

# Prototypes-oriented Transductive Few-shot Learning with Conditional Transport

Long Tian<sup>1</sup>, Jingyi Feng<sup>1</sup>, Xiaoqiang Chai<sup>1</sup>, Wenchao Chen<sup>2,✉</sup>, Liming Wang<sup>1</sup>, Xiyang Liu<sup>1</sup>, Bo Chen<sup>2</sup>  
<sup>1</sup> National Pilot School of Software Engineering, <sup>2</sup> National Key Laboratory of Radar Signal Processing  
<sup>1,2</sup> Xidian University, Xi'an, Shanxi 710071, China

Email: {tianlong, wcchen}\_xidian@163.com

## Abstract

*Transductive Few-Shot Learning (TFSL) has recently attracted increasing attention since it typically outperforms its inductive peer by leveraging statistics of query samples. However, previous TFSL methods usually encode uniform prior that all the classes within query samples are equally likely, which is biased in imbalanced TFSL and causes severe performance degradation. Given this pivotal issue, in this work, we propose a novel Conditional Transport (CT) based imbalanced TFSL model called Prototypes-oriented Unbiased Transfer Model (PUTM) to fully exploit unbiased statistics of imbalanced query samples, which employs forward and backward navigators as transport matrices to balance the prior of query samples per class between uniform and adaptive data-driven distributions. For efficiently transferring statistics learned by CT, we further derive a closed form solution to refine prototypes based on MAP given the learned navigators. The above two steps of discovering and transferring unbiased statistics follow an iterative manner, formulating our EM-based solver. Experimental results on four standard benchmarks including miniImageNet, tiered-ImageNet, CUB, and CIFAR-FS demonstrate superiority of our model in class-imbalanced generalization<sup>1</sup>.*

## 1. Introduction

Deep learning based methods have gained a great success in many real-world applications, such as image classification [10, 25, 36], natural language processing [3, 32, 38], and so on, thanks to its powerful nonlinear representation ability [18, 24]. However, their outstanding performances heavily rely on large-scale annotated training data [6, 16]. When it comes to high cost of collecting and annotating a large amount of data, a major research effort is being dedicated to low data regimes, in which overfitting occurs in optimizing deep learning architectures. Few-shot Learning (FSL) tackles these challenges and has triggered substantial interests

<sup>1</sup>Our code is available at <https://github.com/RashLog/PUTM>

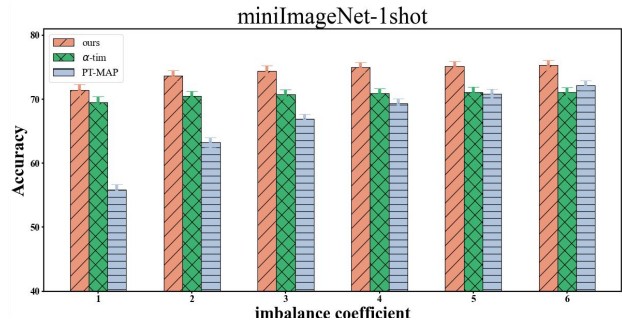


Figure 1: 5-way 1-shot image classification results with miniImageNet [22] dataset on 3000 tasks. Our model achieves the best accuracy in a wide range of imbalanced coefficients compared with PT-MAP [11] and  $\alpha$ -TIM [29].

within the community [8, 20], it aims to train a model from the base classes so as to generalize on the tasks sampled from the novel classes that were never seen during training. There are two major lines of works from the perspective of inference paradigms, namely inductive methods [26, 4] and their transductive counterparts [7, 23]. Both of them adopt well trained models for evaluating new tasks sampled from the novel classes. The former ones only utilize support samples in a supervised way while the later ones leverage both support and query samples in a semi-supervised manner.

Transductive inference is widely used in FSL since it can sufficiently discover statistics of query samples compared with its inductive peer. As a representative example, Hu *et al.* [11] propose PT-MAP for first employing Wide ResNet (WRN) [36] trained with S2M2 [19] on base classes to extract features of new tasks, and then adjusting prototypes from novel classes by aggregating similar query samples per class with the help of transport matrix derived from the Sinkhorn algorithm [5]. Setting marginal distributions of prototypes as uniform encodes strong prior that all the classes within query samples are equally likely, although, it works pretty well in perfectly class-balanced query samples. However, it is proven to be detrimental in class-imbalanced

applications where query samples come with arbitrary and unknown label marginals [21], and classification accuracy drops dramatically, as shown in Fig. 1. To narrow this gap, Veilleux *et al.* [29] propose  $\alpha$ -TIM for handling class-distribution variations effectively. Unfortunately, we have observed the phenomena as shown in Fig. 1 that  $\alpha$ -TIM can not achieve consistent improvement in various class-imbalanced levels, thus restricting its real-world applications. Most recently, other more complicated methods are also be arisen [14, 12]. Hu *et al.* [12] propose a clustering method for capturing the class-wise distributions properly.

To this end, for class-imbalanced TFSL, we develop a novel and simple Prototypes-oriented Unbiased Transfer Model (PUTM). It firstly constructs transferable class-wise prototypes iteratively for capturing the unbiased statistics within class-imbalanced query samples. Then it performs few-shot classification via the similarities between prototypes and query samples when the transfer arrives at convergence. The key to obtain unbiased transferable statistics is estimating class-wise marginal distributions precisely [14]. To achieve this, instead of utilizing doubly stochastic transport matrix constrained with fixed uniform-distributed marginal prior between prototypes and query samples as PT-MAP does, we employ task-adaptive forward and backward transport matrices with learnable marginal priors constrained by minimizing Conditional Transport (CT) distance [39] between prototypes and query samples. CT is a tool for measuring the cost in transporting the mass in one distribution to match another in an exchangeable but asymmetrical manner given a specific point-to-point cost function and an optimizable forward and backward transport matrices, also called navigators in the CT scope. To better calibrate prototypes with the help of learned transport matrices and query samples, we further derive a closed-form solution to refine prototypes based on MAP in an iterative manner. Decisions are finally made using the combination of forward and backward transport matrices when the iteration converges. Experiments on four benchmarks demonstrate that our proposed PUTM can learn task-adaptive prior in various class-imbalanced coefficients, paving a new way to transfer the unbiased statistics of query samples to prototypes.

Our contributions can be summarized as follows:

- We find that the performance of class-imbalanced FSL heavily relies on the accurate estimation of class-wise marginal distribution. Meanwhile, we are the first one to find that there exists a nature connection between CT theory and such marginal distribution estimation.
- We propose a novel and simple PUTM for class-imbalanced FSL based on connections mentioned above. It firstly embeds CT distance to measure the differences between prototypes and query samples for achieving better task-adaptive unbiased statistic migra-

tion elegantly. Then it calibrates prototypes using the transferred statistics in an iterative manner.

- We provide extensive experiments and comparisons on four standard benchmarks for verifying that our model achieves competitive results on both class-balanced and class-imbalanced (in a wide range of imbalanced coefficients) FSL settings.

## 2. Task Formulation

We divide the whole annotated dataset into two disjoint subsets: a base set  $\mathcal{D}_{base}$  with  $\mathcal{B}$  classes and a novel set  $\mathcal{D}_{novel}$  with  $\mathcal{N}$  classes. To build a common N-way K-shot task [30, 26], we first randomly sample  $\mathcal{N}$  classes from  $\mathcal{N}$  novel classes, and then pick  $K$  samples per class for the support set  $\mathcal{S} = \{\mathbf{x}_{i,j}^s, \mathbf{y}_{i,j}^s\}_{i=1, j=1}^{\mathcal{N}, K}$  to train or fine-tune the model. Finally,  $K'_i$  query samples being non-overlapped with support samples per class are randomly picked from the  $\mathcal{N}$  support classes for evaluating the performance of the model, formulating query set as  $\mathcal{Q} = \{\mathbf{x}_{i,j}^q, \mathbf{y}_{i,j}^q\}_{i=1, j=1}^{\mathcal{N}, K'_i}$ . For simplicity, we set  $M = \sum_{i=1}^{\mathcal{N}} K'_i$ .

