Embedding Propagation: Smoother Manifold for Few-Shot Classification

Pau Rodríguez¹, Issam Laradji^{1,2}, Alexandre Drouin¹, and Alexandre Lacoste¹

¹ Element AI, Montreal, Canada
² University of British Columbia, Canada
{pau.rodriguez,issam.laradji,adrouin,allac}@elementai.com

Abstract. Few-shot classification is challenging because the data distribution of the training set can be widely different to the test set as their classes are disjoint. This distribution shift often results in poor generalization. Manifold smoothing has been shown to address the distribution shift problem by extending the decision boundaries and reducing the noise of the class representations. Moreover, manifold smoothness is a key factor for semi-supervised learning and transductive learning algorithms. In this work, we propose to use embedding propagation as an unsupervised non-parametric regularizer for manifold smoothing in few-shot classification. Embedding propagation leverages interpolations between the extracted features of a neural network based on a similarity graph. We empirically show that embedding propagation yields a smoother embedding manifold. We also show that applying embedding propagation to a transductive classifier achieves new stateof-the-art results in *mini*Imagenet, *tiered*Imagenet, Imagenet-FS, and CUB. Furthermore, we show that embedding propagation consistently improves the accuracy of the models in multiple semi-supervised learning scenarios by up to 16% points. The proposed embedding propagation operation can be easily integrated as a non-parametric layer into a neural network. We provide the training code and usage examples at https://github.com/ElementAI/embedding-propagation.

Keywords: few-shot \cdot classification \cdot semi-supervised learning \cdot metalearning

Deep learning methods have achieved state-of-the-art performance in computer vision tasks such as classification [22], semantic segmentation [30], and object detection [38]. However, these methods often need to be trained on a large amount of labeled data. Unfortunately, labeled data is scarce and its collection is expensive for most applications. This has led to the emergence of deep learning methods based on transfer learning [58], few-shot learning (FSL) [9], and semi-supervised learning [5], that address the challenges of learning with limited data.

Few-shot learning methods have the potential to significantly reduce the need for human annotation. This is because such methods learn new tasks with few labeled examples by transferring the knowledge gained across several tasks. Three



Fig. 1: Illustration of the embedding propagation (EP) method. (a) Original decision boundaries for three classes. The color of a region represents the predicted class, and the color of a node represents the node's actual class. (b) Decision boundaries after applying EP, which are smoother than in (a). (c) Predictions after propagating the labels across the graph, leveraging unlabeled points (light gray) to classify a query example (shown in dark gray). *Best viewed in color*.

recent approaches have been successful for few-shot classification (FSC): metric learning, meta learning, and transfer learning. Metric learning approaches [52,44] learn an embedding space where a set of labeled examples (*support set*) is used to predict the classes for unlabeled examples (*query set*). Meta-learning approaches [10,41] learn to infer a set of parameters that can be adapted to new tasks. Transfer learning [6,31] methods aim to learn a general feature representation and then train a classifier for each new task. In this work, we use an approach between metric learning and transfer learning. During training, the model attempts to learn a general feature representation that is fine-tuned using a metric-based classifier.

A key challenge in few-shot classification is training models that generalize well to unseen classes. This requires feature representations that are robust to small changes in the data distribution. This issue has been addressed outside the few-shot learning literature with a number of regularization techniques such as dropout [45], batch normalization [19], and manifold mixup [51]. However, regularization in few-shot learning remains unexplored. In this work, we show that re-framing label propagation to perform manifold smoothing improves the performance of few-shot classifiers, particularly in the transductive and semisupervised settings. Different from manifold mixup [51], the proposed method is unsupervised and captures higher order interactions between the embedding.

We propose an embedding propagation (EP) method that outputs a set of interpolations from the network output features using their similarity in a graph. This graph is constructed with pairwise similarities of the features using the radial basis function (RBF). EP is non-parametric and can be applied on top of any feature extractor. It can be used as part of a network in order to obtain a regularized manifold for both training and testing. We refer to such network as EPNet. For few-shot classification, we empirically show that the proposed regularization improves the performance for transductive and semi-supervised learning. The hypothesis behind this improvement is based on the fact that using interpolated embeddings result in smoother decision boundaries and increased robustness to noise. These properties have been shown to be important for generalization [1,26,51] and semi-supervised learning [5].

For semi-supervised learning (SSL), EPNet takes advantage of an unlabeled set of images at test time in order to make better predictions of the query set. We adapt the SSL approach proposed by Lee *et al.* [24] to the few-shot classification setup. Thus, for each unlabeled image, EPNet selects the class that has the maximum predicted probability as the pseudo label. EPNet then uses these pseudo labels along with the support set to perform label propagation to predict the labels of the query set. This approach achieves significant improvement over previous state-of-the-art in the 1-shot SSL setting. We hypothesize that EPNet is effective in the SSL setting because of the properties of smoother manifolds [5]. Overall, EPNet achieves state-of-the-art results on *mini*Imagenet [52].

*tiered*Imagenet [38], Imagenet-FS [15] and CUB [54] for few-shot classification, and semi-supervised learning scenarios. In our ablation experiments, we evaluate different variations of embedding propagation and their impact on the smoothness of the decision boundaries. We also show that, with EP, we also achieve a clear improvement on the SSL setup compared to the same model without EP.

Our main contributions can be summarized as follows. We show that embedding propagation:

- Regularizes the manifold in an unsupervised manner.
- Leverages embedding interpolations to capture higher order feature interactions.
- Achieves state-of-the-art few-shot classification results for the transductive and semi-supervised learning setups.

1 Related Work

Our work focuses on few-shot classification, but also intersects with manifold regularization, transductive learning, and semi-supervised learning. We describe relevant work for each of these topics and point out their relevance to our method.

Few-shot classification. A common practice for training models for few-shot learning is to use episodic learning [36,52,44]. This training methodology creates episodes that simulate the train and test scenarios of few-shot learning.

