

ReMP: Rectified Metric Propagation for Few-Shot Learning

Yang Zhao¹ Chunyuan Li² Ping Yu¹ Changyou Chen¹
¹University at Buffalo ²Microsoft Research

Abstract

Few-shot learning features the capability of generalizing from a few examples. In this paper, we first identify that a discriminative feature space, namely a rectified metric space, that is learned to maintain the metric consistency from training to testing, is an essential component to the success of metric-based few-shot learning. Numerous analyses indicate that a simple modification of the objective can yield substantial performance gains. The resulting approach, called rectified metric propagation (ReMP), further optimizes an attentive prototype propagation network, and applies a repulsive force to make confident predictions. Extensive experiments demonstrate that the proposed ReMP is effective and efficient, and outperforms the state of the arts on various standard few-shot learning datasets.

1. Introduction

Learning and generalizing from a few examples to effectively make predictions in new domains/tasks is a common scenario in real applications, as the supervised information is often hard to acquire due to some practical considerations, such as high labeling cost, privacy, safety or ethic issues. Naively applying the traditional supervised learning techniques tend to overfit in this scenario [5, 31]. This stimulates the emergence of *few-shot learning (FSL)*, which mimics the human ability of recognizing new data after observing a few instances.

In such a few-shot regime, it is desired to leverage knowledge (e.g. parameters and embeddings) from one task to another. Recent research has shown promising results in exploiting *episodic learning* on this task. In essence, the episodic learning aims to mimic the real testing scenarios where the model generalizes information from a few labeled instances (called *support set*) to predict labels of unlabeled instances (called *query set*) in each episode. Research efforts in solving FSL can be broadly categorized into two types: (i) *Meta-learning* [5, 19] can quickly adapt parameters to a new task after fine-tuning; (ii) *Embedding and metric-based learning* [28, 31] directly performs nearest-neighbor classification given a similarity metric on an em-

bedding space. Our work is falling into this category.

The fundamental task for the metric learning methods is to design an appropriate metric space [20] that can satisfy the following three properties simultaneously: (i) a valid measure (or metric) such as Euclidean distance and Cosine similarity [27, 20, 10]; (ii) a cost function that can fully exploit input information to update the network; (iii) the metric consistency from training to testing, as well as from pretraining to fine-tuning if a pretraining is being applied. A contradictory example that arouses an inconsistency issue is seen in most works as well as recent advanced metric-based FSL frameworks [2, 34]. They adopt an inefficient two-phase training procedure: a *parametric* pre-training with a linear classifier and a *non-parametric* fine-tuning with the nearest-neighbor prediction. In FSL, the pretraining essentially leads to a different metric space from that learnt by fine-tuning. However, the testing is only accessible to use the *non-parametric* nearest neighbor classifier (see more details in the Background section). In the end, the inconsistency is reflected on the very limited performance gain [34].

While the first two properties have largely been individually explored [27, 21, 14], the last one is less studied (if not ignored at all). We observe that this missing ingredient can lead up to 5% absolute recognition accuracy drop when the inconsistent metric exists. In the meantime, however, this kind of parametric linear classifier is essential to explore a discriminative metric space and to encourage faster convergence. So the natural question is: *is there a strategy that can maintain the metric consistency as well as fully utilize the training data information to learn a meaningful metric space?*

To this end, we argue that a suitable interaction of above three properties can yield significant improvements in terms of performance and stability. We study the FSL problem from a unified perspective and propose a rectified metric propagation (ReMP) framework that can progressively make predictions in a discriminative feature space, namely the rectified metric space. The rectified metric is induced from maintaining consistency from training to testing. Specifically, we firstly design a **cooperative learning** objective that considers training-testing consistency both within a global parametric classifier (called **global match-**

ing) and local nearest neighbor prediction among instances (called **local matching**). The framework proceeds in a feed-forward pass without either pre-training or fine-tuning. To further enhance the nearest neighbor prediction confidence, we propose an attention-based contextualized label embedding method to iteratively rectify the prototypes, such that data importance is taken into account when calculating class prototypes. The resulting approach is simpler and more efficient than related recent approaches. The contributions of this paper are threefold:

- We identify the metric inconsistency issue between training and testing in FSL, a long-standing issue of almost all metric-based FSL approaches. To alleviate this problem, we propose a **cooperative label-aligned training** scheme, where the unabridged metric space can be inherited to the testing phase.
- We describe a **contextualized label embedding module containing attentive prototype propagation layers** to take into account data importance to make more confident predictions.
- With the above two novelties, **new state-of-the-art (SoTA)** performances on three standard FSL classification datasets, the *miniImageNet*, *tieredImageNet* and *CIFAR-FS* are achieved.

The rest of the paper is organized as follows. We first summarize the related work and how our proposal is differentiated from those. Then, we describe our contribution in detail and present extensive experiments to demonstrate our justifications. Finally, we conclude our paper and highlight future research directions.

2. Related work

In this section, we introduce related works and make distinction between our approach and related FSL methods. These methods broadly fall into three categories, meta-learning, embedding and metric learning, and transductive learning based approaches.

