

# TAFSSL: Task-Adaptive Feature Sub-Space Learning for few-shot classification

Moshe Lichtenstein<sup>1</sup>, Prasanna Sattigeri<sup>1</sup>, Rogerio Feris<sup>1</sup>,  
Raja Giryes<sup>2</sup>, and Leonid Karlinsky<sup>1</sup>

<sup>1</sup>IBM Research AI, <sup>2</sup>Tel Aviv University

**Abstract.** Recently, Few-Shot Learning (FSL), or learning from very few (typically 1 or 5) examples per novel class (unseen during training), has received a lot of attention and significant performance advances. While number of techniques have been proposed for FSL, several factors have emerged as most important for FSL performance, awarding SOTA even to the simplest of techniques. These are: the backbone architecture (bigger is better), type of pre-training (meta-training vs multi-class), quantity and diversity of the base classes (the more the merrier), and using auxiliary self-supervised tasks (a proxy for increasing the diversity). In this paper we propose TAFSSL, a simple technique for improving the few shot performance in cases when some additional unlabeled data accompanies the few-shot task. TAFSSL is built upon the intuition of reducing the feature and sampling noise inherent to few-shot tasks comprised of novel classes unseen during pre-training. Specifically, we show that on the challenging *mini*ImageNet and *tiered*ImageNet benchmarks, TAFSSL can improve the current state-of-the-art in both transductive and semi-supervised FSL settings by more than 5%, while increasing the benefit of using unlabeled data in FSL to above 10% performance gain.

**Keywords:** Transductive, Semi-supervised, Few-Shot Learning

## 1 Introduction

The great success of Deep Learning (DL) methods to solve complex computer vision problems can be attributed in part to the emergence of large labeled datasets [27, 40] and strong parallel hardware. Yet in many practical situations, the amount of data and/or labels available for training or adapting the DL model to a new target task is prohibitively small. In extreme cases, we might be interested in learning from as little as one example per novel class. This is the typical scenario of Few-Shot Learning (FSL), a very active and exciting research topic of many concurrent works [22, 47, 51, 52]. While many great techniques have been proposed to improve FSL performance, recent studies [3, 15, 52] have shown that there exist a number of important factors that can improve the FSL performance, largely regardless of the model and the learning algorithm used. These include: (i) significant performance gains observed while increasing the size and the number of parameters of the backbone generating the feature representations of the images [3, 52]; (ii) gains while pre-training the FSL model

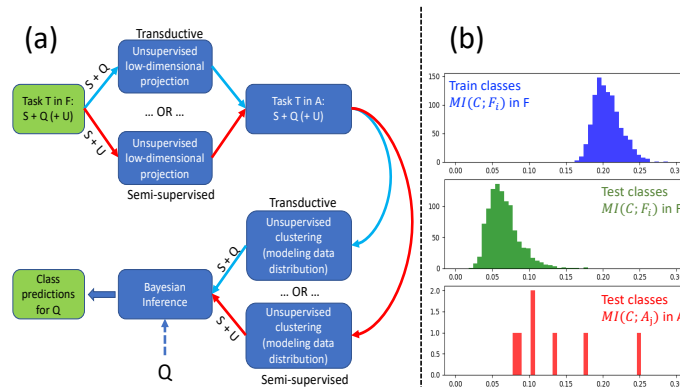


Fig. 1: **(a) TAFSSL overview:** red and blue pathways are for semi-supervised and transductive FSL respectively.  $T$  - few-shot task;  $S$  - support set;  $Q$  - query set;  $U$  - optional set of additional unlabeled examples (semi-supervised FSL);  $F$  - original feature space;  $A$  - task adapted feature sub-space. **(b) Improved SNR in A:** the normalized (by min entropy) Mutual Information (MI) between either train or test classes and the features in  $F$  (of dimension 1024) or in  $A$  (7-dim) provides the motivation to use  $A$  over  $F$ . Computed on *miniImageNet*

on the base classes dataset as a regular multi-class classifier (to all base classes at once) [52], as opposed to the popular meta-training by generating a lot of synthetic few-shot tasks from random small groups of base classes [22, 51]; (iii) gains when pre-training on more (diverse) base classes (e.g. higher empirical FSL performance on seemingly more difficult *tieredImageNet* benchmark than on supposedly simpler *miniImageNet* benchmark [22, 52]; (iv) gains when artificially increasing the diversity and complexity of the base classes dataset by introducing additional self-supervised tasks during pre-training [15]. Correctly using these factors allows the simple Nearest Neighbor classifier to attain state-of-the-art FSL performance [52] improving upon more sophisticated FSL techniques.

All the aforementioned factors and gains concern the base classes pre-training stage of the FSL methods backbones. Much less attention has been given to adapting the feature spaces resulting from these backbones to the novel classes few-shot tasks during test time. It has been shown that some moderate gains can be obtained from using the few training examples (support set) of the novel tasks to fine-tune the backbones (changing the feature spaces slightly), with best gains obtained for higher resolution and higher 'shots' (support examples per class) regimes [31]. Fine-tuning was also shown to be effective for semi-supervised setting [24], where additional unlabeled examples accompany each novel few-shot task. It has also been shown that label propagation and clustering operating in the pre-trained backbone's original feature space provide some gains for transductive FSL (allowing unlabeled queries to be used jointly to predict their labels in a bulk) and semi-supervised FSL [28, 38]. Finally, meta-learned

backbone architecture adaptation mechanics were proposed in [10] allowing for slight backbone architecture transformations adaptive to the few-shot test task.

