

Deep Metric Transfer for Label Propagation with Limited Annotated Data

Bin Liu^{1†} Zhirong Wu^{2†} Han Hu^{2*} Stephen Lin²
¹Tsinghua University ²Microsoft Research Asia

liub18@mails.tsinghua.edu.cn {wuzhiron, hanhu, stevelin}@microsoft.com

Abstract

We study object recognition under the constraint that each object class is only represented by very few observations. Semi-supervised learning, transfer learning, and few-shot recognition all concern with achieving fast generalization with few labeled data. In this paper, we propose a generic framework that utilize unlabeled data to aid generalization for all three tasks. Our approach is to create much more training data through label propagation from the few labeled examples to a vast collection of unannotated images. The main contribution of the paper is that we show such a label propagation scheme can be highly effective when the similarity metric used for propagation is transferred from other related domains. We test various combinations of supervised and unsupervised metric learning methods with various label propagation algorithms. We find that our framework is very generic without being sensitive to any specific techniques. By taking advantage of unlabeled data in this way, we achieve significant improvements on all three tasks. Code is available at <http://github.com/Microsoft/metric-transfer.pytorch>.

1. Introduction

We address the problem of object recognition from a very small amount of labeled data. This problem is of particular importance when limited labels can be collected due to either time or financial constraints. Though this is a difficult challenge, we are encouraged by evidence from cognitive science suggesting that infants can quickly learn new concepts from very few examples [21, 1].

Many recognition problems in computer vision are concerned with learning on few labeled data. Semi-supervised learning, transfer learning, and few-shot recognition all aim to achieve fast generalization from few examples, by leveraging unlabeled data or labeled data from other domains.

The fundamental difficulty of this problem is that naive supervised training with very few examples results in severe

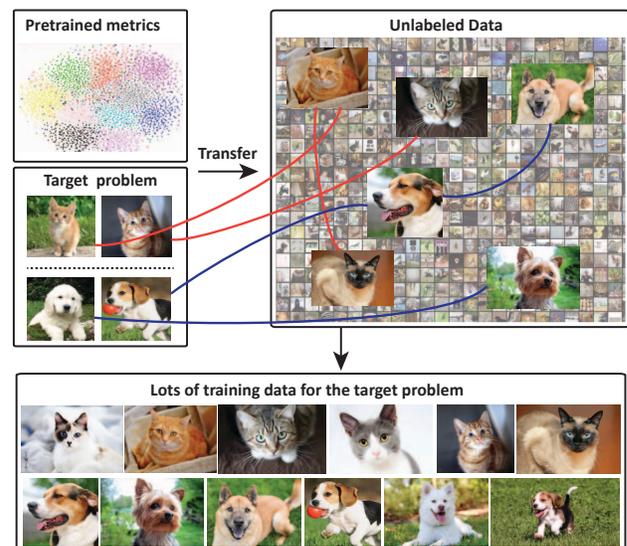


Figure 1: Overview of the approach. Often, object categories are represented by very few images. We transfer a metric learned from another domain and propagate the labels from the few labeled images to a vast collection of unannotated images. We show this can reliably create much more labeled data for the target problem.

over-fitting. Because of this, prior work in semi-supervised learning rely on strong regularizations such as augmentations [10], temporal consistency [20], and adversarial examples [27] to improve performance. Some related works in few-shot learning do not even refine an online classifier. Instead, they simply apply the similarity metric learned from training categories to new categories without adaptation. Meta-learning [8] seeks to optimize an online parametric classifier with few samples, but under the assumption that just a few steps of optimization will lead to effective generalization with less overfitting. These approaches indirectly address the inherent problem of limited training data.

In this paper, we propose a new framework of label propagation via metric transfer to tackle the problem of limited training data. We propagate labels to an unlabeled dataset, so that training a supervised model with great learning ca-

[†]Equal contribution. ^{*}Corresponding author. The work is done when Bin Liu is an intern at MSRA.

capacity no longer faces over-fitting. This approach is related to work on “pseudo-labeling” [22, 31], where the model is bootstrapped from limited data and trained on the new data/label pairs it infers. However, that is unlikely to work well when the labeled data is scarce, since the initial model is likely to be poor. Instead of bootstrapping, our work transfers the metric learned from another related domain, and thus provides a much better generalization ability.

Our approach works with three data domains: a source domain to learn a similarity metric, few labeled examples to define the target problem, and an unlabeled dataset in which to propagate labels. As in Figure 1, we first learn a similarity metric on the source domain, which can be either labeled or unlabeled. Supervised learning or unsupervised (self-supervised) learning is used to learn the metric accordingly. Then, given few observations of the target problem, we propagate the labels from these observations to the unlabeled dataset using the metric learned in the source domain. This creates an abundance of labeled data for learning a classifier. Finally, we train a standard supervised model using the propagated labels.