**Class-imbalanced TFSL:** Veilleux *et al.* [29] figure out that the number of query samples per class should be arbitrary and unknown instead of being strictly equal. That is  $K'$  is a function of class and changes with class indexes. To describe such behavior, they develop a class marginal with Dirichlet distribution for generating query samples. Specifically, the class-wise marginal distribution of query samples per task  $\mathbf{p}$  are determined by:

$$f_{Dir}(\mathbf{p}; \mathbf{a}) = \frac{1}{B(\boldsymbol{\alpha})} \prod_{i=1}^{\mathcal{N}} p_i^{\alpha_i - 1} \quad (1)$$

where  $\mathbf{p} = (p_1, \dots, p_N)$  stands for class-wise distribution,  $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_N)$  denotes Dirichlet-distributed parameter controlling imbalanced ratio of query samples,  $B(\boldsymbol{\alpha}) = \frac{\prod_{i=1}^{\mathcal{N}} \Gamma(\alpha_i)}{\Gamma(\sum_{i=1}^{\mathcal{N}} \alpha_i)}$  represents multivariate Beta function, and  $\Gamma(\cdot)$  is Gamma function. Smaller  $\boldsymbol{\alpha}$  means higher imbalanced level, and vice versa. Taking  $\boldsymbol{\alpha} = 2$  as an example, total number of query samples is 75, then the number of query samples per class could be distributed as  $K' = [10, 5, 2, 35, 18]$ , which is more realistic compared with the class-balanced situation whose  $K'_i = 15, i = 1, \dots, N$ . Besides, it also poses some difficulties such as overfitting to the classes whose query samples are extremely few.

## 3. Our Model

### 3.1. Overall Method

In this work, we propose a novel model called PUTM for class-imbalanced transductive few-shot image classification. Considering the correlations between support and

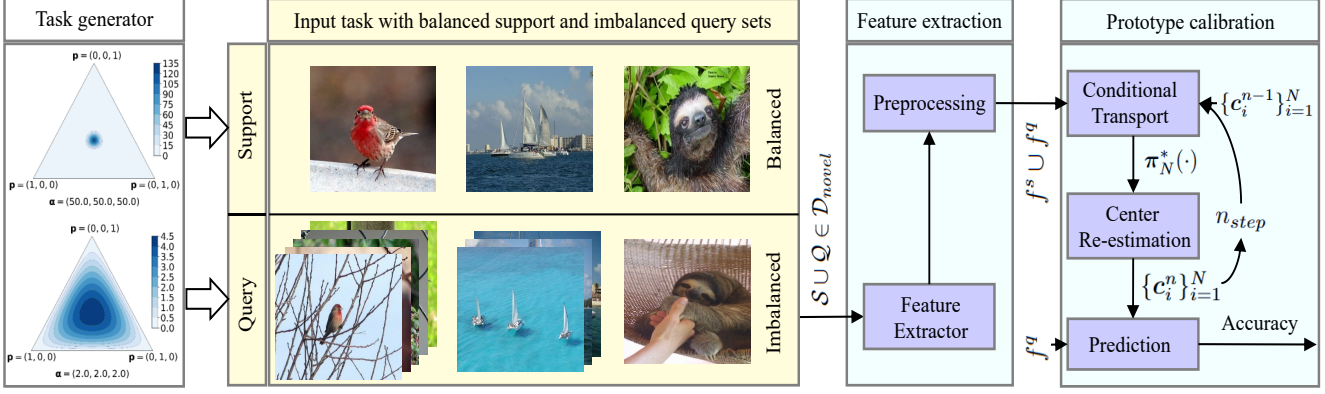


Figure 2: Illustration of our model. Tasks are generated from Dirichlet-distributed generator parameterized by  $\alpha$ . We first adopt a well trained feature extractor [19] and a channel-wise feature adaptation [35] to obtain the input features. Transductive inference is then conducted on the pre-processed features by: i) optimizing the transport matrices using CT; ii) refining the class-wise prototypes via MAP statistics aggregation; iii) making predictions after the above two steps iterate until convergence.

query samples, it is reasonable to refine the representations of class-wise prototypes initialized by support samples with the statistics of query samples. At this moment, the key is how to transfer the unbiased statistics from query samples to prototypes for achieving the best calibration. Moreover, the class-imbalanced setting is more challenging due to its unknown and arbitrary prior distribution of query samples. To address this challenge, we pioneered to depict the important but unknown prior with well-defined CT theory. Specifically, the prior could be measured by the transport mass from query samples to each prototype by summing the combination of forward and backward transport matrices optimized by CT loss along every query sample. Given the well-optimized prior, a prototype refinement algorithm is easy to be derived according to the MAP criterion. Considering the fact that the discriminative prototypes and well-optimized prior are coupled, we further employ an EM-solver to transfer the unbiased statistics efficiently and precisely. The overview of the proposed method can be found in Fig. 2.

### 3.2. CT Theory

CT [39] is a powerful tool developed for measuring distance between any two probability distributions most recently, the effectiveness of which is mostly verified on deep generative models [13, 27] and multi-label classification [15]. In this work, we expand the territory of CT for discrete distributions under class-imbalanced TFSL. Denote  $p(x) = \sum_{i=1}^n a_i \delta_{x_i}$  and  $q(y) = \sum_{j=1}^m b_j \delta_{y_j}$  as source and target discrete distributions with  $n$  and  $m$  points, respectively.  $\mathbf{a} \in \Delta^n$  and  $\mathbf{b} \in \Delta^m$ , where  $\Delta^n$  and  $\Delta^m$  separately denote the probability simplex of  $\mathbb{R}^n$  and  $\mathbb{R}^m$ . The CT distance between  $p(x)$  and  $q(y)$  can be expressed as a combination

of forward CT  $\mathcal{C}_\phi(x \rightarrow y)$  and a backward CT  $\mathcal{C}_\phi(x \leftarrow y)$ :

$$\mathcal{C}_{\phi, \rho}(p, q) = \min_{\phi} \sum_{x \in p, y \in q} \rho \mathcal{C}_\phi(x \rightarrow y) + (1 - \rho) \mathcal{C}_\phi(x \leftarrow y) \quad (2)$$

where  $\phi$  is parameters of CT,  $\rho$  is the trade-off factor controlling the strengths between the forward and backward CTs. To better understand how the CT works when minimizing Eq. 2, we take the forward CT as an example:

$$\mathcal{C}_\phi(x \rightarrow y) = \mathbb{E}_{y_{1:m} \sim q} \mathbb{E}_{x \sim p} [\sum_{j=1}^m c(x, y_j) \hat{\pi}_m(y_j | x, \phi)] \quad (3)$$

where  $c(\cdot, \cdot)$  could be some well defined distance functions, such as Euclidean or Cosine,  $\hat{\pi}_m(y_j | x, \phi)$  is forward transport matrix, also called forward navigator in the scope of CT, satisfying  $\sum_{j=1}^m \hat{\pi}_m(y_j | x, \phi) = 1$ , it describes the possibility of transporting  $x$  to  $y_j$  and can be formulated as:

$$\hat{\pi}_m(y_j | x, \phi) = \frac{e^{-d_\phi(x, y_j)}}{\sum_{j'} e^{-d_\phi(x, y_{j'})}} \quad (4)$$

where  $d_\phi(\cdot, \cdot)$  represents measurement function, such as MLP implemented by Zheng *et al.* [39]. The forward CT can be interpreted as the expected cost to stochastically transport a random source point  $x$  to one of the  $m$  randomly instantiated "anchors" of the target distribution. By minimizing Eq. 3, an unbiased statistical estimation of the two distributions arrives. Proofs can be found in Zheng *et al.* [39]. Hence, it is reasonable to transfer statistics of one distribution to another.

### 3.3. PUTM

Our proposed PUTM aims to exploit discriminative statistics from query samples for reducing the gap between the estimated prototypes and their ground truth prototypes. However, the ground truth prototypes are always intractable in

practice. A reasonable approach to obtain unbiased prototypes is to aggregate discriminative statistics from both support and query samples following a semi-supervised manner [37]. Different from the standard semi-supervised problems, the query samples are usually limited per class, especially in class-imbalanced case. Follow up, we present the importance of unbiased transfer in semi-supervised FSL setting, the algorithm of realizing unbiased exploitation with CT and transductive inference with unbiased transfer in turn.