However, slight adaptation of the backbone’s feature space to a given task, using few iterations of fine-tuning on the support set or other techniques, might not be sufficient to bridge over the generalization gap introduced by the FSL backbone observing completely novel classes unseen during training (as confirmed by the relatively moderate performance gains obtained from these techniques). Intuitively, we could attribute this in part to many of the feature space dimensions (feature vector entries) becoming ‘useless’ for a given set of novel classes in the test few-shot task. Indeed, every feature vector entry can be seen as a certain ‘pattern detector’ which fires strongly when a certain visual pattern is observed on the input image. The SGD (or other) backbone training is making sure all of these patterns are discriminative for the classes used for pre-training. But, due to likely over-fitting, many of these patterns are base classes specific, and do not fire for the novel test classes. Hence, their corresponding feature vector entries will mainly produce ‘noise values’ corresponding to ‘pattern not observed’. In other words, the ratio of feature vector entries that can be used for recognition of novel classes to ones which mainly output ‘noise’ significantly decreases for test few-shot task (Figure 1b). And it is unlikely that small modifications to the feature space recovers a significant portion of the ‘noise producing’ feature entries. The high level of noise in the feature vectors intuitively has significant adverse implications on the performance of the FSL classifier operating on this vector, especially the popular distance based classifiers like nearest-neighbor [52] and Prototypical Networks (PN) [47] are affected. In light of this intuition, we conjecture that for a significant performance boost, we need to concentrate our efforts on the so-called *Task-Adaptive Feature Sub-Space Learning (TAFSSL)* - seeking sub-spaces of the backbone’s feature space that are discriminative for the novel classes of the test few-shot task and which are ‘as noise free as possible’, that is most of the sub-space dimensions indeed ‘find’ the patterns they represent in the images of the novel categories belonging to the task.

In this paper we set to explore TAFSSL under the transductive and the semi-supervised few-shot settings. In many practical applications of FSL, alongside the few labeled training examples (the support set) of the few shot task, additional unlabeled examples containing instances of the target novel classes are available. Such is the situation in transductive FSL which assumes that the query samples arrive in a ‘bulk’ and not one-by-one, and hence we can answer all the queries ‘at once’ while using the query set as unlabeled data. Similar situation exists in semi-supervised FSL, where unlabeled set of images simply accompanies the few-shot task. As can be observed from our experiments, TAFSSL, and especially TAFSSL accompanied by specific (proposed) forms of clustering based approaches, provides very significant boost to FSL performance under the transductive and semi-supervised setting. Specifically, we obtain over 7% and over 10% absolute improvement of the popular 1-shot *miniImageNet* [51] and *tieredImageNet* [38] few-shot benchmarks in the transductive FSL setting, and 13% and over 10% absolute improvement in semi-supervised FSL setting (over

corresponding state-of-the-art while using their respective evaluation protocols). Figure 1a illustrates an overview of the proposed approach.

To summarize, we offer the following contributions: (i) we highlight the Task-Adaptive Feature Sub-Space Learning (TAFSSL) as an important factor for Few-Shot Learning (FSL), we explore several TAFSSL methods and demonstrate significant performance gains obtained using TAFSSL for transductive and semi-supervised FSL; (ii) we propose two variants of clustering that can be used in conjunction with TAFSSL to obtain even greater performance improvements; (iii) we obtain new state-of-the-art transductive and semi-supervised FSL results on two popular benchmarks: *miniImageNet* and *tieredImageNet*; (iv) we offer an extensive ablation study of the various aspects of our approach, including subspace dimension, unlabeled data quantity, effects of out-of-distribution noise in unlabeled data, backbone architectures, and finally - effect of class imbalance (skew) in unlabeled data (so-far unexplored in all previous works).

## 2 Related work

Here we briefly review the modern Few-Shot Learning (FSL) approaches, focusing in more detail on the transductive and semi-supervised FSL methods that leverage unlabeled data. The *meta-learning* methods [51, 47, 50, 23, 57] learn from few-shot tasks (or episodes) rather than from individual labeled samples. Each such task is a small dataset, with few labeled training examples (a.k.a. support), and a few test examples (a.k.a. query). The goal is to learn a model that can adapt to new tasks with novel categories unseen during training. The *gradient-based meta learners* [12, 25, 58, 36, 30, 41, 3, 22, 32] search for models can transfer well to novel few-shot tasks. Typically, in these methods higher order derivatives are used to optimize the loss the model would have after applying few gradient steps. At test time, simple fine-tuning is used. In [11] ensemble methods for few-shot learning are proposed. The *metric learning* based methods [54, 47, 39, 14, 43, 49, 32, 20, 55, 7] a non linear metric is optimized on base classes and applied to unseen few-shot test tasks. In [7, 46] distance to class prototype is replaced by distance to a class sub-space. As opposed to [7] and [46] that try to optimize a sub-space for each class, in TAFSSL we seek a single sub-space optimally adapted to the entire data of the few-shot task - labeled and unlabeled. Notably, in [52] non-meta-learning pre-training was used in combination with 'large' backbones (e.g. DenseNet [19]) and a nearest-neighbor classifier to achieve state-of-the-art results. The *generative and augmentation-based* methods [33, 9, 48, 16, 6, 26, 37, 2, 17, 45, 53, 4, 56, 44, 1] are methods that (learn to) generate more samples from the one or a few training examples in a given few-shot learning task.

**Transductive and semi-supervised FSL:** In many practical applications, in addition to the labeled support set, additional unlabeled data accompanies the few-shot task. Transductive FSL [8, 28, 21, 35] assumes that the set of task queries arrives in a bulk and can be used as a source of unlabeled data, allowing query samples to 'learn' from each other. In [8] the queries are used in fine-tuning in conjunction with entropy minimization loss in order to maximize the certainty

of their predictions. In semi-supervised FSL [24, 38, 46, 28] the unlabeled data comes in addition to the support set and is assumed to have a similar distribution to the target classes (although some unrelated samples noise is also allowed). In the LST [24] self-labeling and soft attention are used on the unlabeled samples intermittently with fine-tuning on the labeled and self-labeled data. Similarly to LST, [38] updates the class prototypes using k-means like iterations initialized from the PN prototypes. Their method also includes down-weighting the potential distractor samples (likely not to belong to the target classes) in the unlabeled data. In [46] unlabeled examples are used through soft-label propagation. In [42] semi-supervised few-shot domain adaptation is considered. In [14, 28, 21] graph neural networks are used for sharing information between labeled and unlabeled examples in semi-supervised [14, 28] and transductive [21] FSL setting. Notably, in [28] a Graph Construction network is used to predict the task specific graph for propagating labels between samples of semi-supervised FSL task.

