

Incremental Few-Shot Learning for Pedestrian Attribute Recognition

Liuyu Xiang¹, Xiaoming Jin^{1*}, Guiguang Ding^{1*}, Jungong Han² and Leida Li³

¹School of Software, Tsinghua University, Beijing, China

²WMG Data Science, University of Warwick, CV4 7AL Coventry, United Kingdom

³School of Artificial Intelligence, Xidian University, Xi'an 710071, China

xiangly17@mails.tsinghua.edu.cn, {xmjin,dingg}@tsinghua.edu.cn

jungonghan77@gmail.com, reader1104@hotmail.com

Abstract

Pedestrian attribute recognition has received increasing attention due to its important role in video surveillance applications. However, most existing methods are designed for a fixed set of attributes. They are unable to handle the incremental few-shot learning scenario, i.e. adapting a well-trained model to newly added attributes with scarce data, which commonly exists in the real world. In this work, we present a meta learning based method to address this issue. The core of our framework is a meta architecture capable of disentangling multiple attribute information and generalizing rapidly to new coming attributes. By conducting extensive experiments on the benchmark dataset PETA and RAP under the incremental few-shot setting, we show that our method is able to perform the task with competitive performances and low resource requirements.

1 Introduction

Pedestrian attribute recognition, aiming at recognizing human part attributes such as age, gender and clothing, is in growing demand for its wide applications in visual surveillance such as person re-identification, social behavioral analysis, and face verification. In the deep learning era, various methods based on either CNN or RNN have been proposed, and have achieved great success in this area. However, in real-world visual surveillance applications, the demand is usually dynamic and evolving, that is, the set of human part attributes that need to be recognized can be changed or enlarged over time.

Consider the example in Figure 1. During phase one, a model is trained (which we call base model) to recognize the existence of several attributes such as 'Customer', 'Suit-up', and 'Long Trousers'. Then during phase two, a few-shot dataset annotated with some new attributes is provided. Our goal is to develop an efficient algorithm that can utilize both the few-shot data and the already well-trained base model, and produce an extended model that can not only successfully recognize the new attributes such as 'hat' and 'Single

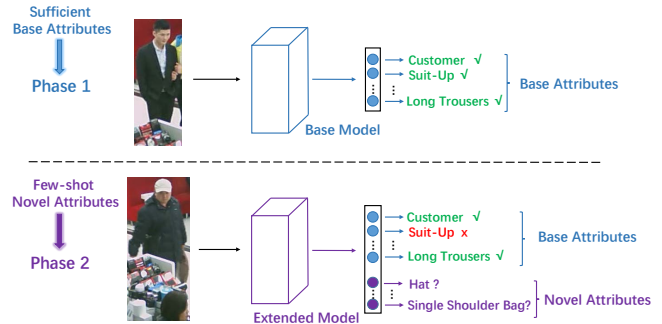


Figure 1: Illustration of Incremental Few-shot Learning for Pedestrian Attribute Recognition

Shoulder Bag', but also retain the capability of classifying old attributes such as 'Customer' and 'Suit-up'. For clarity we denote the previous well-recognized attributes as **base attributes**, and the new coming attributes as **novel attributes**.

Given the scenarios described above, two straight-forward solutions would be finetuning and re-training. However, finetuning, i.e. adding randomly initialized output layers to the previous model and tuning the whole network on the few-shot examples, is prone to overfitting on novel attributes, and may even lead to catastrophic forgetting on base attributes [Goodfellow *et al.*, 2013]. Re-training, i.e. discarding the previous well-trained model and training from scratch using both base and novel data, not only requires to store the previous large-scale training data, but is also time-consuming. Thus it is infeasible to re-train in most applications where the time and memory are limited. Thus, we need to find a solution, that given the previously well-trained model and without storing the previous training data, can efficiently extend the previous model so that it is able to recognize both the base and novel attributes. It is worth to note that this task is challenging, since it involves incremental learning, few-shot learning and multi-label learning. Simply applying or combining any technique from only one of these areas is insufficient to solve this complex problem.

In order to address the challenges described above, we propose to extend the base model with an extra module called Attribute Prototype Generator Module (APGM). The APGM can be regarded as a higher-level meta-learner that extracts the multiple-attribute information from the feature embed-

*Corresponding authors

ding, and produces discriminative Attribute Prototype Embedding (APE). Then the classification weights for the novel attributes are acquired based on APE. Further, if the *base attributes* of the few-shot dataset are also annotated, we put forward another module called Attribute Relationship Module (ARM) that transfers knowledge from the base classification weights to the novel weights based on attributes co-occurrence, to further improve the performances of the generated novel classification weights.