Meta-learning Meta-learning [29, 5], or learning-to-learn, is a framework that is capable of learning a task-specific meta-network. After observing the support set of a new task, the meta-network can quickly adapt to be evaluated on the query set of that new task. [23] proposes to finetune an LSTM-based optimizer besides the meta-learner to maximize the performance. [18] learns to change its inductive bias via fast parameterization. These works include MAML [5], Reptile [19], Meta-SGD [13], Bayesian-MAML [36], Implicit-MAML [22] and LEO [25]. However, the aforementioned approaches often suffer from overfitting and sensitive to architectures, making the performance after fine-tuning on a new task limited. Although

some recent work [1] proposes some modifications, this line of approaches is still unsatisfactory for solving FSL. By contrast, our proposed model can make predictions in a direct and efficient feed-forward manner without the necessity of fine-tuning. Recently, perhaps the most popular approach is model agnostic meta learning (MAML) [5]. MAML learns a meta-parameter-initialization of a meta-network such that it can solve a new task with only a few gradient descent steps.

Embedding and metric learning This class of methods for FSL has drawn more and more attention. The main goal is to learn a transferable and discriminative feature space that preserves the neighborhood structure, *e.g.*, the Matching network [31] and Prototypical network [27]. This means objects belonging to the same class should be consistently closer to each other in the feature space measured by some similarity measures, *e.g.*, the Euclidean distance and Cosine similarity, and vice versa. Two typical works Relational network [28] and TADAM [20] made further improvements. Another line of work proposes to learn global prototypes [6, 21]. Label embedding has been proposed as anchor points to improve text classification in NLP [32]. Different from these approaches, the proposed cooperative learning objective can fully *exploit* the manifold information of both support set and query set in training, and progressively *explore* a discriminative feature space to refine the prototype in testing.

Transductive learning The transductive learning is first introduced in TPN [15] that exploits the manifold structure in the data by learning a graph construction to propagate labels from the support set to the unlabeled query set. It alleviates the small-data problem in FSL and has been shown to outperform the inductive learning (inaccessible to a query set) counterparts. [8] further enhances the graph construction module. FEAT [34], CAN [7], Meta-Fun [33] and EP-Net [24] optimize the embedding manifold to better generalize to unseen classes. DFMN-MCT [10] chooses to propagate prototypes along with assigning confidence scores to all unlabeled queries, leading to SoTA results. However, the pretraining phase of FEAT and the pixel-wise dense classifier of DFMN-MCT lead to an inconsistent pipeline from training to testing. Our work is distinguished from existing works by that we propose to rectify the metric space, which aims to close the gap between training and testing. Furthermore, our approach can achieve new SoTA results with a newly proposed repulsive attention strategy by stacking up more attention layers.

3. Background

3.1. Problem definition

Similar to the supervised-learning setting, a dataset is typically divided into three parts in FSL: a training dataset $\mathcal{D}^{\text{train}}$, a testing dataset $\mathcal{D}^{\text{test}}$ and a validation dataset \mathcal{D}^{val} . The main distinction of FSL is that the three sets have disjoint label spaces. The episodic classification [5, 17, 12, 31, 27, 26] is a common and effective approach to FSL, where the training dataset is exploited to simulate the few-shot learning setting via episode-based training.

Specifically, to characterize generalization, in the training process, one typically randomly samples N classes instances in each episode, containing a *support* set $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N \times K}$ (K samples per class) and a *query* set $\mathcal{Q} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i)\}_{i=1}^{N \times M}$ (M samples per class). *The support set is used to calculate prototypes and then make predictions on the query set, which in turn is used to update the model.* This setting is often abbreviated as N -way K -shot FSL. The number of K is generally very small, e.g. 1 and 5. The goal of FSL is to learn a model to exploit such a low resource data set \mathcal{S} to predict labels for queries in \mathcal{Q} .

3.2. Revisiting the classifiers in the feature space

Current approaches to FSL resort to a two-step procedure, represented as a function decomposition $f \circ g$, where (i) $f : \mathbf{x} \rightarrow \mathbf{z}$ is a feature extractor, which maps an input $\mathbf{x} \in \mathbb{R}^D$ to a feature vector $\mathbf{z} \in \mathbb{R}^d$ in a transferable and discriminative space. It is worth noting that d is not necessarily much smaller than D and it is mostly related to the complexity of f ; (ii) g is a classifier, mapping the feature \mathbf{z} into a vector of logits $g(\mathbf{z})$ and produce a probability distribution $p(\cdot|\mathbf{x})$ over all categories activated by the **Softmax**. Since f and g are sometimes not trained end-to-end, it is crucial to find a good association between a classifier g and the feature representation \mathbf{z} . Typically, there are two general paradigms to design the classifier g . Unfortunately, both paradigms endow a limitation. Our proposed approach would inherit the best of both worlds, as described in the next section.

- **Parametric Methods.** These approaches usually learn a linear classifier $g(\mathbf{z}) = \mathbf{w}^T \mathbf{z} + \mathbf{b}$, $\mathbf{w} \in \mathbb{R}^{d \times N}$, $\mathbf{b} \in \mathbb{R}^N$, thus inherits fast convergence. However, since in FSL, the label space in the testing is disjoint with that in the training phase. This implies that $\{\mathbf{w}, \mathbf{b}\}$ trained on the training data is not ready to be adopted for testing directly. Two solutions can be designed to mitigate this problem: (i) Fine-tuning to adapt the parameters [5, 13, 25] to test sets; (ii) Performing the nearest neighbor search in an embedding space [21, 10, 2] for testing. We only consider the second approach in the following as the fine-tuning approaches are often unsatisfactory and time-consuming. However, the second

solution would lead to **inconsistency between training and testing**. This is because $g(\mathbf{z})$ is computed based on a linear mapping in training; whereas the nearest neighbor is used in testing. In fact, the model is unaware of an appropriate metric in training. Such inconsistency could harm model performance.