Meta-learning approaches make use of this episodic framework. They learn a base network capable of producing parameters for a task-specific network after observing the support set. The task-specific network is then evaluated on the query set and its gradient is used to update the base network. By doing so, the base network learns to use the support set to generate parameters that are suitable for good generalization. This was first introduced in [36]. Perhaps, the most popular meta-learning approach is MAML [10] and other algorithms that derivate from it [32,35], which learn a set of initial weights that are adapted to a specific task in a small amount of gradient steps. However, this choice of architecture, while general, offers limited performance for few-shot image classification. This lead to variants of meta-learning methods more adapted to image classification [33,41,13,34].

Most metric learning approaches are trained using episodes [52,44,29], they can also be seen as meta-learning approaches. Concretely, metric learning approaches are characterized by a classifier learned over a feature space. They focus on learning high-quality and transferable features with a neural network common to all tasks. EPNet leverages the work of Liu *et al.* [29] for learning to propagate labels, and thus falls into this category. Graph-based approaches can also be framed into this category [43,18,56,23]. Gidaris *et al.* [14] proposed to generate classification weights with a graph neural network (GNN) and apply a denoising autoencoder to regularize their representation. EP does also perform a regularization on a graph representation. Set-to-set functions have also been used for embedding adaptation [57]. However, different from GNNs and set-to-set, our graph is unsupervised and non-parametric, its purpose is manifold smoothing, and we show it improves semi-supervised learning approaches.

While metric learning offers a convenient approach to learn transferable features, it has been shown that neural networks trained with conventional supervised learning already learn transferable features [2,6,31]. Hence, to learn a classifier on a new task, it suffices to fine-tune the feature extractor to that task. Also, this approach has shown to learn more discriminative features compared to the episodic scenario. To take advantage of this transfer learning procedure, we use it in our pre-training phase. Thus, EPNet combines a metric-based classifier with the pre-training of transferable features to achieve a more general representation.

Regularization for Generalization. Regularization is a principled approach for improving the generalization performance of deep networks. Commonly used techniques such as dropout [45] and batch normalization [19] attempt to achieve robustness towards input variations. Others are based on regularizing weights [39,42]. Another line of work that is based on manifold regularization [61,51,49,3]. These works propose methods that aim to smooth the decision boundaries and flatten the class representations, which are important factors for generalization [26,1]. Similarly, we attempt to smooth the manifold by incorporating an embedding propagation operation on the extracted features during training. A concurrent work [31] and our work were the first to apply manifold regularization on few-shot classification. However, the method presented in [31] differs from ours in four ways. First, they perform smoothing in an additional training phase. Second, they train linear classifiers at inference time. Third, they use an exemplar self-supervised loss in their training procedure. Fourth, they do not show the efficacy of their method for semi-supervised learning. In the few-shot classification benchmarks, we achieve better classification accuracy on the Imagenet datasets and CUB dataset for the 1-shot, 5-shot, and 10-shot case.

5



Fig. 2: Overview of the EPNet training procedure. EPNet is trained in two phases: a pretraining phase and an episodic fine-tuning phase. (a) First, the model is trained to learn general feature representations using a standard classification loss \mathcal{L}_C and an auxiliar rotation loss \mathcal{L}_R . (b) Then, the model is fine-tuned using episodic learning to learn to generalize to novel classes by minimizing the standard classification loss \mathcal{L}_C and a label propagation loss \mathcal{L}_P . In both phases the features are encoded using a feature extractor followed by our proposed embedding propagation method (Sec. 2).

There are different lines of research showing that perturbing image representations results in better generalization [45,19]. The most closely related to our work are based on feature interpolation. For instance, Zhao and Cho [7] proposed to make predictions based on the interpolation of nearest neighbors to improve adversarial robustness. In *Manifold Mixup* this idea was expanded to smooth the representations of the neural architecture and achieve better generalization [1,26]. *Manifold Mixup* has been applied to FSC architectures as a fine-tuning step to improve their performance [31]. Differently, we propose a novel procedure to smooth the manifold end-to-end. The proposed method is applied only at the output layer and achieves higher classification accuracy than previous approaches. Moreover, and also different from [31], we leverage the properties of smoother manifolds for semi-supervised learning [4], further widening the improvement margin.

Transductive learning (TL). The idea of transduction is to perform predictions only on the test points. In contrast, the goal of inductive learning is to output a prediction function defined on an entire space [50]. Given a small set of labeled examples, transductive learning has been shown to outperform inductive learning [29,20,28]. This makes TL a desirable method for few-shot classification. Liu *et al.* [29] presented one of the few work using TL for FSC. Similar to this work, we use label propagation [62] to predict the labels of the query set. However, they do not incorporate a manifold smoothing method such as the embedding propagation method investigated in this work.

Semi-Supervised learning. While the literature on semi-supervised learning is vast, few works leverage the use of unlabeled data in few-shot image classification. In [37], they develop a soft version of k-means to meta-learn how to use

unlabeled data. Liu *et al.* [29,28] used label propagation to leverage a set of unlabeled data. Their approach works with both semi-supervised and transductive learning and improves results over soft k-means. In this work we use a similar label propagation approach and show that the semi-supervised results can be further improved by using pseudo-labels (labels obtained from the current model). Recently, Sun *et al.* [27] used a meta-learning approach to cherry-pick examples from the unlabeled set and label them with the current model to increase the label set. In contrast, our method does not require cherry-picking which needs an extra learning step.

2 Proposed Method

We propose an embedding propagation network (EPNet) that has the following pipeline. Given an input image, EPNet first extracts image features using a feature extractor. Then, we apply a novel embedding propagation method (described in Sec. 2.1) to map the features to a set of interpolated features that we refer to as embeddings. These embeddings are then used by a classifier to label the images (Sec. 2.3). The goal of embedding propagation is to increase the smoothness of the embedding manifold, which was shown to improve generalization [1,26] and the effectiveness of semi-supervised learning methods [4] (Sec. 2.5). In the following sections, we explain EPNet in more detail.