To verify that our proposed method is capable of performing incremental few-shot learning for pedestrian attribute recognition, we conduct experiments on two benchmark datasets for pedestrian attribute recognition: PETA and RAP. We first transform these datasets into incremental few-shot case, and show by experiments that our method performs superior to the baselines with less time and memory consumption.

Besides, we also discuss the internal relationships between our method and other two contemporary meta learning methods. We explain intuitively the reason why our method is more suitable for the multi-label few-shot case, and demonstrate experimentally that simply extending the multiclass state-of-the-art few-shot meta learning algorithms into their multi-label counterparts will result in poor performances under this context.

In summary, we make the following contributions in this paper:

- We raise a practical issue termed as incremental few-shot (multi-label) attribute learning which commonly exists in real-word applications.
- We put forward a meta learning based method that achieves competitive performances.
- Our method demonstrates the capability of fast adaptation with low annotation and memory requirements which is desirable in low-resource on-device scenarios.

2 Related Work

2.1 Pedestrian Attribute Recognition

Pedestrian attributes have been used as soft biometric traits, and their recognition has been intensively studied due to its wide applications in video surveillance. Traditional methods usually involve extracting hand-crafted features and optimizing classifiers such as Support Vector Machine (SVM) [Prosser *et al.*, 2010]. However, due to the limited representation capability of these hand-designed features, these traditional methods' performances still need to be improved.

With the development of automatic representation learning, deep architectures based on CNN, CNN-RNN and ConvLSTM have been proposed. In particular, [Sudowe *et al.*, 2015] first adopts CNN that is trained end-to-end with independent layers for each single attribute. [Li *et al.*, 2015] proposes DeepMAR with a weighted sigmoid cross-entropy loss to utilize the attribute distribution prior information to handle the class-imbalance problem. Later, [Wang *et al.*, 2017] proposes a CNN-RNN based encoder-decoder framework, and [Zhao *et al.*, 2018] further improves recurrent learning via

attribute grouping. Apart from designing novel network architectures, [Li *et al.*, 2017] exploits a Spatial Transformer Networks (STN) to extract multi-scale context-aware discriminative features, where [Zhu *et al.*, 2017] divides the human body into several parts for more accurate attribute recognition. Efforts have also been made to model the local or hierarchical relationships where techniques such as attention mechanism and automatic gating [Liu *et al.*, 2017; Jin *et al.*, 2018] are exploited.

2.2 Few-Shot Learning and Meta Learning

Meta learning, also known as *learning to learn* [Thrun and Pratt, 2012], has achieved remarkable success in few-shot learning and other applications [Ravi and Larochelle, 2016; Finn *et al.*, 2017; Snell *et al.*, 2017; Xiang *et al.*, 2019].

The key idea of meta learning is to learn a higher-level parameterized network (which we call meta-learner), to guide an ordinary network (which we call learner) to learn. This is often achieved by the so-called episodic training, which enables the meta-network to generalize to few examples by consistently mimicking the few-shot scenarios. Among these methods, [Finn *et al.*, 2017] learns an initialization of a neural network capable of fast adaptation, where [Ravi and Larochelle, 2016] utilizes an LSTM as meta-learner to guide the learner to update parameters using the gradient information. [Snell *et al.*, 2017] and [Sung *et al.*, 2018] approach few-shot learning via metric learning, where the former learns an embedding function followed by KNN, and the latter learns the metric function as well. Dynamical weight methods including [Bertinetto *et al.*, 2016; Qi *et al.*, 2018; Gidaris and Komodakis, 2018] propose to train a meta-network to produce parameters of a base learner, so that the produced parameters (fast weights) can quickly adapt to few examples.

However, most existing meta learning methods are targeted for multi-class few-shot classification, and is not applicable to our **multi-label** setting where complex attribute relationships and co-occurrence exist.

2.3 Incremental Learning

The problem we address here is also related to (but distinct from) incremental learning, also known as lifelong learning or continual learning, where the model is given a sequence of tasks to learn. The main issue of incremental learning is to overcome catastrophic forgetting [Goodfellow *et al.*, 2013], a phenomenon where the model performs poorly on the previous tasks after learning the new tasks.

