

# Few-shot Metric Learning: Online Adaptation of Embedding for Retrieval

Deunsol Jung, Dahyun Kang, Suha Kwak, and Minsu Cho

Pohang University of Science and Technology (POSTECH), South Korea  
`{deunsol.jung,dahyun.kang,suha.kwak,mscho}@postech.ac.kr`

**Abstract.** Metric learning aims to build a distance metric typically by learning an effective embedding function that maps similar objects into nearby points in its embedding space. Despite recent advances in deep metric learning, it remains challenging for the learned metric to generalize to unseen classes with a substantial domain gap. To tackle the issue, we explore a new problem of few-shot metric learning that aims to adapt the embedding function to the target domain with only a few annotated data. We introduce three few-shot metric learning baselines and propose the *Channel-Rectifier Meta-Learning* (CRML), which effectively adapts the metric space online by adjusting channels of intermediate layers. Experimental analyses on *miniImageNet*, CUB-200-2011, MPII, as well as a new dataset, *miniDeepFashion*, demonstrate that our method consistently improves the learned metric by adapting it to target classes and achieves a greater gain in image retrieval when the domain gap from the source classes is larger.

## 1 Introduction

The ability of measuring a reliable distance between objects is crucial for a variety of problems in the fields of artificial intelligence. Metric learning aims to learn such a distance metric for a type of input data, *e.g.*, images or texts, that conforms to semantic distance measures between the data instances. It is typically achieved by learning an embedding function that maps similar instances to nearby points on a manifold in the embedding space and dissimilar instances apart from each other. Along with the recent advance in deep neural networks, deep metric learning has evolved and applied to a variety of tasks such as image retrieval [31], person re-identification [6] and visual tracking [34]. In contrast to conventional classification approaches, which learn category-specific concepts using explicit instance-level labels for predefined classes, metric learning learns the general concept of distance metrics using relational labels between samples in the form of pairs or triplets. This type of learning is natural for information retrieval, *e.g.*, image search, where the goal is to return instances that are most similar to a query, and is also a powerful tool for open-set problems where we match or classify instances of totally new classes based on the learned metric. For this reason, metric learning has focused on generalization to unseen classes, that have never been observed during training [15, 31, 44]. Despite recent progress of




Fig. 1: **Few-shot classification vs. few-shot metric learning.** (a) A few-shot classifier learns to construct a decision boundary between support examples of different classes (red and blue). The query (white) is correctly classified as the red class, but fails to be placed away from blue samples. (b) A *few-shot metric learner* learns improved distance relations by adapting the embedding space online using the support examples, where the contour line represents the points with the same distance from the query.

deep metric learning, however, it remains exceedingly challenging for the learned embedding to generalize to unseen classes with a substantial domain gap.

To bridge the generalization gap in metric learning, we investigate the problem of few-shot metric learning that aims to *adapt an embedding function on the fly* to target classes with only a few annotated data. While the problem of few-shot learning has been actively studied for classification [21, 30], the problem for metric learning has never been directly investigated so far to the best of our knowledge. Few-shot metric learning and classification problems share the goal of adapting to scarce labeled data, but diverge in terms of their training objectives and evaluation protocols, thus requiring different approaches. As illustrated in Fig. 1, few-shot classification focuses on forming a decision boundary between samples of different classes and often fails to measure the relative distance relations between the samples, which are crucial for a retrieval task. Our analysis also shows that the improvement of classification accuracy does not necessarily lead to that of retrieval accuracy (Table 7). While one of the main approaches to few-shot classification is to learn a metric space for nearest-neighbor classification to generalize to unseen classes [21, 38, 39, 33, 45–47, 24, 7, 18], it exploits the learned metric in testing without adapting the embedding space to the few-shot instances online. In this sense, metric-based few-shot classification is very different from few-shot metric learning that we investigate in this paper.

In this work we introduce three baselines for few-shot metric learning by adapting existing methods and propose the Channel-Rectifier Meta-Learning (CRML), which effectively adapts the metric space online by adjusting channels of intermediate layers. We compare them to conventional metric learning as well as few-shot classification counterparts on *mini*ImageNet [30], CUB-200-2011 [40], and MPII [23]. We also introduce a new multi-attribute dataset for image retrieval, dubbed *mini*DeepFashion, where two instances may have an opposite

similarity relationship over multiple attributes (perspectives), and thus embedding functions are required to adapt to target attributes. Experiments show that CRML as well as the three baselines significantly improve image retrieval quality by adapting the model to target classes. Notably, such improvement is significant when the gap between training and testing domains arises, which conventional metric learning often fails to bridge.

## 2 Related work

### 2.1 Metric learning

Metric learning has been studied extensively in the past decades [32], and has shown a great success recently using deep embedding networks trained with new losses. One of the widely studied losses is the pair-wise loss [8, 31, 35, 43] which minimizes the distance between two instances that have the same label and separates them otherwise. Such losses include contrastive loss [4, 8, 12], triplet loss [31, 41], lifted structured loss [35], and multi-similarity loss [43]. Unlike pair-based losses, proxy-based losses [28, 2, 29, 19] associate proxy embeddings with each training class as a part of learnable parameters and learn semantic distance between an instance and the proxies. These previous methods, which we refer to as *conventional* metric learning, emphasize the generalization performance on unseen classes that have never been observed during training. However, they often suffer from a significant gap between source and target classes [27]. Although it is very practical to utilize a few labeled data from the target classes on the fly, online adaptation of the metric has never been explored so far to the best of our knowledge. Recently, Milbich *et al.* [27] showcase the effect of few-shot adaptation as a mean of out-of-distribution deep metric learning, their work does not present a problem formulation and a method dedicated to few-shot learning while our work does both of them. In many practical applications of metric learning, a metric can also be learned with continuous labels of similarity, which are more informative but costly to annotate [36, 11, 23, 20]. Online adaptation of metric learning may be particularly useful for such a scenario where we need to adapt the metric to the target classes with their few yet expensive labels.