2.1 Embedding propagation

Embedding propagation takes a set of feature vectors $\mathbf{z}_i \in \mathbb{R}^m$, obtained from applying a feature extractor (CNN) to the samples of an episode. Then, it outputs a set of embeddings $\tilde{\mathbf{z}}_i \in \mathbb{R}^m$ through the following two steps. First, for each pair of features (i, j), the model computes the distance as $d_{ij}^2 = \|\mathbf{z}_i - \mathbf{z}_j\|_2^2$ and the adjacency as $A_{ij} = \exp(-d_{ij}^2/\sigma^2)$, where σ^2 is a scaling factor and $A_{ii} = 0, \forall i$, as done in TPN [29]. We chose $\sigma^2 = \operatorname{Var}(d_{ij}^2)$ which we found to stabilize training.

Next we compute the Laplacian of the adjacency matrix,

$$L = D^{-\frac{1}{2}}AD^{-\frac{1}{2}}, \quad D_{ii} = \sum_{j} A_{ij}.$$
 (1)

Finally, using the label propagation formula described in [62], we obtain the propagator matrix P as,

$$P = (I - \alpha L)^{-1},$$
 (2)

where $\alpha \in \mathbb{R}$ is a scaling factor, and I is the identity matrix. Then, the embeddings are obtained as follows,

$$\widetilde{\mathbf{z}}_i = \sum_j P_{ij} \mathbf{z}_j. \tag{3}$$

Since the $\tilde{\mathbf{z}}_i$ are now a weighted sum of their neighbors, embedding propagation has the effect of removing undesired noise from the feature vectors. Note that this operation is simple to implement and compatible with a wide range of feature extractors and classifiers. Further, note that the computational complexity of Eq. 2 is negligible for few-shot episodes [29] since the size of the episode is small.

2.2 Few-shot classification setup

Following the common few-shot setups [52,37], we are given three datasets: a base dataset (\mathcal{D}_b) , a novel dataset (\mathcal{D}_n) , and a validation dataset (\mathcal{D}_v) . The base dataset is composed of a large amount of labeled images $\mathcal{D}_b = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_{base}}$, where each image \mathbf{x}_i is labeled with class $y_i \in \mathcal{Y}_{base}$. The novel dataset $\mathcal{D}_n = \{(\mathbf{x}_j, y_j)\}_{j=1}^{N_n ovel}$, where \mathbf{x}_j comes from previously unseen classes $y_j \in \mathcal{Y}_{novel}$, such that $\mathcal{Y}_{base} \cap \mathcal{Y}_{novel} = \emptyset$, is used to evaluate the transfer learning capabilities of the model. The validation dataset \mathcal{D}_v contains classes not present in \mathcal{D}_b and \mathcal{D}_n and is used to conduct hyperparameter search.

Furthermore, we have access to episodes. Each episode consists of n classes sampled uniformly without replacement from the set of all classes, a support set S (k examples per class) and a query set Q (q examples per class). This is referred to as n-way k-shot learning.

2.3 Inference Phase

Given an episode, we perform inference by extracting features of an input image, applying embedding propagation on those features, then applying label propagation. More formally, this is performed as follows. Let $\widetilde{Z} \in \mathbb{R}^{(k+q)\times m}$ be the matrix of propagated embeddings obtained by jointly applying Eq. 1-3 to the support and query sets. Let $P_{\widetilde{Z}}$ be the corresponding propagator matrix. Further, let $Y_S \in \mathbb{R}^{k \times n}$ be a one-hot encoding of the labels in the support set and $\mathbf{0} \in \mathbb{R}^{q \times n}$ a matrix of zeros. We compute the logits for the query set (\hat{Y}_Q) by performing label propagation as described in [62].

2.4 Training procedure

EPNet is trained in two phases as illustrated in Fig. 2. First, the model is trained on \mathcal{D}_b using the common pretraining procedure for few-shot classification [41] in order to learn a general feature representation. Second, the model is finetuned using episodes in order to learn to generalize to novel classes. Episodes are drawn from the same dataset \mathcal{D}_b . In both phases, EPNet uses the same feature extractor $f_{\theta}(\mathbf{x})$ parametrized by θ to obtain the features $\tilde{\mathbf{z}}$ extracted for a given input image \mathbf{x} . However, each phase relies on a different objective.

Pre-training phase. As shown in Fig. 2a, we train f_{θ} using two linear classifiers, which are linear layers with softmax activations parametrized by W_l and W_r , respectively. The first classifier is trained to predict the class labels of examples in \mathcal{D}_b . It is optimized by minimizing the cross-entropy loss,

$$\mathcal{L}_c(\mathbf{x}_i, y_i; W_l, \theta) = -\ln p(y_i | \widetilde{\mathbf{z}}_i, W_l), \tag{4}$$

where $y_i \in \mathcal{Y}_b$ and the probabilities are obtained by applying softmax to the logits provided by the neural network.

For fair comparison with recent literature, we also add a self-supervision loss [31,12] to obtain more robust feature representations. Hence, we use the second classifier to predict image rotations and use the following loss,

$$\mathcal{L}_r(\mathbf{x}_i, r_j; W_r, \theta) = -\ln p(r_j | \mathbf{\tilde{z}}_i, W_r),$$
(5)

where $r_j \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$, and $p(r_j | \tilde{\mathbf{z}}_i, W_r)$ is the probability of the input being rotated by r_j as predicted by a softmax classifier with weights W_r .

