

FisherMatch: Semi-Supervised Rotation Regression via Entropy-based Filtering

Yingda Yin Yingcheng Cai He Wang[†] Baoquan Chen[†]
 Peking University

Abstract

Estimating the 3DoF rotation from a single RGB image is an important yet challenging problem. Recent works achieve good performance relying on a large amount of expensive-to-obtain labeled data. To reduce the amount of supervision, we for the first time propose a general framework, FisherMatch, for semi-supervised rotation regression, without assuming any domain-specific knowledge or paired data. Inspired by the popular semi-supervised approach, FixMatch, we propose to leverage pseudo label filtering to facilitate the information flow from labeled data to unlabeled data in a teacher-student mutual learning framework. However, incorporating the pseudo label filtering mechanism into semi-supervised rotation regression is highly non-trivial, mainly due to the lack of a reliable confidence measure for rotation prediction. In this work, we propose to leverage matrix Fisher distribution to build a probabilistic model of rotation and devise a matrix Fisher-based regressor for jointly predicting rotation along with its prediction uncertainty. We then propose to use the entropy of the predicted distribution as a confidence measure, which enables us to perform pseudo label filtering for rotation regression. For supervising such distribution-like pseudo labels, we further investigate the problem of how to enforce loss between two matrix Fisher distributions. Our extensive experiments show that our method can work well even under very low labeled data ratios on different benchmarks, achieving significant and consistent performance improvement over supervised learning and other semi-supervised learning baselines. Our project page is at <https://yd-yin.github.io/FisherMatch>.

1. Introduction

Incorporating deep neural networks to perform rotation regression is exerting an ever-important influence in computer vision, graphics and robotics. This is now one of the key technology in enabling a multitude of applications such as camera relocalization and visual odometry [6, 11], ob-

ject pose estimation and tracking [41, 44], and 6DoF robot grasping [5, 15]. One of the major obstacles to improving rotation regression is expensive rotation annotations. Though many large-scale image datasets have been curated with sufficient semantic annotations, obtaining a large-scale real dataset with rotation annotations can be extremely laborious, expensive and error-prone [43]. With the amount of labeled data being the bottleneck, there is a demand for methods that can leverage unlabeled data.

Regarding training models with fewer labels, semi-supervised learning (SSL) has been a powerful approach, mitigating the requirement for labeled data by providing a means of leveraging unlabeled data and thus attracting more and more attention.

Recent years have witnessed many processes in semi-supervised classification [4, 12, 20, 36, 37], semi-supervised object detection [25, 40], and semi-supervised human and hand pose estimation [17, 32]. However, only few works address semi-supervised rotation regression, with most of them leveraging domain-specific knowledge, *e.g.*, temporal smoothness of object pose [24] and strong assumptions, *e.g.*, paired images from different viewpoints [27].

The underlying reason for little work in this field is that rotation regression is very unique and challenging. First, it is undesirable to turn the rotation regression into a classification problem. Given that 3D rotation space is continuous, discretizing the space into a small number of bins will lead to limited accuracy, which is intolerable for many applications involving rotation estimation. Also, rotation regression is even not a standard regression problem. Given that rotation space $SO(3)$ is a non-Euclidean manifold [51], a general regression algorithm needs to be tailored, taking the nonlinear structure of the rotation space into account. This further makes semi-supervised rotation regression a more challenging and less studied topic.

In this work, for the first time, we propose a general framework, namely *FisherMatch*, for semi-supervised rotation regression. The problem we tackle is very general: *using a neural network to regress rotation from a single RGB image*. Inspired by a popular semi-supervised learning approach, FixMatch [36], initially developed for classification tasks, we attempt to process rotation regression problems in

[†]He Wang and Baoquan Chen are the corresponding authors ({hewang, baoquan}@pku.edu.cn).

a similar flavor.

The key idea to the success of FixMatch is *to filter out the pseudo labels with low classification confidence and only supervise the model outputs with highly confident labels*. This mechanism ensures the quality of pseudo labels and thus significantly improves the performance of semi-supervised learning. The underlying assumption is that the more confident a pseudo label is, the more closed this label is to the ground truth. Or, in other words, this system needs to predict a confidence that can well indicate the correctness of its prediction. Fortunately, a classification output naturally carries the information: the probability of its prediction can be used as its prediction confidence. We argue that the availability of such a reliable confidence measure is crucial to the success of FixMatch on semi-supervised classification tasks. Similarly, when adopting FixMatch to 3D object detection, 3DIoUMatch [40] constructs a separate branch to predict the 3D IoU between the predicted bounding box and the ground truth bounding box as a localization confidence to filter out poor predictions. Although 3D IoU estimation is a regression task, 3DIoUMatch can move around the predicted bounding boxes as an augmentation trick, thus creating an infinite amount of training data for this confidence estimation module. This augmentation is crucial for such a confidence estimation module since the confidence estimation modules can only be trained using labeled data and must work on unlabeled data.

However, we argue that adopting FixMatch for rotation estimation is highly non-trivial. The biggest obstacle is how to estimate the prediction confidence for rotation regression. For rotation regression, we don't have the largest probability from the bins as our confidence; also, for rotation estimation from a single RGB image, we can't perform such augmentation to change our rotation prediction; yet, we still need this uncertainty estimation module to work on unlabeled data with only training on a small set of labeled data.

As pointed out by [35], probabilistic modeling of rotation is the correct way to model the uncertainty of rotation regression. Parametric statistical methods for orientation statistics have long been established [9, 16, 19, 33]. In order to better resort to $SO(3)$ manifold which has a different topology than unconstrained values in \mathbb{R}^N , Deng *et al.* [8] and Mohlin *et al.* [28] incorporate Bingham distribution and matrix Fisher distribution respectively to automatically learn uncertainties along with predictions, without further supervision. Thus, such networks can provide valuable information about the quality of the prediction. We prefer matrix Fisher distribution to Bingham distribution, since its rotation representation is continuous and its loss is convex with bounded gradient magnitudes, resulting in a stable training for neural networks [22, 28].