### 3 Method

Here we derive the formal definition of TAFSSL and examine several approaches for it. In addition, we propose several ways to combine TAFSSL with clustering, which is shown to be very beneficial to the performance in section 4.

#### 3.1 FSSL and TAFSSL

Let a CNN backbone  $\mathcal{B}$  (e.g. ResNet [18] or DenseNet [19]) pre-trained for FSL on a (large) dataset  $\mathcal{D}_b$  with a set of base (training) classes  $\mathcal{C}_b$ . Here for simplicity, we equally refer to all different forms of pre-training proposed for FSL in the literature, be it meta-training [51] or 'regular' training of a multi-class classifier for all the classes  $\mathcal{C}_b$  [52]. Denote by  $\mathcal{B}(x) \in \mathcal{F} \subset \mathbf{R}^m$  to be a feature vector corresponding to an input image  $x$  represented in the feature space  $\mathcal{F}$  by the backbone  $\mathcal{B}$ . Under this notation, we define the goal of linear Feature Sub-Space Learning (FSSL) to find an 'optimal' (for a certain task) linear sub-space  $\mathcal{A}$  of  $\mathcal{F}$  and a linear mapping  $W$  of size  $r \times m$  (typically with  $r \ll m$ ) such that:

$$\mathbf{R}^r \supset \mathcal{A} \ni A = W \cdot \mathcal{B}(x) \quad (1)$$

is the new representation of an input image  $x$  as a vector  $A$  in the feature sub-space  $\mathcal{A}$  (spanned by rows of  $W$ ).

Now, consider an  $n$ -way +  $k$ -shot few-shot test task  $\mathcal{T}$  with a query set  $Q$ , and a support set:  $S = \{s_i^j | 1 \leq i \leq n, 1 \leq j \leq k, \mathcal{L}(s_i^j) = i\}$ , where  $\mathcal{L}(x)$  is the class label of image  $x$ , so in  $S$  we have  $k$  training samples (shots) for each of the  $n$  classes in the task  $\mathcal{T}$ . Using the PN [47] paradigm we assume  $k = 1$  (otherwise support examples of the same class are averaged to a single class prototype) and that each  $q \in Q$  is classified using Nearest Neighbor (NN) in  $\mathcal{F}$ :

$$CLS(q) = \underset{i}{\operatorname{argmin}} \|\mathcal{B}(s_i^1) - \mathcal{B}(q)\|^2 \quad (2)$$

Then, in the context of this given task  $\mathcal{T}$ , we can define linear Task-Adaptive FSSL (TAFSSL) as a search for a linear sub-space  $\mathcal{A}_{\mathcal{T}}$  of the feature space  $\mathcal{F}$  defined by a  $\mathcal{T}$ -specific projection matrix  $W_{\mathcal{T}}$ , such that the probability:

$$\frac{\exp(-\tau \cdot \|W_{\mathcal{T}} \cdot (\mathcal{B}(s_{\mathcal{L}(q)}^1) - \mathcal{B}(q))\|^2)}{\sum_i \exp(-\tau \cdot \|W_{\mathcal{T}} \cdot (\mathcal{B}(s_i^1) - \mathcal{B}(q))\|^2)} \quad (3)$$

of predicting  $q$  to belong to the same class as the 'correct' support  $s_{\mathcal{L}(q)}^1$  is maximized, while of course the true label  $\mathcal{L}(q)$  is unknown at test time (here  $\tau$  in eq. 3 is a temperature parameter, we used  $\tau = 1$ ).

**Discussion.** Using the 'pattern detectors' intuition from section 1, lets consider the activations of each dimension  $F_d$  of  $F \in \mathcal{F}$  as a random variable with a Mixture of (two) Gaussians (MoG) distribution:

$$F_d \sim P_d = \rho_n \cdot N(\mu_n, \sigma_n) + \rho_s \cdot N(\mu_s, \sigma_s) \quad (4)$$

where  $(\mu_n, \sigma_n)$  and  $(\mu_s, \sigma_s)$  are the expectation and variance of the  $F_d$ 's distribution of activations when  $F_d$  does not detect (noise) or detects (signal) the pattern respectively. The  $\rho_n$  and  $\rho_s$  are the noise and the signal prior probabilities respectively ( $\rho_n + \rho_s = 1$ ). For brevity, we drop the index  $d$  from the distribution parameters. Naturally, for the training classes  $C_b$ , for most dimensions  $F_d$  the  $\rho_s \gg 0$  implying that the dimension is 'useful' and does not produce only noise (Figure 1b, top). However, for the new (unseen during training) classes of a test task  $\mathcal{T}$  this is no longer the case, and it is likely that for the majority of dimensions  $\rho_s^{\mathcal{T}} \approx 0$  (Figure 1b, middle). Assuming (for the time being) that  $F_d$  are conditionally independent, the square Euclidean distance could be seen as an aggregation of votes for the 'still useful' (for the classes of  $\mathcal{T}$ ) patterns, and a sum of squares of i.i.d (zero mean) Gaussian samples for the patterns that are 'noise only' on the classes of  $\mathcal{T}$ . The latter 'noise dimensions' randomly increase the distance on the expected order of  $N^{\mathcal{T}, \mathcal{F}} \cdot \sigma_n^2$ , where  $N^{\mathcal{T}, \mathcal{F}}$  is the number of *noise features* of the feature space  $\mathcal{F}$  for the classes of task  $\mathcal{T}$ . Using this intuition, if we could find such a TAFSSL sub-space  $\mathcal{A}_{\mathcal{T}}$  adapted to the task  $\mathcal{T}$  so that  $N^{\mathcal{T}, \mathcal{S}\tau}$  is reduced (Figure 1b, bottom), we would improve the performance of the NN classifier on  $\mathcal{T}$ . With only few labeled samples in the support set  $S$ , we cannot expect to effectively learn the  $W_{\mathcal{T}}$  projection to the sub-space  $\mathcal{A}_{\mathcal{T}}$  using SGD on  $S$ . Yet, when unlabeled data accompanies the task  $\mathcal{T}$  ( $Q$  in transductive FSL, or an additional set of unlabeled samples  $U$  in semi-supervised FSL), we can use this data to find such  $W_{\mathcal{T}}$  that: (a) the dimensions of  $\mathcal{A}_{\mathcal{T}}$  are 'disentangled', meaning their pairwise independence is maximized; (b) after the 'disentanglement' we choose the dimensions that are expected to 'exhibit the least noise' or in our previous MoG notation have the largest  $\rho_s$  values.