The main contribution of this work is the metric transfer approach for label propagation. By studying different combinations of metric pretraining methods (*e.g.* unsupervised, supervised) and label propagation algorithms (*e.g.* nearest neighbors, spectral clustering), we find that our metric transfer approach on unlabeled data is general enough to work effectively for many settings. For semi-supervised learning on CIFAR10 and ImageNet, we obtain an absolute 20% improvement over the state-of-the-art when labeled data is limited (5 – 10 labels per category). We also achieve a 10% improvement on transferring representations from ImageNet to CIFAR10 for transfer learning, and 6% improved performance for few-shot recognition on the mini-ImageNet benchmark.

Due to this generic framework, our work also brings individual insights into the respective tasks we studied: 1) for semi-supervised learning, algorithms may better develop from unsupervised learning, as opposed to using unlabeled data for regularization. 2) for transfer learning, we propose an alternative method for transferring knowledge other than the dominant finetuning approach. 3) for few-shot recognition, in certain scenarios, unlabeled data in the target domain is more beneficial than labeled data in the source domain.

2. Related Work

Large-scale Recognition. To solve a computer vision problem, it has become a common practice to build a large-scale dataset [6, 3] and train deep neural networks [19, 34] on it. This philosophy has achieved unprecedented success on many important computer vision problems [6, 24, 32].

However, constructing a large-scale dataset is often time-consuming and expensive, and this has motivated work on unsupervised learning and problems defined on few labeled samples.

Semi-supervised Learning. Semi-supervised learning [39] is a problem that lies in between supervised learning and unsupervised learning. It aims to make more accurate predictions by leveraging a large amount of unlabeled data than by relying on the labeled data alone. In the era of deep learning, one line of work leverages unlabeled data through deep generative models [17, 29]. However, training of generative models is often unstable, making it tricky to work with recognition tasks. Recent efforts on semi-supervised learning focus on regularization by self-ensembling through consistency loss, such as temporal ensembling [20], adversarial ensembling [27], teacher-student distillation [36], and cross-view ensembling [2]. The pseudo-labeling approach [22, 31] initializes a model on a small labeled dataset and bootstraps on the new data it predicts. This tends to fail when the labeled set is small.

Our work is most closely related to the transductive approaches [16, 45]. Prior work [7] in computer vision shows that label propagation can work well with handcrafted GIST descriptors. We bring it to the context of deep learning, and demonstrate that metric transfer may further improve the accuracy of label propagation.

Few-shot Recognition. Given some training data in training categories, few-shot recognition [1] requires the classifier to generalize to new categories from observing very few examples, often 1-shot or 5-shot. A body of work approaches this problem by offline metric learning [37, 35, 40], where a generic similarity metric is learned on the training data and directly transferred to the new categories using simple nearest neighbor classifiers without further adaptation. Recent works on meta-learning [8, 23, 26] take a learning-to-learn approach using online algorithms. In order not to overfit to the few examples, they develop meta-learners to find a common embedding space, which can be further finetuned with fast convergence to the target problem. Recent works [30, 9] using meta-learning consider the combined problem of semi-supervised learning and few-shot recognition, by allowing access to unlabeled data in few-shot recognition. This drives few-shot recognition into more realistic scenarios. We follow this setting as we study few-shot recognition.

Transfer Learning. Since the inception of the ImageNet challenge [32], transfer learning has emerged almost everywhere in visual recognition, such as in object detection [11] and semantic segmentation [25], by simply transferring the network weights learned on ImageNet classification and finetuning on the target task. When the pretraining task and the target task are closely related, this tends to generalize

much better than training from scratch on the target task alone. Domain adaptation seeks to address a much more difficult scenario where there is a large gap between the inputs of the source and target domains [14], for example, between real images and synthetic images. What we study in this paper is metric transfer. Different from prior work [42] that employ metric transfer just to reduce the distribution divergence of different domains, we use metric transfer to propagate labels. Through this, we show that metric propagation is an effective method for learning with small data.

3. Approach

To deal with the shortage of labeled data, our approach is to enlarge it by propagating labels from annotated images to unlabeled data using the similarity metric between data pairs. The creation of much more labeled data enables us to train deep neural networks to their full learning capacity.

Our framework works on three data domains: the source domain \mathcal{S} , the target domain \mathcal{T} , and additional unlabeled data \mathcal{U} . The source domain \mathcal{S} can be labeled or unlabeled with abundant data, and it is used to learn a generic similarity metric between data pairs. The target domain \mathcal{T} only has few labeled data, but it defines the problem we want to optimize. The unlabeled data \mathcal{U} is the resource in which to propagate labels, and may potentially contain similar classes to the task defined in \mathcal{T} . It may or may not have overlapping classes with \mathcal{S} .

The approach we propose in this paper is very general, suggesting that a spectrum of metric pretraining and label propagation algorithms can all work well in this framework. Below we introduce our method in details, and overview several metric learning and label propagation methods we used for our experiments.