### 2.2 Few-shot classification

Few-shot learning has been actively investigated for classification problems, and recent work related to ours is roughly categorized into three types: metric-based, optimization-based, and transfer-learning methods. The key idea of metric-based methods [21, 38, 39, 33, 45–47, 24, 7, 18] is to learn an embedding space via episodic training so that the class membership of a query is determined based on its nearest class representations in the embedding space. Although the metric-based few-shot classification and few-shot metric learning share the terminology “metric”, they clearly differ from each other in terms of their learning objectives and practical aspects. While metric-based few-shot classifiers construct a decision

boundary using a learned metric without online adaptation of embeddings, few-shot metric learners learn improved metric function by adapting the embeddings online. In this aspect,  $N$ -way 1-shot retrieval task, proposed in Triantafillou *et al.* [38], is different from our task, few-shot metric learning. The  $N$ -way 1-shot retrieval task in [38] does not perform any online adaptation of embedding in the inference time; thus, it is exactly the same as the conventional deep metric learning. In this work, we focus on instance retrieval problems on an adaptive embedding space of a few examples, while class discrimination is out of our interest.

The optimization-based few-shot classification methods [30, 10, 37] learn how to learn a base-learner using a few annotated data. While two aforementioned lines of work follow meta-learning frameworks, recent studies suggest that the standard transfer learning is a strong baseline for few-shot classification [5, 42, 16, 25, 9]. Such transfer learning methods pre-train a model using all available training classes and leverage the model for testing.

The contribution of this paper is four-fold:

- We introduce a new problem of *few-shot metric learning* that aims to adapt an embedding function to target classes with only a few annotated data.
- We present three few-shot metric learning baselines and a new method, Channel-Rectifier Meta-Learning (CRML), which tackles the limitations of the baselines.
- We extensively evaluate them on standard benchmarks and demonstrate that the proposed methods outperform the conventional metric learning and few-shot classification approaches by a large margin.
- We introduce *miniDeepFashion*, which is a challenging multi-attribute retrieval dataset for few-shot metric learning.

### 3 Few-shot metric learning

The goal of few-shot metric learning is to learn an embedding function for target classes with a limited number of labeled instances. In this section, we first revisit conventional metric learning, and then introduce the problem formulation and setup of few-shot metric learning.

#### 3.1 Metric learning revisited

Let us assume data  $\mathcal{X}$  of our interest, *e.g.*, a collection of images. Given an instance  $x \in \mathcal{X}$ , we can sample its positive example  $x^+$ , which is from the same class with  $x$ , and its negative example  $x^-$ , which belongs to a different class from  $x$ . The task of metric learning is to learn a distance function  $d$  such that

$$\forall(x, x^+, x^-), d(x, x^+; \theta) < d(x, x^-; \theta). \quad (1)$$

Deep metric learning solves the problem by learning a deep embedding function  $f(\cdot, \theta)$ , parameterized by  $\theta$ , that projects instances to a space where the




Fig. 2: The problem formulation of few-shot metric learning as an episodic training.

Euclidean distance is typically used as a distance function  $d$ :

$$d(x, x'; \theta) = \|f(x; \theta) - f(x'; \theta)\|_2^2. \quad (2)$$

Note that metric learning focuses on unseen class generalization. The conventional setup for metric learning [35] assumes a set of training classes  $\mathcal{C}^{\text{tr}}$  and its dataset  $\mathcal{D}^{\text{tr}} = \{(x_t, y_t) | y_t \in \mathcal{C}^{\text{tr}}\}_t$  that contains labeled instances  $x_t$  of the training classes. The task of metric learning then is to learn an embedding model using the dataset  $\mathcal{D}^{\text{tr}}$  so that it generalizes to a dataset of *unseen classes*  $\mathcal{D}^{\text{un}} = \{(x_u, y_u) | y_u \in \mathcal{C}^{\text{un}}\}_u$ , which contains instances from the classes not observed in training, *i.e.*,  $\mathcal{C}^{\text{tr}} \cap \mathcal{C}^{\text{un}} = \emptyset$ .

### 3.2 Problem formulation of few-shot metric learning

In contrast to the conventional metric learning, few-shot metric learning uses annotated instances of the target classes (thus, not totally unseen any more) but only a few instances per class, which is reasonable in most real-world scenarios. For simplicity and fair comparison, we applied  $N$ -way  $K$ -shot setting. Let us assume a target class set  $\mathcal{C}$ , which contains  $N$  classes of our interest, and its *support set*  $\mathcal{S} = \{(x_s, y_s) | y_s \in \mathcal{C}\}_{s=1}^{NK}$ , which contains  $K$  labeled instances for each of the target classes;  $K$  is supposed to be small. The task of  $N$ -way  $K$ -shot metric learning is to learn an embedding model using the support set  $\mathcal{S}$  so that it generalizes to a *prediction set*  $\mathcal{P} = \{(x_p, y_p) | y_p \in \mathcal{C}\}_{p=1}^M$ , which contains  $M$  unseen instances, *i.e.*,  $\mathcal{S} \cap \mathcal{P} = \emptyset$ , from the target classes. The generalization performance is evaluated using instance retrieval on  $\mathcal{P}$ .