Two representative works of overcoming catastrophic forgetting are LwF [Li and Hoiem, 2018] and EWC [Kirkpatrick *et al.*, 2017]. LwF aims to alleviate the catastrophic forgetting by exploiting knowledge distillation on the old output nodes of the model. EWC, on the other hand, evaluates each neurons importance for previous tasks using fisher information matrix and keep the important neurons close to their original values. The incremental few-shot problem we are facing here differs from the standard incremental setting where the learning of the novel classes is harder and catastrophic forgetting is less severe.

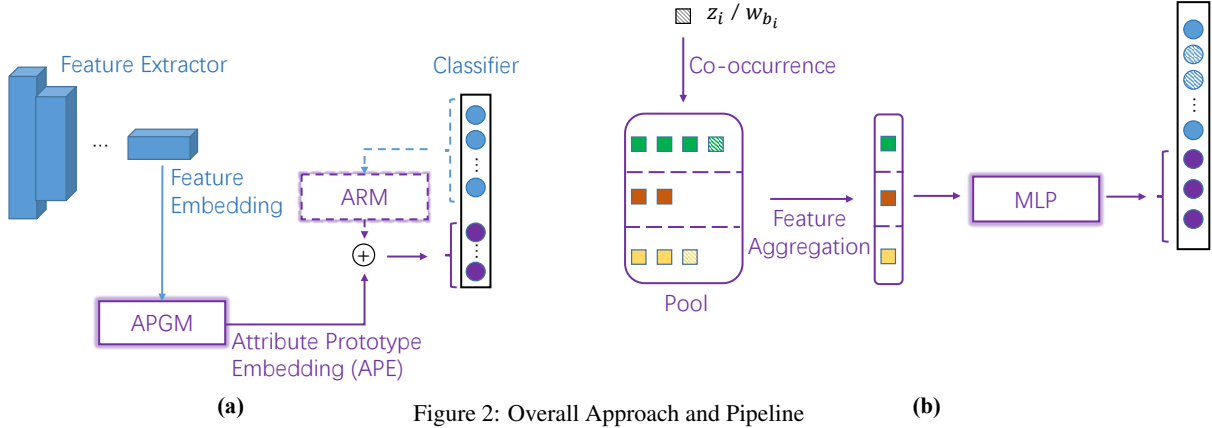


Figure 2: Overall Approach and Pipeline

3 Problem Setup

Incremental few-shot attribute learning is essentially a **multi-label few-shot learning** problem. We assume that in phase one we are given a large-scale pedestrian attribute dataset $\mathcal{D}_{base} = \{(x_1, y_1), \dots, (x_{K_{base}}, y_{K_{base}})\}$, $y_i \in \{0, 1\}^{N_{base}}$, with attributes $\mathcal{A}_{base} = \{a_1, \dots, a_{N_{base}}\}$, together with a model \mathcal{M}_{base} which is well-trained on it.

Then in phase two, we are provided with a new dataset \mathcal{D}_{novel} , in which novel attributes $\mathcal{A}_{novel} = \{a_1, \dots, a_{N_{novel}}\}$ appear. The \mathcal{D}_{novel} contains K_{novel} examples annotated with N_{novel} novel attributes where K_{novel} is usually small, and the entire learning process is termed as N -way K -shot learning.

Our goal is to extend the \mathcal{M}_{base} to \mathcal{M}_{ext} in the second phase so that it is able to recognize both \mathcal{A}_{base} and \mathcal{A}_{novel} . It is worth noting that \mathcal{A}_{novel} can be either absent or not annotated in \mathcal{D}_{base} . Unlike the normal multiclass incremental few-shot learning setting which assumes $\mathcal{D}_{base} \cap \mathcal{D}_{novel} = \phi$, our setting only requires $\mathcal{A}_{base} \cap \mathcal{A}_{novel} = \phi$ and \mathcal{A}_{novel} will not appear in \mathcal{D}_{base} , in which case the same example may appear both in \mathcal{D}_{base} and \mathcal{D}_{novel} , but are annotated with \mathcal{A}_{base} and \mathcal{A}_{novel} respectively. In other words, it is more of **class incremental** rather than the **data incremental** setting.