We thus devise a matrix Fisher-based rotation regressor that takes input a single RGB image and outputs the param-

eter of a matrix-Fisher distribution. Given the predicted distribution, we propose to use the entropy of this distribution as a confidence measure for pseudo label filtering. Basically, only pseudo labels with high confidence, *i.e.* lower entropy than a threshold τ_{entropy} , will pass the filtering and be used for supervising the model under training. Our experiment consistently proves that entropy is an efficient indicator of the prediction performance, not only in the case of 100 percent labeled data, but also in low data ratio cases down to 5 percent. Since FisherMatch outputs a distribution rather than a single rotation, our pseudo labels become a distribution, which requires research into the unsupervised loss enforced between two distributions. In this work, we investigate cross entropy loss and negative log likelihood loss, draw a connection between them, and find their proper usage in our experiments.

On common benchmark datasets of object rotation estimation from RGB images (ModelNet10-SO(3) and Pascal3D+) under various labeled data ratios, our experiment demonstrates a significant and consistent performance improvement over supervised learning and other semi-supervised learning baselines.

2. Related Work

Rotation regression The choice of rotation representation is one of the core issues concerning rotation regression. The commonly used representations include Euler angles, axis-angles, unit quaternions, *etc.* However, Euler angles suffer from gimbal lock, and quaternions have a double embedding giving rise to the existence of two disconnected local minima. Moreover, [51] argues that representations less than 4 dimensions are bound to have discontinuities and are difficult for neural networks to learn. To this end, the continuous 6D representation with Gram-Schmidt orthogonalization [51] and 9D representation with SVD orthogonalization [22] have been proposed respectively, leading to superior performance in rotation regression.

Several works propose to use probability distributions over rotations to further model prediction uncertainties along with rotation regression. In Prokudin *et al.* [35], parameters of a mixture of Von Mises distribution using a biternion network are estimated. Deng *et al.* [8] uses Bingham distribution over unit quaternions to jointly predict the rotation as well as the uncertainty. Estimation with matrix Fisher distribution [28] learns to build the probability distribution over rotation matrices with unconstrained parameters. To further express arbitrary rotation distributions and better tackle rotation regression for symmetry objects, Implicit-PDF [29] chooses to represent the distributions implicitly by neural networks, instead of distribution parameters, where the $SO(3)$ space is uniformly discretized with the help of Hopf fibration [48].

Semi-supervised classification Semi-supervised learning is a long-studied field with a diversity of approaches, many in the field of classification. Consistency regularization and pseudo labeling are two measures with in-depth exploration. Consistency regularization was first proposed in [2] which enforces the model to predict consistently across multiple perturbations [18,20,37,45,46]. Pseudo labels [21] are artificial labels generated by the model itself and are used to further train the model, often applied along with a confidence-based thresholding to ensure the pseudo label quality. Mixmatch [4], ReMixmatch [3] and FixMatch [36] are holistic methods utilizing various augmentation and label sharpening strategies.

More recently, SIMPLE [13] proposes the paired loss minimizing the statistical distance between confident and similar pseudo labels. SemCo [31] considers label semantics to prevent the degradation of pseudo label quality for visually similar classes in a co-training manner. Dash [47] and FlexMatch [49] propose dynamic and adaptive pseudo label filtering, better suited for the training process.

Semi-supervised regression Semi-supervised regression is a less-touched field compared with classification, where most of the works deal with regressing Euclidean variables, e.g., Parkinson’s disease rating scales from multiple telemonitoring data in UCI repository [1]. Early work of CoReg [52] utilizes multiple k-nearest neighbor regressors with different distance metrics and leverages the predictions of one regressor to label the other regressors in a co-training manner. SSDKL [14] leverages the unlabeled data by minimizing predictive variance in the posterior regularization framework through the composition of neural networks and the probabilistic modeling of Gaussian processes.

Self-/semi-supervised rotation estimation Several works tackle rotation estimation in a self-supervised manner. Mustikovela *et al.* [30] leverages the analysis-by-synthesis technique that requires a lot of extra images for training a generative model. ViewNet [27] assumes the availability of paired data (same object, different poses). The most relevant semi-supervised learning work is NVSM [39], which shares the same assumptions on data and labels with us. In contrary to *regression*, NVSM builds a category-level 3D cuboid mesh with feature vectors and estimates the object rotation in a render-and-compare technique through the distance-based rotation retrieval. Less literature has been seen in the field of semi-supervised rotation regression. Mariotti *et al.* [26] requires paired images of an object and enforces cross-reconstruction in an analysis-and-synthesis manner via rotating the encoded neural latent variables.

Our work draws insight from both the orientation statistics and the semi-supervised learning techniques introduced above, dedicated to correlating the techniques in two well-

explored fields to tackle the problem in the general setting of semi-supervised rotation regression.

3. Method

In this work, we tackle the problem of learning to predict 3D object rotation from single RGB images under a semi-supervised setting, where we have only a (small) set of labeled data $\{\mathbf{x}_i^l, \mathbf{y}_i^l\}_{i=1}^{N_l}$ and a larger set of unlabeled data $\{\mathbf{x}_i^u\}_{i=1}^{N_u}$. Here, \mathbf{x}^l and \mathbf{x}^u represent the labeled and unlabeled RGB image respectively, and \mathbf{y}^l represents the ground-truth rotation in $SO(3)$ for a labeled data; N_l and N_u are the number of labeled and unlabeled images, respectively.