Our goal is to learn such a few-shot metric learning model using a set of labeled instances from *non-target* classes, which can be viewed as a meta-learning problem. In training a few-shot metric learning model, we thus adopt the episodic training setup of meta-learning [39, 30, 10, 33] as illustrated in Fig. 2. In this setup, we are given a meta-train set  $\mathcal{D}^{\text{mtr}}$  and a meta-validation set  $\mathcal{D}^{\text{mvl}}$ . They both consist of labeled instances from non-target classes but their class sets,  $\mathcal{C}^{\text{mtr}}$  and  $\mathcal{C}^{\text{mvl}}$ , are disjoint, *i.e.*,  $\mathcal{C} \cap \mathcal{C}^{\text{mtr}} = \mathcal{C} \cap \mathcal{C}^{\text{mvl}} = \mathcal{C}^{\text{mtr}} \cap \mathcal{C}^{\text{mvl}} = \emptyset$ .

A meta-train episode is constructed from  $\mathcal{D}^{\text{mtr}}$  by simulating a support set and its prediction set; (1) a support set  $\mathcal{S}^{\text{mtr}}$  is simulated by sampling  $N$  classes from  $\mathcal{C}^{\text{mtr}}$  then  $K$  instances for each of the  $N$  classes, and (2) a prediction set  $\mathcal{P}^{\text{mtr}}$ , which is disjoint from  $\mathcal{S}^{\text{mtr}}$ , is simulated by sampling other  $K'$  instances

for each of the  $N$  classes. A meta-validation episode is constructed from  $\mathcal{D}^{\text{mvl}}$  likewise. The meta-trained few-shot metric learning model is tested on meta-test set  $\mathcal{E} = \{(\mathcal{S}_m, \mathcal{P}_m)\}_m$  from the target classes  $\mathcal{C}$ .

## 4 Methods

We introduce three baselines for few-shot metric learning (Sec. 4.1-4.3) by adapting existing few-shot learning methods, representative for few-shot classification and appropriate for few-shot metric learning in the sense that it adapts the embedding spaces online. We then discuss the limitations of the baselines and propose our method that overcomes the limitations (Sec. 4.4).

### 4.1 Simple Fine-Tuning (SFT)

As the first simple baseline for few-shot metric learning, we use the standard procedure of inductive transfer learning. In training, the embedding model  $f$  is trained on meta-training set  $\mathcal{D}^{\text{mtr}}$  from scratch with a metric learning loss; this is done in the same way with conventional metric learning, not involving any episodic training. In testing with an episode  $(\mathcal{S}, \mathcal{P})$ , the trained model  $f(x; \theta_0)$  is simply fine-tuned using the target support set  $\mathcal{S}$  by computing gradients with respect to a metric learning loss on  $\mathcal{S}$ :

$$\theta' = \theta_0 - \alpha \nabla_{\theta} \mathcal{L}(\mathcal{S}; \theta_0). \quad (3)$$

After fine-tuning, the model is tested on  $\mathcal{P}$ . We choose the number of updates that shows the highest performance on the (meta-)validation set.

### 4.2 Model-Agnostic Meta-Learning (MAML)

As the second baseline for few-shot metric learning, we employ MAML [10], which meta-learns a good initialization for few-shot adaptation via episodic training. Given a meta-train episode  $(\mathcal{S}_k^{\text{mtr}}, \mathcal{P}_k^{\text{mtr}})$ , the meta-training process consists of inner and outer loops. In the inner loop, the parameters of a *base-learner* are updated using a *meta-learner* and the support set  $\mathcal{S}_k^{\text{mtr}}$ . In MAML, the base-learner corresponds to the embedding model  $f$ , and the meta-learner corresponds to the initializer  $\theta_0$  for the embedding model. The inner loop updates the base learner using  $\mathcal{S}_k^{\text{mtr}}$  by a few gradient descent steps:

$$\theta_1 = \theta_0 - \alpha \nabla_{\theta} \mathcal{L}(\mathcal{S}_k^{\text{mtr}}; \theta_0), \quad (4)$$

where  $\alpha$  is the step size of inner-loop updates. In the outer loop, the meta-learner is updated using the loss of the updated base-learner on the prediction set  $\mathcal{P}_k^{\text{mtr}}$ :

$$\theta'_0 = \theta_0 - \eta \nabla_{\theta} \mathcal{L}(\mathcal{P}_k^{\text{mtr}}; \theta_1), \quad (5)$$

where  $\eta$  is the step size of outer-loop updates. This meta-optimization with episodic training seeks to learn the initialization of the embedding model that

generalizes well to unseen classes. For computational efficiency, we use the first-order approximation of meta-gradient update as in [10].

The meta-testing procedure with an episode  $(\mathcal{S}, \mathcal{P})$  is the same as that of the inner loop; the meta-trained model  $f(x; \theta'_0)$  is evaluated by fine-tuning it on  $\mathcal{S}$  with a few gradient steps and testing it on  $\mathcal{P}$ . We choose the number of steps that shows the highest performance on the  $\mathcal{D}^{\text{mvl}}$ .