Luckily, simple classical methods can be used for TAFSSL approximating the requirements (a) and (b). Both Principle Component Analysis (PCA) [34] and Independent Component Analysis (ICA) [5] applied in  $\mathcal{F}$  on the set of samples:  $S \cup Q$  (transductive FSL) or  $S \cup U$  (semi-supervised FSL) can approximate (a).

PCA under the approximate joint Gaussianity assumption of  $\mathcal{F}$ , and ICA under approximate non-Gaussianity assumption. In addition, if after the PCA rotation we subtract the mean, the variance of the (zero-mean) MoG mixtures for the transformed (independent) dimensions would be:

$$\rho_n \cdot (\mu_n^2 + \sigma_n^2) + \rho_s \cdot (\mu_s^2 + \sigma_s^2) \quad (5)$$

Then assuming  $\mu_n$  and  $\sigma_n$  are roughly the same for all dimensions (which is reasonable due to heavy use of Batch Normalization (BN) in the modern backbones), choosing the dimensions with higher variance in PCA would lead to larger  $\rho_s$ ,  $\mu_s$ , and  $\sigma_s$  - all of which are likely to increase the signal-to-noise ratio of the NN classifier. Larger  $\mu_s$  leads to patterns with stronger 'votes', larger  $\sigma_s$  means wider range of values that may better discriminate multiple classes, and larger  $\rho_s$  means patterns that are more frequent for classes of  $\mathcal{T}$ . Similarly, the dimensions with bigger  $\rho_s$  exhibit stronger departure from Gaussianity and hence would be chosen by ICA.

**TAFSSL summary.** To summarize, following the discussion above, both PCA and ICA are good simple approximations for TAFSSL using unlabeled data and therefore we simply use them to perform the 'unsupervised low-dimensional projection' in the first step of our proposed approach (Figure 1a). As we show in the Results section 4, even on their own (when directly followed by an NN classifier) they lead to significant FSL performance boosts (Tables 1 and 2).

### 3.2 Clustering

It was shown that clustering is a useful tool for transductive and semi-supervised FSL [38]. There, it was assumed that modes of the task  $\mathcal{T}$  data distribution (including both labeled and unlabeled image samples) correspond classes. However, in the presence of feature 'noise' in  $\mathcal{F}$ , as discussed in section 3.1, the 'class' modes may become mixed with the noise distribution modes, that may blur the class modes boundaries or swallow the class modes altogether. Indeed, the performance gains in [38] were not very high.

In contrast, after applying PCA or ICA based TAFSSL, the feature noise levels are usually significantly reduced (Figure 1b) making the task-adapted feature sub-space  $\mathcal{A}_{\mathcal{T}}$  of the original feature space  $\mathcal{F}$  to be much more effective for clustering. We propose two clustering-based algorithms, the Bayesian K-Means (BKM) and Mean-Shift Propagation (MSP). In the Results section 4 we show that following PCA or ICA based TAFSSL, these clustering techniques add about 5% to the performance. They are used to perform the 'unsupervised clustering' + 'bayesian inference' steps of our approach (Figure 1a).

The **BKM** is a soft k-means [29] variant accompanied with Bayesian inference for computing class probabilities for the queries. In BKM, each k-means cluster, obtained for the entire set of (labeled + unlabeled) task  $\mathcal{T}$  data, is treated as a Gaussian mixture distribution with a mode for each class. The BKM directly computes the class probability for each query  $q \in Q$  by averaging

the posterior of  $q$  in each of the mixtures using weights being the probability of  $q$  to belong to each cluster. BKM details are provided in Algorithm 1.

---

**Algorithm 1** Bayesian K-Means (BKM)
 

---

Cluster the samples of task  $\mathcal{T}$  ( $Q \cup S$  or  $U \cup S$  in transductive or semi-supervised FSL respectively) into  $k$  clusters, associating each to  $c_k$  - the centroid of cluster  $k$ .

**for each**  $s \in S$ ,  $q \in Q$ , and  $k$  **do**

$$P(\text{cluster}(q) = k) = \frac{\exp(-\|q - c_k\|^2)}{\sum_j \exp(-\|q - c_j\|^2)}$$

$$P(\text{cluster}(s) = k) = \frac{\exp(-\|s - c_k\|^2)}{\sum_j \exp(-\|s - c_j\|^2)}$$

$$P(\mathcal{L}(q) = i | \text{cluster}(q) = k) = \sum_{\mathcal{L}(s)=i} \frac{\exp(-\|q - s\|^2) \cdot P(\text{cluster}(s)=k)}{\sum_{t \in S} \exp(-\|q - t\|^2) \cdot P(\text{cluster}(t)=k)}$$

$$P(\mathcal{L}(q) = i) = \sum_k P(\mathcal{L}(q) = i | \text{cluster}(q) = k) \cdot P(\text{cluster}(q) = k)$$


---

The **MSP** is a mean-shift [13] based approach, that is used to update the prototype of each class. In MSP we perform a number of mean-shift like iterations on the prototypes [47] of the classes taken within the distribution of all the (labeled and unlabeled) samples of  $\mathcal{T}$ . In each iteration, for each prototype  $p_i$  (of class  $i$ ), we compute a set of  $K$  most confident samples within a certain confidence radius and use the mean of this set as the next prototype (of class  $i$ ). The  $K$  itself is balanced among the classes. The details of MSP are summarized in Algorithm 2 box. Following MSP, the updated prototypes are used in standard NN classifier fashion to obtain the class probabilities.