Following a popular semi-supervised learning approach, FixMatch [36], we adopt the teacher-student mutual learning framework, which we summarize in Section 3.1. In Section 3.2, we make use of two probabilistic models of rotation for depicting the uncertainty in rotation prediction, namely Bingham distribution and matrix Fisher distribution [8,28], and propose to use the entropy of the predicted matrix Fisher distribution as the prediction confidence for pseudo label filtering; In Section 3.4, for the purpose of enforcing loss between the teacher and the student, we construct two loss functions between pseudo labels and predicted distributions; Finally, in Section 3.5, we introduce our training protocol in detail.

3.1. Revisit FixMatch

The teacher-student mutual learning framework is a popular approach for semi-supervised learning. Mean Teacher [37] proposes the first version, containing two jointly learned models - a *teacher* and a *student*. The parameters of the teacher model are the exponential moving average (EMA) of the student model parameters that are updated by the stochastic gradient descent. The student model is trained by the ground-truth labels for the labeled data, and for the unlabeled data, the predictions of the teacher model serve as the *pseudo labels* and are used to supervise the student network, through which, a history consistency is enforced between the two models.

FixMatch [36] further develops this approach by proposing two strategies: asymmetric data augmentation and confidence-based pseudo label filtering. Asymmetric data augmentation means that the teacher model is fed by weakly augmented unlabeled samples while the student model takes strongly augmented unlabeled samples which contributes to the performance gap between the teacher and the student, facilitating correct information flow to the student.

Arguably, the most important contribution of FixMatch is to demonstrate the effectiveness of confidence-based pseudo label filtering. For a non-trivial semi-supervised learning task, previous works recognize that the pseudo la-

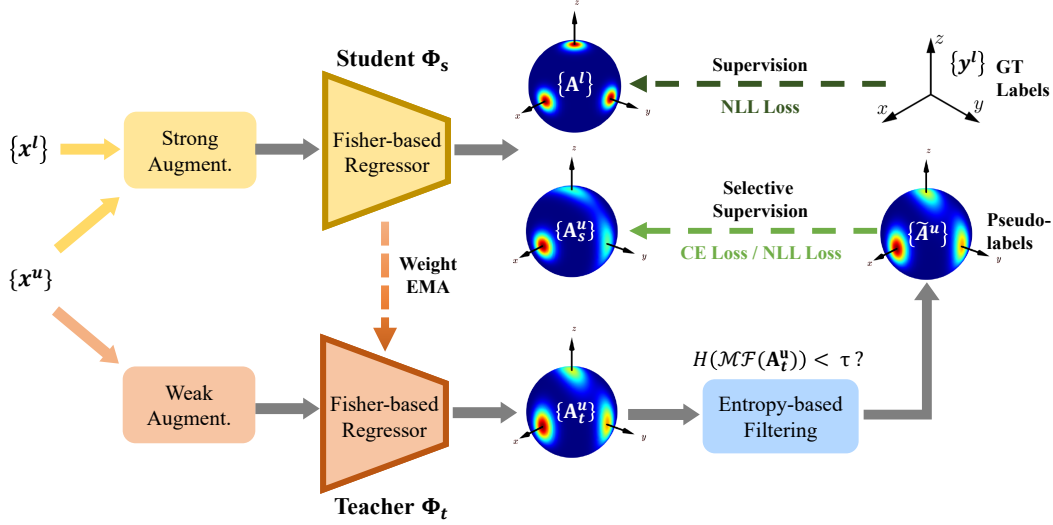


Figure 1. **Pipeline overview.** Our matrix Fisher based-rotation regressor Φ takes an RGB image \mathbf{x} as input and outputs the parameter \mathbf{A} of the predicted matrix Fisher distribution. We leverage a teacher-student mutual learning framework composed of a learnable student model and an exponential-moving-average (EMA) teacher model. On labeled data, the student network is trained by the ground-truth labels with the supervised loss; while on unlabeled data, the student model takes the pseudo labels from the EMA teacher. We leverage an entropy-based filtering technique to filter out noisy teacher predictions. The distribution visualization is borrowed from [28] where x , y and z shown in black axes correspond to the standard basis of \mathbb{R}^3 , and the pdf is shown on the sphere with a *jet* color coding. See supplementary material Section D for details of the visualization method.

bels generated by the teacher output suffer from significant noises [36, 40]. To this end, FixMatch proposes to filter out low-quality predictions and only supervise the student model with predictions with high confidence. This strategy avoids wrong supervision to the student model and has been proved to be very effective for challenging tasks, *e.g.*, object detection [25, 40]. Given the difficulty of rotation regression, we further propose to leverage FixMatch as the basis of our framework for the rotation regression task.

3.2. Probabilistic Modeling of Rotation

To model the uncertainty of rotation estimation, we leverage *matrix Fisher* distribution to build a probabilistic model of rotation prediction, following Mohlin *et al.* [28].

Matrix Fisher distribution [19, 34] $\mathcal{MF}(\mathbf{R}; \mathbf{A})$ is a probability distribution over $\text{SO}(3)$ for rotation matrices, whose probability density function is in the form of

$$p(\mathbf{R}) = \mathcal{MF}(\mathbf{R}; \mathbf{A}) = \frac{1}{F(\mathbf{A})} \exp(\text{tr}(\mathbf{A}^T \mathbf{R})) \quad (1)$$

where parameter $\mathbf{A} \in \mathbb{R}^{3 \times 3}$ is an arbitrary 3×3 matrix and $F(\mathbf{A})$ is the normalizing constant. The mode and dispersion of the distribution can be computed from computing singular value decomposition of the parameter \mathbf{A} . Assume $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ and the singular values are sorted in descending order, the mode of the distribution is computed as

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

and the singular values $\mathbf{S} = \text{diag}(s_1, s_2, s_3)$ indicates the strength of concentration. The larger a singular value s_i is, the more concentrated the distribution is along the corresponding axis (the i -th column of mode $\hat{\mathbf{R}}$).