3.1. Metric Pretraining

The source domain \mathcal{S} is used for pretraining a similarity metric between data pairs. Ideally, we desire the metric to capture the inherent structure in the target domain \mathcal{T} , so that transferring labels from \mathcal{T} is reliable and useful. For this to happen, we usually hold some prior knowledge about the source \mathcal{S} and the target \mathcal{T} . For example, the source domain is sampled from the same distribution as the target domain, but is completely unannotated, or the source domain is annotated with a different task but is closely related to the target. Formally, a similarity metric s_{ij} between data x_i and x_j can be defined as

$$s_{ij} = f(x_i, x_j), \quad (1)$$

where f is the similarity function to be learned. In this work, we use deep neural networks as a parametric model of this similarity function. The metric can be trained with either supervised or unsupervised methods, depending on

whether labels are given in the source domain \mathcal{S} . We briefly review the training algorithms as follows.

Unsupervised Metric Pretraining

Recently, there has been growing interest in unsupervised learning and self-supervised learning. Different algorithms are based on different data properties (e.g. color [44], context [4], motion [46]) and thus may vary in performance on the target task we may want to transfer. However, it is not our intent to give a comprehensive comparison over various methods and choose the best one. Instead, we show that general unsupervised transfer is beneficial for label propagation and leads to improved performance.

In this work, we utilize two unsupervised learning methods: instance discrimination [41] and colorization [44]. For instance discrimination, we treat each instance as a class, and maximize the probability of each example belonging to the class of itself,

$$P(i|x_i) = \frac{\exp(s_{ii})}{\sum_{j=1}^n \exp(s_{ij})}. \quad (2)$$

For colorization, the idea is to learn a mapping from grayscale images to colorful ones. Following the original paper [44], instead of predicting raw pixel colors, we quantize the color space into soft bins q , and use the cross-entropy loss on the soft bins,

$$\mathcal{L}_{\text{color}} = - \sum_{h,w} q_{h,w} \log(q_{h,w}), \quad (3)$$

where h, w are spatial indices. We follow previous work [5] for applying ResNet to colorization, where we use a base network to map inputs to features, and a head network of three convolutional layers to convert features to colors. Since colorization does not automatically output a metric, we use the Euclidean distance on the features from the base network to measure similarity.

Supervised Metric Pretraining

In some scenarios, we have access to a labeled dataset, such as PASCAL VOC and ImageNet, having commonalities with the target task. Traditional metric learning with supervision minimizes the intra-class distance and maximizes the inter-class distance of the labeled samples. For this purpose, many types of loss functions such as contrastive loss, triplet loss [13], and neighborhood analysis [12] have been proposed. In this work, we use neighborhood analysis [12] to learn our metric. Concretely, we maximize the likelihood of each example being supported by other examples belonging to the same category,

$$P(y_i|x_i) = \frac{\sum_{y_k=y_i} \exp(s_{ik})}{\sum_{j=1}^n \exp(s_{ij})}. \quad (4)$$

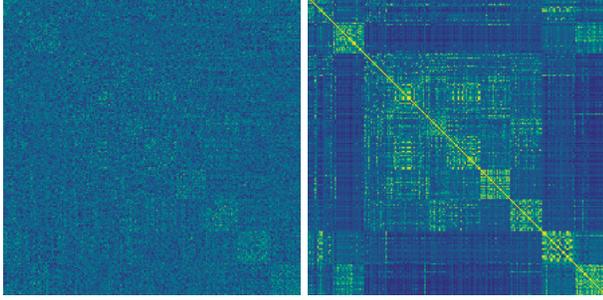


Figure 2: Left: raw similarity matrix. Right: similarity matrix by spectral embedding. Through spectral embedding, sparse similarities are propagated to distant areas to reveal *global* structure. Samples are sorted by their class id for better visualization.

3.2. Label Propagation

Given a target \mathcal{T} represented by a small number of labeled examples, and an unlabeled set \mathcal{U} , we propagate labels from \mathcal{T} to \mathcal{U} using the similarity function $f(\cdot)$ learned from \mathcal{S} . Suppose $\mathcal{T} = \{(x_1, y_1), (x_2, y_2), \dots, (x_{n_t}, y_{n_t})\}$, and $\mathcal{U} = \{x_{n_t+1}, x_{n_t+2}, \dots, x_{n_t+n_u}\}$, where n_t and n_u are the number of images in \mathcal{T} , \mathcal{U} respectively. Label y_i is represented as a vector with the ground-truth class element set to 1 and the others set to -1 . We consider two propagation algorithms.

Naive Nearest Neighbors

A straightforward propagation approach is to vote for the class of an unlabeled sample based on its similarity to each of the exemplars in the target set \mathcal{T} . For an unlabeled example $x_u \in \mathcal{U}$, we calculate its logits $z_{u,c}$ for every class c ,

$$z_{u,c} = \frac{1}{n_{t,c}} \sum_{i=1}^{n_t} I(y_{i,c} = 1) \cdot W_{i,u}, \quad (5)$$

where $I(\cdot)$ is the indicator function, $W_{i,u} = \exp(f(x_i, x_u))$ denotes the similarity between example i and u , and $n_{t,c}$ is the number of labeled images available for class c .

