deep directional statistics: Pose estimation with uncertainty quantification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 534–551, 2018.
- [21] Stephen Se, David Lowe, and Jim Little. Vision-based mobile robot localization and mapping using scale-invariant features. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164)*, volume 2, pages 2051–2058. IEEE, 2001.
- [22] Hao Su, Charles R Qi, Yangyan Li, and Leonidas J Guibas. 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*, pages 2686–2694, 2015.
- [23] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. *arXiv preprint arXiv:1809.10790*, 2018.
- [24] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1510–1519, 2015.
- [25] Ulrich Weidenbacher, Georg Layher, Pierre Bayerl, and Heiko Neumann. Detection of head pose and gaze direction for human-computer interaction. In *Perception and Interactive Technologies*, pages 9–19. Springer Berlin Heidelberg, 2006.
- [26] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.
- [27] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2014.
- [28] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *Robotics: Science and Systems (RSS)*, 2018.
- [29] Xucong Zhang, Yusuke Sugano, and Andreas Bulling. Revisiting data normalization for appearance-based gaze estimation. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, pages 1–9, 2018.
- [30] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019.