Another important probabilistic model for rotation is *Bingham* distribution on \mathcal{S}^3 for unit quaternions. The probability density function is defined as

$$\mathcal{B}(\mathbf{q}; \mathbf{M}, \mathbf{Z}) = \frac{1}{F(\mathbf{Z})} \exp(\mathbf{q}^T \mathbf{M} \mathbf{Z} \mathbf{M}^T \mathbf{q}) \quad (3)$$

where $\mathbf{M} \in \text{O}(4)$ is a 4×4 orthogonal matrix and $\mathbf{Z} = \text{diag}(0, z_1, z_2, z_3)$ is a 4×4 diagonal matrix with $0 \geq z_1 \geq z_2 \geq z_3$. The first column of parameter \mathbf{M} indicates the mode and the remaining columns describe the orientation of dispersion while the corresponding z_i , ($i \in 1, 2, 3$) describe the strength of the dispersion. $F(\mathbf{Z})$ is the normalizing constant.

It is well recognized that rotation matrix \mathbf{R} and quaternion \mathbf{q} are two different representations of rotation. Similarly, as discussed in [34], matrix Fisher distribution and Bingham distribution are equivalent to each other differing only in parameterizations and rotation representations. However, given that quaternion is not a continuous representation of rotation [51], using matrix representation to learn a deep rotation estimation model has an intrinsic advantage and usually yields better performance. [28] further shows that matrix Fisher distribution has a bounded gradient, which is favored by deep neural networks. Therefore, 9D rotation matrix is chosen as our representation, and ma-

trix Fisher distribution is used for building our probabilistic rotation model.

3.3. Entropy-based Pseudo Label Filtering

Inspired by FixMatch, we only want the accurate predictions from the teacher model to “teach” the student model. Otherwise, noisy pseudo labels may slow down the training procedure, or even do harm to the whole process.

For depicting the confidence of a predicted distribution, we propose to use *entropy*, which is widely used in statistics acting as the degree of disorder or randomness in the system, as a measure of *uncertainty*. A lower entropy generally indicates a more *peaked* distribution which exhibits less uncertainty and higher confidence.

In this work, we propose an entropy-based filtering mechanism leveraging the probabilistic modeling of the rotation estimation over $SO(3)$. We devise a rotation regressor Φ that takes a single RGB image \mathbf{x} and outputs the parameter $\mathbf{A} \in \mathbb{R}^{3 \times 3}$ of a matrix Fisher distribution

$$\mathbf{A} = \Phi(\mathbf{x}), \quad (4)$$

which not only contains a predicted rotation as the *mode* of this distribution, but also encode the information of the distribution *concentration*. We then compute the entropy of this predicted distribution (see Equation 9).

For *pseudo label filtering*, we set a fixed entropy threshold τ , and only reserve the prediction as a pseudo label if its entropy is lower than the threshold. Specifically, for unlabeled data \mathbf{x}^u , assume $p_t = \mathcal{MF}(\mathbf{A}_t^u)$ is the teacher output with $\mathbf{A}_t^u = \Phi_t(\mathbf{x}^u)$ and $p_s = \mathcal{MF}(\mathbf{A}_s^u)$ is the student output with $\mathbf{A}_s^u = \Phi_s(\mathbf{x}^u)$, the loss on unlabeled data is therefore:

$$L_u(\mathbf{x}^u) = \mathbb{1}(H(p_t) \leq \tau) L(p_t, p_s) \quad (5)$$

We discuss the loss function enforced between two distribution $L(p_t, p_s)$ in Section 3.4.

3.4. Loss Function between Distributions

For the labeled set $\{\mathbf{x}_i^l, \mathbf{y}_i^l\}_{i=1}^{N_l}$, we adopt the most common loss function, negative log likelihood (NLL) loss, to learn the probabilistic model of rotation, as in [8, 28]. This loss minimizes the negative log likelihood of the ground-truth rotation in the predicted distributions, as shown below:

$$L_l(\mathbf{x}^l, \mathbf{y}^l) = -\log(\mathcal{MF}(\mathbf{y}^l; \mathbf{A}^l)) \quad (6)$$

where $\mathbf{A}(\mathbf{x}^l)$ denotes the network output fed with input \mathbf{x}^l .

For unlabeled data, both our network predictions and pseudo labels are distributions, and thus we need to enforce loss between two distributions, which is rarely the case for a regression problem. We investigate two types of losses, *i.e.*, negative log likelihood (NLL) loss and cross entropy (CE) loss.

Cross Entropy Loss L^{CE} In classification problems, a widely-used loss function between two discrete distributions is cross entropy loss L^{CE} , whose gradient is equivalent to the gradient of KL divergence between two distribution [10]. We thus extend cross entropy loss L^{CE} so as to enforce the consistency between pseudo labels and the student outputs:

$$L^{\text{CE}}(p_t, p_s) = H(p_t, p_s) \quad (7)$$

To compute L^{CE} between two continuous distribution on $SO(3)$, we derive the analytical formula for the cross entropy between two matrix Fisher distributions $f \sim \mathcal{MF}(\mathbf{A}_f)$ and $g \sim \mathcal{MF}(\mathbf{A}_g)$, as shown below:

Assume $\mathbf{A}_f = \mathbf{U}_f \mathbf{S}_f \mathbf{V}_f^T$, $\mathbf{A}_g = \mathbf{U}_g \mathbf{S}_g \mathbf{V}_g^T$, γ is the standard transform from unit quaternion to rotation matrix, \mathbf{e}_i is the i -th column of \mathbf{I}_4 , and $\mathbf{E}_i = \gamma(\mathbf{e}_i)$, then we can derive

$$H(f, g) = \log F_g - \sum_{i=1}^4 z_{gi} \left(b_i^2 + \sum_{j=1}^4 (a_{ij}^2 - b_i^2) \frac{1}{F_f} \frac{\partial F_f}{\partial z_{fj}} \right) \quad (8)$$

$$\begin{aligned} \text{where } z_{gi} &= \text{tr}(\mathbf{E}_i^T \mathbf{S}_g) & z_{fj} &= \text{tr}(\mathbf{E}_j^T \mathbf{S}_f) \\ a_{ij} &= \gamma^{-1}(\mathbf{U}_f \mathbf{E}_i \mathbf{V}_f^T) \cdot \gamma^{-1}(\mathbf{U}_g \mathbf{E}_j \mathbf{V}_g^T) \\ b_i &= \gamma^{-1}(\mathbf{U}_f \mathbf{E}_i \mathbf{V}_f^T) \cdot \gamma^{-1}(\mathbf{U}_g \mathbf{E}_i \mathbf{V}_g^T) \end{aligned}$$

and F_f and F_g are constant wrt. parameter \mathbf{Z} . See supplementary material Section B for the derivation. Note that when $f = g$, we can also get the entropy $H(f)$ for matrix Fisher distribution, as shown below:

$$H(f) = \log F_f - \sum_{i=1}^4 \left(z_{fi} \frac{1}{F_f} \frac{\partial F_f}{\partial z_{fi}} \right) \quad (9)$$

NLL Loss L^{NLL} Another option of the loss is to consider the negative log likelihood of the mode predicted by the teacher in the distribution predicted by the student, which is basically the NLL loss treating the teacher prediction as ground truth, as in the case of labeled data.

$$L^{\text{NLL}}(p_t, p_s) = -\log p_s(\mathbf{y}_t^u), \quad (10)$$

where \mathbf{y}_t^u is the mode predicted by the teacher and can be computed by SVD of \mathbf{A}_t^u (see Section 3.2).

Relationship between L^{NLL} and L^{CE} Here we intend to make connection between L^{NLL} and L^{CE} . We find that L^{CE} becomes L^{NLL} when we decreases the dispersion of the distribution p_t to a Dirac distribution $\delta(\mathbf{R}; \mathbf{y}_t^u)$ with its mode located at \mathbf{y}_t^u . We give a brief proof as below:

$$\begin{aligned} L^{\text{CE}}(\text{Dirac}(p_t), p_s) &= H(\delta(\mathbf{y}_t^u), p_s) \\ &= - \int_{SO(3)} \delta(\mathbf{y}_t^u) \log p_s d\mathbf{R} \\ &= -\log p_s(\mathbf{y}_t^u) = L^{\text{NLL}}(p_t, p_s). \end{aligned}$$

This exactly resembles the label sharpening technique used in semi-supervised classification [4,36], where the teacher’s output is either sharpened or turned into a hard label. To be specific, when we turn a predicted distribution $\mathcal{MF}(\mathbf{A}_t^u)$ into a hard label \mathbf{y}_t^u , L^{CE} becomes L^{NLL} . We use L^{CE} in the experiments and investigate the different behavior of these two losses in Section 4.4.

3.5. Training Protocol

Our training is composed of two stages: a pre-training stage, where we train our rotation regressor on the labeled data, followed by an SSL stage where both the labeled and the unlabeled data are utilized. Our matrix Fisher-based rotation regressor is fed with an RGB image \mathbf{x} and outputs a 3×3 matrix \mathbf{A} as the predicted parameter of the matrix Fisher distribution. We take the mode of the distribution as the predicted value.

Pre-training We start with a supervised training procedure on the labeled set with the supervised loss as Eq. 6. We clone the rotation regressor to obtain a pair of teacher and student networks with the same initialization, once converged.

Semi-supervised training In SSL stage, we utilize both the labeled data and the unlabeled data. A training batch contains a mixture of $\{\mathbf{x}_i^l\}_{i=1}^{B_l}$ labeled samples and $\{\mathbf{x}_i^u\}_{i=1}^{B_u}$ unlabeled samples. The loss function is composed of the supervised loss applied to the labeled samples and the unsupervised loss for the unlabeled samples

$$L = L_l(\mathbf{x}^l, \mathbf{y}^l) + \lambda_u L_u(\mathbf{x}^u) \quad (11)$$

where L_l is computed as Eq. 6, L_u is as Eq. 7, and λ_u is the unsupervised loss weight.

In this stage, We adopt asymmetric augmentation and an exponential-moving-average teacher as stated in Sec. 3.1.

4. Experiment

4.1. Datasets

ModelNet10-SO(3) [23] is created by rendering 3D models of ModelNet-10 [42] that are rotated by uniformly sampled random rotations in $\text{SO}(3)$. Following [7,28], we focus on the `chair` and `sofa` category which exhibit the least rotational symmetries in the dataset. In the experiments, we set the ratio of labeled data as 5% and 10% of the training set.

Pascal3D+ [43] contains real images from Pascal VOC and ImageNet of 12 rigid object classes. Following NVSM [39], we evaluate 6 vehicle categories (`aeroplane`, `bicycle`, `boat`, `bus`, `car`, `motorbike`) which have relatively evenly distributed poses in azimuth angles, and set the number of labeled images as 7, 20 and 50 for each category respectively. We share the same selected 7 images as NVSM such that they are spread around the pose space.

We follow the original train-test split and further divide the training split into the labeled set with ground truth and the unlabeled set without ground truth.

4.2. Evaluation setup

Baselines To the best of our knowledge, we are the first to tackle semi-supervised rotation regression in this setting, hence the comparisons are made with self-made baselines. **Supervised-L1** uses a normal regressor and only trains on the labeled set with L1 loss with the 9D-SVD [22] rotation representation, while **Supervised-Fisher** uses our matrix Fisher regressor and also only go through the pretraining stage. As an SSL baseline, **SSL-L1-Consistency** refers to adopting FixMatch into the task with the EMA teacher and asymmetric data augmentation preserved, but only applying L1 loss as the consistency supervision between the student and teacher predictions without filtering, due to lacking the confidence measure. Here, for non-Fisher regressors, we choose L1 instead of L2 loss, as [8] points out that L1 outperforms L2 for rotation regression.