The nearest neighbor propagation method is essentially a one-step random walk where the similarity metric acts as the transition matrix and the indicator function acts as the initial distribution. The effectiveness of such one-step propagation depends heavily on the quality of the similarity metric.

In general, it is hard to learn such a metric well, especially when limited supervision is available, because of the visual diversity of images. Figure 2 (left) shows a typical similarity matrix computed from unsupervised features. Data points in the similarity matrix are sparsely connected, thus limiting the one-step label propagation approach.

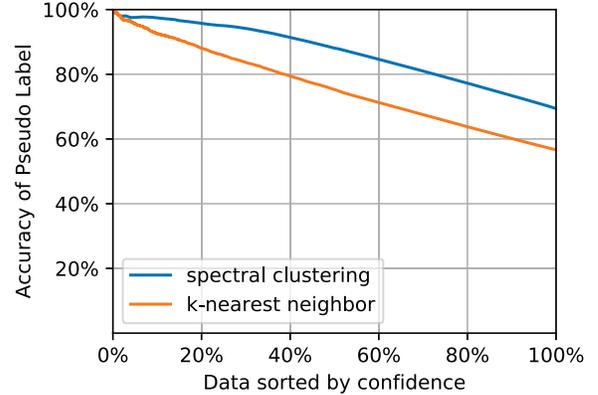


Figure 3: The accumulated accuracy of the pseudo labels on the validation data sorted by the confidence measure.

Constrained Spectral Clustering

Constrained spectral clustering [15, 7] may potentially relieve such a problem. Instead of propagating labels by one step as in the naive nearest neighbor approach, constrained spectral clustering propagates labels through multiple steps by taking advantage of structure within the unlabeled dataset. It computes a spectral embedding [33, 38] from the original similarity metric, which is then used as the new metric for label propagation. The spectral embedding is formulated as

$$W' = \sum_{j=2}^{\eta} \frac{1}{\lambda_j} e_j e_j^T, \quad (6)$$

where λ_j and e_j are the eigenvalues and eigenvectors of the normalized Laplacian in ascending order. The Laplacian matrix L_{sym} is derived from the original similarity metric as $L_{\text{sym}} = I - D^{-1/2} W D^{-1/2}$, with degree matrix $D = \text{diag}(d)$ and $d_i = \sum_j W_{ij}$. Parameter η is the total number of eigen components used.

Due to its globalized nature, spectral clustering is able to pass messages between distant areas, which is in contrast to the *local* behavior of the naive nearest neighbors approach. The embedded metric is usually densely connected and better aligned with object classes, as illustrated in Figure 2 (right). Using the same voting approach as in Eqn (5), labeled propagation can be more accurate than using the original raw similarity metric.

Constrained spectral clustering is also efficient. By following the common practice of using k -nearest neighbors to build the similarity graph [38], propagating labels to $50k$ images takes about 10 seconds on a regular GPU.

3.3. Confidence Weighted Supervised Training

Given the logits z_i , the pseudo label \hat{y}_i is estimated as

$$\hat{y}_i = \underset{c}{\text{argmax}} z_{i,c}. \quad (7)$$

Table 1: Ablation study of the mean average precision (mAP) of pseudo labels on CIFAR10.

Metric pretraining	Propagation method	50	100	250	500	1000	2000	4000	8000
Bootstrapping	Nearest neighbor	22.03	25.74	48.35	68.03	77.57	77.28	87.77	90.88
	Spectral	23.49	28.88	54.46	70.02	80.94	87.77	93.94	96.23
Colorization [44]	Nearest neighbor	57.32	67.61	75.48	79.34	80.70	82.14	83.66	84.79
	Spectral	60.85	67.34	76.31	80.04	81.78	81.89	82.93	82.03
Instance [41]	Nearest neighbor	54.82	62.99	77.08	84.90	88.68	91.34	92.72	93.67
	Spectral	72.59	79.21	86.64	90.01	91.04	91.57	91.77	91.94

Table 2: Ablation study of semi-supervised performance on CIFAR10.

Metric pretraining	Propagation method	50	100	250	500	1000	2000	4000	8000
No	No	20.95	25.35	41.63	54.06	65.08	73.22	81.44	86.23
Bootstrapping	Nearest neighbor	21.79	25.37	42.70	54.14	68.08	75.17	83.30	87.68
	Spectral	22.78	27.95	47.28	60.73	72.60	78.20	85.10	88.26
Colorization [44]	No	49.57	55.41	64.65	68.81	73.40	77.93	82.17	86.25
	Nearest neighbor	49.96	52.69	65.63	65.88	70.88	76.36	80.16	84.64
	Spectral	53.47	55.08	68.40	71.15	72.38	76.50	80.31	84.03
	No	35.27	37.87	62.46	71.04	75.96	80.12	83.90	87.82
Instance [41]	Nearest neighbor	46.68	54.45	66.93	74.16	79.17	82.24	84.56	87.92
	Spectral	56.34	63.53	71.26	74.77	79.38	82.34	84.52	87.48

With the estimated pseudo labels on the unlabeled data, we have considerably more data for training a classifier. However, the pseudo labels may not be accurate, and directly using these labels may lead to degraded performance. For example, not all the data in the unlabeled set are related to the target problem. Here, we devise a simple weighting mechanism to compensate for inaccurate labels.