4 Proposed Method

4.1 Overview

Our overall approach for incremental few-shot attribute learning is illustrated in Figure 2(a). We first start with a well-trained base model \mathcal{M}_{base} (the blue part) with conceptually two parts: feature extractor \mathcal{F} and classifier \mathcal{W} . We then introduce the APGM module which has no interference with the base model so that the performances on \mathcal{A}_{base} will not be hampered. The APGM takes the feature embedding that contains the encoded information from multiple attributes as inputs, and produces the separated APE embedding for each attribute. Then the classification weights for novel attributes are produced using the APEs. Moreover, when \mathcal{D}_{novel} is also annotated with base attributes, we propose the ARM, which takes the base attributes' classification weights as inputs, and works in a similar way as APGM to produce the novel classification weights. In order to facilitate few-shot learning, we

train these modules with episodic training on \mathcal{D}_{base} , which will be discussed in the next section.

4.2 Episodic Training

The essence of meta learning involves two levels of learning. During the low-level training, tasks $\{\mathcal{T}_t\}$ (such as an N -way K -shot classification) are sampled from \mathcal{D}_{base} as support examples $\mathcal{D}_{support}$, on which the learner is guided by the meta-learner to learn. Then during the high-level training, the learner is evaluated on \mathcal{D}_{query} (which is also from \mathcal{D}_{base} and held out from $\mathcal{D}_{support}$) with some meta-objective to find out how well the meta-learner is capable of guiding the learner to generalize to new tasks, and then the meta-learner is optimized using the meta-objective (which is cross-entropy loss in our case). This whole process is called an *episode* [Vinyals *et al.*, 2016]. By repetitively episodic training, which is essentially mimicking the expected few-shot setting, the meta-learner will gradually acquire the ability to generalize on few-shot task \mathcal{T}' from \mathcal{D}_{novel} held-out from $\{\mathcal{T}_t\}$.

Following the idea of episodic training, we first sample N -way K -shot tasks from \mathcal{D}_{base} as $\mathcal{D}_{support}$, and treat the sampled attributes from \mathcal{A}_{base} as *fake novel attributes* $\mathcal{A}_{fake.novel}$. We then produce the classification weights $\hat{W}_{fake.novel}$ for $\mathcal{A}_{fake.novel}$ under the guidance of APGM (and ARM). Then we evaluate the classification performances of $\hat{W}_{fake.novel}$ on \mathcal{D}_{query} , and let the gradient backpropagated to update the APGM and ARM.

The whole episodic training procedure on \mathcal{D}_{base} is shown in Algorithm 1.

4.3 Attribute Prototype Generator Module

The basic pipeline of APGM is shown in Figure 2(b). Note that different from the multiclass classification case, where each feature embedding only belongs to a single class, here each feature embedding contains fused information of multiple attributes, and the goal of APGM is to make efficient transformation to separate the fused information so that the generated APE is discriminative enough for classification.

The APGM mainly consists of two components: a pool and an MLP. The pool contains N_{novel} slots, and each slot stores all the relevant feature embedding for one novel attribute. The

Algorithm 1: N -way K -shot Episodic Training

Input : Base dataset \mathcal{D}_{base} , Base model \mathcal{M}_{base} , Base attributes set \mathcal{A}_{base} , Number of sampled fake novel attributes N_{fake_novel} , Number of support examples $K_{support}$, Number of episodes E .

Output: \mathcal{M}_{base} with trained APGM, ARM modules

while episode $\leq E$ **do**

```

 $\mathcal{A}_{fake\_novel} = \text{Sample}(\mathcal{A}_{base}, N_{fake\_novel})$ 
 $\mathcal{D}_{fake\_novel}, \mathcal{D}_{fake\_base} = \text{Split}(\mathcal{D}_{base})$ 
// where  $\mathcal{D}_{fake\_novel}$  contains at least one
// attribute from  $\mathcal{A}_{fake\_novel}$ 
 $\mathcal{D}_{support} = \text{Sample}(\mathcal{D}_{fake\_novel}, K_{support})$ 
 $\mathcal{D}_{query} = \mathcal{D}_{base} \setminus \mathcal{D}_{support}$ 
for  $(x_i, y_i)$  in  $\mathcal{D}_{support}$  do
     $z_i = \text{FeatureExtract}(x_i)$ 
     $Z_{support}.append(z_i)$ 
end
 $\hat{W}_{APE} = \text{APGM}(Z_{support}, \mathcal{D}_{support})$ 
 $\hat{W}_{ARM} = \text{ARM}(W_{fake\_base}, \mathcal{D}_{support})$ 
 $\hat{W}_{fake\_novel} = \hat{W}_{APE} + \hat{W}_{ARM}$ 
 $\hat{W} = \text{concat}(W_{fake\_base}, \hat{W}_{fake\_novel})$ 
// low-level learning
for  $(x_j, y_j)$  in  $\mathcal{D}_{query}$  do
     $score = \text{FeedForward}(x_j, [\mathcal{F}, \hat{W}])$ 
    update AGPM and ARM using meta-objective
     $\mathcal{L}(score, y_j)$ 
    // high-level update
end
end
return APGM, ARM module
    
```