- **Non-parametric Methods.** In training and testing, one consistently applies a metric-based similarity measure $g(\mathbf{z}) = \kappa(\mathbf{z}, \mathbf{c}_n)$, where \mathbf{c}_n is the class prototype (prototype), usually defined as the intra-class mean over support embeddings; $\kappa(\cdot, \cdot)$ is often defined by Cosine similarity [20], negative Euclidean distance [27, 20, 10] and kernel-based functions [33, 15]. However, it has been found that these approaches **are hard to explore and learn a discriminative metric space**¹ [27, 20, 28], thus limiting its potential to reach a good convergence point (see Tab. 3).

4. The Proposed Method

Our goal is to overcome the two aforementioned limitations, and design a model to close the gap between training and testing, and at the same time, to learn a discriminative metric space. To this end, we first define a cooperatively label-aligned learning objective. A contextualized label embedding module with self-attention is then further formulated to iteratively propagate labels on $\mathcal{S} \cup \mathcal{Q}$, and to rectify prototypes such that more confident predictions can be made. The overall framework is illustrated in Figure 1.

4.1. Label-Aligned Cooperative Learning

We firstly define a label-aligned learning objective to optimize the feature extractor f and the classifier g in training. It completes the cooperative learning on two levels, *global matching* on the global prototype of $\mathcal{D}^{\text{train}}$ and *local matching* on the local prototype of the episodic $\mathcal{S} \cup \mathcal{Q}$.

Notations To facilitate the reading, notations are clarified. The number of classes in $\mathcal{D}^{\text{train}}$ is N^{train} . The i -th pair of \mathcal{Q} is $(\tilde{\mathbf{x}}_i, \tilde{y}_i)$. For the N -way classification protocol on the instance $\tilde{\mathbf{x}}_i$, $(\tilde{y}_i^g = \tilde{y}_i, \hat{y}_i^g)$ are the true label and predicted label ranging from 1 to N^{train} on the global level; whereas $(\tilde{y}_i^l, \hat{y}_i^l)$ denotes the true label and predicted label ranging from 1 to N on the local level. Note \tilde{y}_i^g and \tilde{y}_i^l are different random variables corresponding to the same data instance $\tilde{\mathbf{x}}_i$ in a FSL setting.

Global matching To encourage fast convergence, we first define a global matching loss on $\mathcal{D}^{\text{train}}$. The global matching

¹A metric space means an embedding space, where one can compute similarity given a metric. It is different from the definition in mathematics, e.g., the Cosine similarity, which can be negative.

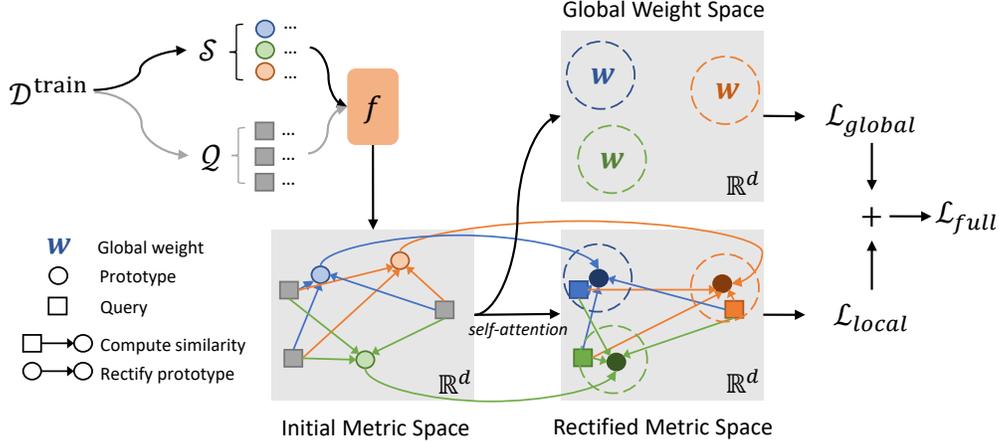


Figure 1. Illustration of the proposed label-aligned learning. The training considers matching in both a parametric global weight space and a non-parametric rectified metric space. In testing, predictions are only made through local matching in a feed-forward manner, where rectified prototypes help to generate a better decision boundary based on self-attention.

is a parametric model, aiming to make predictions matched with the global-level label \hat{y}_i^g . To this end, we first introduce a metric $\kappa(\cdot, \cdot)$ to measure the similarity between two vectors. The metric can be flexible. We adopt the Cosine similarity and Euclidean distance in our experiments. The goal is to match the feature of a data instance with a set of global learnable weight vectors $\{\mathbf{w}_k \in \mathbb{R}^d\}_{k=1}^{N^{\text{train}}}$, each representing a global class prototype, as illustrated in Figure 1. To this end, the global likelihood on $\tilde{\mathbf{x}}_i$ is:

$$p(\hat{y}_i^g = k | \tilde{\mathbf{x}}_i) = \frac{\exp(\kappa(f(\tilde{\mathbf{x}}_i), \mathbf{w}_k))}{\sum_{k=1}^{N^{\text{train}}} \exp(\kappa(f(\tilde{\mathbf{x}}_i), \mathbf{w}_k))}. \quad (1)$$