---

**Algorithm 2** Mean-Shift Propagation (MSP)
 

---

**Initialize:**

Compute prototypes:  $\{p_i = \frac{1}{k} \cdot \sum_{s \in S, \mathcal{L}(s)=i} s\}$ , where  $k$  is # of shots in task  $\mathcal{T}$

**for**  $N$  times **do**

$$\text{Compute } P(\mathcal{L}(x) = i) = \frac{\exp(-\|x - p_i\|^2)}{\sum_j \exp(-\|x - p_j\|^2)}, \forall x \in Q \cup S \text{ (or } x \in U \cup S)$$

Compute predictions  $c(x) = \text{argmax}_i P(\mathcal{L}(x) = i)$

$K_i = \sum_x \mathbb{1}_{(c(x)=i) \wedge (P(\mathcal{L}(x)=i) > T)}$ , where  $T$  is a threshold parameter

$K = \min_i \{K_i\}$

Compute the new prototypes:  $\{p_i = \frac{1}{K} \cdot \sum_{x \in \hat{S}_i} x\}$ , where  $\hat{S}_i$  are the top  $K$  samples that have  $c(x) = i$  sorted in decreasing order of  $P(\mathcal{L}(x) = i)$

**return** labels  $c(q), \forall q \in Q$

---

### 3.3 Implementation details

All the proposed TAFSSL approaches were implemented in PyTorch. We have used the PyTorch native version of SVD for PCA, and FastICA from sklearn for ICA. The k-means from sklearn was used for BKM. The sub-space dimensions



were 4 for PCA based TAFSSL, and 10 for ICA based TAFSSL. These were set using validation (section 4.4). The  $T = 0.3$  and  $N = 4$  were used for MSP, and  $k = 5$  for BKM, all set using validation. We used the backbones implementations from [52]. Unless otherwise specified, DenseNet backbone was used (for backbones ablation, please see section 4.4). The most time consuming of the proposed TAFSSL approaches (ICA + BKM) runs in below 0.05 seconds (CPU) for a typical 1-shot and 5-way episode with 15 queries per class.

## 4 Results

We have evaluated our approach on the popular few-shot classification benchmarks, namely *miniImageNet* [51] and *tieredImageNet* [38], used in all transductive and semi-supervised FSL works [8, 28, 21, 35, 24, 38, 46, 28]. On these benchmarks, we used the standard evaluation protocols, exactly as in corresponding (compared) works. The results of the transductive and semi-supervised FSL evaluation, together with comparison to previous methods, are summarized in tables 1 and 2 respectively and are detailed and discussed in the following sections. All the performance numbers are given in accuracy % and the 0.95 confidence intervals are reported. The tests are performed on 10,000 random 5-way episodes, with 1 or 5 shots (number of support examples per class), and with 15 queries per episode (unless otherwise specified). For each dataset, the standard train / validation / test splits were used. For each dataset, training subset was used to pre-train the backbone (from scratch) as a regular multi-class classifier to all the train classes, same as in [52]; the validation data was used to select the best model along the training epochs and to choose the hyper-parameters; and episodes generated from the test data (with test categories unseen during training and validation) were used for meta-testing to obtain the final performance. In all experiments not involving BKM, the class probabilities were computed using the NN classifier to the class prototypes.

### 4.1 FSL benchmarks used in our experiments

The ***miniImageNet benchmark (Mini)*** [51] is a standard benchmark for few-shot image classification, that has 100 randomly chosen classes from ILSVRC-2012 [40]. They are randomly split into disjoint subsets of 64 meta-training, 16 meta-validation, and 20 meta-testing classes. Each class has 600 images of size  $84 \times 84$ . We use the same splits as [22] and prior works.

The ***tieredImageNet benchmark (Tiered)*** [38] is a larger subset of ILSVRC-2012 [40], consisted of 608 classes grouped into 34 high-level classes. These are divided into disjoint 20 meta-training high-level classes, 6 meta-validation classes, and 8 meta-testing classes. This corresponds to 351, 97, and 160 classes for meta-training, meta-validation, and meta-testing respectively. Splitting using higher level classes effectively minimizes the semantic similarity between classes belonging to the different splits. All images are of size  $84 \times 84$ .

Table 1: Transductive setting

	Mini 1-shot	Mini 5-shot	Tiered 1-shot	Tiered 5-shot
<b>Simple shot [52]</b>	64.30 $\pm$ 0.20	81.48 $\pm$ 0.14	71.26 $\pm$ 0.21	86.59 $\pm$ 0.15
<b>TPN [28]</b>	55.51 $\pm$ 0.86	69.86 $\pm$ 0.65	59.91 $\pm$ 0.94	73.30 $\pm$ 0.75
<b>TEAM [35]</b>	60.07 $\pm$ N.A.	75.90 $\pm$ N.A.	-	-
<b>EGNN + trans. [21]</b>	-	76.37 $\pm$ N.A.	-	80.15 $\pm$ N.A.
<b>Trans. Fine-Tuning [8]</b>	65.73 $\pm$ 0.68	78.40 $\pm$ 0.52	73.34 $\pm$ 0.71	85.50 $\pm$ 0.50
<b>Trans-mean-sub</b>	65.58 $\pm$ 0.20	81.45 $\pm$ 0.14	73.49 $\pm$ 0.21	86.56 $\pm$ 0.15
<b>Trans-mean-sub(*)</b>	65.88 $\pm$ 0.20	82.20 $\pm$ 0.14	73.75 $\pm$ 0.21	87.16 $\pm$ 0.15
<b>PCA</b>	70.53 $\pm$ 0.25	80.71 $\pm$ 0.16	80.07 $\pm$ 0.25	86.42 $\pm$ 0.17
<b>ICA</b>	72.10 $\pm$ 0.25	81.85 $\pm$ 0.16	80.82 $\pm$ 0.25	86.97 $\pm$ 0.17
<b>BKM</b>	72.05 $\pm$ 0.24	80.34 $\pm$ 0.17	79.82 $\pm$ 0.25	85.67 $\pm$ 0.18
<b>PCA + BKM</b>	75.11 $\pm$ 0.26	82.24 $\pm$ 0.17	83.19 $\pm$ 0.25	87.83 $\pm$ 0.17
<b>ICA + BKM</b>	75.79 $\pm$ 0.26	82.83 $\pm$ 0.16	83.39 $\pm$ 0.25	88.00 $\pm$ 0.17
<b>MSP</b>	71.39 $\pm$ 0.27	82.67 $\pm$ 0.15	76.01 $\pm$ 0.27	87.13 $\pm$ 0.15
<b>PCA + MSP</b>	76.31 $\pm$ 0.26	84.54 $\pm$ 0.14	84.06 $\pm$ 0.25	89.13 $\pm$ 0.15
<b>ICA + MSP</b>	<b>77.06 <math>\pm</math> 0.26</b>	<b>84.99 <math>\pm</math> 0.14</b>	<b>84.29 <math>\pm</math> 0.25</b>	<b>89.31 <math>\pm</math> 0.15</b>