Overall, we use stochastic gradient descent (SGD) with batches of size 128 and 4 rotations per image to optimize the following loss,

$$\underset{\theta, W_l, W_r}{\operatorname{argmin}} \sum_{i=1}^{128} \sum_{j=1}^{4} \mathcal{L}_c(\mathbf{x}_i, y_i; W_l, \theta) + \mathcal{L}_r(\mathbf{x}_i, r_j; W_r, \theta).$$
(6)

Episodic Learning phase. As shown in Fig. 2b, after the pre-training phase, we use episodic training to learn to recognize new classes. In this phase, we also optimize EPNet using two classifiers. The first classifier is based on label propagation. It computes class probabilities by applying a softmax to the query set logits \hat{Y}_Q defined in Sec. 2.3, i.e.,

$$\mathcal{L}_p(\mathbf{x}_i, y_i; \theta) = -\ln p(y_i | \widetilde{\mathbf{z}}_i, \widetilde{Z}, Y_S).$$
(7)

The second classifier is identical to the W_l -based classifier used in pretraining. It is included to preserve a discriminative feature representation. Hence, we minimize the following loss:

$$\underset{\theta, W_l}{\operatorname{argmin}} \left[\frac{1}{|Q|} \sum_{(\mathbf{x}_i, y_i) \in Q} \mathcal{L}_p(\mathbf{x}_i, y_i; \theta) + \frac{1}{|S \cup Q|} \sum_{(\mathbf{x}_i, y_i) \in S \cup Q} \frac{1}{2} \mathcal{L}_c(\mathbf{x}_i, y_i; W_l, \theta) \right].$$
(8)

2.5 Semi-supervised learning

In the semi-supervised learning scenario, we also have access to an unlabeled set of images U. We use the unlabeled set as follows. First, we use the inference procedure described in Sec. 2.3 to predict the labels \hat{c}_U for the unlabeled set as pseudo-labels. Then, we augment the support set with U using their pseudo-labels as the true labels. Finally, we use the inference procedure in Sec. 2.3 on the new support set to predict the labels for the query set.

We also consider the semi-supervised scenario proposed by Garcia and Bruna [11]. In this scenario the model is trained to perform 5-shot 5-way classification but only 20% to 60% of the support set is labeled.

As shown by Lee *et al.* [24], this procedure is equivalent to entropy regularization, an effective method for semi-supervised learning. Entropy regularization is particularly effective in cases where the decision boundary lies in low-density regions. With embedding propagation we achieve a similar decision boundary by smoothing the manifold.

3 Experiments

In this section, we present the results on three standard FSC datasets, miniImagenet [52], tieredImagenet [37], CUB [54], and Imagenet-FS [15]. We also provide ablation experiments to illustrate the properties of embedding propagation. As common procedure, we averaged accuracies on \mathcal{D}_n over 1000 episodes [52].

3.1 Datasets

*mini*Imagenet [36]. A subset of the Imagenet dataset [40] consisting of 100 classes with 600 images per class. Classes are divided in three disjoint sets of 64 base classes, 16 for validation and 20 novel classes.

*tiered*Imagenet [37]. A more challenging subset of the Imagenet dataset [40] where class subsets are chosen from supersets of the wordnet hierarchy. The top hierarchy has 34 super-classes, which are divided into 20 base (351 classes), 6 validation (97 classes) and 8 novel (160 classes) categories.

Imagenet-FS [15]. A large-scale version of ImageNet. It is split into 389 base classes and 611 novel classes. The training set consists of 193 of the base classes. Validation consists of 193 of the base classes plus 300 novel classes. The test set consists of the remaining 196 base classes and the remaining 311 novel classes.

CUB [6,16]. A fine-grained dataset based on CUB200 [54] composed of 200 classes and 11,788 images split in 100 base, 50 validation, and 50 novel classes.

3.2 Implementation Details

For fair comparison with previous work, we used three common feature extractors: (i) a 4-layer convnet [52,44] with 64 channels per layer, (ii) a 12-layer resnet [33], and (iii) a wide residual network (WRN-28-10) [41,60]. For *mini*, *tiered*Imagenet, and CUB, images are resized to 84×84 . For Imagenet-FS, as described in [12,13] we use a resnet-18 and images are resized to 224×224 .

In the training stage, the models are optimized using SGD with learning rate of 0.1 for 100 epochs. The learning rate is reduced by a factor of 10 every time the model reached a plateau, in which case the validation loss had not improved for 10 epochs. α is cross-validated on the 4-layer convnet.

In the episodic fine-tuning stage we randomly sample 5 classes per episode, where in each class k instances are selected for the support set and 15 for the query set. Similar to the training stage, the model is optimized with SGD with learning rate 0.001 reduced by a factor of 10 on plateau.

For training the wide residual networks (WRN), we apply the standard data augmentation methods mentioned by Szegedy *et al.* [48,41]. For the other architectures, we do not apply data augmentation.

For Imagenet-FS, we add EP right after the denoising autoencoder of wDAE-GNN [14]. We use the original code provided by the authors.³. Different from

³ github.com/gidariss/wDAE_GNN_FewShot

Table 1: Comparison of test accuracy against state-of-the art methods for 1shot and 5-shot classification using *mini*Imagenet and *tiered*Imagenet. The second column shows the number of parameters of each model in thousands (K). **Robust* - 20 uses an 18-layer residual network. --Net is identical to EPNet but without EP. Gray colored results are obtained using 224×224 pixels instead of the standard 84×84 pixel images.