### 3.3.1 Unbiased Statistics Transfer Matters in TFSL

Given a few-shot task, unbiased prototypes arrives if and only if the aggregated query samples are noiseless, which should heavily rely on the unbiased statistics transfer. Leveraging similarities between query samples and prototypes is a common way to realize statistics transfer in transductive few shot image classification. Given the transport matrix  $\pi^n \in \mathbb{R}^{M \times N}$  at the  $n$ -th iteration, it measures the closeness of prototypes  $\{c_i^n\}_{i=1}^N$  and query samples  $\{x_{i,j}^q\}_{j=1}^{K'_i}$ . For simplicity, we rewrite all query samples as  $\{x_r^q\}_{r=1}^M$ . Thus, the refined prototypes  $\{c_i^{n+1}\}_{i=1}^N$  can be expressed as:

$$c_i^{n+1} = \frac{\sum_{r=1}^M \pi_{r,i}^n x_r^q + \sum_{j=1}^K x_{i,j}^s}{\sum_{r'=1}^M \pi_{r',i}^n + K}, \quad i = 1, \dots, N \quad (5)$$

where  $\pi_{r,i}^n = \frac{e^{-d_\phi(x_r^q, c_i^n)}}{\sum_{r'=1}^M e^{-d_\phi(x_{r'}^q, c_i^n)}}$ ,  $d_\phi(\cdot, \cdot)$  denotes the well defined measurement function, and we denote optimal transport matrix as  $\pi^*$ . To guarantee the unbiased refinement of prototypes,  $\pi^*$  need to be precise, which is usually difficult to be satisfied in practice. Fortunately, partial noisy is tolerable. For example, given 15 query samples per class, a few wrong transferring such as 1 or 2 samples is tolerable in calibrating prototypes. As for class-imbalanced case, a tiny mistake may result in disastrous consequence on prototypes calibration. For the classes having only 1 or 2 query samples, the tolerances of error transmissions should be extremely low. Therefore, to improve the robustness of calibrated prototypes in class-imbalanced TFSL, transferring unbiased statistics of query samples with  $\pi^*$  becomes crucial.

There are mainly two lines of works in defining  $\pi^*$ : i) Liu *et al.* [17] employ Cosine function [17]; ii) Hu *et al.* [11] use Optimal Transport (OT). Cosine function may not powerful enough to fully reflect the relationships between prototypes and query samples since it ignores various importance between each two points. OT describes a doubly stochastic transport matrix such that  $\Pi(p, q) = \{\pi^* | \sum_{r=1}^M \pi_{r,i}^* = a_i, \sum_{i=1}^N \pi_{r,i}^* = b_r\}$ . Unfortunately, we usually do not have any priors on distributions of  $p$  and  $q$ , thus, we usually have  $a_i = \frac{1}{N}, i = 1, \dots, N$  and  $b_r = \frac{1}{M}, r = 1, \dots, M$ , which should be contradictory to class-imbalanced setting. For example, once we employ Dirichlet distribution with  $\alpha = \mathbf{2}_N$  to mimic the generation of query samples,  $a_i = \frac{1}{N}$  will never

fit the data. According to the analysis of CT, it replaces the uniform constraint on prior distribution of classes derived from query samples with a data-driven manner, thus ensuring to transfer the unbiased statistics more flexible and efficient.

### 3.3.2 Unbiased Statistics Exploitation with CT

As we can see, CT has an appealing property of exploiting unbiased statistics thanks to its task-adaptive prior rather than uniform-distributed one as OT. Here after, we present details of achieving such unbiased exploitation with CT.

As shown in Fig. 2, given a class-imbalanced few-shot task in Sec. 2, we first employ a well-trained feature extractor  $f_{\theta^*}(\cdot)$  to get features of support and query samples as  $f_{i,j}^s = f_{\theta^*}(x_{i,j}^s)$  and  $f_r^q = f_{\theta^*}(x_r^q)$ , where  $\theta^*$  is the set of optimized parameters, being fixed after training on the base classes following Mangla *et al.* [19]. Then, we conduct channel attention with preprocessing developed by Yang *et al.* [35]. Finally, unbiased statistics exploitation is realized by optimizing transport matrices constrained by CT theory.

Specifically, the initial prototypes are averaged on the features of support set as  $c_i^0 = \frac{1}{K} \sum_{j=1}^K f_{i,j}^s, i = 1, \dots, N$ , where the superscript of  $\{c_i^0\}_{i=1}^N$  describes the generation of prototypes is zero. They will be updated subsequently, detailed in Sec. 3.3.3. At this moment, we treat  $\{c_i^0\}_{i=1}^N$  and  $\{f_r^q\}_{r=1}^M$  as samples from the two discrete distributions, and utilizing CT to learn the transport matrices for exploiting unbiased statistics. According to the analysis of Eq. 2, an unbiased sample estimate of CT given query features  $\{f_r^q\}_{r=1}^M$  and prototypes  $\{c_i^0\}_{i=1}^N$  can be expressed as:

$$\begin{aligned} \mathcal{L}_{\phi, \theta^*, \rho}(f_{1:M}^q, c_{1:N}^0) &= \sum_{i=1}^N \sum_{r=1}^M c(c_i^0, f_r^q) \\ &= \left( \frac{\rho}{M} \hat{\pi}_N(c_i^0 | f_r^q, \phi) + \frac{1-\rho}{N} \hat{\pi}_M(f_r^q | c_i^0, \phi) \right) \end{aligned} \quad (6)$$

where  $\hat{\pi}_N$  and  $\hat{\pi}_M$  separately denote the forward and backward transport matrices. Again substituting query features  $\{f_r^q\}_{r=1}^M$  and prototypes  $\{c_i^0\}_{i=1}^N$  into Eq. 4, we have:

$$\hat{\pi}_N(c_i^0 | f_r^q, \phi) = \frac{e^{-d_\phi(f_r^q, c_i^0)}}{\sum_{i'=1}^N e^{-d_\phi(f_r^q, c_{i'}^0)}} \quad (7)$$

$$\hat{\pi}_M(f_r^q | c_i^0, \phi) = \frac{e^{-d_\phi(f_r^q, c_i^0)}}{\sum_{r'=1}^M e^{-d_\phi(f_{r'}^q, c_i^0)}} \quad (8)$$

The forward transport matrix represents the probability of query samples belonging to the corresponding support classes. The two transport matrices with the optimal parameter  $\phi^*$  can be obtained by minimizing the approximated CT estimation of Eq. 6. Once  $\hat{\pi}_N^*$  and  $\hat{\pi}_M^*$  with the optimal  $\phi^*$  are obtained, the transport matrix described in Sec. 3.3.1 for unbiased statistics exploitation can be naturally realized. To better understand this, we depict the mechanisms of the

---

**Algorithm 1** Work flow of PUTM for one task.
 

---

**Input:** Support samples  $\{\mathbf{x}_{i,j}^s\}_{i,j=1}^{N,K}$  and class-imbalanced query samples  $\{\mathbf{x}_r^q\}_{r=1}^M$  generated from Dirichlet distribution with parameter  $\alpha$ , a well-trained feature extractor  $f_{\theta^*}(\cdot)$ ;

**Parameters:** Distance measurement function  $d_\phi(\cdot, \cdot)$ ;

**Initialization** Hyper-parameters  $\alpha, \rho, \lambda$ , and learning schedule for updating  $\phi$ ; Initialize prototypes  $\mathbf{c}_i^0 = \frac{1}{K} \sum_{j=1}^K f_{i,j}^s, i = 1, \dots, N$ ;

**for**  $n=1: n_{steps}$  **do**

  Initialize parameters  $\phi$ ;

$\phi^* \leftarrow \operatorname{argmin}_{\phi} \mathcal{L}_{\phi, \theta^*, \rho}(f_{1:M}^q, \mathbf{c}_{1:N}^n)$ ;