We find the most relevant work to ours is NVSM [39], which, though not regression-based, tackles the same task as ours and leverages a render-and-compare scheme through distance-based rotation retrieval. We borrow NVSM and their developed baselines as our compared baselines, including two supervised rotation estimation works (**StarMap** [50] and **NeMo** [38]) and two standard classification networks (**Res50-Gene** and **Res50-Spec**), adapted into semi-supervised learning, respectively. Due to the unavailability of the training code, we exactly follow the experiment settings of NVSM and evaluate on Pascal3D+ dataset. See supplementary Section A for more details.

Evaluation metrics We evaluate the experiments by the mean error, the median error (in degrees) and the accuracy within 30° between the prediction and the ground truth.

4.3. Results

Result comparison Table 1 shows the results of our method compared with baselines on ModelNet10-SO(3) under different labeled data ratios. We can see that the results of supervised learning with the labeled data only perform similarly, regardless of using a normal or a Fisher regressor. Since these models are in fact the pre-trained models for the SSL methods in the SSL stage, their similar performance sets a common basis for a fair comparison in the SSL stage. For methods that undergo a second SSL stage, our proposed **FisherMatch** method consistently outperforms the baseline SSL method **SSL-L1-Consistency**, which demonstrates the importance of performing pseudo label filtering.

The experiment results on Pascal3D+ dataset are shown in Table 2. The results illustrate that, with the effective teacher-student mutual learning framework as well as the

Table 1. Comparing our proposed FisherMatch with the baselines on ModelNet10-SO(3) under different ratios of labeled data.

Category	Method	5%		10%	
		Mean↓	Med.↓	Mean↓	Med.↓
Sofa	Sup.-L1 [22]	44.64	11.42	32.65	9.03
	Sup.-Fisher [28]	45.19	13.16	32.92	8.83
	SSL-L1-Consist.	36.86	8.65	25.94	6.81
	SSL-FisherMatch	32.02	7.78	21.29	5.25
	Full Sup.	18.62	5.77	18.62	5.77
Chair	Sup.-L1 [22]	40.41	16.09	29.02	10.64
	Sup.-Fisher [28]	39.34	16.79	28.58	10.84
	SSL-L1-Consist.	31.20	11.29	23.59	8.10
	SSL-FisherMatch	26.69	9.42	20.06	7.44
	Full Sup.	17.38	6.78	17.38	6.78

Table 2. Comparing our proposed FisherMatch with the baselines on the 6 categories of Pascal3D+ dataset with few annotations (7, 20, 50 images). The results are averaged on 6 categories.

Method	7		20		50	
	Med.↓	Acc _{30°} ↑	Med.↓	Acc _{30°} ↑	Med.↓	Acc _{30°} ↑
Res50-Gene	39.1	36.1	26.3	45.2	20.2	54.6
Res50-Spec	46.5	29.6	29.4	42.8	23.0	50.4
StarMap [50]	49.6	30.7	46.4	35.6	27.9	53.8
NeMo [38]	60.0	38.4	33.3	51.7	22.1	69.3
NVSM [39]	37.5	53.8	28.7	61.7	24.2	65.6
FisherMatch	28.3	56.8	23.8	63.6	16.1	75.7
Full Sup.	8.1	89.6	8.1	89.6	8.1	89.6



Figure 2. The visualization of the training process of SSL-FisherMatch on ModelNet10-SO(3) Sofa dataset with 5% labeled data. The four plots, from left to right and from top to bottom, show the mean errors of the predictions, the pseudo label coverage, the pseudo label quality represented by the mean errors of the pseudo labels, and the mean errors between the student model and the corresponding pseudo labels, in the process of training. All the errors are measured in degrees.

entropy-based pseudo label filtering scheme, our algorithm significantly outperforms the state-of-the-art baselines under all different numbers of labeled images.

Training process analysis Here we show how our SSL method works during the training. In Fig. 2, the upper left plot shows that the performance of the unlabeled data

increases together with the test data, which indicates the increasing quality of the teacher predictions. We can also note that the performance on the unlabeled data is slightly better than that of the test data, which is sometimes referred to as *transductive semi-supervised learning*.

We also show the changes over the training process of the pseudo label coverage, the pseudo label quality, and the error between the student predictions and the corresponding pseudo labels, respectively. Here, we refer to *pseudo labels* as the teacher predictions that pass the entropy threshold. The pseudo label coverage means the percentage of teacher predictions that pass the confidence threshold. The pseudo label quality simply means the error of the pseudo labels to the ground truth.

As shown in the curves, as the SSL goes on, the improving model leads to more confident predictions indicated by the decreasing entropy and increasing pseudo label coverage, which in return fuels the learning process. The coverage of pseudo labels increases by a large margin from 40% to the final 70%, while the pseudo label quality still keeps stable with a shaking around 2.5°. This indicates that entropy always acts as a good indicator of performance during the whole process. The error of the student model to the pseudo labels keeps decreasing, which further proves the effectiveness of our unsupervised loss.

4.4. Ablation Study

Effect of Different Unsupervised Loss and Entropy Threshold Here, we analyze the performance of our Fish-




Figure 3. **The performance of FisherMatch with CE or NLL unsupervised losses with different entropy thresholds.** The experiments are done on ModelNet10-SO(3) Sofa dataset with 5% labeled data.




Figure 4. **Comparison of the pseudo label coverage over training process with CE or NLL unsupervised losses and entropy thresholds.** The experiments are done on ModelNet10-SO(3) Sofa dataset with 5% labeled data.