### 4.3 Meta-Transfer Learning (MTL)

As the third baseline, we adapt MTL [37] for few-shot metric learning. MTL first pre-trains an embedding model  $f$  on the meta-training set  $\mathcal{D}^{\text{mtr}}$  and then conducts an optimization-based meta-learning process while freezing all the layers except for the last fully-connected one. Here, the base-learner is the last fully-connected layer  $\psi$  and the meta-learner consists of two groups: a set of channel scaling/shifting parameters  $\Phi = \{(\gamma^l, \beta^l)\}_l$  and the initialization of the last fully-connected layer  $\psi_0$ . The channel-scaling/shifting parameters are applied to the frozen convolutional layers of the embedding model  $f$  so that the parameters of each conv layer are scaled by  $\gamma^l$  and shifted by  $\beta^l$ :

$$\text{Conv}_{\text{CSS}}^l(X; W^l, b^l, \gamma^l, \beta^l) = (W^l \odot \gamma^l) * X + (b^l + \beta^l), \quad (6)$$

where  $W^l$  and  $b^l$  are the weight and bias of each convolution layer,  $\odot$  is channel-wise multiplication, and  $*$  is convolution. For example, if the  $3 \times 3$  convolution kernel is size of  $128 \times 64 \times 3 \times 3$ , then channel-scaling parameter is size of  $128 \times 64 \times 1 \times 1$ .

In the inner loop of the meta-training, the last fully-connected layer is fine-tuned on  $\mathcal{S}_k^{\text{mtr}}$  from the initialization  $\psi_0$ . In the outer loop, the set of channel-scaling/shifting parameters  $\Phi$  and the initialization of the last fully-connected layer  $\psi_0$  are meta-updated using the prediction loss.

In meta-testing with an episode  $(\mathcal{S}, \mathcal{P})$ , following the process of the inner loop, only the last layer  $\psi$  is updated via a few gradient steps from the meta-learned initialization  $\psi_0$  using  $\mathcal{S}$ , and the fine-tuned model is tested on  $\mathcal{P}$ .

### 4.4 Channel-Rectifier Meta-Learning (CRML)

In this subsection we discuss limitations of the aforementioned baselines and then propose a simple yet effective method for few-shot metric learning.

The main challenge in few-shot metric learning is how to effectively adapt the vast metric space using only a few examples while avoiding the danger of overfitting; as expected, the issue of overfitting is particularly critical in few-shot learning since only a few annotated examples are given for adaptation online. Updating all learnable parameters during meta-testing, as typically done in simple fine-tuning (Sec. 4.1) and MAML (Sec. 4.2), often causes quick over-fitting. A reasonable alternative is to fine-tune a part of the network only, *e.g.*, the output layer as in MTL (Sec. 4.3), with all the other parts frozen. This partial update strategy is shown to be effective for classification problems where class decision

**Algorithm 1** Channel-Rectifier Meta-Learning

---

**Input:** Meta training set  $\mathcal{D}^{\text{mtr}}$ , learning rate  $\alpha, \eta$ , pre-trained embedding model  $f$   
**Output:** The set of initialization of channel-scaling/shifting parameter  $\Phi_0 = \{(\gamma_0^l, \beta_0^l)\}_l$

- 1: Initialize  $\gamma^l \leftarrow \mathbf{1}, \beta^l \leftarrow \mathbf{0}$
- 2: **for**  $(\mathcal{S}_k^{\text{mtr}}, \mathcal{P}_k^{\text{mtr}}) \in \mathcal{D}^{\text{mtr}}$  **do**
- 3:    $\Phi_1 \leftarrow \Phi_0 - \alpha \nabla_{\Phi} \mathcal{L}(\mathcal{S}_k^{\text{mtr}}; f, \Phi_0)$  // Inner loop
- 4:    $\Phi_0 \leftarrow \Phi_0 - \eta \nabla_{\Phi} \mathcal{L}(\mathcal{P}_k^{\text{mtr}}; f, \Phi_1)$  // Outer loop
- 5: **end for**

---

boundaries can be easily affected by the specific part. In metric learning, however, fine-tuning the output layer or other specific layer online turns to hardly change the metric space, *i.e.*, distance relations among embeddings (Sec. 5.2).

To tackle the issue, we propose the *Channel-Rectifier Meta-Learning* (CRML) that meta-learns *how to rectify* channels of intermediate feature layers. The main idea is inspired by channel scaling/shifting of MTL (Eq. (6)), which is shown to be effective in adapting pretrained layers but never used in online adaptation. Unlike MTL, we propose to leverage the channel scaling/shifting module, dubbed *channel rectifier*, for online adaptation. In other words, we set the channel rectifier  $\Phi = \{(\gamma^l, \beta^l)\}_l$  as a base-learner and its initialization  $\Phi_0$  as a meta-learner. In this setup, we pre-train an embedding model  $f$  on the meta-training set  $\mathcal{D}^{\text{mtr}}$  and all the pre-trained parameters are frozen during the subsequent meta-learning process. Instead, in the meta-training stage, we update the channel rectifier  $\Phi$  in the inner loop (Eq. (4)) while updating the initialization of the channel rectifier  $\Phi_0$  in the outer loop. This meta-learning process of CRML is summarized in Alg. 1, where we describe a single-step inner loop with meta-batch of size 1 for the sake of simplicity.

In meta-testing with an episode  $(\mathcal{S}, \mathcal{P})$ , only the channel rectifier  $\Phi$  is fine-tuned by the support set  $\mathcal{S}$  with a few gradient steps from the learned initialization  $\Phi_0$ , and the fine-tuned model is tested on  $\mathcal{P}$ . Note that the CRML allows the channel rectifier to effectively exploit the support set to adapt the embedding function online.