MLP serves as the parameterized meta-learner, which is optimized during the high-level learning. A well-trained APGM works as follows. Suppose that we have N_{novel} novel attributes \mathcal{A}_{novel} , where in Figure 2(b) $N_{novel} = 3$. First, the i^{th} example $(x_i, y_i) \in \mathcal{D}_{novel}$ are fed into the feature extractor to obtain the feature embedding z_i . With its annotated novel attributes $\{a_{i1}, \dots, a_{il_i}\} \subseteq \mathcal{A}_{novel}$, we push z_i into the slots corresponding to a_{i1}, \dots, a_{il_i} , where in the figure, z_i is pushed into the first and the third slots. After forwarding the whole few-shot dataset \mathcal{D}_{novel} , we aggregate features over each slot, and feed the aggregated embedding for each novel attribute into the MLP to produce the APE. The aggregation operation can have multiple choices, and we choose summation in our experiment. The whole forward process of APGM module during meta training is shown in Algorithm 2.

4.4 Attribute Relationship Module

During the generation of attribute prototype generation, only the novel attributes \mathcal{A}_{novel} in \mathcal{D}_{novel} is needed. If we are further provided with the base attributes \mathcal{A}_{base} on \mathcal{D}_{novel} , then we can exploit the well-trained base classification weights (the output node in blue) and attributes relationships for further improvements.

The ARM module shares a similar structure with APGM module with a disjoint set of parameters where the input to

Algorithm 2: APGM Module

Input : Support set $\mathcal{D}_{support} = \{(x_i, y_i)\}$, support features $Z_{support} = \{z_i\}$, fake novel attributes \mathcal{A}_{fake_novel} .

Output: Attribute Prototype Embedding \hat{W}_{APE}

```

for  $z_i$  in  $Z_{support}$  do
    Get  $x_i$ 's attributes  $[a_{i1}, a_{i2}, \dots, a_{il_i}]$ 
    for  $a_{ij}$  in  $[a_{i1}, a_{i2}, \dots, a_{il_i}]$  do
        |  $\text{slot}[a_{ij}].append(z_i)$ 
    end
end
for  $a_k$  in  $\mathcal{A}_{fake\_novel}$  do
    |  $\hat{W}_{APE}[a_k] = \text{MLP}(\text{Aggregate}(\text{slot}[a_k]))$ 
end
return  $\hat{W}_{APE}$ 
    
```

the ARM is the base classification weights W_{base} (the output node in blue). When forwarding an example (x_i, y_i) , whose attributes are denoted as $\{a_{i1}, \dots, a_{il_i}\}$, for each co-occurred base/novel attribute pair (a_{ij}, a_{ik}) , $a_{ij} \in \mathcal{A}_{base}$ and $a_{ik} \in \mathcal{A}_{novel}$, we push $w_{a_{ij}} \in W_{base}$ into a_{ik} 's corresponding slots. Then all the weights in the slots are aggregated and feed into the MLP. This process is in essence a weighted attention over W_{base} determined by how often two attributes will co-appear on the same instance.

Finally, the output classification weights for novel attributes are the addition of APGM and ARM's outputs. During practical implementation, we further normalize the classification weights to be unit vector following [Qi *et al.*, 2018], in which case the softmax classifier is equivalent to nearest neighbour with cosine distance, which will help stabilizing the meta-training.

4.5 Inference

During the second phase when the actual novel attributes come, we fix the meta-learner, and feed forward the whole \mathcal{D}_{novel} through APGM (and ARM if possible) to produce the classification weights W_{novel} for \mathcal{A}_{novel} . Then a complete model \mathcal{M}_{ext} is obtained by concatenating W_{base} and W_{novel} , where the feature extractor \mathcal{F} remains unchanged. It is worth noting that unlike re-training or finetuning which entails several steps of backpropagation and optimization, our method **only requires forwarding \mathcal{D}_{novel} once to generate the \mathcal{M}_{ext}** , which could be valuable for applications where models need to update on-device.