The global loss is then defined as the standard cross-entropy loss:

$$\mathcal{L}_{\text{global}} = - \sum_{i=1}^{NM} \sum_{k=1}^{N^{\text{train}}} \mathbb{I}(\tilde{y}_i^g = k) \log p(\hat{y}_i^g = k | \tilde{\mathbf{x}}_i), \quad (2)$$

where $\mathbb{I}(\cdot)$ is an indicator function. In each episode, all global weights are being optimized with interactive information sharing, which could significantly encourage the model's convergence. Note that, the proposed approach is different from other types of inconsistent global matching, e.g., the linear classifier in [7] and the dense matching in [10], where the embedding has a different dimension from the global weight.

Local matching Note the above global loss $\mathcal{L}_{\text{global}}$ is parametric, which still could cause training-testing inconsistency. To close the gap, we further incorporate a local loss on episodic $\mathcal{S} \cup \mathcal{Q}$ with N classes. Instead of using a parametric model, we define a non-parametric model to directly match predictions with local prototypes. We first

define the local likelihood on $\tilde{\mathbf{x}}_i$ as:

$$p(\hat{y}_i^l = n | \tilde{\mathbf{x}}_i) = \frac{\exp(\kappa(f(\tilde{\mathbf{x}}_i), \mathbf{c}_n))}{\sum_{n=1}^N \exp(\kappa(f(\tilde{\mathbf{x}}_i), \mathbf{c}_n))} \quad (3)$$

where \mathbf{c}_n represents the prototype for local class n , calculated as

$$\mathbf{c}_n = \frac{1}{K} \sum_{i=1}^{NK} \mathbb{I}(y_i = n) f(\mathbf{x}_i), \text{ for } (\mathbf{x}_i, y_i) \in \mathcal{S}. \quad (4)$$

We then define the local loss as:

$$\mathcal{L}_{\text{local}} = - \sum_{i=1}^{NM} \sum_{n=1}^N \mathbb{I}(\tilde{y}_i^l = n) \log p(\hat{y}_i^l = n | \tilde{\mathbf{x}}_i). \quad (5)$$

It is clear that $\mathcal{L}_{\text{local}}$ can leverage episodic \mathcal{S} and \mathcal{Q} for similarity comparison both in training and testing, without consulting the global parametric classifier.

Label-aligned objective To align labels in the local matching with that in the global metric space, we propose a full objective that combines Eq. (2) and Eq. (5):

$$\mathcal{L}_{\text{full}} = \mathcal{L}_{\text{global}} + \alpha \mathcal{L}_{\text{local}}, \quad (6)$$

where α is a balance hyper-parameter. At testing, predictions will only be made following Eq. (3) such that the local matching is consistent for both training and testing. We claim that:

Remark 1 A good metric space is all we need in metric-based FSL. On the premise that the embedding \mathbf{z} and the global weight \mathbf{w} are jointly learned via Eq. (6), global matching helps accelerate convergence and explore a space that is endowed with a strong transferability; meanwhile local matching preserves the metric inheritance from training to testing.

Besides, the metric in Equation (2) and Equation (5) don't need to be the same since the global weight vectors are learned rather than computed by the episodic instances. This claim is essentially different from previous methods, e.g., CAN [7] and MCT-DFMN [10]. We will give detailed justifications in the experiments.

4.2. Contextualized Label Embeddings for Prototype Rectification

In FSL, the training set $\mathcal{D}^{\text{train}}$ and test set $\mathcal{D}^{\text{test}}$ have disjoint label space. This makes the resulting embedding space trained on the training dataset not discriminative enough for making predictions on the testing dataset. Consequently, the prototypes computed on the support set \mathcal{S} with Eq. (4) might not reflect the ground true. Following the transductive-learning setting [15], we propose to incorporate the query set \mathcal{Q} to refine the prototype progressively, so that one can make more accurate predictions. To this end, we propose a self-attentive prototype rectification process, where contextualized label embeddings are propagated from \mathcal{S} to \mathcal{Q} . The overall procedure is illustrated in Figure 2.

Embedding We first use the feature extractor f to map the input $\mathcal{S} \cup \mathcal{Q}$ to an embedding matrix $\mathbf{Z} = [\mathbf{Z}^{\mathcal{S}}; \mathbf{Z}^{\mathcal{Q}}] \in \mathbb{R}^{(NK+NM) \times d}$, where $\mathcal{S} \xrightarrow{f} \mathbf{Z}^{\mathcal{S}} \in \mathbb{R}^{NK \times d}$ and $\mathcal{Q} \xrightarrow{f} \mathbf{Z}^{\mathcal{Q}} \in \mathbb{R}^{NM \times d}$. Meanwhile, the prototype matrix $\mathbf{C} = [\mathbf{c}_1; \dots; \mathbf{c}_n; \dots] \in \mathbb{R}^{N \times d}$ is initialized from $\mathbf{Z}^{\mathcal{S}}$ with Eq. (4).