Given the logits z_i produced by the label propagation algorithm, we first normalize it into a probabilistic distribution,

$$\bar{z}_{i,c} = \frac{\exp(z_{i,c}/\tau)}{\sum_j \exp(z_{i,j}/\tau)}, \quad (8)$$

where c indexes the dimension of categories, and the temperature τ controls the sharpness of the distribution. We then define the confidence measure α_i of the pseudo label as the difference between the maximum response and the second largest response,

$$\alpha_i = \max_j \bar{z}_{i,j} - \max_{c \neq \arg\max_j \bar{z}_{i,j}} \bar{z}_{i,c}. \quad (9)$$

A high value of α_i indicates a confident estimate of the pseudo label, and a low value of α_i indicates an ambiguous estimate. In Figure 3, we measure the accumulated accuracy of pseudo labels on validation data sorted by this confidence. It can be seen that our confidence measure gives a good indication of the quality of pseudo labels.

Our final training criterion is given by

$$\mathcal{L} = -\frac{1}{N} \sum_i \alpha_i \cdot \log p_{\hat{y}_i} \quad (10)$$

where \hat{y}_i is the pseudo label for example i , and $p(\cdot)$ is the softmax probability output of the classification network.

In practice, since some pseudo labels have relatively low confidence, e.g. $\alpha < 0.01$, and thus contribute negligibly to the overall learning criterion, we may safely discard those examples to speed up learning.

4. Experiments

Through experiments, we show that, with unlabeled data, metric propagation is able to effectively label lots of data when little labeled data is given. We verify our approach on semi-supervised learning, where an unsupervised metric is transferred, and on transfer learning, where supervised metrics generalize across different data distributions, and on few-shot recognition, where the metric can generalize across open-set object categories. While studying few-shot recognition, we leverage an extra unlabeled data for label propagation, which is also known as semi-supervised few-shot recognition [30].

Our approach has two major hyper-parameters: the number of the eigenvectors η for spectral clustering and the temperature σ controlling the confidence distribution. Different parameter settings may slightly change the performance. We use $\eta = 200$ and $\sigma = 40$ across the experiments. A detailed analysis is provided in the supplementary materials.

4.1. Semi-Supervised Learning

We follow a recent evaluation paper [28], which gives a comprehensive benchmark for state-of-the-art semi-

Table 3: Scalability to large network architectures on CIFAR10.

Methods	Network architectures	50	100	250	500	1000	2000	4000	8000
Mean Teacher	WideResNet-28-2	29.66	36.62	45.49	57.19	65.07	79.26	84.38	87.55
Ours		56.34	63.53	71.26	74.77	79.38	82.34	84.52	87.48
Mean Teacher	WideResNet-28-10	27.35	38.83	49.44	59.45	70.03	82.62	86.71	89.38
Ours		73.13	75.87	80.30	81.76	84.97	86.82	88.70	91.01

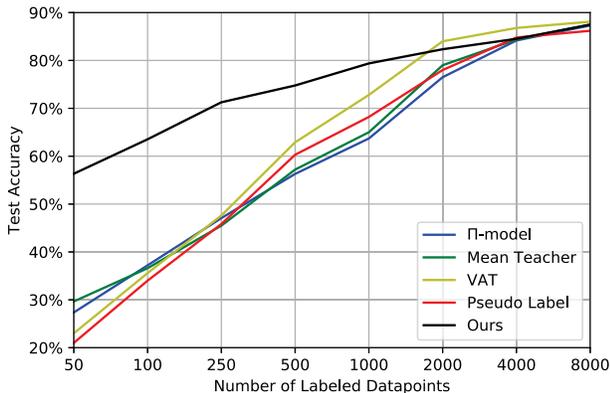


Figure 4: Comparisons to the state-of-the-art on CIFAR10.

supervised learning approaches. A majority of our ablation studies are conducted on CIFAR10 [18], while we also test our method on ImageNet. On CIFAR10, we use the same Wide-ResNet [43] architecture with 28 layers and a width factor of 2. We report performance varying the number of labeled examples from 50 to 8,000 of 50,000 examples.

For training our model, we pretrain the metric using the unlabeled split, and propagate labels to the same unlabeled set. This means $\mathcal{S} = \mathcal{U}$ in our framework. We use SGD for optimization with an initial learning rate of 0.01 and a cosine decay schedule. We fix the total number of optimization iterations to 200K as opposed to fixing optimization epochs, because it gives more consistent comparisons when the number of labeled data varies.