4.6 Discussions: Relationships with Other Meta Learning Methods

In this section, we discuss the relationships among our method and two state-of-the-art meta learning based methods [Snell *et al.*, 2017; Gidaris and Komodakis, 2018]. First we analyze by decoupling the whole network \mathcal{M} into feature extractor \mathcal{F} and classifier \mathcal{W} . In [Snell *et al.*, 2017] the Prototypical Network adopts the entire feature embedding encoder \mathcal{F} as meta-learner, where the classifier \mathcal{W} is the non-parametric KNN, which indicates that the prototype is the feature embedding itself. Both [Gidaris and Komodakis,

2018] and our method use an extra parameterized weight generator as meta-learner, which outputs the classifier weights \mathcal{W} . However, the meta-learner in [Gidaris and Komodakis, 2018] is a Hadamard production with learnable parameters that transforms the feature embedding into \mathcal{W} , where our meta-learner is an Multi-Layer Perceptron (MLP).

As we will illustrate in the experimental section, that simply modifying [Snell *et al.*, 2017; Gidaris and Komodakis, 2018] to their multi-label versions perform poorly in our setting. For Prototypical Network, we argue that choosing \mathcal{F} as the meta-learner in multi-label setting will not be able to learn class prototypes that generalize and transfer as they do in the multiclass case. This is because the same prototype can be attributed to multiple labels, thus it is almost impossible to represent prototype information from multiple labels within a single class prototype. Moreover, updating the \mathcal{F} during episodic training will cause catastrophic forgetting on the base attributes. By comparison, our method and [Gidaris and Komodakis, 2018] avoid this issue by **decoupling the meta-learner and the feature extractor** \mathcal{F} , so that \mathcal{F} can remain unchanged during the second phase, thus keeping the performances on the base attributes. However, [Gidaris and Komodakis, 2018] uses Hadamard production with learnable parameters as meta network, which indicates that **the same set of coefficients is used for all attribute prototypes generation** and we argue that it is not discriminative enough to disentangle the fused information and produce the classification weights. Our meta network, however, is an MLP that provides different sets of coefficients for each class prototype generation, thus leading to more discriminative transformation and superior performances.

5 Experimental Results

5.1 Data Preparation

We evaluate on two pedestrian attribute benchmark datasets PETA and RAP. We first modify these two datasets into incremental few-shot setting. To be specific, for N -way K -shot incremental few-shot learning, we first sample N attribute groups (see [Zhao *et al.*, 2018]), then randomly choose one attribute from each of the selected N groups to form \mathcal{A}_{novel} , and treat the remaining attributes as \mathcal{A}_{base} . We follow [Li *et al.*, 2015] and [Li *et al.*, 2016], and divide the PETA and RAP into train/val/test set with 5 random partitions. All the reported results are the average results of these 5 partitions. We then further divide each the train/val/test set into two parts: base and novel. An example is included in the novel sets if it contains at least one novel attribute, and is included in the base sets following the same rule. We then randomly keep K examples in the train_{novel} set and discard the rest of the examples. During phase one, the model is trained on $\text{train}_{base}/\text{val}_{base}$, and test with test_{base} . During the second phase, the model is provided with the K -shot $\text{train}_{novel}/\text{val}_{novel}$, and test on the whole test set containing both base and novel attributes.

5.2 Baselines

Finetune with Different Learning Rate is the most straight-forward method in dealing with few-shot settings.

We fix the learning rate of classifier \mathcal{W} to 0.001, and without specification, the lr refers to the learning rate of the feature extractor \mathcal{F} .

Prototypical Network [Snell *et al.*, 2017] is a simple yet effective meta learning method. During the experiment, we change the metric from Euclidean distance to cosine similarity to cater for the multi-label setting.

DFWF [Gidaris and Komodakis, 2018] is a recent state-of-the-art meta learning method. During the experiment, we also modify the original DFWF for the multi-label setting.

5.3 Evaluation Metric

In the experiment, we mainly focus on four metrics: (1) Class-centric mA: the mean average of each attribute’s accuracy of positive and negative example. Specifically, we mainly focus on mA_{base} and mA_{novel} as they indicate how well the model can adapt to the novel attributes as well as remembering the old attributes in our class-incremental setting. (2) Instance-centric: the Precision, Recall and F1 calculated on each instance.




Figure 3: Selected Novel Attributes Distribution of 5-way 50-shot Learning on PETA and RAP

5.4 Implementation Details

Base Network. Since we focus on the learning algorithm which is agnostic to the specific model, we adopt DeepMAR[Li *et al.*, 2015] with ResNet50 as backbone as our base model \mathcal{M}_{base} for its competitive performances and simplicity.