		miniIn	tieredImagenet										
	Params	1-shot	5-shot	1-shot	5-shot								
	CONV-4												
Matching [52]	112K	43.56 ± 0.84	55.31 ± 0.73	-	-								
MAML [29]	112K	48.70 ± 1.84	63.11 ± 0.92	51.67 ± 1.81	$70.30{\scriptstyle~\pm0.08}$								
ProtoNet [44]	112K	$49.42{\scriptstyle~\pm 0.78}$	68.20 ± 0.66	53.31 ± 0.89	$72.69{\scriptstyle~\pm 0.74}$								
ReNet $[47]$	223K	50.44 ± 0.82	$65.32{\scriptstyle~\pm 0.70}$	54.48 ± 0.92	$71.32{\scriptstyle~\pm 0.78}$								
GNN [11]	1619K	50.33 ± 0.36	66.41 ± 0.63	-	-								
TPN [29]	171K	53.75 ± 0.86	69.43 ± 0.67	57.53 ± 0.96	72.85 ± 0.74								
CC+rot [12]	112K	54.83 ± 0.43	71.86 ± 0.33	-	-								
EGNN [21]	5068K	-	76.37	-	80.15								
Net (ours)	112K	57.18 ± 0.83	72.57 ± 0.66	57.60 ± 0.93	73.30 ± 0.74								
EPNet (ours)	112K	$59.32 \hspace{0.1 in} \pm 0.88$	72.95 ± 0.64	59.97 ± 0.95	73.91 ± 0.75								
RESNET-12													
ProtoNets++ [55]	7989K	56.52 ± 0.45	74.28 ± 0.20	58.47 ± 0.64	78.41 ± 0.41								
TADAM [33]	7989K	58.50 ± 0.30	76.70 ± 0.30	-	-								
MetaOpt-SVM [25]	12415K	62.64 ± 0.61	78.60 ± 0.46	65.99 ± 0.72	81.56 ± 0.53								
TPN [29]	8284K	59.46	75.65	-	-								
* Robust-20++ [8]	11174K	58.11 ± 0.64	75.24 ± 0.49	70.44 ± 0.32	85.43 ± 0.21								
MTL [46]	8286K	61.20 ± 1.80	75.50 ± 0.80	-	-								
CAN [17]	8026K	67.19 ± 0.55	80.64 ± 0.35	73.21 ± 0.58	84.93 ± 0.38								
Net (ours)	7989K	65.66 ± 0.85	$81.28 \hspace{0.1 cm} \pm 0.62$	72.60 ± 0.91	85.69 ± 0.65								
EPNet (ours)	7989K	66.50 ± 0.89	$81.06 \hspace{0.1 in} \pm 0.60$	76.53 ± 0.87	$87.32 \hspace{0.1 in} \pm 0.64$								
		WRN-28	8-10										
LEO [41]	37582K	61.76 ± 0.08	77.59 ± 0.12	66.33 ± 0.05	81.44 ±0.09								
Robust- $20 + [8]$	37582K	62.80 ± 0.62	$80.85{\scriptstyle~\pm 0.43}$	-	-								
wDAE-GNN [14]	48855K	62.96 ± 0.15	78.85 ± 0.10	68.18 ± 0.16	83.09 ± 0.12								
CC+rot [12]	37582K	62.93 ± 0.45	$79.87{\scriptstyle~\pm 0.33}$	70.53 ± 0.51	84.98 ± 0.36								
Manifold mixup [31]	37582K	64.93 ± 0.48	83.18 ± 0.72	-	-								
Net (ours)	37582K	65.98 ± 0.85	82.22 ± 0.66	74.04 ± 0.93	86.03 ± 0.63								
EPNet (ours)	37582K	$\textbf{70.74} \pm 0.85$	84.34 ± 0.53	78.50 ± 0.91	88.36 ± 0.57								

the other datasets, in this one evaluation is performed on the 311 test classes at the same time (311-way), with the number of supports $k \in \{1, 2, 5, 10, 20\}$.

We evaluate 3 variations of our method: (i) EPNet as described in Sec. 2; (ii) --Net, which is identical to EPNet but without EP; and (iii) EPNet_{SSL} (semi-supervised learning) as described in Sec. 2.5.

Table 2: Comparison with the state of the art on CUB-200-2011. *Robust-20++ uses an 18-layer residual network, and Accuracies obtained with 224×224 images appear in gray.

	backbone	1-shot	5-shot
*Robust-20++ $[8]$ EPNet (ours)	RESNET-18 RESNET-12	$ \begin{array}{c} 68.68 \pm 0.69 \\ \textbf{82.85} \pm 0.81 \end{array} $	$83.21 \pm 0.44 \\ \textbf{91.32} \pm 0.41$
Manifold mixup [31] EPNet (ours)	WRN-28-10 WRN-28-10	$ \begin{vmatrix} 80.68 \\ \pm 0.81 \\ 87.75 \\ \pm 0.70 \end{vmatrix} $	$\begin{array}{c} 90.85 \\ \pm 0.44 \\ \textbf{94.03} \\ \pm 0.33 \end{array}$

Table 3: Top-5 test accuracy on Imagenet-FS.

			Novel Classe	5	All classes					
Approach	K=1	2	5	10	20	K=1	2	5	10	20
Batch SGM [15]	-	-	-	-	-	49.3	60.5	71.4	75.8	78.5
PMN [53]	45.8	57.8	69.0	74.3	77.4	57.6	64.7	71.9	75.2	77.5
LwoF [13]	46.2	57.5	69.2	74.8	78.1	58.2	65.2	72.7	76.5	78.7
CC+ Rot [12]	46.43 ± 0.24	57.80 ± 0.16	69.67 ± 0.09	74.64 ± 0.06	77.31 ± 0.05	57.88 ± 0.15	64.76 ± 0.10	72.29 ± 0.07	75.63 ± 0.04	77.40 ± 0.03
wDAE-GNN [14]	48.00 ± 0.21	59.70 ± 0.15	70.30 ± 0.08	75.00 ± 0.06	77.80 ± 0.05	59.10 ± 0.13	66.30 ± 0.10	73.20 ± 0.05	76.10 ± 0.04	77.50 ± 0.03
wDAE-GNN + EP (ours)	50.07 ± 0.27	62.16 ± 0.16	72.89 ± 0.11	77.25 ± 0.07	79.48 ± 0.05	60.87 ± 0.16	$68.53 {\scriptstyle \pm 0.10}$	75.56 ± 0.07	78.28 ± 0.04	$\textbf{78.89} \pm 0.03$