  Refining prototypes  $\mathbf{c}_{1:N}^{n+1}$  using Eq. 11 based on  $\mathbf{c}_{1:N}^n$ ;

**end for**

Evaluating performance using Eq. 12;

**Output** Parameters  $\{\phi^*, \mathbf{c}_{1:N}^{n_{steps}}\}$ .

---

transport matrices developed above from three perspectives: 1) the distance function in transport matrices is implemented with DNNs, hence, they are able to learn a comprehensive distance metric beyond a simple Cosine function; 2) the CT distance is proven to be unbiased given samples of two distributions  $p$  and  $q$  in [39]; 3) The forward transport matrix holds free for sum along query samples given a specific prototype, namely,  $\sum_{r=1}^M \hat{\pi}_N(\mathbf{c}_i^0 | f_r^q, \phi^*) = C$ , and  $C$  is not restricted to be equal for each class, thus embedding task-adaptive prior instead of uniform distribution as PT-MAP [11] does. On the other hand, the backward transport matrix holds tight for sum along query samples given a specific prototype, that is,  $\sum_{r=1}^M \hat{\pi}_M(f_r^q | \mathbf{c}_i^0, \phi^*) = 1$ , hence, uniform-distributed prior holds, and it is proven to be underfitting in class-imbalanced case [29]. In practice, we employ a trade-off to balance the strengths of the two transport matrices by summing them together with a soft coefficient  $\rho$ . For simplicity, we only use the forward transport matrix in the rest of the work.

### 3.3.3 Unbiased Statistics Transfer with EM-Solver

Once the forward transport matrix  $\hat{\pi}_N^*$  is obtained by minimizing Eq. 6, it is straightforward to derive the refined prototypes by substituting  $\hat{\pi}_N^*$  into Eq. 5:

$$\mathbf{c}_i' = \frac{\sum_{r=1}^M \hat{\pi}_{N,r,i}^* f_r^q + \sum_{j=1}^K f_{i,j}^s}{K + \sum_{r=1}^M \hat{\pi}_{N,r,i}^*}, \quad i = 1, \dots, N \quad (9)$$

It is worth noting that once the optimal parameters  $\phi^*$  are arrived, the only parameters left for optimizing are prototypes. Recalling our objective in Eq. 6, it is easy to find that the optimal prototypes at this moment could also be derived by  $\frac{\partial \mathcal{L}_{\phi^*, \theta^*, \rho}}{\partial \mathbf{c}_{1:N}^0}$ , thereby Eq. 9 could be derived. Instead of directly calibrating class-wise prototypes, We employ an inertia parameter  $\alpha$  to update prototypes proportionally:

$$\mathbf{c}_i^1 = \mathbf{c}_i^0 + \alpha(\mathbf{c}_i' - \mathbf{c}_i^0), \quad i = 1, \dots, N \quad (10)$$

Right now, recursive expression emerges for updating prototypes according to Eq. 10:

$$\mathbf{c}_i^{n+1} = \mathbf{c}_i^n + \alpha(\mathbf{c}_i' - \mathbf{c}_i^n), \quad i = 1, \dots, N \quad (11)$$

As we can see: i) updating the forward transport matrix  $\hat{\pi}_N^*$  and prototypes  $\mathbf{c}_{1:N}^n$  are coupled; ii)  $\hat{\pi}_N^*$  is a random variable determined by  $\phi$  while  $\mathbf{c}_{1:N}^n$  is an unknown but deterministic variable. These two facts meet conditions of implementing EM optimization, and we formulate our EM-Solver as follows: i) E Step, updating  $\hat{\pi}_N^*$  by minimizing Eq. 6 while keeping  $\mathbf{c}_{1:N}^n$  fixed; ii) M Step, updating  $\mathbf{c}_{1:N}^n$  by minimizing Eq. 6 while keeping  $\hat{\pi}_N^*$  fixed. Besides, it is worth noting that at each iteration,  $\phi^*$  need to be re-initialized. The pseudo code is provided in Alg. 1.

After  $n_{step}$  loops between estimating  $\hat{\pi}_N^*$  and refining  $\mathbf{c}_{1:N}^n$ , the decision per sample is finally made by:

$$y_r^q = \max_{i \in [1:N]} \hat{\pi}_N(\mathbf{c}_i^{n_{step}} | f_r^q, \phi^*), \quad r = 1, \dots, M \quad (12)$$

Statistic results of average accuracy are usually conducted using query samples on various of different tasks.

## 3.4. Complexity Analysis

$\phi$  is implemented with MLP in our work. Taking a two-layers MLP for example, namely,  $d_f \rightarrow d_1^\phi \rightarrow d_2^\phi$ . For  $T$  tasks, computational complexity originally takes  $\mathcal{O}(T \times \max\{NM \times d_f \times d_1^\phi, NM \times d_1^\phi \times d_2^\phi\})$ . In practice,  $T$  is usually set to be large enough for reliable results, thus resulting in time-consuming evaluation of our model. To mitigate the real-time issue, we trade space for time by paralleling our proposed PUTM among all tasks. Therefore, the computational complexity is reduced to  $\mathcal{O}(\max\{NM \times d_f \times d_1^\phi, NM \times d_1^\phi \times d_2^\phi\})$ . Besides, we can further control the size of MLP, that is  $\{d_i^\phi\}_{i=1}^2$ , to achieve a trade-off between accuracy and real-time efficiency.

## 4. Experiments

### 4.1. Experimental Setup

**Baselines:** We mainly compare our model with some recently proposed State-Of-The-Art (SOTA) methods on one Nvidia RTX 3090 GPU. For inductive methods, we adopt Baseline/Baseline(++) [4], ProtoNet [26], SimpleShot [33], and HOT [9] in inductive. For transductive methods, we adopt PT-MAP [11], LaplacianShot [40], TIM[2], BD-CSPN [17],  $\alpha$ -TIM [29], iLPC [14], and BAVAR [12].

**Datasets:** We evaluate our model on four benchmarks including miniImageNet [22], tieredImageNet [23], CUB [31], and CIFAR-FS [1]. MiniImageNet is randomly chosen from ILSVRC-12 dataset [6] and it contains 100 classes with 600 images per class, the resolution of which is  $84 \times 84 \times 3$ . Following previous works [22, 29, 11, 35], the dataset is splitted into 64 base classes, 16 validation classes, and



Table 1: Classification accuracy (%) of SOTA methods on class-imbalanced and class-balanced settings with different datasets and backbones. For imbalanced case, query samples are generated following a Dirichlet distribution with  $\alpha = \mathbf{2}_N$ . Results are averaged over 3000 tasks with 95% confidence intervals. Results with superscript † are our re-implementations based on the code provided by Veilleux *et al.* [29]. The best and the second best scores are marked in bold and blue, respectively.