Study of different pretrained metrics.

Our label propagation algorithm needs a pretrained similarity metric to guide it. The pretrained metric can be learned by supervised methods using limited labeled data, or by unsupervised methods using large-scale unlabeled data. Here, we consider three metric pretraining methods:

1. supervised bootstrapping on limited labeled data.
2. self-supervised learning by image colorization [44].
3. unsupervised learning by instance discrimination [41].

We train the models using the optimal parameters for each pretraining method. Then we use cosine similarity in the feature space for propagating labels to the unlabeled data.

In Table 1, we evaluate the quality of pseudo labels as the mean average precision (mAP) sorted by the confidence as in Figure 3. Table 2 lists the final semi-supervised recogni-

tion accuracy. We can see that both unsupervised methods generalize much better than the supervised bootstrapping method most of the time, until the labeled set is relatively large with 4000 labels. This confirms our claim that unsupervised transfer is the key for label propagation. For the unsupervised methods, non-parametric metric learning performs better than colorization, probably because it explicitly learns a similarity metric. We also include the result of the naive baseline which trains from scratch using limited labeled data without label propagation.

Study of different label propagation schemes.

Given the pretrained metrics, there are various ways to transfer the metrics. We consider three possible solutions:

1. no propagation, only transfer network weights.
2. nearest neighbor metric transfer.
3. spectral metric transfer.

The first baseline is a common practice, which basically transfers the network weights and then finetunes on the labeled data. The second is much weaker than the third because it only considers one-hop distances, without taking into account the similarities between unlabeled pairs.

The results are summarized in Table 1 and Table 2. Compared to the state-of-the-art performance in Table 4, even a simple finetuning approach outperforms the state-of-the-arts when the labeled data is small. For example, by finetuning from instance discrimination, we achieve 62.46% with 250 labeled data, significantly outperforming the state-of-the-art result of 47.07%. This suggests that unsupervised pretraining generally improves semi-supervised learning.

When unlabeled data is used for label propagation, metric transfer can be much stronger than just weight transfer, improving the performance to 71.26% with 250 labeled data. It is also evident that the spectral clustering method performs better than weighted nearest neighbors because of its globalization behavior.

Scalability to large network architectures.

In contrast to prior methods which face over-fitting issues, our approach can easily scale to larger network architectures. Here, we keep all the learning parameters unchanged, and experiment with a wider version of Wide-ResNet-28 with a width factor of 10. We consider a state-of-the-art method mean-teacher [36] for comparison. In Table 3, mean-teacher only shows a limited improvement of about 2 – 3%. Our method enjoys consistently significant

Table 4: Ours is complementary to all prior state-of-the-art methods on CIFAR10.

Num Labeled	250	4000
Ours	71.26	84.52
Pi Model [20]	47.07	84.17
+ Ours	74.90	85.32
Mean Teacher [36]	45.49	84.38
+ Ours	74.54	85.45
VAT [27]	44.83	86.79
+ Ours	78.34	86.93
VAT+EM [27]	46.29	86.96
+ Ours	78.63	87.20

Table 5: Semi-supervised classification results on the ImageNet dataset.

Num Labeled	1%	2%	4%
Scratch	22.4	40.2	58.2
Finetune	39.2	52.8	65.2
Ours	58.6	66.3	72.4

gains from a larger network on all the testing scenarios. It achieves an unprecedented 73.13% accuracy using only 50 labels with Wide-ResNet-28-10.

Comparison to the state-of-the-art on CIFAR10.

We compare our approach to state of the art methods in Figure 4. Ours is particularly stronger when the labeled set is small, but this advantage diminishes as the labeled set grows. However, as most prior approaches focus on self-ensembling, ours is orthogonal to them. We examine the complementarity of our method by combining it with each of the prior approaches. To do so, we generate our most confident $10K$ pseudo labels (about 20% of the full data), and use it as ground-truth for the other algorithms. For fair comparisons, we run public code¹ with our generated pseudo labels. In Table 4, combining our approach leads to improved performance for all of the methods.

Comparison to the state-of-the-art on ImageNet.

We notice that few literature report semi-supervised classification performance on ImageNet consistently. In this paper, we consider finetuning from an unsupervised model trained with instance discrimination [41] as our baseline. We vary the number of labeled examples from 1% to 4% of the entire ImageNet. We use ResNet-50 to pretrain the unsupervised model, and split the dataset into 10 chunks for spectral clustering to speed up the computation. In Table 5, finetuning from unsupervised model significantly improves upon training from scratch, and our label propagation approach outperforms the finetuning approach. Notable, ours is 18% better when 1% labeled data is available.

¹<https://github.com/brain-research/realistic-ssl-evaluation>

4.2. Transfer Learning

We also examine whether the proposed metric transfer can work across different data distributions. We pretrain the metric on the source \mathcal{S} ImageNet, and transfer it to the unlabeled \mathcal{U} CIFAR10. For this, we study supervised and unsupervised pretraining for transfer learning.