Figure 5. **Visualization of the indication ability of the distribution entropy wrt. the performance.** The horizontal axis is the distribution entropy and the vertical axis is the number of data points, color coded by the errors (in degrees). The experiments are done on ModelNet10-SO(3) Sofa dataset with 10% labeled data.

erMatch with different unsupervised losses, L^{CE} (Eq. 7) and L^{NLL} (Eq. 10), and how they are dependent on the entropy threshold τ by sweeping the parameter τ . Shown in Fig. 3 and 4, the CE loss performs slightly better with a more tolerant threshold, while the NLL loss encourages a higher confidence of the network. The results verify that the NLL loss is a sharpened version of CE loss, where all the pseudo labels passing the threshold are seen as *absolute confident* regardless of the actually predicted uncertainty. This behavior results in a more- but maybe over- confident network, especially with a tolerant entropy threshold. On the other hand, since pseudo labels already exhibit much confidence as they pass the threshold, further sharpening does not lead to additional performance gains. Thus we believe CE loss is a better choice in our task with broader compatibility.

Indication Ability of Distribution Entropy To clearly exhibit the indication ability of the distribution entropy wrt. the performance, we plot the relationship between the error of the prediction and the corresponding distribution entropy on test set in Fig. 5. The figure shows that the entropy effectively captures the prediction error, even under a low

Table 3. **Semi-supervised learning experiment based on the Bingham distribution on ModelNet10-SO(3) Sofa dataset with 10% labeled data.**

Method	Mean↓	Med.↓
Sup.-Bingham	39.61	12.68
Sup.-Fisher	32.92	8.83
SSL-BinghamMatch	27.01	6.77
SSL-FisherMatch	21.29	5.25

labeled data ratio.

Comparison with Bingham-based Regressor Our designed algorithm is agnostic to the choice of the rotation representation as well as the distribution model. We further test our framework based on the Bingham distribution and report the results in Table 3.

As shown in the table, the Bingham-based framework is also able to utilize the unlabeled data and significantly improve the performance of rotation estimation. However, for both its supervised and semi-supervised version, its rotation errors are in general larger than those of matrix Fisher-based framework, since its rotation representation, quaternion, is not a continuous rotation representation, as pointed in [51], thus leading to inferior performance. See supplementary material Section A for detailed settings for SSL-BinghamMatch.

5. Conclusion and Limitations

In this paper, we tackle the problem of semi-supervised rotation regression from single RGB images in a general way. Without requiring any domain-specific knowledge or paired images, we leverage the teacher-student mutual learning framework and propose an entropy-based pseudo label filtering strategy based on the probabilistic modeling of SO(3). Our experiments demonstrate the effectiveness and advantage of our method on both ModelNet10-SO(3) and Pascal3D+ datasets.

The performance of our method may degrade when both the numbers of labeled and unlabeled data are not sufficient. In this case, the uncertainty predicted by our network can be under-estimated due to over-fitting in the small labeled data, leading to reduced effectiveness in the pseudo label filtering and thus the mutual learning.

Acknowledgements

We thank the anonymous reviewers for the insightful feedback. We would like to credit Jiangran Lv from DUT for the fruitful discussions and valuable help in experiments and Yang Wang from PKU for the help in the derivation of maths. This work is supported in part by grants from the Joint NSFC-ISF Research Grant (62161146002).

References

- [1] Arthur Asuncion and David Newman. Uci machine learning repository, 2007. [3](#)
- [2] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. *Advances in neural information processing systems*, 27, 2014. [3](#)
- [3] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. [3](#)
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 32, 2019. [1](#), [3](#), [6](#)
- [5] Michel Breyer, Jen Jen Chung, Lionel Ott, Siegwart Roland, and Nieto Juan. Volumetric grasping network: Real-time 6 dof grasp detection in clutter. In *Conference on Robot Learning*, 2020. [1](#)
- [6] Mai Bui, Tolga Birdal, Haowen Deng, Shadi Albarqouni, Leonidas Guibas, Slobodan Ilic, and Nassir Navab. 6d camera relocalization in ambiguous scenes via continuous multimodal inference. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 139–157. Springer, 2020. [1](#)
- [7] Jiayi Chen, Yingda Yin, Tolga Birdal, Baoquan Chen, Leonidas Guibas, and He Wang. Projective manifold gradient layer for deep rotation regression. *arXiv preprint arXiv:2110.11657*, 2021. [6](#)
- [8] Haowen Deng, Mai Bui, Nassir Navab, Leonidas Guibas, Slobodan Ilic, and Tolga Birdal. Deep bingham networks: Dealing with uncertainty and ambiguity in pose estimation. *arXiv preprint arXiv:2012.11002*, 2020. [2](#), [3](#), [5](#), [6](#)
- [9] Thomas D Downs. Orientation statistics. *Biometrika*, 59(3):665–676, 1972. [2](#)
- [10] Jared Marshall Glover. *The quaternion Bingham distribution, 3D object detection, and dynamic manipulation*. PhD thesis, Massachusetts Institute of Technology, 2014. [5](#)
- [11] Zan Gojcic, Caifa Zhou, Jan D Wegner, Leonidas J Guibas, and Tolga Birdal. Learning multiview 3d point cloud registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1759–1769, 2020. [1](#)
- [12] Chengyue Gong, Dilin Wang, and Qiang Liu. Alphamatch: Improving consistency for semi-supervised learning with alpha-divergence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13683–13692, 2021. [1](#)
- [13] Zijian Hu, Zhengyu Yang, Xuefeng Hu, and Ram Nevatia. Simple: Similar pseudo label exploitation for semi-supervised classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15099–15108, 2021. [3](#)
- [14] Neal Jean, Sang Michael Xie, and Stefano Ermon. Semi-supervised deep kernel learning: Regression with unlabeled data by minimizing predictive variance. *Advances in Neural Information Processing Systems*, 31, 2018. [3](#)
- [15] Zhenyu Jiang, Yifeng Zhu, Maxwell Svetlik, Kuan Fang, and Yuke Zhu. Synergies between affordance and geometry: 6-dof grasp detection via implicit representations. *Robotics: science and systems*, 2021. [1](#)
- [16] Peter E Jupp and Kanti V Mardia. Maximum likelihood estimators for the matrix von mises-fisher and bingham distributions. *The Annals of Statistics*, 7(3):599–606, 1979. [2](#)
- [17] Atul Kanaujia, Cristian Sminchisescu, and Dimitris Metaxas. Semi-supervised hierarchical models for 3d human pose reconstruction. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [1](#)
- [18] Zhanghan Ke, Daoye Wang, Qiong Yan, Jimmy Ren, and Rynson WH Lau. Dual student: Breaking the limits of the teacher in semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6728–6736, 2019. [3](#)
- [19] CG Khatri and Kanti V Mardia. The von mises–fisher matrix distribution in orientation statistics. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):95–106, 1977. [2](#), [4](#)
- [20] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. [1](#), [3](#)
- [21] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. [3](#)
- [22] Jake Levinson, Carlos Esteves, Kefan Chen, Noah Snaveley, Angjoo Kanazawa, Afshin Rostamizadeh, and Ameesh Makadia. An analysis of svd for deep rotation estimation. *Advances in Neural Information Processing Systems*, 33:22554–22565, 2020. [2](#), [6](#), [7](#)
- [23] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 9759–9767, 2019. [6](#)
- [24] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3d hand-object poses estimation with interactions in time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14687–14697, 2021. [1](#)
- [25] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. [1](#), [4](#)
- [26] Octave Mariotti and Hakan Bilen. Semi-supervised viewpoint estimation with geometry-aware conditional generation. In *European Conference on Computer Vision*, pages 631–647. Springer, 2020. [3](#)
- [27] Octave Mariotti, Oisín Mac Aodha, and Hakan Bilen. Viewnet: Unsupervised viewpoint estimation from conditional generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10418–10428, 2021. [1](#), [3](#)