Setting	Method	mini-ImageNet (WRN/WRN <sup>§</sup> )			CUB (RN18/WRN <sup>§</sup> )			tiered-ImageNet (WRN/WRN <sup>§</sup> )			CIFAR-FS (WRN <sup>§</sup> )			
		1-shot	3-shot	5-shot	1-shot	3-shot	5-shot	1-shot	3-shot	5-shot	1-shot	3-shot	5-shot	
Induct. (%)	Imbalanced & Balanced	Baseline (ICLR'19)	62.2/62.8 <sup>†</sup>	76.9 <sup>†</sup> /77.3 <sup>†</sup>	81.9/82.9 <sup>†</sup>	64.6/70.7 <sup>†</sup>	82.3 <sup>†</sup> /87.3 <sup>†</sup>	86.9/90.0 <sup>†</sup>	64.6/63.6 <sup>†</sup>	80.0 <sup>†</sup> /83.0 <sup>†</sup>	84.9/87.2 <sup>†</sup>	69.5 <sup>†</sup>	83.9 <sup>†</sup>	86.7 <sup>†</sup>
		Baseline++ (ICLR'19)	64.5/64.8 <sup>†</sup>	77.5 <sup>†</sup> /78.4 <sup>†</sup>	82.1/82.7 <sup>†</sup>	69.4/81.7 <sup>†</sup>	83.3 <sup>†</sup> /89.4 <sup>†</sup>	87.1/91.2 <sup>†</sup>	68.7/66.9 <sup>†</sup>	80.9 <sup>†</sup> /78.7 <sup>†</sup>	85.4/82.2 <sup>†</sup>	74.9 <sup>†</sup>	85.0 <sup>†</sup>	87.4 <sup>†</sup>
		ProtoNet (NeurIPS'17)	61.2/62.5 <sup>†</sup>	76.9 <sup>†</sup> /77.2 <sup>†</sup>	81.1/81.4 <sup>†</sup>	63.1 <sup>†</sup> /69.6 <sup>†</sup>	82.3 <sup>†</sup> /85.6 <sup>†</sup>	85.4 <sup>†</sup> /88.8 <sup>†</sup>	63.5 <sup>†</sup> /64.2 <sup>†</sup>	79.9 <sup>†</sup> /80.1 <sup>†</sup>	84.1 <sup>†</sup> /84.0 <sup>†</sup>	69.5 <sup>†</sup>	83.8 <sup>†</sup>	86.7 <sup>†</sup>
		SimpleShot (ARXIV)	66.2/67.1 <sup>†</sup>	77.6 <sup>†</sup> /78.2 <sup>†</sup>	82.4/83.3 <sup>†</sup>	70.6/81.7 <sup>†</sup>	83.6 <sup>†</sup> /89.5 <sup>†</sup>	87.5/91.4 <sup>†</sup>	70.7/71.4 <sup>†</sup>	81.3 <sup>†</sup> /84.2 <sup>†</sup>	85.9/87.5 <sup>†</sup>	75.2 <sup>†</sup>	84.7 <sup>†</sup>	87.4 <sup>†</sup>
		HOT (NeurIPS'22)	66.3 <sup>†</sup> /69.1	78.5 <sup>†</sup> /79.8 <sup>†</sup>	82.8 <sup>†</sup> /84.4	74.5 <sup>†</sup> /81.2	85.4 <sup>†</sup> /88.2 <sup>†</sup>	87.8 <sup>†</sup> /91.5	74.1 <sup>†</sup> /75.9	82.6 <sup>†</sup> /83.9 <sup>†</sup>	85.4 <sup>†</sup> /87.3	75.4	85.3 <sup>†</sup>	87.5
		PT-MAP (ARXIV)	60.6/63.2 <sup>†</sup>	65.9 <sup>†</sup> /68.0 <sup>†</sup>	66.8/69.3 <sup>†</sup>	65.1/69.8 <sup>†</sup>	71.2 <sup>†</sup> /73.2 <sup>†</sup>	71.3/74.3 <sup>†</sup>	65.1/67.8 <sup>†</sup>	69.1 <sup>†</sup> /72.0 <sup>†</sup>	71.0/73.2 <sup>†</sup>	65.9 <sup>†</sup>	69.3 <sup>†</sup>	70.3 <sup>†</sup>
Imbalanced	LaplacianShot (ICML'20)	68.1/69.8 <sup>†</sup>	78.8 <sup>†</sup> /79.2 <sup>†</sup>	83.2/83.6 <sup>†</sup>	73.7/84.2 <sup>†</sup>	85.7 <sup>†</sup> /90.4 <sup>†</sup>	87.7/91.4 <sup>†</sup>	73.5/72.6 <sup>†</sup>	82.6 <sup>†</sup> /85.0 <sup>†</sup>	86.8/87.7 <sup>†</sup>	78.3 <sup>†</sup>	85.1 <sup>†</sup>	87.3 <sup>†</sup>	
	TIM (NeurIPS'20)	69.8/70.9 <sup>†</sup>	79.7 <sup>†</sup> /80.7 <sup>†</sup>	81.6/83.2 <sup>†</sup>	74.8/82.8 <sup>†</sup>	84.8 <sup>†</sup> /89.1 <sup>†</sup>	86.9/90.5 <sup>†</sup>	75.8/77.8 <sup>†</sup>	82.5 <sup>†</sup> /86.3 <sup>†</sup>	85.4/88.3 <sup>†</sup>	77.0 <sup>†</sup>	84.6 <sup>†</sup>	86.2 <sup>†</sup>	
	BD-CSPN (ECCV'20)	70.4/71.2 <sup>†</sup>	77.1 <sup>†</sup> /78.6 <sup>†</sup>	82.3/83.4 <sup>†</sup>	74.5/85.1 <sup>†</sup>	85.7 <sup>†</sup> /90.5 <sup>†</sup>	87.1/91.5 <sup>†</sup>	75.4/76.3 <sup>†</sup>	82.5 <sup>†</sup> /85.0 <sup>†</sup>	85.9/87.2 <sup>†</sup>	79.1 <sup>†</sup>	86.0 <sup>†</sup>	87.3 <sup>†</sup>	
	$\alpha$ -TIM (NeurIPS'2021)	69.8/70.9 <sup>†</sup>	80.4 <sup>†</sup> /81.1 <sup>†</sup>	<b>84.8/84.9<sup>†</sup></b>	75.7/83.6 <sup>†</sup>	87.3 <sup>†</sup> /91.2 <sup>†</sup>	<b>89.8/92.8<sup>†</sup></b>	76.0/77.7 <sup>†</sup>	<b>83.6<sup>†</sup>/86.6<sup>†</sup></b>	<b>87.8/89.9<sup>†</sup></b>	78.1 <sup>†</sup>	<b>87.2<sup>†</sup></b>	<b>88.9<sup>†</sup></b>	
	iLPC (ICCV'21)	<b>71.5<sup>†</sup>/74.3<sup>†</sup></b>	78.6 <sup>†</sup> /81.1 <sup>†</sup>	81.0 <sup>†</sup> /83.3 <sup>†</sup>	78.7 <sup>†</sup> /86.4 <sup>†</sup>	85.6 <sup>†</sup> /90.1 <sup>†</sup>	86.9 <sup>†</sup> /90.8 <sup>†</sup>	<b>76.6<sup>†</sup>/81.6<sup>†</sup></b>	83.2 <sup>†</sup> /87.5 <sup>†</sup>	85.1 <sup>†</sup> /89.0 <sup>†</sup>	80.8 <sup>†</sup>	85.4 <sup>†</sup>	86.5 <sup>†</sup>	
	BAVAR (AISTATS'23)	<b>74.1/74.2<sup>†</sup></b>	<b>82.6<sup>†</sup>/82.6<sup>†</sup></b>	<b>85.5/85.6<sup>†</sup></b>	<b>82.0/85.7<sup>†</sup></b>	<b>89.2<sup>†</sup>/89.8<sup>†</sup></b>	<b>90.7/90.8<sup>†</sup></b>	<b>77.5/80.6<sup>†</sup></b>	<b>84.8<sup>†</sup>/87.2<sup>†</sup></b>	<b>87.5/89.3<sup>†</sup></b>	<b>81.8<sup>†</sup></b>	87.0 <sup>†</sup>	87.9 <sup>†</sup>	
Transduct. (%)	PUTM (Ours)	70.9/73.8	<b>80.5/82.9</b>	84.6/85.7	<b>78.9/86.9</b>	<b>87.8/91.3</b>	89.3/92.2	75.0/81.1	<b>83.6/88.4</b>	86.0/90.4	<b>81.4</b>	<b>87.3</b>	<b>88.6</b>	
		PT-MAP (ARXIV)	<b>78.9/82.5<sup>†</sup></b>	<b>86.3<sup>†</sup>/87.5<sup>†</sup></b>	86.6/88.6 <sup>†</sup>	<b>85.5/92.0<sup>†</sup></b>	<b>90.5<sup>†</sup>/94.0<sup>†</sup></b>	<b>91.3/94.3<sup>†</sup></b>	<b>84.6/88.1<sup>†</sup></b>	<b>88.2<sup>†</sup>/91.8<sup>†</sup></b>	<b>90.0/92.5<sup>†</sup></b>	<b>87.6<sup>†</sup></b>	<b>90.5<sup>†</sup></b>	<b>90.9<sup>†</sup></b>
		LaplacianShot (ICML'20)	72.9/73.4 <sup>†</sup>	79.0 <sup>†</sup> /80.5 <sup>†</sup>	83.8/84.2 <sup>†</sup>	78.9/87.5 <sup>†</sup>	86.7 <sup>†</sup> /90.4 <sup>†</sup>	88.8/92.3 <sup>†</sup>	78.8/78.7 <sup>†</sup>	84.6 <sup>†</sup> /85.6 <sup>†</sup>	87.3/87.0 <sup>†</sup>	82.1 <sup>†</sup>	86.3 <sup>†</sup>	86.8 <sup>†</sup>
		TIM (NeurIPS'20)	74.6/77.6 <sup>†</sup>	85.5 <sup>†</sup> /86.0 <sup>†</sup>	85.9/88.0 <sup>†</sup>	80.3/88.1 <sup>†</sup>	89.8 <sup>†</sup> /93.0 <sup>†</sup>	<b>90.5/94.0<sup>†</sup></b>	80.3/83.8 <sup>†</sup>	<b>87.7<sup>†</sup>/90.6<sup>†</sup></b>	88.9/92.1 <sup>†</sup>	83.0 <sup>†</sup>	89.5 <sup>†</sup>	<b>90.5<sup>†</sup></b>
		BD-CSPN (ECCV'20)	72.5/73.8 <sup>†</sup>	79.2 <sup>†</sup> /80.7 <sup>†</sup>	83.7/85.3 <sup>†</sup>	77.9/87.3 <sup>†</sup>	86.9 <sup>†</sup> /91.4 <sup>†</sup>	88.9/92.4 <sup>†</sup>	77.7/78.7 <sup>†</sup>	84.5 <sup>†</sup> /86.4 <sup>†</sup>	87.4/88.3 <sup>†</sup>	81.6 <sup>†</sup>	87.4 <sup>†</sup>	88.6 <sup>†</sup>
		$\alpha$ -TIM (NeurIPS'2021)	71.0/71.5 <sup>†</sup>	76.8 <sup>†</sup> /77.7 <sup>†</sup>	82.4/83.3 <sup>†</sup>	77.2 <sup>†</sup> /84.9 <sup>†</sup>	85.1 <sup>†</sup> /90.2 <sup>†</sup>	87.6 <sup>†</sup> /92.1 <sup>†</sup>	75.1 <sup>†</sup> /78.9 <sup>†</sup>	80.6 <sup>†</sup> /84.4 <sup>†</sup>	85.3 <sup>†</sup> /88.2 <sup>†</sup>	79.7 <sup>†</sup>	85.7 <sup>†</sup>	87.9 <sup>†</sup>
Balanced	iLPC (ICCV'21)	75.5 <sup>†</sup> /77.8 <sup>†</sup>	82.0 <sup>†</sup> /83.9 <sup>†</sup>	83.9 <sup>†</sup> /85.7 <sup>†</sup>	82.4 <sup>†</sup> /89.6 <sup>†</sup>	87.9 <sup>†</sup> /92.3 <sup>†</sup>	88.9 <sup>†</sup> /93.0 <sup>†</sup>	81.0 <sup>†</sup> /85.5 <sup>†</sup>	86.4 <sup>†</sup> /89.8 <sup>†</sup>	87.5 <sup>†</sup> /90.9 <sup>†</sup>	84.8 <sup>†</sup>	88.6 <sup>†</sup>	89.2 <sup>†</sup>	
	BAVAR (AISTATS'23)	<b>78.5/78.5<sup>†</sup></b>	84.8 <sup>†</sup> /84.9 <sup>†</sup>	<b>87.4/87.7<sup>†</sup></b>	<b>85.6/89.0<sup>†</sup></b>	<b>90.8<sup>†</sup>/91.7<sup>†</sup></b>	<b>91.4/92.1<sup>†</sup></b>	81.5/84.7 <sup>†</sup>	87.4 <sup>†</sup> /89.3 <sup>†</sup>	88.3/90.3 <sup>†</sup>	85.0 <sup>†</sup>	88.4 <sup>†</sup>	88.9 <sup>†</sup>	
	PUTM(Ours)	78.4/81.4	<b>84.9/86.7</b>	86.5/88.1	84.9/91.2	89.2/93.3	90.1/93.7	<b>81.6/86.8</b>	86.9/90.7	<b>88.6/92.2</b>	<b>86.4</b>	<b>89.5</b>	90.4	