## 5 Experiments

### 5.1 Experimental Settings

**Datasets and scenarios.** We evaluate CRML and three baselines on two standard few-shot learning datasets, *mini*ImageNet [30] and CUB-200-2011 [40]. The *mini*ImageNet dataset is a subset of the ImageNet [22] and consists of 60,000 images categorized into 100 classes with 600 images each. We use the splits divided into 64/16/20 classes for (meta-)training, (meta-)validation, and (meta-)testing, which has been introduced by [30]. The CUB-200-2011 (CUB) is a fine-grained classification dataset of bird species. It consists of 200 classes with 11,788 images in total. Following the evaluation protocol of [14], the split is divided into 100/50/50 species for (meta-)training, (meta-)validation, and (meta-)testing. We



also use the MPII dataset, a human pose dataset, to train the metric learning model with continuous labels. Following the split of [23], 22,285 full-body pose images are divided into 12,366 for (meta-)training and 9,919 for (meta-)testing. Following [20], the label for a pose pair is defined as a pose distance, the sum of Euclidean distances between body-joint locations. To verify the influence of the domain gap between the source and the target classes, we conduct cross-domain experiments [5], which is designed to have a significant domain gap between training and evaluation; (meta-)training set consists of all samples in *mini*ImageNet where each (meta-)validation and (meta-)test set consists of 50 classes from CUB. Lastly, we propose a new multi-attribute dataset for few-shot metric learning, dubbed *mini*DeepFashion. The *mini*DeepFashion is built on DeepFashion [26] dataset, which is a multi-label classification dataset with six fashion attributes. It consists of 491 classes and the number of all instances in the dataset amounts to 33,841. The details of *mini*DeepFashion is in Sec. 5.5.

**Evaluation metrics.** We use two standard evaluation metrics, mAP (mean value of average precision) [3] and Recall@ $k$  [17], to measure image retrieval performances. Recall@ $k$  evaluates the retrieval quality beyond  $k$  nearest neighbors while mAP evaluates the full ranking in retrieval. Since the MPII dataset for the human pose retrieval is labeled with continuous real values, we employ two metrics defined on continuous labels following [20]: mean pose distance (mPD) and a modified version of normalized discounted cumulative gain (nDCG). The mPD $_k$  evaluates the mean pose distance between a query and  $k$  nearest images. The modified nDCG $_k$  evaluates the rank of the  $k$  nearest images and their relevance scores. The details about evaluation metrics are specified in the supplementary material.

**Implementation details.** We use ResNet-18 [13] for the main backbone from scratch. We append a fully-connected layer with the embedding size of 128 followed by  $l_2$  normalization on top of the backbone. We use the multi-similarity loss [43] for training all the baselines and ours, CRML. For the human pose retrieval task on the MPII, we use ResNet-34 with the embedding size of 128, which is pre-trained on ImageNet [22] for a fair comparison with [20]. We fine-tune the network with log-ratio loss [20] using from 25 to 300 pairs out of all possible  $\binom{12366}{2} \approx 7.6 \times 10^7$  pairs. Complete implementation details are specified in the supplementary material.

## 5.2 Effectiveness of few-shot metric learning

**Few-shot metric learning is effective on discrete-label benchmarks.** We first compare few-shot metric learning to conventional metric learning (DML) in Tables 1, 2, and 3. All the few-shot metric learning methods consistently outperform DML not only on the 5-way 5-shot setting, which is standard for few-shot learning but also on the full-way 5-shot setting, which is standard for image retrieval. The result also shows that only five shots for each class is enough

Table 1: Performance on *miniImageNet*. Table 2: Performance on CUB-200-2011.

Method	5-way 5-shot				20-way 5-shot			
	mAP	R@1	R@2	R@4	mAP	R@1	R@2	R@4
DML	46.9	73.5	84.1	91.2	21.0	49.2	62.6	74.2
FSML <sub>SFT</sub>	63.3	78.6	85.7	91.8	29.9	52.6	65.1	75.3
FSML <sub>MAML</sub>	65.9	79.8	87.5	92.2	28.6	52.7	65.6	76.3
FSML <sub>MTL</sub>	56.5	77.2	86.7	92.6	24.0	50.0	64.6	76.2
CRML	<b>69.2</b>	<b>83.2</b>	<b>89.9</b>	<b>93.8</b>	<b>30.7</b>	<b>56.3</b>	<b>68.6</b>	<b>78.5</b>

Method	5-way 5-shot				50-way 5-shot			
	mAP	R@1	R@2	R@4	mAP	R@1	R@2	R@4
DML	57.8	81.4	88.6	93.6	26.3	51.8	62.5	72.2
FSML <sub>SFT</sub>	79.9	87.7	91.6	93.8	31.1	55.3	66.2	74.8
FSML <sub>MAML</sub>	82.0	89.5	93.2	95.4	33.2	54.6	66.3	75.7
FSML <sub>MTL</sub>	71.9	86.2	91.4	94.5	30.2	54.3	65.2	73.6
CRML	<b>82.7</b>	<b>90.0</b>	<b>93.5</b>	<b>95.5</b>	<b>33.9</b>	<b>58.1</b>	<b>68.4</b>	<b>76.5</b>

Table 3: Performance on cross-domain. Table 4: Performance on *miniDeepFashion*.