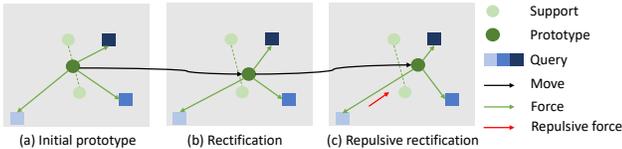


Figure 2. Prototype rectification. The color depth of a query represents the attention score to the prototype. (a) The initial prototype is the mean embedding of the two instances in the support set. (b) Rectification will take the neighbor information into consideration to refine the prototype. (c) Repulsive rectification further imposes a repulsive force to get a more representative prototype.

Prototype rectification Our idea is to refine the prototypes by considering the similarity information of the support set \mathcal{S} and query set \mathcal{Q} . To achieve this, we modify the popular attention-based method [30] to be suited to FSL.

Specifically, the prototype rectification aims to aggregate information from weighted embeddings of \mathbf{Z} , where the attentive weight $\mathbf{A} \in \mathbb{R}^{N \times (NK+NM)}$ is given by the similarity between the current prototype \mathbf{C} and the embedding \mathbf{Z} . As a result, the output \mathbf{C}^* (the rectified prototype) is:

$$\mathbf{C}^* = \mathbf{AZ} = \text{Softmax}(\kappa(\mathbf{C}, \mathbf{Z}))\mathbf{Z}. \quad (7)$$

Considering the labels of \mathcal{S} have been defined in advance in the local level, those labels shouldn't be ambiguous even if a new prototype is computed. So the way to compute the attention is not straightforward as shown in Eq. (7), otherwise the attention matrix will be problematic as presented in [35]. Given that, we split the attention matrix \mathbf{A} into two parts $\mathbf{A} = [\mathbf{A}^{\mathcal{S}}, \mathbf{A}^{\mathcal{Q}}]$ with

$$\mathbf{A}^{\mathcal{S}} = \text{Softmax}(\kappa(\mathbf{C}, \mathbf{Z}^{\mathcal{S}})) \in \mathbb{R}^{N \times NK}, \quad (8)$$

$$\mathbf{A}^{\mathcal{Q}} = \text{Softmax}(\kappa(\mathbf{C}, \mathbf{Z}^{\mathcal{Q}})) \in \mathbb{R}^{N \times NM}.$$

where the $\text{Softmax}(\cdot)$ is performed on each column. To preserve the predefined label information in the support set, we re-define the metric $\kappa(\cdot, \cdot)$ in $\mathbf{A}^{\mathcal{S}}$. Recall the label of the current data instance is n . The metric is defined as

$$\kappa(\mathbf{c}_n, \mathbf{z}_s) = \begin{cases} +\infty & \text{if } n = y_s, \\ -\infty & \text{if } n \neq y_s. \end{cases} \quad (9)$$

where $-\infty$ and $+\infty$ indicate attention scores are 0 and 1. In this way, the attention weights in $\mathbf{A}^{\mathcal{S}}$ are non-zero only if the labels are the same. In other words, $\mathbf{A}^{\mathcal{S}}$ can be written as

$$\mathbf{A}^{\mathcal{S}} = [\mathbf{v}(1); \dots; \mathbf{v}(n); \dots], \text{ for } n \in \{1, 2, \dots, N\} \quad (10)$$

where $\mathbf{v}(n)$ is a zero-vector except for elements from $K \times (n-1)$ to $K \times n$ with value 1. With $\mathbf{A}^{\mathcal{S}}$ being hard-coded, the label of each support instance won't be ambiguous even if the prototype is iteratively refined. Eventually, when concatenating the hard-coded $\mathbf{A}^{\mathcal{S}}$ with $\mathbf{A}^{\mathcal{Q}}$, we re-normalize the resulting matrix per row.

To further explore the metric space and ensure the embedding is not drifted away, a residual projection layer $h: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is stacked on top of the above attention layer:

$$\mathbf{C} = h(\mathbf{C}^*) + \mathbf{C}, \mathbf{Z} = h(\mathbf{Z}) + \mathbf{Z} \quad (11)$$

Repulsive attention In an episode of the standard N -way FSL classification scenario, each of the N classes is associated with M queries. With the attention mechanism defined above, the prototype rectification for a specific class might be affected by the other $N(M-1)$ queries beyond this class. Intuitively, queries too far away from the current instance is not expected to interact with the current instance. As a result, we propose a repulsive self-attention mechanism to refine the attention scores. Specifically, given a threshold β , we refine the attention score as:

$$\mathbf{A}[\text{mask}] = -\min(\mathbf{A}), \text{ where } \text{mask} = (\mathbf{A} < \beta). \quad (12)$$

Also, a larger β tends to enforce weaker repulsive force. The enforced negative attention score, which corresponds to some dissimilar queries, can induce a repulsive force