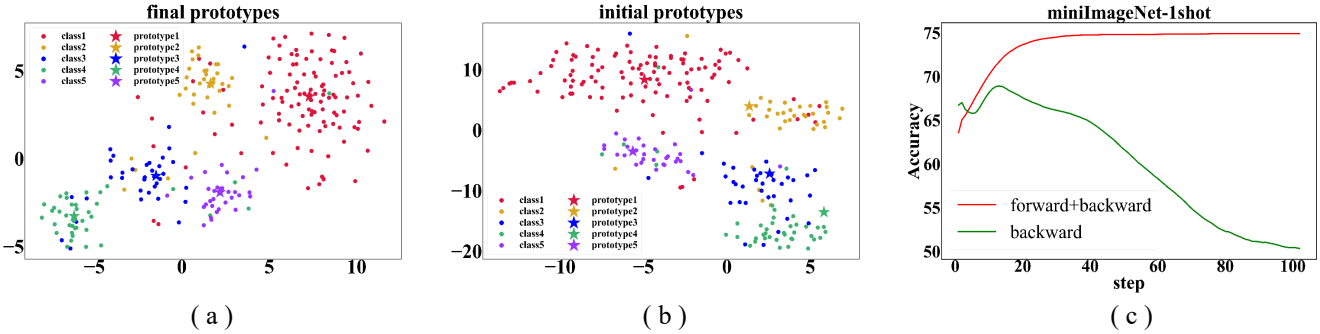


Figure 3: (a)-(b): The t-SNE [28] visualizations before (middle) and after (left) calibration on a 5-way 1-shot miniImageNet task. Pentagrams are prototypes while solid dots are query features. Classes change and best view in colors. (c): Impact of forward and backward transport matrices in calibrating prototypes on miniImageNet over 3000 5-way 1-shot tasks.

20 novel classes. TieredImageNet [23] is also a subset of ILSVRC-12 dataset, which is larger than miniImageNet. It contains 608 classes sampled from a hierarchical category structure with around 1281 images per class. Following the previous work [23], we split the dataset into 351 base classes, 97 validation classes, and 160 novel classes. The resolution is also  $84 \times 84 \times 3$ . CUB is a fine-grained dataset, containing 200 classes of birds with around 60 images per class. And the resolution is still  $84 \times 84 \times 3$ . We adopt 100 base classes, 50 validation classes, and 50 novel classes following [4]. CIFAR-FS consists of all 100 classes from CIFAR-100 [34], whose classes are randomly split into 64 base classes, 16 classes, and 20 novel classes. Each class contains 600 images of size  $32 \times 32 \times 3$ .

**Implementation details:** We use pre-trained backbones for extracting visual features of four benchmarks, as shown in Table. 1. WRN and RN18 are separately WideResNet-28-

10 and ResNet-18 pre-trained with standard classification loss [29]. WRN<sup>§</sup> is WideResNet-28-10 pre-trained with self-supervised loss [19]. Besides, we employ log transform, a channel-wise feature adaptation, introduced in Yang *et al.* [35] to pre-process the pre-extracted features. The distance function  $\phi$  is instantiated as a two-layers MLP activated with LeakyReLU whose architecture is  $[d_f, 128, 64, 1]$ . As for hyper-parameters, we mainly set  $\beta = 0.5$ ,  $\alpha = 0.1$ ,  $\rho = 0.2$ . The reported results are averaged over 3000 tasks.