## 4.2 Transductive FSL setting

In these experiments we consider the transductive FSL setting, where the set of queries is used as the source of the unlabeled data. This setting is typical for cases when an FSL classifier is submitted a bulk of query data for offline evaluation. In Table 1 we report the performance of our proposed TAFSSL (PCA, ICA), clustering (BKM, MSP), and TAFSSL+clustering (PCA/ICA + BKM/MSP) approaches and compare them to a set of baselines and state-of-the-art (SOTA) transductive FSL methods from the literature: TPN [28] and Transductive Fine-Tuning [8]. We also compare to SOTA regular FSL result of [52] in order to highlight the effect of using the unlabeled queries for prediction. As baselines, we try to maximally adapt the method of [52] to the transductive FSL setting. These are the so-called "trans-mean-sub" that on each test episode subtracts the mean of all the samples ( $S \cup Q$ ) from all the samples followed by L2 normalization (in order reduce the episode bias); and the "trans-mean-sub(\*)" where we do the same but computing and subtracting the means of the  $S$  and  $Q$  sample sets separately (in order to better align their distributions). As can be seen from Table 1, on both the Mini and the Tiered transductive FSL benchmarks, the top performing of our proposed TAFSSL based approaches (ICA+MSP) consistently outperforms all the previous (transductive and non-transductive) SOTA and the baselines by more than 10% in the more challenging 1-shot setting and by more than 2% in the 5-shot setting, underlining the benefits of using the transductive setting, and the importance of TAFSSL to this setting. In the following section, we only evaluate the ICA based TAFSSL variants as it was found to consistently outperform the PCA based variant under all settings.

## 4.3 Semi-supervised FSL setting

Table 2: Semi supervised setting

	# Unlabeled	Mini 1-shot	Mini 5-shot	Tiered 1-shot	Tiered 5-shot
TPN [28]	360	52.78 $\pm$ 0.27	66.42 $\pm$ 0.21	-	-
PSN [46]	100	-	68.12 $\pm$ 0.67	-	71.15 $\pm$ 0.67
TPN [28]	1170	-	-	55.74 $\pm$ 0.29	71.01 $\pm$ 0.23
SKM [38]	100	62.10 $\pm$ N.A.	73.60 $\pm$ N.A.	68.60 $\pm$ N.A.	81.00 $\pm$ N.A.
TPN [28]	100	62.70 $\pm$ N.A.	74.20 $\pm$ N.A.	72.10 $\pm$ N.A.	83.30 $\pm$ N.A.
LST [24]	30	65.00 $\pm$ 1.90	-	75.40 $\pm$ 1.60	-
LST [24]	50	-	77.80 $\pm$ 0.80	-	83.40 $\pm$ 0.80
LST [24]	100	70.10 $\pm$ 1.90	78.70 $\pm$ 0.80	77.70 $\pm$ 1.60	85.20 $\pm$ 0.80
ICA	30	72.00 $\pm$ 0.24	81.31 $\pm$ 0.16	80.24 $\pm$ 0.24	86.57 $\pm$ 0.17
ICA	50	72.66 $\pm$ 0.24	81.96 $\pm$ 0.16	80.86 $\pm$ 0.24	87.03 $\pm$ 0.17
ICA	100	72.80 $\pm$ 0.24	82.27 $\pm$ 0.16	80.91 $\pm$ 0.25	87.14 $\pm$ 0.17
ICA + BKM	30	75.70 $\pm$ 0.22	83.59 $\pm$ 0.14	82.97 $\pm$ 0.23	88.34 $\pm$ 0.15
ICA + BKM	50	76.46 $\pm$ 0.22	84.36 $\pm$ 0.14	83.51 $\pm$ 0.22	88.81 $\pm$ 0.15
ICA + BKM	100	76.83 $\pm$ 0.22	84.83 $\pm$ 0.14	83.73 $\pm$ 0.22	88.95 $\pm$ 0.15
ICA + MSP	30	78.55 $\pm$ 0.25	84.84 $\pm$ 0.14	85.04 $\pm$ 0.24	88.94 $\pm$ 0.15
ICA + MSP	50	79.58 $\pm$ 0.25	85.41 $\pm$ 0.13	85.75 $\pm$ 0.24	89.32 $\pm$ 0.15
ICA + MSP	100	80.11 $\pm$ 0.25	85.78 $\pm$ 0.13	86.00 $\pm$ 0.23	89.39 $\pm$ 0.15

In this section we evaluate our proposed approaches in the semi-supervised FSL setting. In this setting, we have an additional set of unlabeled samples  $U$  that accompanies the test task  $\mathcal{T}$ . In  $U$  we usually expect to have additional samples from the  $\mathcal{T}$ 's target classes distribution, possibly mixed with additional unrelated samples from some number of distracting classes (please see section 4.4 for an ablation on this). In Table 2

we summarize the performance of our proposed TAFSSL based approaches, and compare them to the SOTA semi-supervised FSL methods of [38, 28, 24, 46]. In addition, we also present results for varying number of additional unlabeled samples in  $U$  (semi-supervised). As can be seen from Table 2 the TAFSSL-based approaches outperform all competing methods by a large margins of over 8% and 4% accuracy gain in both the Mini and the Tiered benchmarks in 1-shot and 5-shot settings respectively. Interestingly, same as for the transductive FSL, for the semi-supervised FSL the ICA+MSP approach is the best performing.