Models	Backbone	FT	<i>miniImageNet</i>		<i>tieredImageNet</i>	
			1-shot	5-shot	1-shot	5-shot
Prototypical Net [27]	ConvNet	✗	49.42 \pm 0.78	68.20 \pm 0.66	53.31 \pm 0.89	72.69 \pm 0.74
Relation Net [28]	ConvNet	✗	50.44 \pm 0.82	65.32 \pm 0.70	54.48 \pm 0.93	71.32 \pm 0.78
TPN [15]	ConvNet	✗	55.51 \pm 0.86	69.86 \pm 0.65	59.91 \pm 0.94	73.30 \pm 0.75
FEAT [35]	ConvNet	✓	55.75	72.17	–	–
EPNet [24]	ConvNet	✓	59.32 \pm 0.84	72.95 \pm 0.63	60.70 \pm 0.97	73.91 \pm 0.74
DFMN-MCT [10]	ConvNet	✗	64.65 \pm 0.89	75.96 \pm 0.54	65.66 \pm 0.98	75.72 \pm 0.61
ReMP	ConvNet	✗	66.21\pm0.36	76.50\pm0.29	67.12\pm0.87	76.43\pm0.50
TADAM [20]	ResNet	✗	58.50 \pm 0.30	76.70 \pm 0.30	–	–
TPN [15]	ResNet	✗	59.46	75.65	–	–
FEAT [35]	ResNet	✓	62.60	78.06	–	–
CAN [7]	ResNet	✗	67.19 \pm 0.55	80.64 \pm 0.35	73.21 \pm 0.58	84.93 \pm 0.38
EPNet [24]	ResNet	✓	70.74 \pm 0.86	81.52 \pm 0.58	78.50 \pm 0.86	87.48 \pm 0.69
DFMN-MCT [10]	ResNet	✗	78.30 \pm 0.81	86.48 \pm 0.42	80.89 \pm 0.84	87.30 \pm 0.49
ReMP	ResNet	✗	79.25\pm0.31	87.01\pm0.29	82.01\pm0.71	87.92\pm0.38

Table 1. Comparison with state-of-the-art methods on 5-way classification (FT: fine-tuning).

that prevents a prototype from being pulled away by these queries, as demonstrated in Figure 2.

In our experiment, we stack L layers of the rectification process described above, corresponding to an L -layer attention mechanism, which has been demonstrated effective in natural language processing (NLP) tasks [30, 11].

5. Experiments

In this section, we evaluate the proposed ReMP to seek answers for the following questions: **Q1**: How does ReMP perform on standard benchmarks compared to SoTA? **Q2**: Does ReMP indeed reduce the inconsistency between training and testing? **Q3**: How much does each component of ReMP contribute to the performance?

5.1. Experimental Settings

Datasets In the main part, we consider three standard FSL datasets, *miniImageNet*, *tieredImageNet* and *CIFAR-FS*.

- *miniImageNet* is a subset of ILSVRC-12 [3] proposed in [31] for FSL evaluation. It consists of 100 classes with 600 images per class. We follow the standard protocol to divide the dataset into three subsets: $\mathcal{D}^{\text{train}}$ of 64 classes, $\mathcal{D}^{\text{test}}$ of 20 classes and \mathcal{D}^{val} of 16 classes.
- *tieredImageNet* is an alternative subset of ILSVRC-12 prepared for more challenging FSL evaluation. It has a hierarchical structure of 34 coarse categories with fine-grained classes. There are three subsets: $\mathcal{D}^{\text{train}}$ with 20 categories and 351 classes, $\mathcal{D}^{\text{test}}$ with 8 categories and 160 classes, and \mathcal{D}^{val} with 6 categories and 97 classes.
- *CIFAR-FS* is based on CIFAR-100 [9]. It is split into

64 training classes, 16 validation classes and 20 testing classes.

Implementation details All experiments are conducted in PyTorch-1.3 on NVIDIA TITAN XP (12GB) platform. Following [10, 28, 12], we consider two standard embedding modules (f) as backbones: a 4-layer ConvNet with 64 channels per layer and a 12-layer ResNet, as well as standard data augmentation (including random resized crop and horizontal flip). Following [28, 10], in each episode, the number of queries in each class is set to $M = 15$. Cosine similarity is used in (1) and negative Euclidean distance is used in (3). All models are trained with a Stochastic Gradient Descent (SGD) optimizer with momentum 0.9. The initial learning rate is set to 0.1 and the weight decay is $5e^{-3}$. We decay the learning rate by a factor of 10 every 25000 iterations until convergence. We stack $L = 2$ layers in training and $L = 10$ layers in testing. The default loss balance factor $\alpha = 0.1$. All reported accuracy results are averaged over 600 test episodes with 95% confidence intervals.

Baseline Methods We compare our approach with seven baselines. All baselines, except for Prototypical Net [27], Relation Net [28] and TADAM [20], are transductive-learning based approaches. Specifically, FEAT [35] and CAN [7] use embedding propagation while TPN [15], EPNet [24] and DFMN-MCT [10] are based on label propagation. We also note that, (i) TADAM introduces a co-training method that takes the parametric linear classifier as an auxiliary task whose importance is decayed, however, it is still decoupled from the primary FSL task; (ii) FEAT borrows

the entire Transformer architecture from [30]. This tends to break both the pretrained embedding space and ignore the predefined label information (if used in a transductive manner), resulting in a limited performance gain; (iii) CAN [7] also combines the local and global losses, both of which are defined with cosine similarity. However, we demonstrate that it is not necessary to use the same metric as the global loss and the local loss don't necessarily share the same "prototypes".

5.2. Comparison with SoTA methods

The main results are shown in Table 1. As is seen, ReMP obtains the highest FSL classification accuracy as well as the lowest deviations in most settings by a large margin. Table 2 shows the 1-shot and 5-shot accuracy results on *CIFAR-FS*.