## 4.2. Main Results

**Quantitative results:** The classification results on mini-ImageNet, tieredImageNet, CUB, and CIFAR-FS are reported in Table 1, according to which five-fold observations can be drawn: 1) inductive methods are not bothered by class-imbalanced problem since they do not leverage transferable statistics of query samples; 2) most transductive methods

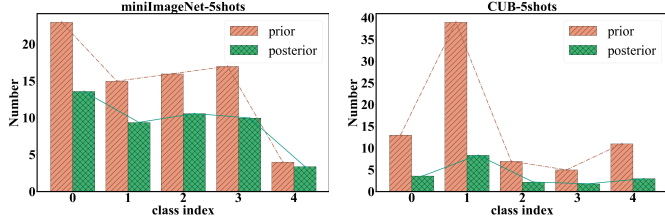


Figure 4: Case studies for 5-way 5-shots tasks on CUB [31] (left) and miniImageNet (right). Comparisons between unnormalized prior distribution (true number of query samples counting for each class) and unnormalized posterior distribution (estimated number of query samples counting by the well optimized forward transport matrix for each class).

still outperform the inductive peers in class-imbalanced setting, which implies that there exists a trade-off between extra statistics conveyed by class-imbalanced query samples and algorithm bias; 3) results of features extracted by WRN§ usually outperforms those by WRN or RN18, it indicates that self-supervised learning is helpful in obtaining generalized representations in class-imbalanced and class-balanced FSL applications; 4) most existed transductive inference methods suffer from performance cliff since their mismatch between data distribution and prior distribution. For example, TIM implicitly introduces uniform prior on query samples by using entropy regularization since minimizing entropy-based regularization is equivalent to minimizing the KL divergence between the predicted marginal distribution and uniform distribution. Hence, data distribution and model bias are unsuitable, leading to poor accuracy. Fortunately, such kind of mismatch can be eased by re-weighting in [29]. On the contrary, PT-MAP holds a strong but unmodified uniformly distributed prior on its transport matrix. Therefore, results of TIM are robuster than those of PT-MAP; 5) our proposed method outperforms most of the transductive methods and achieves competitive results with some most recently proposed methods, such as  $\alpha$ -TIM, iLPC and BAVAR. It indicates that prototype refinement benefits from the well optimized transport matrices when shot number is low. Moreover, compared with iLPC and BAVAR, on the one hand, our model is easy to implement. On the other hand, BAVAR heavily relies on prior of imbalanced ratio  $\alpha$  in class-imbalanced case while iLPC depends on uniform prior in class-balanced case [12], importantly however, our model is free of these priors, making it suitable for real-world applications when we have no idea on class distributions.

**Qualitatively analysis:** To qualitatively illustrate the evolution of prototypes of our model, we map prototypes and query features of one task into a 2-D subspace with t-distributed Stochastic Neighbor Embedding (t-SNE) [28], it is a non-linear dimensional reduction technique well-suited for embedding high-dimensional data into a low-dimensional

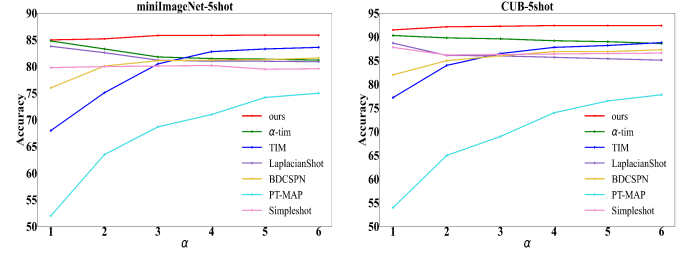


Figure 5: 5-way 5-shot classification accuracy on miniImageNet and CUB of transductive methods versus imbalance level, lower  $\alpha$  corresponds to severer class imbalance.

Table 2: End-to-end real-time efficiencies ( $s$ ) of 5-way 1-shot on CUB with two backbones averaged on 100 tasks.

Backbone	$\alpha$ - TIM (NeurIPS'21)	iLPC (ICCV'21)	BAVAR (AISTATS'23)	Ours
RN18	5.6	4.6	3.4	2.9
WRN	6.2	4.7	3.6	3.7

space. As we can see from Fig. 3, due to class-imbalanced issue, although prototypes can not match their corresponding class centers well at beginning, they become consistent with the ground truth centers after refinement using our proposed method, verifying the effectiveness of our PUTM in dealing with unknown and arbitrary prior situation.

### 4.3. Model Analysis

**Forward transport V.S. backward transport:** As discussed in Sec. 3.3.2, the forward and backward transport matrices put different priors of  $\{a_i\}_{i=1}^N$ . Namely, we do not restrict uniform distribution in the forward transport and enable it to be adaptive to data-driven class distribution. On the contrary, prior in the backward transport is always assumed to be  $\frac{1}{N}$ . Therefore, we employ the forward transport for predictions to meet the class-imbalanced essence, the efficiency is illustrated in Fig. 4, it depicts that the predicted distribution is consistent with the true distribution. Moreover, we conduct an extra experiment to verify the opposite direction of the two transport matrices, results are illustrated in Fig. 3, where we separately use the forward transport and the backward transport to calibrate the prototypes over 3000 tasks. According to the results, we find that, from a statistical point of view, the backward transport does harm the efficient migration from the class-imbalanced query samples.

**Robustness of imbalanced coefficients:** An important point in this work is that our model is robust for various class-imbalanced coefficients, which is vital and crucial for real-world applications since we usually do not know whether the few-shot tasks are balanced or not in advance, thus, we could not know what kind of methods are suitable for describing the samples distribution of tasks. Once they are

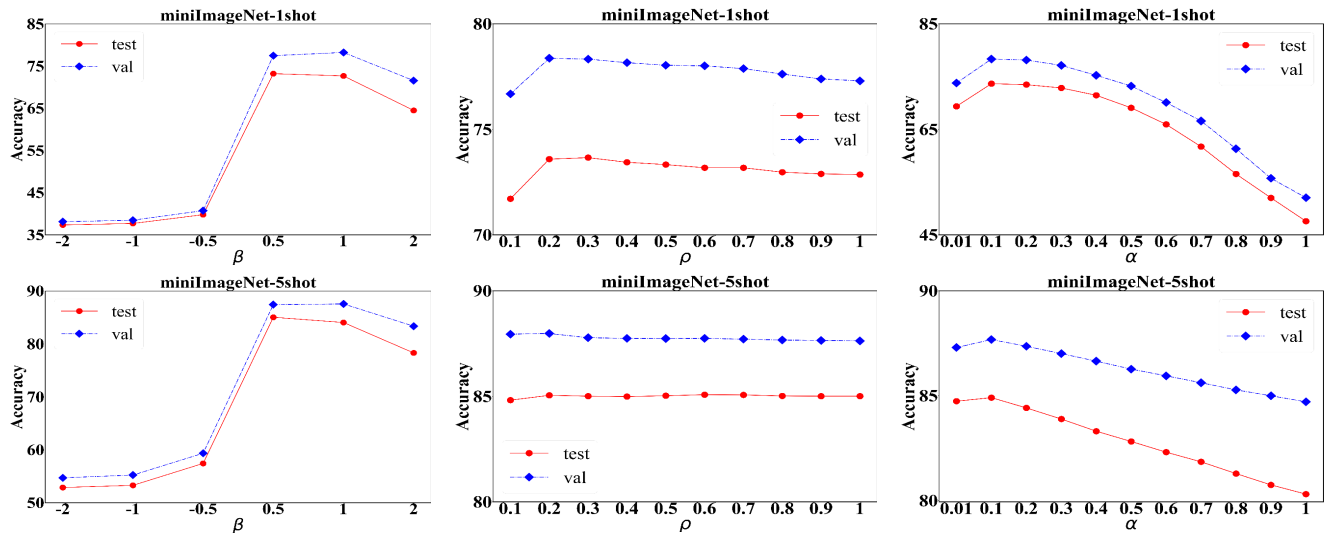


Figure 6: Sensitivity analysis of our model to hyper-parameters including log transform  $\beta$ , CT coefficient  $\rho$ , and inertia  $\alpha$ . Results are obtained with WRN $\xi$  and are reported on both validation and test sets. Best viewed in color.

unmatched, horrible performance degradation would happen, such as PT-MAP does. We verify the robustness of our model with a wide range of imbalanced coefficients as shown in Fig. 5, where  $\alpha$  changes from 1 to 6. The performance of most methods decreases as  $\alpha$  decreasing because smaller  $\alpha$  indicates more imbalanced distribution of query samples, thus making tasks harder. On the contrary, performances of LaplacianShot and  $\alpha$ -TIM drop obviously as  $\alpha$  increases. Our model achieves consistent improvement and reasonable trend in a wide range of  $\alpha$ , we attribute this to the introduction of adaptive transport matrices, which enables PUTM to fit the class distributions precisely and flexibly.