Fig. 2: **Num. queries (transductive)**. (a) *miniImageNet*; (b) *tieredImageNet*

#### 4.4 Ablation study

**Number of queries in transductive FSL.** Since the unlabelled data in transductive FSL is comprised entirely from the query samples, the size of the query set  $Q$  in the meta-testing episodes affects the performance. To test this we have evaluated the proposed TAFSSL ICA-based methods, as well as two baselines, namely SimpleShot [52], and its adaptation to transductive setting "trans-mean-sub\*" (sub). All methods were tested varying the number of queries from 2 to 50. The results of this ablation on both the Tiered and Mini benchmarks are shown on figure 2. As can be seen from the figure, already for as little as 5 queries a substantial gap can be observed (for both the benchmarks) between the proposed best performing ICA+MSP technique and the best of the baselines.

#### Out of distribution noise (distraction classes) in unlabeled data.

In many applications, the unlabeled data may become contaminated with samples "unrelated" to the few-shot task  $\mathcal{T}$  target classes. This situation is most likely to arise in the semi-supervised FSL setting, as in transductive FSL the unlabeled samples are the queries and unless we are interested in open-set FSL mode (to the best of our knowledge not explored yet), these are commonly expected to belong only to the target classes distribution. In the semi-supervised

FSL literature [38, 28, 24], this type of noise is evaluated using additional random samples from random "distracting" classes added to the unlabeled set. In figure 3 we compare our proposed ICA-based TAFSSL approaches to SOTA semi-supervised FSL methods [38, 28, 24]. By varying the number of distracting classes from 0 to 7, we see that about 8% accuracy gap is maintained between top TAFSSL method and the top baseline across all the tested noise levels.

**The number of TAFSSL sub-space dimensions.** An important parameter for TAFSSL is the number of the dimensions of the sub-space selected by the TAFSSL approach. In figure 4 we explore the effect of the number of chosen dimensions in ICA-based TAFSSL on both the Mini and the Tiered benchmarks. As can be seen from the figure, the optimal number of dimensions for ICA-based TAFSSL approaches is 10, which is consistent between both test and validation sets. Interestingly, the same number 10 is consistent between the two




Fig. 3: **The affect of the unlabeled data noise.** Plots for LST [24], TPN [28], and SKM [38] are from original papers. (a) *miniImageNet*; (b) *tieredImageNet*

Table 3: Backbones comparison. The 1-shot transductive FSL setting for *miniImageNet* (Mini) and *tieredImageNet* (Tiered) was used for this comparison

	Backbone	Mini 1-shot	Tiered 1-shot
<b>TPN [28]</b>	Conv-4	55.51 $\pm$ 0.86	59.91 $\pm$ 0.94
<b>TPN [28]</b>	ResNet-12	59.46 $\pm$ N.A.	-
<b>Transductive Fine-Tuning [8]</b>	WRN	65.73 $\pm$ 0.68	73.34 $\pm$ 0.71
<b>PCA + MSP</b>	Conv-4	56.63 $\pm$ 0.27	60.27 $\pm$ 0.29
<b>PCA + MSP</b>	ResNet-10	70.93 $\pm$ 0.28	76.27 $\pm$ 0.28
<b>PCA + MSP</b>	ResNet-18	73.73 $\pm$ 0.27	80.60 $\pm$ 0.27
<b>PCA + MSP</b>	WRN	73.72 $\pm$ 0.27	81.61 $\pm$ 0.26
<b>PCA + MSP</b>	DenseNet	76.31 $\pm$ 0.26	84.06 $\pm$ 0.25

benchmarks. Similarly, using validation, the optimal dimension for PCA-based TAFSSL was found to be 4 (also consistently on the two benchmarks).

### Backbone architectures.

The choice of backbone turned out to be an important factor for FSL methods performance [3, 52]. In Table 3 we evaluate the performance of one of the proposed TAFSSL approaches, namely PCA+MSP while using different backbones pre-trained on the training set to compute the base feature space  $\mathcal{F}$ . We used the 1-shot transductive FSL setting on both Mini and Tiered benchmarks for this evaluation. As can be seen from the table, larger backbones produce

better performance for the TAFSSL approach. In addition, we list the reported performance of the competing SOTA transductive FSL methods in the same table for direct comparison using the same backbones. As can be seen, above 8% accuracy advantage is maintained by our proposed approach above the top previous method using the corresponding WRN architecture.

### Unbalanced (long-tail) test classes distribution in unlabeled data.

In all previous transductive FSL works, balanced test tasks (in terms of the number of queries for each class) were used. While this is fine for experimental evaluation, in practical applications there is no guarantee that the bulk of queries sent for offline evaluation will be class-balanced. In fact, class-skew (lack of balance) is much more likely. To test the effect of query set skew, we have evaluated the proposed ICA-based TAFSSL approaches, as well as the Simple-Shot [52] and its




Fig. 4: **ICA dim. vs accuracy.** (a) *miniImageNet* (b) *tieredImageNet*

transductive adaptation ”trans-mean-sub\*” (sub) baselines, under varying levels of this skew. The level of skew was controlled through the so-called ”unbalanced factor” parameter  $R$ : in each test episode, for each class  $15 + uni([0, R])$  query samples were randomly chosen (here  $uni$  refers to a uniform distribution). Figure 5 shows the effect of varying  $R$  from 10 to 50, while at the extreme setting (50) above factor 4 skew is possible between the classes in terms of the number of associated queries. Nevertheless, as can be seen from the figure, the effect of lack of balance on the performance of the TAFSSL based approaches is minimal, leading to at most 1% performance loss at  $R = 50$ . Since no prior work offered a similar evaluation design, we believe that the proposed protocol may become an additional important tool for evaluating transductive FSL methods under lack of query set balance in the future.