Models	Backbone	1-shot	5-shot
ProtoNet	ConvNet	55.50 \pm 0.70	72.00 \pm 0.60
MetaOpt-SVM	ResNet	72.00 \pm 0.70	84.20 \pm 0.50
DFMN-MCT	Resnet	87.51 \pm 0.48	90.23 \pm 0.63
ReMP	ResNet	87.83\pm0.31	90.17\pm0.50

Table 2. Comparison with various baselines on *CIFAR-FS*.

***N*-way *M*-query learning** We conduct further analysis to determine the effect of the number of queries per class and the number of classes per testing episodes. Decreasing the number of queries and increasing the number of classes could both affect the model's prediction accuracy. The related work DFMN-MCT, which holds the current state-of-the-arts, is chosen as a strong baseline. ReMP distinguishes from DFMN-MCT in a rectified metric objective and an efficient repulsive attention-based propagation layer. As seen in Figure 3, ReMP consistently outperforms DFMN-MCT under two sets of scenarios, ending both higher accuracy and lower deviation.

5.3. Why does ReMP improve performance?

On the rectification of learning objective We present four scenarios covering the impact of training schedules and metric options for the demonstration: (i) Cooperative training - apply \mathcal{L}_{full} (6); (ii) Pretrain and finetune - apply \mathcal{L}_{global} (2) then use \mathcal{L}_{local} (5); (iii) Local matching - only apply \mathcal{L}_{local} ; and (iv) Global matching - only apply \mathcal{L}_{global} . The main results are shown in Tab. 3, from which we can conclude:

- Global matching stimulates a more discriminative metric space, as it is observed that the accuracy in (b) is always higher than the corresponding one in (c), regardless of the metrics.

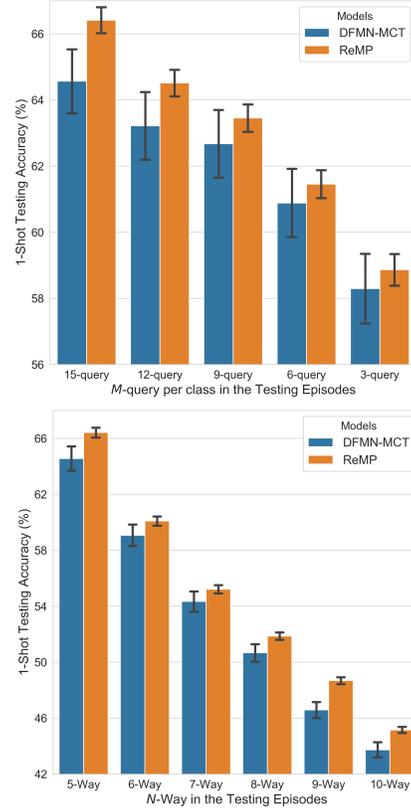


Figure 3. Comparison showing the effect of: (Top) the number of test queries per class and (Bottom) the number of classes per testing episode for 1-shot protocol on the *miniImageNet*.

- Local matching mitigates the discrepancy between training and testing as seen in (a) and (d). The reason lies in the improved transferability from training to testing, as they are being optimized towards exactly the same goal in (3).

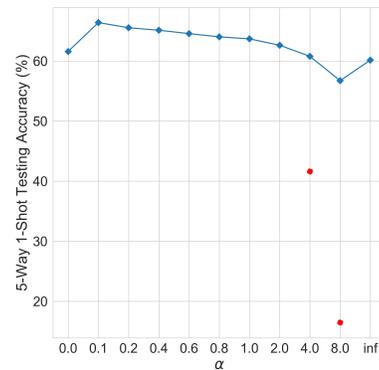


Figure 4. The impact of α under Euclidean distance.

In addition, comparing (a) with (b) and (d), we find that cooperative learning surprisingly obtains the best

performance when Cosine similarity and Euclidean distance are applied at the same time, which nearly closes the gap between the two measures. We hypothesize that the metric-based FSL problem is fundamentally difficult without inductive biases on both the models and the data. The most suitable metric space can not be easily identified. Figure 5 in Appendix B shows the t-SNE [16] visualization of embedding spaces. The observation is consistent with Tab. 3. Above findings provide strong support for Remark 1 to answer Q2. We also visualize the impact of the balance

	\mathcal{L}_{local}	Cosine	Euclidean
\mathcal{L}_{global}			
Cosine		64.69	65.55
Euclidean		65.43	65.15

(a) Cooperative training with $\alpha = 0.1$

	Test	Cosine	Euclidean
Train			
Cosine		61.89	62.41
Euclidean		61.68	62.79

(b) Pretrain \mathcal{L}_{global} , finetune \mathcal{L}_{local} .

	Test	Cosine	Euclidean
Train			
Cosine		59.67	58.53
Euclidean		57.26	59.45

(c) Training with \mathcal{L}_{local} .

	Test	Cosine	Euclidean
Train			
Cosine		56.34	60.62
Euclidean		56.31	61.31

(d) Training with \mathcal{L}_{global} .

Table 3. The impact of similarity $\kappa(\cdot, \cdot)$ and training schedule. Results are averaged over 10 independent runs of 6000 episodes.

factor α in Figure 4, where two red dots appeared as the unstable model training tries (about 6 out of 10). When $\alpha = 0$, global matching itself will cause inconsistency issues. As α gets a bit larger, global and local prototypes try to reach a consensus to induce an optimal model. If α is very large, instability issues will be triggered where global and local prototypes struggle to fit each other. When $\alpha \rightarrow \text{inf}$, the local matching itself won't find a discriminative space as the global matching does.