Method	5-way 5-shot				50-way 5-shot			
	mAP	R@1	R@2	R@4	mAP	R@1	R@2	R@4
DML	36.2	57.5	73.2	86.1	6.1	19.6	29.0	41.1
FSML <sub>SFT</sub>	49.4	65.2	77.1	86.0	9.8	24.1	35.2	47.8
FSML <sub>MAML</sub>	51.5	67.0	78.8	87.2	10.0	23.7	35.4	48.8
FSML <sub>MTL</sub>	40.3	62.7	28.7	41.1	6.8	20.0	29.9	42.8
CRML	<b>56.4</b>	<b>71.0</b>	<b>81.5</b>	<b>88.9</b>	<b>10.9</b>	<b>27.0</b>	<b>38.3</b>	<b>51.1</b>

Method	5-way 5-shot				20-way 5-shot			
	mAP	R@1	R@2	R@4	mAP	R@1	R@2	R@4
DML	31.8	50.3	65.8	80.2	11.3	26.1	37.3	50.2
FSML <sub>SFT</sub>	35.2	51.3	66.1	79.8	12.5	26.4	37.8	50.4
FSML <sub>MAML</sub>	38.2	<b>53.5</b>	<b>67.6</b>	80.2	<b>13.3</b>	27.7	<b>39.1</b>	<b>52.1</b>
FSML <sub>MTL</sub>	35.2	52.2	66.8	79.9	12.3	26.8	38.0	50.6
CRML	<b>38.3</b>	50.7	66.3	<b>80.2</b>	13.0	<b>27.8</b>	<b>39.1</b>	52.0

to boost the image retrieval quality regardless of the number of classes to retrieve. Such improvement is clear not only in Recall@ $k$ , *i.e.*, the measurement of top- $k$  nearest neighbors, but also in mAP, *i.e.*, the quality of all distance ranks. More importantly, the proposed CRML outperforms all baselines in the most settings, improving over MAML and SFT baselines by a large margin on the *miniImageNet* and the cross domain setting. CRML is trained to *rectify* the base feature maps by learning a small set of channel scaling and shifting parameters and thus effectively avoids overfitting to the few-shot support set from an unseen domain. In contrast, both the MAML and the SFT baselines update all parameters in the embedding functions online, thus being vulnerable to overfitting to the small number of support set. Note that the worst model is the MTL baseline, which fine-tunes the last fully-connected layer while all the other layers frozen, suggesting that fine-tuning only a single output layer in the embedding function is insufficient for online adaptation. Note that simple fine-tuning (SFT) often performs comparable to meta-learning baselines (MAML and MTL) as recently reported in transfer learning based few-shot learning work [5].

#### Few-shot metric learning is effective on continuous-label benchmarks.

We also evaluate few-shot metric learning on the human pose retrieval on MPII to demonstrate its applications and show its effectiveness. Given a human image with a certain pose, the goal of the human pose retrieval is to retrieve the most similar image of a human pose, where the supervisions in MPII consist of a continuous value on a pair-wise similarity. Since such labels are expensive to collect, few-shot learning is a practical solution for this problem, while the standard image classification approach is unlikely to applied due to the pair-based form of supervisions. Table 5 summarizes the retrieval performances with increasing numbers of the pair-wise supervisions. The retrieval performance gradually

Table 5: Performance on MPII [1]. For Table 6: Adaptation growth rate<sup>1</sup>(%) nDCG<sub>1</sub>, the higher the better, and for of CRML on three datasets. The models mPD<sub>1</sub>, the lower the better. † and ‡ de- are trained and evaluated in the 5-way notes the performances from [23] and 5-shot setting on each dataset. [20].

Metric	#(pair labels used)								Metric →	mAP (%)			recall@1 (%)		
	0	25	50	100	200	300	76M <sup>†</sup>	76M <sup>‡</sup>	Dataset →	CUB mini cross			CUB mini cross		
before adapt.										72.2	58.1	41.8	87.5	80.7	65.3
after adapt.										82.7	69.2	56.4	90.0	83.2	71.0
growth rate (%)										14.5	19.0	<b>34.8</b>	2.8	3.1	<b>8.7</b>
nDCG <sub>1</sub>	40.4	41.2	41.5	41.9	42.4	43.0	70.8	74.2							
mPD <sub>1</sub>	31.5	30.9	30.8	30.7	30.3	29.9	17.5	16.5							

improves as the number of labels increase, although the source classes (object classes in ImageNet [22]) used for pre-training deviate from the target classes of human poses. Few-shot adaptation achieves 7.7% of the state-of-the-art performance [20] trained with full supervisions using only 0.00039% of supervisions.

Please refer to the supplementary material for qualitative visualizations of our method and more experiments about effectiveness of few-shot metric learning with 1) 10 shots, 2) additional metric learning losses, 3) qualitative results.

### 5.3 Influence of domain gap between source and target

To verify the influence of the domain gap between the source and the target classes, we conduct cross-domain experiments. Table 3 shows the results of the 5-way and 50-way 5-shot experiments on the cross-domain setting. Due to the substantial domain gap between the source and target classes, the performances are much lower than those on CUB experiment in Table 2. However, the performance improvement is between 1.5 and 2 times higher than that of CUB. We observe that CRML results in remarkable improvement, which implies CRML is learned to properly rectify the base features adapted to given a support set.

We investigate the correlation between domain gap and effects of few-shot adaptation. For *mini*ImageNet, we randomly sample 60 instances from each target class in to match the prediction set size equal to that of CUB for a fair comparison. Note that the CUB is fine-grained thus has the smallest domain gap, while the cross-domain setting has the biggest. For each dataset, we measure the ratio of performance improvement from online adaptation and refer to it as adaptation growth rate. As shown in Table 6, the growth rate increases as the domain gap arises. It implies that few-shot metric learning is more effective when the target classes diverge more from the source classes.