Table 4: SSL results with 100 unlabeled samples. --Net is identical to EPNet but without embedding propagation. *Re-implementation [59]

		miniIm	tieredImagenet			
	Backbone	1-shot	5-shot	1-shot	5-shot	
TPN_{SSL} [29]	CONV-4	52.78	66.42	55.74	71.01	
k-Means _{masked,soft} [37]	CONV-4	50.41 ± 0.31	64.39 ± 0.24	-	-	
EPNet (ours)	CONV-4	59.32 ± 0.88	72.95 ± 0.64	59.97 ± 0.95	73.91 ± 0.75	
Net_{SSL} (ours)	CONV-4	63.74 ± 0.97	75.30 ± 0.67	65.01 ± 1.04	74.24 ± 0.80	
$EPNet_{SSL}$ (ours)	CONV-4	$65.13 \pm 0.97 $	75.42 ± 0.64	$66.63 \hspace{0.1 in} \pm 1.04$	$75.70 \hspace{0.1 in} \pm 0.74$	
LST [27]	RESNET-12	70.10 ± 1.90	78.70 ± 0.80	77.70 ± 1.60	85.20 ±0.80	
EPNet (ours)	RESNET-12	66.50 ± 0.89	81.06 ± 0.60	76.53 ± 0.87	87.32 ± 0.64	
Net_{SSL} (ours)	RESNET-12	73.42 ± 0.94	83.17 ± 0.58	80.26 ± 0.96	88.06 ± 0.59	
$EPNet_{SSL}$ (ours)	RESNET-12	$75.36 \hspace{0.1 in} \pm 1.01$	84.07 ± 0.60	$81.79 \hspace{0.1 in} \pm 0.97$	88.45 ± 0.61	
*k-Means _{masked,soft} [37]	WRN-28-10	52.78 ± 0.27	66.42 ± 0.21	-	-	
TransMatch [59]	WRN-28-10	63.02 ± 1.07	$81.19{\scriptstyle~\pm 0.59}$	-	-	
EPNet (ours)	WRN-28-10	70.74 ± 0.85	84.34 ± 0.53	78.50 ± 0.91	88.36 ± 0.57	
Net_{SSL} (ours)	WRN-28-10	77.70 ± 0.96	86.30 ± 0.50	82.03 ± 1.03	88.20 ± 0.61	
$EPNet_{SSL}$ (ours)	WRN-28-10	$79.22 \hspace{0.1 in} \pm 0.92$	$88.05 \hspace{0.1 in} \pm 0.51$	$83.69 \hspace{0.1 cm} \pm 0.99$	$89.34 {\ \pm 0.59}$	

3.3 Experimental Results

In this section we compare EPNet with previous methods in the standard fewshot classification scenario, and semi-supervised learning scenarios.

Main Results. We first report the results for the methods that do not use an unlabeled set. As seen in Tables 1 and 2, EPNet obtains state-of-the-art accuracies on miniImagenet, tieredImagenet, and CUB-200-2011 for the 1-shot and

Table 5: SSL results for the 5-shot 5-way scenario with different amounts of unlabeled data. The percentages refer to the amount of supports that are labeled in a set of 5 images per class.

	Params	20%	40%	60%	100%
GNN [11]	112K	52.45	58.76	-	66.41
EGNN [21]	5068K	63.62	64.32	66.37	76.37
Net_{SSL} (ours)	112K	58.52 ± 0.97	64.46 ± 0.79	67.81 ± 0.74	$57.18{\scriptstyle~\pm 0.83}$
$EPNet_{SSL}$ (ours)	112K	60.66 ± 0.97	67.08 ± 0.80	$68.74 \hspace{0.1cm} \pm 0.74$	$59.32{\scriptstyle~\pm 0.88}$

5-shot benchmarks even when compared with models that use more parameters or higher resolution images. It can also be observed the effectiveness of EP in isolation when comparing EPNet with an identical model without EP (--Net). Higher-way and 10-shot results can be found in the supplementary material. Note that EGNN [21] uses $\times 45$ parameters. On the large-scale Imagenet-FS, EP improves all benchmarks by approximately 2% accuracy, see Table 3. These results demonstrate the scalability of our method and the orthogonality with other embedding transformations such as denoising autoencoders [14].

Semi-supervised learning. We evaluate EPNet on the SSL setting where 100 additional unlabeled samples are available [37,29] (EPNet_{SSL}). We observe in Table 4 that including unlabeled samples increases the accuracy of EPNet for all settings, surpassing the state of the art by a wide margin of up to 16% accuracy points for the 1-shot WRN-28-10. Similar to previous experiments, removing EP to EPNet (--Net) is detrimental for the model, supporting our hypotheses. Following the same setting described in [11,21], we trained our model in the 5-shot 5-way scenario where the support samples are partially labeled. In Table 5, we report the test accuracy with conv-4 when labeling 20%, 40%, 60% and 100% of the support set. EPNet obtains up to 2.7% improvement over previous state-of-the-art when 40% of the support are labeled. Moreover, EPNet also outperforms EGNN [21] in the 40% and 60% scenarios, although EPNet has $45 \times$ less parameters.

3.4 Ablation Studies

In this section we investigate the different hyperparameters of our method and the properties derived from embedding propagation. Additional ablations are provided in the supplementary material.