Repulsive attention We empirically define a layer-wise threshold $\beta_l = 1.5/N(L - l)$ throughout our experiments, where l is the layer number of a stacked propagation. β_l is expected to gradually increase as we rectify the prototype

with more layers. We visualize the rectification process using attention heatmap in Figure 5. Apparently, the iterative rectification process leads to a better decision boundary. And generally, the boundary can be further optimized by the repulsive rectification as seen in Figure 2(c).

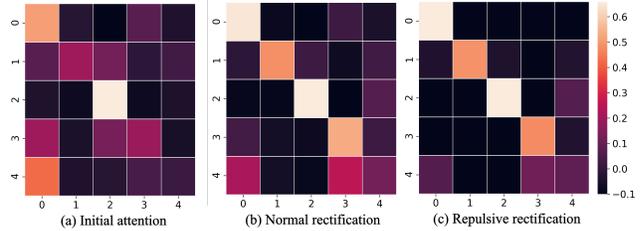


Figure 5. Visualization of the rectification process for 5-way 1-shot classification on *minilImageNet*. In a heatmap, the x -axis and y -axis represent class types for prototypes and queries, respectively. For example, the value in (row 1, column 2) represents the similarity between query of class 1 and prototype of class 2. For better classification, the diagonal values should be the highest in each row.

Ablate the component To answer Q3, we conduct an ablation study to examine the contribution of each component to our model. As it can be seen in Table 4, the transductive FSL obviously outperforms the inductive implementation. More importantly, the global matching indeed encourages a more discriminative metric space than the local matching does. In summary, the cooperative learning can give the best performance.

Ablation	Accuracy
ReMP	66.21 \pm 0.36
ReMP w/o repulsive attention	65.55 \pm 0.36
ReMP w/o \mathcal{L}_{local}	60.62 \pm 0.37
ReMP w/o \mathcal{L}_{global}	58.53 \pm 0.36
ReMP w/o contextualization (inductive)	56.13 \pm 0.78

Table 4. Ablation study results with 95% confidence intervals on 1-shot 5-way classification on *minilImageNet* (backbone: ConvNet-64).

6. Conclusions

In this paper, we have presented a rectified metric propagation (ReMP) framework for few-shot learning. To cope with the inconsistency issues between training and testing in metric-based FSL, we propose a label-aligned learning objective to close the gap, where a discriminative and transferable embedding space is induced. To further enhance the quality of the prototype in the testing phase, we propose a contextualized label embedding module, which progressively propagates labels to the query set basis in a repulsive self-attention mechanism. Empirically, ReMP achieves state-of-the-art results on various FSL datasets.

References

- [1] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. *arXiv preprint arXiv:1810.09502*, 2018.
- [2] Yinbo Chen, Xiaolong Wang, Zhuang Liu, Huijuan Xu, and Trevor Darrell. A new meta-baseline for few-shot learning. *arXiv preprint arXiv:2003.04390*, 2020.
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [4] Guneet S Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. *arXiv preprint arXiv:1909.02729*, 2019.
- [5] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [6] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018.
- [7] Ruibing Hou, Hong Chang, MA Bingpeng, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. In *Advances in Neural Information Processing Systems*, pages 4005–4016, 2019.
- [8] Jongmin Kim, Taesup Kim, Sungwoong Kim, and Chang D Yoo. Edge-labeling graph neural network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11–20, 2019.
- [9] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [10] Seong Min Kye, Hae Beom Lee, Hoirin Kim, and Sung Ju Hwang. Transductive few-shot learning with meta-learned confidence. *arXiv preprint arXiv:2002.12017*, 2020.
- [11] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [12] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10657–10665, 2019.
- [13] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-grad: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- [14] Yann Lifchitz, Yannis Avrithis, Sylvaine Picard, and Andrei Bursuc. Dense classification and implanting for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9258–9267, 2019.
- [15] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. *arXiv preprint arXiv:1805.10002*, 2018.
- [16] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [17] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- [18] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2554–2563. JMLR. org, 2017.
- [19] Alex Nichol and John Schulman. Reptile: a scalable meta-learning algorithm. *arXiv preprint arXiv:1803.02999*, 2:2, 2018.
- [20] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, pages 721–731, 2018.
- [21] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5822–5830, 2018.
- [22] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Advances in Neural Information Processing Systems*, pages 113–124, 2019.
- [23] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- [24] Pau Rodríguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. *arXiv preprint arXiv:2003.04151*, 2020.
- [25] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018.
- [26] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850, 2016.
- [27] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017.
- [28] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- [29] Sebastian Thrun. Lifelong learning algorithms. In *Learning to learn*, pages 181–209. Springer, 1998.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [31] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.

- [32] Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. Joint embedding of words and labels for text classification. *ACL*, 2018.
- [33] Jin Xu, Jean-Francois Ton, Hyunjik Kim, Adam R Kosiorek, and Yee Whye Teh. Metafun: Meta-learning with iterative functional updates. *arXiv preprint arXiv:1912.02738*, 2019.
- [34] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions, 2018.
- [35] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Learning embedding adaptation for few-shot learning. *arXiv preprint arXiv:1812.03664*, 2018.
- [36] Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. Bayesian model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, pages 7332–7342, 2018.