### 5.4 Few-shot metric learning vs. few-shot classification

To verify the differences between few-shot metric learning and few-shot classification, we compare them both in image retrieval and classification. We evaluate different types of few-shot classification methods for comparison: transfer-based (Baseline [5], Baseline++ [5]), optimization-based (MAML [10]), and metric-based (MatchingNet [39], ProtoNet [33], FEAT [45]) methods. Note that as

Table 7: Classification and image retrieval performances of few-shot classification and few-shot metric learning methods on *miniImageNet* in a 5-way 5-shot setting. TTA stands for test-time adaptation via gradient descents.

Method	Classification	Image retrieval	
	Accuracy	Recall@1	mAP
Transfer-based few-shot classification			
Baseline [5]	62.21	74.04	53.14
Baseline++ [5]	74.88	78.97	65.22
Metric-based few-shot classification			
MatchingNet [39]	67.42	69.70	51.10
+TTA	71.50	71.05	52.51
ProtoNet [33]	74.46	71.04	51.73
+TTA	77.15	71.89	51.78
FEAT [45]	80.37	79.15	49.73
+TTA	<b>80.59</b>	79.31	51.72
Optimization-based few-shot classification			
MAML [10]	68.80	75.81	57.03
Few-shot metric learning			
FSML <sub>SFT</sub>	69.92	79.14	65.22
FSML <sub>MAML</sub>	72.69	79.77	65.86
FSML <sub>MTL</sub>	70.34	77.19	56.50
CRML	76.64	<b>83.22</b>	<b>69.15</b>

mentioned earlier, unlike the transfer-based and optimization-based methods, the metric-based ones in their original forms do not use online adaptation to the given support on test time. For a fair comparison, we thus perform add-on online adaptation, which is denoted by TTA, for the metric-based methods by a few steps of gradient descent using the support set.

We observe that there is little correlation between the classification accuracy and the image retrieval performances as shown in Table 7. All the few-shot metric learning methods outperform few-shot classification methods on image retrieval in terms of Recall@1 and mAP, while their classification accuracies are lower than those of few-shot classification methods. Interestingly, only Baseline++ shows competitive results with few-shot metric learning on image retrieval; we believe it is because the learned vectors in Baseline++ behave similarly to proxies in proxy-based metric learning methods. The results imply that few-shot classification learning and few-shot metric learning are distinct and result in different effects indeed. Note that even metric-based few-shot classification is not adequate for organizing the overall metric space, and additional test-time adaptation contributes insignificantly to improving the image retrieval performances.

## 5.5 Results on *miniDeepFashion*

Our design principle for the *miniDeepFashion* is to make a binary (similar/dis-similar) relationship between two instances inconsistent across the non-target

Table 8: The split of *miniDeep Fashion*.

splits	attribute (#classes)	#instances
$\mathcal{C}^{\text{mtr}}$	fabric (99)	27,079
	part (95)	
	style (127)	
$\mathcal{C}^{\text{mvl}}$	shape (76)	7,600
$\mathcal{C}$	category (37)	8,685
	texture (57)	
total		33,841

Belted trench coat	Open-shoulder shift dress	Cami trapeze dress	Dotted cami dress
			
attribute: class	attribute: class	attribute: class	attribute: class
part: belted	part: belted	shape: cami	shape: cami
category: coat	category: dress	texture: watercolor	texture: dots

Fig. 3: Two example pairs sharing attributes. Blue rows denote positive pairs on one attribute, and red rows denote negative pairs on the other.

and target classes. We thus split six attributes, each of which indicates a semantic perspective that categorizes instance into (meta-)training, (meta-)validation, and (meta-)testing as shown in Table 8. For example, a trench coat and a shift dress are either a positive pair in terms of parts they share, or a negative pair in terms of their categories as shown in Fig. 3. As existing few-shot learning datasets [30, 40] do not assume such semantic switch, one embedding space is enough for a global information. In contrast, the assumption is no longer valid on *miniDeepFashion*, thus online adaptation is inevitable; this feature is well-aligned with the goal of few-shot metric learning. *miniDeepFashion* is built on DeepFashion [26] dataset, which is a multi-label classification dataset with six fashion attributes. We construct *miniDeepFashion* by randomly sampling 100 instances from each attribute class. The number of classes in each attribute and the number of instances in each split is shown in Table 8. Also, the class configuration for each attribute is in the supplementary materials.

Table 4 shows the results on *miniDeepFashion*. We observe that it is exceptionally challenging to reexamine distance ranking between instances online when the context of target class similarity switches from that of (meta-)training class. CRML and the baselines result in moderate performance growth from DML, opening the door for future work. Figure 4 shows the qualitative retrieval results of DML and CRML on the *miniDeepFashion*. The leftmost images are queries and the right eight images are top-eight nearest neighbors. As shown in Fig. 4, DML is misled by similar colors or shapes without adapting to target attributes, texture and category, only retrieving images of common patterns. In contrast, CRML adapts to attribute-specific data, thus retrieving correct images. For example, when the query is blue chinos, while DML only retrieves the blue pants regardless of the category, CRML retrieves the chinos successfully irrespective of the color (Fig. 4 (a)). Note that *miniDeepFashion* benchmark has an original characteristics that requires online adaptation to a certain attribute of given a support set, thus this benchmark makes more sense for evaluating few-shot metric learning in comparison to prevalent few-shot classification benchmarks.



Fig. 4: Retrieval results on the (a) texture and (b) category attribute in *mini*-DeepFashion. The leftmost images are queries, and the right images are top-eight nearest neighbors. Green and red boxes are positive and negative images.