Transferring from labeled ImageNet. We resize ImageNet images to a resolution of 32×32 and pretrain the metric on them by supervised learning. We keep the network architecture WideResNet-28-2 for meaningful comparison with the semi-supervised settings in Sec 4.1. This obtains an accuracy of 42% on the ImageNet validation set. Then we transfer the metric to CIFAR10. This transfer is conducted by network finetuning and by metric propagation. In Table 6, we can see that simple network finetuning can reach the best results obtained in the semi-supervised settings of the previous subsection. By using label propagation with spectral clustering, we can observe a large improvement, yielding 86.07% accuracy with just 250 labeled images. This illustrates the generality of our metric transfer approach, where supervised transfer can also take advantage of unlabeled data to improve generalization.

Transferring from unlabeled ImageNet. Instead of supervised training which encodes prior knowledge about object categories, we treat ImageNet images as unlabeled and repeat the previous experiment. Different from the earlier unsupervised experiments, this setting involves substantially more unlabeled data, which could potentially lead to a better unsupervised metric. However, our results suggest otherwise. When propagating to CIFAR10, the unsupervised metric learned from ImageNet is inferior to the metric learned from CIFAR10. This is possibly due to the data distribution gap between CIFAR10 and ImageNet. Nevertheless, our unsupervised transfer from ImageNet still surpasses the state-of-the-art in the semi-supervised setting when labeled samples are limited.

4.3. Few-Shot Recognition

Few-shot recognition targets a more challenging scenario, the generalization across object categories (*a.k.a.* open-set recognition). Originally, the problem is defined with numerous labeled examples in a source dataset, and few examples in the target categories. Recent works [30, 9] also explore the scenario where extra unlabeled data is available for this problem. This fits into our framework for studying label propagation via metric transfer.

We follow the protocols in [30] for conducting the experiments, because it introduces distractor categories in the unlabeled set. The experiments are evaluated on the mini-ImageNet dataset, consisting of a total of 100 categories, with 64 for training, 16 for validation and 20 for testing. Images in each category are split into 40% as labeled, and

Table 6: Transfer learning from ImageNet to CIFAR10.

Metric pretraining	Transfer method	50	100	250	500	1000	2000	4000	8000
Unsupervised	Network finetuning	28.92	34.56	57.14	67.54	76.20	80.92	85.01	88.74
	Spectral	44.30	46.51	61.29	68.31	72.61	77.86	84.00	88.19
Supervised	Network finetuning	54.95	61.88	73.01	78.43	84.52	88.79	91.44	93.05
	Spectral	77.71	85.34	86.07	86.91	88.27	89.93	91.22	93.49

Table 7: Few-shot recognition on Mini-ImageNet dataset.

Method	Fintune	Unlabel data	5-way Classification	
			1-shot	5-shot
NN baseline [37]	No	No	41.1±0.7	51.0±0.7
MAML [8]	Yes	No	48.7±0.7	63.2±0.9
Meta-SGD [23]	No	No	50.5±1.9	64.0±0.9
Matching net [37]	Yes	No	46.6±0.8	60.0±0.7
Prototypical [35]	No	No	49.4±0.8	68.2±0.7
SNCA [40]	No	No	50.3±0.7	64.1±0.8
Soft k-means [30]	Yes	Yes	50.4±0.3	64.4±0.2
Our supervised	Yes	Yes	56.1±0.6	70.7±0.5
Our unsupervised	Yes	Yes	50.8±0.6	66.0±0.5

60% as unlabeled. Training uses only the labeled split in the training categories. During evaluation, a testing episode is constructed by sampling few-shot labeled observations from the labeled split in the testing categories, and all of the unlabeled images in all the testing categories. A testing episode requires the model to find useful information in the unlabeled set to aid recognition from the few-shot observations. Unlike [30], which includes five distractor categories in the unlabeled set, we consider all 20 categories in the testing set, which better reflects practical scenarios. We test 300 episodes and report the results.

We follow prior work [37] by using a shallow architecture with four convolutional layers and a final fully connected layer. Each convolutional layer has 64 channels, interleaved with ReLU, subsampling and a batch normalization layer. Images are resized to 84×84 to train the model. We use the spectral embedding approach for label propagation. During online training, we use an initial learning rate of 0.01 with a total of 30 epochs and decrease the learning rate to be 5 times smaller after 20 epochs.

Transfer from supervised models. We use a recent supervised metric learning approach SNCA [40] as the baseline. After label propagation and finetuning on the new data, our supervised propagation obtains a significant boost of 6% over SNCA. Prior work [30] improves upon its baselines, but fails to make further improvement because of limited training data. In Figure 5, we visualize the top retrievals from the unlabeled set in the one-shot scenario. These retrievals not only belong to the same class as the ground



Figure 5: Visualizations of top ranked retrievals from the unlabeled set given one-shot observations.

truth, but their diversity facilitates a strong classifier.