**Sensitivity of hyper-parameters:** To further verify the robustness of our model, we validate the sensitivity of PUTM to different hyper-parameters introduced in Sec. 4.1, and results are reported in Fig. 6. Generally speaking, the tendencies on validation and test sets are same. Besides there are three-fold observations: i) CT coefficient is less sensitive than log transform and inertia; ii) log transform controls the strength of channel-wise attention, which is proven to be important especially in few-shot case, thus it affects performance mostly; iii) inertia reflects the bottleneck of information flow from the aggregated prototypes to the refined prototypes. Bigger or smaller value of inertia is helpless for robust transfer since refinement of prototypes is associated with transport matrices. Namely, if the matrices are not well-optimized, a big inertia may cause algorithm to fall into a local optimum and vice versa.

**Real-time efficiency:** To verify the efficiency of PUTM, we conduct parallel versions of  $\alpha$ -TIM, iLPC, and BAVAR based on their official implementations with the same machine. Results are reported in Table. 2, as we can see that our

model consumes less time using RN18 and achieves competitive results using WRN. And we attribute this to the fact that : i) both  $\alpha$ -TIM and iLPC employ Sinkhorn algorithm to solve the Optimal Transport (OT) problem involving thousands of inner loop, and it is a kind of time consuming; ; ii) our EM-Solver might be more efficient for dealing with the class-imbalanced TFSL since it imposes potential capability for modeling such unknown distribution with CT theory.

## 5. Conclusion and Future Works

We propose a novel model called PUTM with EM-Solver based on CT theory for both class-imbalanced and class-balanced TFSL in this work according to an important observation that CT has a natural connection with unknown and arbitrary class prior problem. Specifically, PUTM exploits transferable statistics with CT for refining biased prototypes solved by EM to capture the unbiased statistics within imbalanced query samples. Extensive experiments on standard benchmarks verify the effectiveness of our model in both class-imbalanced and class-balanced few-shot applications. In the future, we will explore the applicability of PUTM on more problem settings and applications, such as multi-domain few-shot classification and few-shot object detection.

## 6. Acknowledgement

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant U21B2006; in part by Shaanxi Youth Innovation Team Project; in part by the National NSFC under Grant 82172860; in part by the National NSFC under Grant 6220010437; in part by the Foundation of Aerospace under Grant SAST2021-



012. We also thanks Dandan Guo and He Zhao for their constructive comments on transductive few-shot learning, we thanks Zhisong Zhang for his very initial exploration on class-imbalanced few-shot image classification, we thanks Dongsheng Wang for his passion discussions on CT.

## References

- [1] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- [2] Malik Boudiaf, Imtiaz Ziko, Jérôme Rony, José Dolz, Pablo Piantanida, and Ismail Ben Ayed. Information maximization for few-shot learning. *Advances in Neural Information Processing Systems*, 33:2445–2457, 2020.
- [3] wenchao chen, Chaojie Wang, Bo Chen, Yicheng Liu, Hao Zhang, and Mingyuan Zhou. Bidirectional convolutional poisson gamma dynamical systems. 33, 2020.
- [4] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019.
- [5] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- [6] 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.
- [7] Guneet S Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. *arXiv preprint arXiv:1909.02729*, 2019.
- [8] Li Fei-Fei, Robert Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006.
- [9] Dandan Guo, Long Tian, He Zhao, Mingyuan Zhou, and Hongyuan Zha. Adaptive distribution calibration for few-shot learning with hierarchical optimal transport. *Advances in Neural Information Processing Systems*, 35:6996–7010, 2022.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [11] Yuqing Hu, Vincent Gripon, and Stéphane Pateux. Leveraging the feature distribution in transfer-based few-shot learning. In *Artificial Neural Networks and Machine Learning–ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part II 30*, pages 487–499. Springer, 2021.
- [12] Yuqing Hu, Stéphane Pateux, and Vincent Gripon. Adaptive dimension reduction and variational inference for transductive few-shot classification. In *International Conference on Artificial Intelligence and Statistics*, pages 5899–5917. PMLR, 2023.
- [13] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- [14] Michalis Lazarou, Tania Stathaki, and Yannis Avrithis. Iterative label cleaning for transductive and semi-supervised few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8751–8760, 2021.
- [15] Miaoge Li, Dongsheng Wang, Xinyang Liu, Zequn Zeng, Ruiying Lu, Bo Chen, and Mingyuan Zhou. Patchct: Aligning patch set and label set with conditional transport for multi-label image classification. *arXiv preprint arXiv:2307.09066*, 2023.
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [17] Jinlu Liu, Liang Song, and Yongqiang Qin. Prototype rectification for few-shot learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 741–756. Springer, 2020.
- [18] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
- [19] Puneet Mangla, Nupur Kumari, Abhishek Sinha, Mayank Singh, Balaji Krishnamurthy, and Vineeth N Balasubramanian. Charting the right manifold: Manifold mixup for few-shot learning. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2218–2227, 2020.
- [20] Erik G Miller, Nicholas E Matsakis, and Paul A Viola. Learning from one example through shared densities on transforms. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, volume 1, pages 464–471. IEEE, 2000.
- [21] Mateusz Ochal, Massimiliano Patacchiola, Amos Storkey, Jose Vazquez, and Sen Wang. Few-shot learning with class imbalance. *arXiv preprint arXiv:2101.02523*, 2021.
- [22] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International conference on learning representations*, 2017.
- [23] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018.
- [24] Franco Scarselli and Ah Chung Tsoi. Universal approximation using feedforward neural networks: A survey of some existing methods, and some new results. *Neural networks*, 11(1):15–37, 1998.
- [25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [26] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.

- [27] Korawat Tanwisuth, Shujian Zhang, Huangjie Zheng, Pengcheng He, and Mingyuan Zhou. Pouf: Prompt-oriented unsupervised fine-tuning for large pre-trained models. *arXiv preprint arXiv:2305.00350*, 2023.
- [28] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [29] Olivier Veilleux, Malik Boudiaf, Pablo Piantanida, and Ismail Ben Ayed. Realistic evaluation of transductive few-shot learning. *Advances in Neural Information Processing Systems*, 34:9290–9302, 2021.
- [30] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.
- [31] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [32] Chaojie Wang, Hao Zhang, Bo Chen, Dongsheng Wang, Zhengjue Wang, and Mingyuan Zhou. Deep relational topic modeling via graph poisson gamma belief network. *In NeurIPS*, 33, 2020.
- [33] Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019.
- [34] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.
- [35] Shuo Yang, Lu Liu, and Min Xu. Free lunch for few-shot learning: Distribution calibration. *arXiv preprint arXiv:2101.06395*, 2021.
- [36] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [37] Haipeng Zhang, Zhong Cao, Ziang Yan, and Changshui Zhang. Sill-net: Feature augmentation with separated illumination representation. *arXiv preprint arXiv:2102.03539*, 2021.
- [38] Hao Zhang, Bo Chen, Dandan Guo, and Mingyuan Zhou. Whai: Weibull hybrid autoencoding inference for deep topic modeling. *In ICLR*, 2018.
- [39] Huangjie Zheng and Mingyuan Zhou. Exploiting chain rule and bayes’ theorem to compare probability distributions. *Advances in Neural Information Processing Systems*, 34:14993–15006, 2021.
- [40] Imtiaz Ziko, Jose Dolz, Eric Granger, and Ismail Ben Ayed. Laplacian regularized few-shot learning. In *International conference on machine learning*, pages 11660–11670. PMLR, 2020.