## 6 Conclusion

We have presented a few-shot metric learning framework and proposed an effective method, CRML, as well as other baseline methods. All the few-shot metric learning methods consistently outperform the conventional metric learning approach, demonstrating that they effectively adapt the learned embedding using a few annotations from target classes. Moreover, few-shot metric learning is more effective than classification approaches on relational tasks such as learning with continuous labels and multi-attribute image retrieval tasks. For this direction, we have introduced a challenging multi-attribute image retrieval dataset, *mini*DeepFashion. We believe few-shot metric learning is a new promising direction for metric learning, which effectively bridges the generalization gap of conventional metric learning.

**Acknowledgements** This work was supported by Samsung Electronics Co., Ltd. (IO201208-07822-01), the IITP grants (2022-0-00959: Few-Shot Learning of Causal Inference in Vision and Language (30%), 2019-0-01906: AI Graduate School Program at POSTECH (20%)) and NRF grant (NRF-2021R1A2C3012728 (50%)) funded by the Korea government (MSIT).

## References

1. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014)
2. Azieri, N., Todorovic, S.: Ensemble deep manifold similarity learning using hard proxies. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
3. Boyd, K., Eng, K.H., Page, C.D.: Area under the precision-recall curve: point estimates and confidence intervals. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD) (2013)
4. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a " siamese" time delay neural network. In: Advances in Neural Information Processing Systems (NerulPS) (1994)
5. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A closer look at few-shot classification. In: International Conference on Learning Representations (ICLR) (2019)
6. Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
7. Chikontwe, P., Kim, S., Park, S.H.: Cad: Co-adapting discriminative features for improved few-shot classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14554–14563 (2022)
8. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2005)
9. Chowdhury, A., Jiang, M., Chaudhuri, S., Jermaine, C.: Few-shot image classification: Just use a library of pre-trained feature extractors and a simple classifier. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9445–9454 (2021)
10. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning (ICML) (2017)
11. Gordo, A., Larlus, D.: Beyond instance-level image retrieval: Leveraging captions to learn a global visual representation for semantic retrieval. In: The IEEE conference on computer vision and pattern recognition (CVPR) (2017)
12. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2006)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
14. 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)
15. Hoffer, E., Ailon, N.: Deep metric learning using triplet network. In: International Workshop on Similarity-Based Pattern Recognition (SIMBAD) (2015)
16. Hu, S.X., Li, D., Stühmer, J., Kim, M., Hospedales, T.M.: Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9068–9077 (2022)

17. Jegou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **33**(1), 117–128 (2010)
18. Kang, D., Kwon, H., Min, J., Cho, M.: Relational embedding for few-shot classification. In: *The IEEE International Conference on Computer Vision (ICCV)* (2021)
19. Kim, S., Kim, D., Cho, M., Kwak, S.: Proxy anchor loss for deep metric learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3238–3247 (2020)
20. Kim, S., Seo, M., Laptev, I., Cho, M., Kwak, S.: Deep metric learning beyond binary supervision. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
21. Koch, G., Zemel, R., Salakhutdinov, R.: Siamese neural networks for one-shot image recognition. In: *International Conference on Machine Learning (ICML) Workshop on Deep Learning* (2015)
22. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems (NeurIPS)* (2012)
23. Kwak, S., Cho, M., Laptev, I.: Thin-slicing for pose: Learning to understand pose without explicit pose estimation. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
24. Li, W.H., Liu, X., Bilen, H.: Universal representation learning from multiple domains for few-shot classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9526–9535 (2021)
25. Liu, B., Cao, Y., Lin, Y., Li, Q., Zhang, Z., Long, M., Hu, H.: Negative margin matters: Understanding margin in few-shot classification (2020)
26. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
27. Milbich, T., Roth, K., Sinha, S., Schmidt, L., Ghassemi, M., Ommer, B.: Characterizing generalization under out-of-distribution shifts in deep metric learning. *Advances in Neural Information Processing Systems* **34**, 25006–25018 (2021)
28. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: *The IEEE International Conference on Computer Vision (ICCV)* (2017)
29. Qian, Q., Shang, L., Sun, B., Hu, J., Li, H., Jin, R.: Softtriple loss: Deep metric learning without triplet sampling. In: *The IEEE International Conference on Computer Vision (ICCV)* (2019)
30. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: *International Conference on Learning Representations (ICLR)* (2017)
31. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
32. Short, R., Fukunaga, K.: The optimal distance measure for nearest neighbor classification. *IEEE Transactions on Information Theory (TIT)* **27**(5), 622–627 (1981)
33. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
34. Son, J., Baek, M., Cho, M., Han, B.: Multi-object tracking with quadruplet convolutional neural networks. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)



35. Song, H.O., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
36. Sumer, O., Dencker, T., Ommer, B.: Self-supervised learning of pose embeddings from spatiotemporal relations in videos. In: The IEEE International Conference on Computer Vision (ICCV) (2017)
37. Sun, Q., Liu, Y., Chua, T.S., Schiele, B.: Meta-transfer learning for few-shot learning. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
38. Triantafillou, E., Zemel, R., Urtasun, R.: Few-shot learning through an information retrieval lens. *Advances in neural information processing systems* **30** (2017)
39. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2016)
40. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
41. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014)
42. Wang, X., Huang, T.E., Darrell, T., Gonzalez, J.E., Yu, F.: Frustratingly simple few-shot object detection (July 2020)
43. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
44. Wohlhart, P., Lepetit, V.: Learning descriptors for object recognition and 3d pose estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
45. Ye, H.J., Hu, H., Zhan, D.C., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8808–8817 (2020)
46. Ye, H.J., Hu, H., Zhan, D.C., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
47. Zhang, C., Cai, Y., Lin, G., Shen, C.: Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)