Hyperparameters. In order to train our meta network, we adopt SGD and a small learning rate of 10^{-5} , with a momentum of 0.9 and weight decay of 0.0005. We choose the batch size to be 32, train for 800 episodes and choose the sampled N_{fake_novel} to be twice the expected N_{novel} . During the meta training, we re-sample the support set if any a_{fake_novel} is absent in the $D_{support}$.

Alleviating Overfitting. In order to alleviate overfitting to a certain sampled task, we apply early stopping when the loss stops to decrease within each episode. During inference, we choose the model with the smallest training loss at the end of each episode as our final extended model.

5.5 Results on 5-way 50-shot Learning

We first experiment on PETA and RAP with 5-way 50-shot learning with two sets of randomly selected novel attributes, and the result is shown in Table 1. Here only novel attributes are annotated on \mathcal{D}_{novel} . It is also worth to mention that although the total number of examples is 50, these attributes

Metric \ Dataset	PETA					RAP				
	Method	mA_{Base}	mA_{Novel}	Precision	Recall	F1	mA_{Base}	mA_{Novel}	Precision	Recall
Ft (lr=0)	↓ 0.15	65.66	83.22	82.22	82.71	↓ 0.47	67.57	77.18	72.86	74.95
Ft (lr=0.01)	↓ 2.36	61.84	81.12	81.06	81.08	↓ 1.19	62.79	76.34	73.58	74.92
Ft (lr=0.001)	↓ 0.24	64.12	83.10	82.05	82.57	↓ 0.39	66.38	77.19	72.97	75.02
Ft (lr=0.0001)	↓ 0.12	65.69	83.24	82.11	82.67	↓ 0.44	67.47	77.24	72.98	75.05
Ft (lr=0.00001)	↓ 0.13	65.52	83.23	82.28	82.75	↓ 0.46	67.67	77.22	72.93	75.01
ProtoNet	49.63	50.00	27.55	75.37	40.26	49.66	50.00	18.13	91.65	30.27
DFWF	↓ 0	54.46	70.42	82.51	75.78	↓ 0	57.80	70.96	77.14	65.82
APGM (Ours)	↓ 0.0	68.13	74.05	84.53	78.93	↓ 0.0	73.06	71.31	76.95	74.02

Table 1: 5-way 50-shot learning results on PETA and RAP benchmark. Ft (lr) refers to finetuning, and lr indicates the learning rate of \mathcal{F} . ↓ means decrease in mA_{base} as compared to the base model.

are highly imbalanced, where some attributes only contain less than 5 positive examples. The attribute distribution of the first five selected novel attributes is shown in Figure 3. Without specification, the rest of the results are reported on these five novel attributes. The results from Table 1 show that compared to finetuning (Ft) with different learning rates, our method outperforms significantly in mA_{novel} on both datasets, and holds comparable F1 on RAP while suffers a decrease in F1 on PETA. The drop in F1 mainly comes from the drop in Precision, and further regularization might be considered in future work. Besides, finetuning comes at a cost of forgetting on the base attributes where our method is able to retain the base model’s performances. From the class incremental perspective, the class-centric metric mA_{base}/mA_{novel} demonstrates our method’s superiority in learning novel attributes on such few-shot learning setting. Moreover, finetuning involves several backpropagation and updating while our method requires much less time consumption. Finally, the results also accord with the discussions in Section 4.6, that the current meta learning method is not applicable in our setting where the ProtoNet acts almost like random guess, and DFWF also yields poor performances.




Figure 4: Effect of Number of Shots and Ways

5.6 Effect of ARM Module

If we have access to the *base attributes* of the *novel data*, then we can apply ARM for further improvements. We experiment on PETA with 5-way 50-shot, and RAP with 8-way 80-shot where ARM can bring improvements in mA_{novel} from 67.16%/71.90% to 68.05%/72.88%.

5.7 Effect of Number of Ways and Shots

We compare our method against finetuning (lr=0.0001) with different number of shots and ways. The results in Figure 4 show that our method is able to consistently outperform the baseline on mA_{novel} . With different number of ways, our method is still able to generalize well to the novel attributes. However, the gap gets smaller as the number of shots increases, which implies that finetuning becomes a stronger baseline as the number of examples increases, and our method’s superiority lies mostly on the few-shot scenarios.