Transfer from unsupervised models. We also investigate pretraining the metric without labels, using instance discrimination [41] for learning the metric. Surprisingly, in Table 7, our unsupervised propagation obtains better performance than the offline metric learning approach with annotations [40], by 0.5% in 1-shot recognition and 2% for 5-shot. This suggests that leveraging unlabeled data in the target problem may possibly be more beneficial than using labeled samples in the source domain.

5. Discussions

- The effectiveness of label propagation depends heavily on the learned metric, so advances in metric learning should lead to improved results. Since the prevalent pretraining methods in deep learning use softmax classification, we hope to draw more attention to pretraining networks with metric learning.
- Currently, we study metric pretraining and label propagation separately. It may be beneficial to formulate them jointly in an end-to-end framework.
- Our algorithm takes advantage of the unlabeled dataset \mathcal{U} to create more training data. The overall performance is affected by the relevance of image content in the unlabeled set \mathcal{U} to that of the target \mathcal{T} , as this impacts the ability to effectively propagate labels.

References

- [1] S. Carey and E. Bartlett. Acquiring a single new word. 1978. 1, 2
- [2] K. Clark, T. Luong, and Q. V. Le. Cross-view training for semi-supervised learning. In *ICLR*, 2019. 2
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*. Ieee, 2009. 2
- [4] C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 3
- [5] C. Doersch and A. Zisserman. Multi-task self-supervised visual learning. In *ICCV*, 2017. 3
- [6] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 2
- [7] R. Fergus, Y. Weiss, and A. Torralba. Semi-supervised learning in gigantic image collections. In *NIPS*, 2009. 2, 4
- [8] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017. 1, 2, 8
- [9] V. Garcia and J. Bruna. Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043*, 2017. 2, 7
- [10] X. Gastaldi. Shake-shake regularization. *arXiv preprint arXiv:1705.07485*, 2017. 1
- [11] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 2
- [12] J. Goldberger, G. E. Hinton, S. T. Roweis, and R. R. Salakhutdinov. Neighbourhood components analysis. In *Advances in neural information processing systems*, 2005. 3
- [13] E. Hoffer and N. Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*. Springer, 2015. 3
- [14] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017. 3
- [15] H. Hu, J. Feng, C. Yu, and J. Zhou. Multi-class constrained normalized cut with hard, soft, unary and pairwise priors and its applications to object segmentation. *TIP*, 2013. 4
- [16] T. Joachims. Transductive learning via spectral graph partitioning. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 290–297, 2003. 2
- [17] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *NIPS*, 2014. 2
- [18] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009. 6
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 2012. 2
- [20] S. Laine and T. Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. 1, 2, 7
- [21] B. Lake, R. Salakhutdinov, J. Gross, and J. Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2011. 1
- [22] D.-H. Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, 2013. 2
- [23] Z. Li, F. Zhou, F. Chen, and H. Li. Meta-sgd: Learning to learn quickly for few shot learning. *arXiv preprint arXiv:1707.09835*, 2017. 2, 8
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*. Springer, 2014. 2
- [25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 2
- [26] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. A simple neural attentive meta-learner. 2018. 2
- [27] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. 2017. 1, 2, 7
- [28] A. Oliver, A. Odena, C. Raffel, E. D. Cubuk, and I. J. Goodfellow. Realistic evaluation of semi-supervised learning algorithms. 2018. 5
- [29] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-supervised learning with ladder networks. In *NIPS*, 2015. 2
- [30] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018. 2, 5, 7, 8
- [31] C. Rosenberg, M. Hebert, and H. Schneiderman. Semi-supervised self-training of object detection models. In *WACV/MOTION*, 2005. 2
- [32] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 2
- [33] J. Shi and J. Malik. Normalized cuts and image segmentation. *TPAMI*, 2000. 4
- [34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [35] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017. 2, 8
- [36] A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, 2017. 2, 6, 7
- [37] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot learning. In *NIPS*, 2016. 2, 8
- [38] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 2007. 4
- [39] J. Weston, F. Ratle, H. Mobahi, and R. Collobert. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*. Springer, 2012. 2
- [40] Z. Wu, A. A. Efros, and S. X. Yu. Improving generalization via scalable neighborhood component analysis. *arXiv preprint arXiv:1808.04699*, 2018. 2, 8
- [41] Z. Wu, Y. Xiong, X. Y. Stella, and D. Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 3, 5, 6, 7, 8

- [42] Y. Xu, S. J. Pan, H. Xiong, Q. Wu, R. Luo, H. Min, and H. Song. A unified framework for metric transfer learning. *IEEE Trans. Knowl. Data Eng.*, 2017. [3](#)
- [43] S. Zagoruyko and N. Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [6](#)
- [44] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *ECCV*. Springer, 2016. [3](#), [5](#), [6](#)
- [45] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In *Advances in neural information processing systems*, pages 321–328, 2004. [2](#)
- [46] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. [3](#)