Algorithm ablation. In Table 6, we investigate the impact of the rotation loss (ROT), embedding fine-tuning (EFT), label propagation (LP), and embedding propagation (EP) on the 1-shot *mini*Imagenet accuracy. When label propagation is deactivated, we substitute it with a prototypical classifier. Interestingly, it can be seen that the improvement is larger when using LP in combination

Table 6: Algorithm ablation with conv-4 on 1-shot *mini*Imagenet. EFT: Episodic Fine-tuning, ROT: Rotation loss, LP: Label Propagation, EP: Embedding Propagation

0																
EXP	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
EFT									\checkmark	\checkmark	\checkmark	 ✓ 	√	 ✓ 	√	 ✓
ROT					√	\checkmark	\checkmark	\checkmark					\checkmark	 ✓ 	\checkmark	 ✓
LP			 ✓ 	\checkmark			\checkmark	\checkmark			\checkmark	 ✓ 			\checkmark	 ✓
EP		\checkmark		\checkmark		\checkmark		\checkmark		\checkmark		\checkmark		 ✓ 		√
ACC	49.57	52.83	53.40	55.75	50.83	53.63	53.38	55.55	54.29	56.38	56.93	58.35	54.92	56.46	57.35	58.85



Fig. 3: Interpolation of embedding pairs of different classes vs probability of belonging to the first of these two classes. The top row shows the class probability for resnet-12 embeddings extracted from EPNet, and the second(--Net) from the same network trained without embedding propagation. The scalar α controls the weight of the first embedding in the linear interpolation.

with EP (Table 6; columns 2-4, and 10-12). This finding is in accordance with the hypothesis that EP performs manifold smoothing, and this is beneficial for transductive and SSL algorithms. We included a rotation loss for fair comparison with other SotA [12,31], however, we see that the main improvement is due to the combination of EP with LP. We also find that episodic fine-tuning successfully adapts our model to the episodic scenario (Table 6; line 2).

Embedding propagation on manifold smoothness. We explore whether embedding propagation helps the classifier to attain smoother decision boundaries. We use EPNet to obtain image embeddings and select a set of random pairs $\mathbf{z}_i, \mathbf{z}_j$ that belong to different classes y_i, y_j . We then interpolate between each pair as $\tilde{\mathbf{z}} = \alpha \cdot \mathbf{z}_i + (1-\alpha)\mathbf{z}_j$ where $\alpha \in [0..1]$, and plot this value against $p(y_i|\tilde{\mathbf{z}})$ (Sec. 7) in Fig. 3. We also plot $p(y_i|\hat{\mathbf{z}})$ where embeddings were obtained using EPNet without embedding propagation (--Net). We observe that EPNet has significantly smoother probability transitions than --Net as the embedding $\tilde{\mathbf{z}}$ changes from \mathbf{z}_i to \mathbf{z}_j . In contrast, --Net yields sudden probability transitions. This suggests that embedding propagation encourages smoother decision boundaries.

In Figure 4, we show the effect of embedding propagation on a toy dataset. The dataset consists of embeddings that are arranged in two disjoint moons. The embeddings in the top moon belong to first class, and the other to the second class. Figure 4a) illustrates the effect of batch sampling during the pre-training



Fig. 4: Visualization of embedding propagation on the two moons dataset. *The embeddings* are shown on the same scale.

phase. Gray points correspond to the extracted embeddings when no embedding propagation was applied. Each unique color shows multiple projections of the same gray point. Each projection is performed using a different batch of the data. This suggests that the projections of the same embedding fill a wide space in the manifold. As a result, the density and smoothness at the inter-class boundary increases. Figure 4b) shows the result of applying embedding propagation on all the gray points in Figure 4a). The blue and red colors correspond to the two-moon classes. We observe that the propagated manifold is denser and more compact than the original one, possibly reducing the noise of the representations.

4 Conclusion

We have presented EPNet, an approach to address the problem of distribution shift few-shot learning. EPNet introduces a simple embedding propagation step to regularize the feature representations. Empirically, we have shown that embedding propagation smooths decision boundaries, which is an important factor for generalization. By leveraging the properties of smooth manifolds, we have shown significant improvements on the transductive, and semi-supervised learning setup compared to methods that do not use embedding propagation. As a result, EPNet achieves state-of-the-art results on *mini*Imagenet, *tiered*Imagenet, and CUB for the standard and semi-supervised scenarios. Further, we have shown that EP scales to the larger Imagenet-FS dataset, improving by more than 2% the accuracy of the state-of-the-art wDAE-GNN [14] in all setups. We have compared EPNet with a non-smooth version of the same model (--Net), showing that smoothing alone accounts for 4.8% accuracy improvement on 1-shot *mini*Imagenet with Wide Residual Networks.

References

- 1. Bartlett, P., Shawe-Taylor, J.: Generalization performance of support vector machines and other pattern classifiers. Advances in Kernel methods—support vector learning (1999)
- Bauer, M., Rojas-Carulla, M., Światkowski, J.B., Schölkopf, B., Turner, R.E.: Discriminative k-shot learning using probabilistic models. arXiv preprint arXiv:1706.00326 (2017)
- Belkin, M., Niyogi, P., Sindhwani, V.: Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. JMLR 7(Nov), 2399–2434 (2006)
- 4. Chapelle, O., Scholkopf, B., Zien, A.: Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. IEEE Transactions on Neural Networks (2009)
- Chapelle, O., Zien, A.: Semi-supervised classification by low density separation. In: AISTATS. vol. 2005, pp. 57–64 (2005)
- Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C.F., Huang, J.B.: A closer look at few-shot classification. In: ICLR (2019)
- Cho, K., et al.: Retrieval-augmented convolutional neural networks against adversarial examples. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11563–11571 (2019)
- Dvornik, N., Schmid, C., Mairal, J.: Diversity with cooperation: Ensemble methods for few-shot classification. In: ICCV. pp. 3723–3731 (2019)
- Fei-Fei, L., Fergus, R., Perona, P.: One-shot learning of object categories. TPAMI 28(4), 594–611 (2006)
- Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML. pp. 1126–1135. JMLR. org (2017)
- 11. Garcia, V., Bruna, J.: Few-shot learning with graph neural networks. ICLR (2017)
- Gidaris, S., Bursuc, A., Komodakis, N., Pérez, P., Cord, M.: Boosting few-shot visual learning with self-supervision. In: CVPR. pp. 8059–8068 (2019)
- Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: CVPR. pp. 4367–4375 (2018)
- Gidaris, S., Komodakis, N.: Generating classification weights with gnn denoising autoencoders for few-shot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 21–30 (2019)
- Hariharan, B., Girshick, R.: Low-shot visual recognition by shrinking and hallucinating features. In: ICCV. pp. 3018–3027 (2017)
- Hilliard, N., Phillips, L., Howland, S., Yankov, A., Corley, C.D., Hodas, N.O.: Few-shot learning with metric-agnostic conditional embeddings. arXiv preprint arXiv:1802.04376 (2018)
- Hou, R., Chang, H., Bingpeng, M., Shan, S., Chen, X.: Cross attention network for few-shot classification. In: Advances in Neural Information Processing Systems. pp. 4005–4016 (2019)
- Hu, Y., Gripon, V., Pateux, S.: Exploiting unsupervised inputs for accurate fewshot classification (2020)
- Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML. pp. 448–456 (2015)
- Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Label propagation for deep semisupervised learning. In: CVPR. pp. 5070–5079 (2019)
- Kim, J., Kim, T., Kim, S., Yoo, C.D.: Edge-labeling graph neural network for fewshot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11–20 (2019)