- [28] David Mohlin, Josephine Sullivan, and Gérald Bianchi. Probabilistic orientation estimation with matrix fisher distributions. *Advances in Neural Information Processing Systems*, 33:4884–4893, 2020. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [29] Kieran Murphy, Carlos Esteves, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia. Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold. *arXiv preprint arXiv:2106.05965*, 2021. [2](#)
- [30] Siva Karthik Mustikovela, Varun Jampani, Shalini De Mello, Sifei Liu, Umar Iqbal, Carsten Rother, and Jan Kautz. Self-supervised viewpoint learning from image collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3971–3981, 2020. [3](#)
- [31] Islam Nassar, Samitha Herath, Ehsan Abbasnejad, Wray Buntine, and Gholamreza Haffari. All labels are not created equal: Enhancing semi-supervision via label grouping and co-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7241–7250, 2021. [3](#)
- [32] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7753–7762, 2019. [1](#)
- [33] MJ Prentice. On invariant tests of uniformity for directions and orientations. *The Annals of Statistics*, pages 169–176, 1978. [2](#)
- [34] Michael J Prentice. Orientation statistics without parametric assumptions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 48(2):214–222, 1986. [4](#)
- [35] 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. [2](#)
- [36] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33:596–608, 2020. [1](#), [3](#), [4](#), [6](#)
- [37] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. [1](#), [3](#)
- [38] Angtian Wang, Adam Kortylewski, and Alan Yuille. Nemo: Neural mesh models of contrastive features for robust 3d pose estimation. *arXiv preprint arXiv:2101.12378*, 2021. [6](#), [7](#)
- [39] Angtian Wang, Shenxiao Mei, Alan L Yuille, and Adam Kortylewski. Neural view synthesis and matching for semi-supervised few-shot learning of 3d pose. *Advances in Neural Information Processing Systems*, 34, 2021. [3](#), [6](#), [7](#)
- [40] He Wang, Yezhen Cong, Or Litany, Yue Gao, and Leonidas J Guibas. 3dioumatch: Leveraging iou prediction for semi-supervised 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14615–14624, 2021. [1](#), [2](#), [4](#)
- [41] Yijia Weng, He Wang, Qiang Zhou, Yuzhe Qin, Yueqi Duan, Qingnan Fan, Baoquan Chen, Hao Su, and Leonidas J. Guibas. Captra: Category-level pose tracking for rigid and articulated objects from point clouds. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 13209–13218, October 2021. [1](#)
- [42] 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. [6](#)
- [43] 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*, pages 75–82. IEEE, 2014. [1](#), [6](#)
- [44] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. [1](#)
- [45] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268, 2020. [3](#)
- [46] Rongchang Xie, Chunyu Wang, Wenjun Zeng, and Yizhou Wang. An empirical study of the collapsing problem in semi-supervised 2d human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11240–11249, 2021. [3](#)
- [47] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning*, pages 11525–11536. PMLR, 2021. [3](#)
- [48] Anna Yershova, Swati Jain, Steven M Lavalley, and Julie C Mitchell. Generating uniform incremental grids on so(3) using the hopf fibration. *The International journal of robotics research*, 29(7):801–812, 2010. [2](#)
- [49] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34, 2021. [3](#)
- [50] Xingyi Zhou, Arjun Karapur, Linjie Luo, and Qixing Huang. Starmap for category-agnostic keypoint and viewpoint estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018. [6](#), [7](#)
- [51] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. [1](#), [2](#), [4](#), [8](#)
- [52] Zhi-Hua Zhou, Ming Li, et al. Semi-supervised regression with co-training. In *IJCAI*, volume 5, pages 908–913, 2005. [3](#)