During the experiments, we also empirically find that sampling more fake novel attributes than the expected number of ways, will yield a mA_{novel} gain of $\sim 1\%$. This training strategy is firstly adopted in [Snell *et al.*, 2017], which can be explained intuitively that making the sampled task more difficult (higher ways) during meta training is beneficial to the future generalization in the actual test case.

6 Conclusion and Future Work

In this paper, we raise the practical issue of incremental few-shot learning for pedestrian attribute recognition, a commonly existing problem which has gained little attention by the academic community. We tackle this challenging task by proposing a meta learning based framework, which only requires annotations on the novel attributes and one feed forward time to produce the adapted model. We evaluate our method against finetuning and other state-of-the-art meta learning methods, and argue that our method is able to rapidly adapt to the novel few-shot data with low resource consumption and yield competitive performances. We believe that this task is far from fully solved, and future work may involve designing meta-learner with stronger representation capability such as graph neural networks.

Acknowledgements

The work is supported by National Key R&D Program of China (2018YFC0806900), National Natural Science Foundation of China (61571269) and National Basic Research Program of China (2015CB352300).

References

- [Bertinetto *et al.*, 2016] Luca Bertinetto, João F Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In *Advances in Neural Information Processing Systems*, pages 523–531, 2016.
- [Finn *et al.*, 2017] 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.
- [Gidaris and Komodakis, 2018] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4367–4375. IEEE, 2018.
- [Goodfellow *et al.*, 2013] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- [Jin *et al.*, 2018] Xiaoming Jin, Tao He, Cheng Wan, Lan Yi, Guiguang Ding, and Dou Shen. Automatic gating of attributes in deep structure. In *IJCAI*, pages 2305–2311, 2018.
- [Kirkpatrick *et al.*, 2017] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [Li and Hoiem, 2018] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018.
- [Li *et al.*, 2015] Dangwei Li, Xiaotang Chen, and Kaiqi Huang. Multi-attribute learning for pedestrian attribute recognition in surveillance scenarios. In *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 111–115. IEEE, 2015.
- [Li *et al.*, 2016] Dangwei Li, Zhang Zhang, Xiaotang Chen, Haibin Ling, and Kaiqi Huang. A richly annotated dataset for pedestrian attribute recognition. *arXiv preprint arXiv:1603.07054*, 2016.
- [Li *et al.*, 2017] Dangwei Li, Xiaotang Chen, Zhang Zhang, and Kaiqi Huang. Learning deep context-aware features over body and latent parts for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 384–393, 2017.
- [Liu *et al.*, 2017] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Shuai Yi, Junjie Yan, and Xiaogang Wang. Hydraplus-net: Attentive deep features for pedestrian analysis. In *Proceedings of the IEEE international conference on computer vision*, pages 350–359, 2017.
- [Prosser *et al.*, 2010] Bryan James Prosser, Wei-Shi Zheng, Shaogang Gong, Tao Xiang, and Q Mary. Person re-identification by support vector ranking. In *BMVC*, volume 2, page 6, 2010.
- [Qi *et al.*, 2018] 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.
- [Ravi and Larochelle, 2016] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- [Snell *et al.*, 2017] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [Sudowe *et al.*, 2015] Patrick Sudowe, Hannah Spitzer, and Bastian Leibe. Person attribute recognition with a jointly-trained holistic cnn model. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 87–95, 2015.
- [Sung *et al.*, 2018] 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.
- [Thrun and Pratt, 2012] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012.
- [Vinyals *et al.*, 2016] 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.
- [Wang *et al.*, 2017] Jingya Wang, Xiatian Zhu, Shaogang Gong, and Wei Li. Attribute recognition by joint recurrent learning of context and correlation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 531–540, 2017.
- [Xiang *et al.*, 2019] Liuyu Xiang, Xiaoming Jin, Lan Yi, and Guiguang Ding. Adaptive region embedding for text classification. In *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.
- [Zhao *et al.*, 2018] Xin Zhao, Liufang Sang, Guiguang Ding, Yuchen Guo, and Xiaoming Jin. Grouping attribute recognition for pedestrian with joint recurrent learning. In *IJCAI*, pages 3177–3183, 2018.
- [Zhu *et al.*, 2017] Jianqing Zhu, Shengcai Liao, Zhen Lei, and Stan Z Li. Multi-label convolutional neural network based pedestrian attribute classification. *Image and Vision Computing*, 58:224–229, 2017.