- 16 P. Rodríguez et al.
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NeurIPS. pp. 1097–1105 (2012)
- Kye, S.M., Lee, H.B., Kim, H., Hwang, S.J.: Transductive few-shot learning with meta-learned confidence. arXiv preprint arXiv:2002.12017 (2020)
- Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on challenges in representation learning, ICML. vol. 3, p. 2 (2013)
- Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: CVPR. pp. 10657–10665 (2019)
- Lee, W.S., Bartlett, P.L., Williamson, R.C.: Lower bounds on the vc dimension of smoothly parameterized function classes. Neural Computation 7(5), 1040–1053 (1995)
- Li, X., Sun, Q., Liu, Y., Zhou, Q., Zheng, S., Chua, T.S., Schiele, B.: Learning to self-train for semi-supervised few-shot classification. In: NeurIPS. pp. 10276–10286 (2019)
- Liu, B., Wu, Z., Hu, H., Lin, S.: Deep metric transfer for label propagation with limited annotated data. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 0–0 (2019)
- Liu, Y., Lee, J., Park, M., Kim, S., Yang, E., Hwang, S., Yang, Y.: Learning to propagate labels: transductive propagation network for few-shot learning. In: ICLR (2019)
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440 (2015)
- Mangla, P., Singh, M., Sinha, A., Kumari, N., Balasubramanian, V.N., Krishnamurthy, B.: Charting the right manifold: Manifold mixup for few-shot learning. arXiv preprint arXiv:1907.12087 (2019)
- Nichol, A., Schulman, J.: Reptile: a scalable metalearning algorithm. arXiv preprint arXiv:1803.02999 (2018)
- Oreshkin, B., López, P.R., Lacoste, A.: Tadam: Task dependent adaptive metric for improved few-shot learning. In: NeurIPS. pp. 721–731 (2018)
- Qiao, S., Liu, C., Shen, W., Yuille, A.L.: Few-shot image recognition by predicting parameters from activations. In: CVPR. pp. 7229–7238 (2018)
- Rajeswaran, A., Finn, C., Kakade, S.M., Levine, S.: Meta-learning with implicit gradients. In: NeurIPS. pp. 113–124 (2019)
- Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2016)
- Ren, M., Ravi, S., Triantafillou, E., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: ICLR (2018)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS. pp. 91–99 (2015)
- Rodríguez, P., Gonzalez, J., Cucurull, G., Gonfaus, J.M., Roca, X.: Regularizing cnns with locally constrained decorrelations. ICLR (2016)
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. IJCC (2015)
- Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: ICLR (2018)
- 42. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In: NeurIPS. pp. 901–909 (2016)

- 43. Satorras, V.G., Estrach, J.B.: Few-shot learning with graph neural networks. In: ICLR (2018)
- 44. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: NeurIPS. pp. 4077–4087 (2017)
- 45. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. JMLR 15(1), 1929–1958 (2014)
- Sun, Q., Liu, Y., Chua, T.S., Schiele, B.: Meta-transfer learning for few-shot learning. In: CVPR. pp. 403–412 (2019)
- Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR. pp. 1199–1208 (2018)
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR. pp. 1– 9 (2015)
- Tokozume, Y., Ushiku, Y., Harada, T.: Between-class learning for image classification. In: CVPR. pp. 5486–5494 (2018)
- 50. Vapnik, V.N.: An overview of statistical learning theory. IEEE transactions on neural networks (1999)
- Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., Bengio, Y.: Manifold mixup: Better representations by interpolating hidden states. In: ICML. pp. 6438–6447 (2019)
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: NeurIPS. pp. 3630–3638 (2016)
- Wang, Y.X., Girshick, R., Hebert, M., Hariharan, B.: Low-shot learning from imaginary data. In: CVPR. pp. 7278–7286 (2018)
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-UCSD Birds 200. Tech. Rep. CNS-TR-2010-001, California Institute of Technology (2010)
- 55. Xing, C., Rostamzadeh, N., Oreshkin, B., Pinheiro, P.O.: Adaptive cross-modal few-shot learning. In: NeurIPS. pp. 4848–4858 (2019)
- Yang, L., Li, L., Zhang, Z., Zhou, X., Zhou, E., Liu, Y.: Dpgn: Distribution propagation graph network for few-shot learning. In: CVPR. pp. 13390–13399 (2020)
- 57. Ye, H.J., Hu, H., Zhan, D.C., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: CVPR. pp. 8808–8817 (2020)
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? In: NeurIPS. pp. 3320–3328 (2014)
- Yu, Z., Chen, L., Cheng, Z., Luo, J.: Transmatch: A transfer-learning scheme for semi-supervised few-shot learning. In: CVPR. pp. 12856–12864 (2020)
- Zagoruyko, S., Komodakis, N.: Wide residual networks. In: BMVC. pp. 87.1–87.12. BMVA Press (September 2016). https://doi.org/10.5244/C.30.87
- Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: ICLR (2018)
- Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Schölkopf, B.: Learning with local and global consistency. In: NeurIPS. pp. 321–328 (2004)