## 5 Summary and Conclusions

In this paper we have shown that the Feature Sub-Space Learning (FSSL), and specifically it’s Task Adaptive variant (TAFSSL), has a significant positive effect on FSL performance. We have explored different methods for benefiting from TAFSSL in FSL and have shown great promise for this kind of techniques by achieving large margin improvements over transductive and semi-supervised FSL state-of-the-art, as well as over methods that do not use additional unlabeled

data, thus highlighting the benefit of the latter. Potential future work directions include using TAFSSL with meta-training, e.g. by propagating training episodes gradients through pyTorch PCA/ICA implementations, and the proposed clustering techniques BKM/MSP; exploring non-linear TAFSSL variants (e.g. kernel or DNN TAFSSL); further exploring the effect of TAFSSL in *any-shot* learning and the significance of the *way* parameter of the task; exploring the benefits of TAFSSL in cross-domain few-shot learning where the FSL backbone pre-training occurs in different visual domain from the one test classes are sampled from.

**Acknowledgment.** This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. FA8750-19-C-1001. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA. Raja Giryes is supported by ERC-StG grant no. 757497 (SPADE).




Fig. 5: **Unbalanced.** (a) *miniImageNet* (b) *tieredImageNet*

## References

1. Alfassy, A., Karlinsky, L., Aides, A., Shtok, J., Harary, S., Feris, R., Giryas, R., Bronstein, A.M.: LaSO: Label-Set Operations networks for multi-label few-shot learning. In: CVPR (2019)
2. Antoniou, A., Storkey, A., Edwards, H.: Data Augmentation Generative Adversarial Networks. arXiv:1711.04340 (2017), <https://arxiv.org/pdf/1711.04340.pdf>
3. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A Closer Look At Few-Shot Classification. In: ICLR (2019)
4. Chen, Z., Fu, Y., Zhang, Y., Jiang, Y.G., Xue, X., Sigal, L.: Multi-Level Semantic Feature Augmentation for One-Shot Learning. IEEE Transactions on Image Processing **28**(9), 4594–4605 (2019). <https://doi.org/10.1109/tip.2019.2910052>
5. Comon, P.: Independent component analysis, A new concept? Tech. rep. (1994)
6. Cubuk, E.D., Zoph, B., Mané, D., Vasudevan, V., Le, Q.V.: AutoAugment: Learning Augmentation Policies from Data <https://arxiv.org/pdf/1805.09501v1.pdf>
7. Devos, A., Grossglauser, M.: Subspace Networks for Few-shot Classification. Tech. rep. (2019)
8. Dhillon, G.S., Chaudhari, P., Ravichandran, A., Soatto, S.: A Baseline For Few-Shot Image Classification. Tech. rep. (2019)
9. Dosovitskiy, A., Springenberg, J.T., Tatarchenko, M., Brox, T.: Learning to Generate Chairs, Tables and Cars with Convolutional Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence **39**(4), 692–705 (2017). <https://doi.org/10.1109/TPAMI.2016.2567384>
10. Doveh, S., Schwartz, E., Xue, C., Feris, R., Bronstein, A., Giryas, R., Karlinsky, L.: MetAdapt: Meta-Learned Task-Adaptive Architecture for Few-Shot Classification. Tech. rep. (2019)
11. Dvornik, N., Schmid, C., Mairal, J.: Diversity with Cooperation: Ensemble Methods for Few-Shot Classification. The IEEE International Conference on Computer Vision (ICCV) (2019), <http://arxiv.org/abs/1903.11341>
12. Finn, C., Abbeel, P., Levine, S.: Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. arXiv:1703.03400 (2017), <http://arxiv.org/abs/1703.03400>
13. Fukunaga, K., Hostetler, L.D.: The Estimation of the Gradient of a Density Function, with Applications in Pattern Recognition. IEEE Transactions on Information Theory **21**(1), 32–40 (1975). <https://doi.org/10.1109/TIT.1975.1055330>
14. Garcia, V., Bruna, J.: Few-Shot Learning with Graph Neural Networks. arXiv:1711.04043 pp. 1–13 (2017), <http://arxiv.org/abs/1711.04043>
15. Gidaris, S., Bursuc, A., Komodakis, N., Pérez, P., Cord, M.: Boosting Few-Shot Visual Learning with Self-Supervision (6 2019), <http://arxiv.org/abs/1906.05186>
16. Guu, K., Hashimoto, T.B., Oren, Y., Liang, P.: Generating Sentences by Editing Prototypes. Arxiv:1709.08878 (2017), <https://arxiv.org/pdf/1709.08878.pdf>
17. Hariharan, B., Girshick, R.: Low-shot Visual Recognition by Shrinking and Hallucinating Features. IEEE International Conference on Computer Vision (ICCV) (2017), <https://arxiv.org/pdf/1606.02819.pdf>
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. arXiv:1512.03385 (2015), <https://arxiv.org/pdf/1512.03385.pdf>
19. Huang, G., Liu, Z., v. d. Maaten, L., Weinberger, K.Q.: Densely Connected Convolutional Networks. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 2261–2269 (2017). <https://doi.org/10.1109/CVPR.2017.243>, <https://arxiv.org/pdf/1608.06993.pdf>

20. Jiang, X., Havaei, M., Varno, F., Chartrand, G.: Learning To Learn With Conditional Class Dependencies pp. 1–11 (2019)
21. Kim, J., Kim, T., Kim, S., Yoo, C.D.: Edge-Labeling Graph Neural Network for Few-shot Learning. Tech. rep.
22. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-Learning with Differentiable Convex Optimization. In: CVPR (2019), <https://github.com/kjunelee/MetaOptNet>
23. Li, H., Eigen, D., Dodge, S., Zeiler, M., Wang, X.: Finding Task-Relevant Features for Few-Shot Learning by Category Traversal **1** (2019), <http://arxiv.org/abs/1905.11116>
24. Li, X., Sun, Q., Liu, Y., Zheng, S., Zhou, Q., Chua, T.S., Schiele, B.: Learning to Self-Train for Semi-Supervised Few-Shot Classification (6 2019), <http://arxiv.org/abs/1906.00562>
25. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-SGD: Learning to Learn Quickly for Few-Shot Learning. arXiv:1707.09835 (2017), <http://arxiv.org/abs/1707.09835>
26. Lim, S., Kim, I., Kim, T., Kim, C., Brain, K., Kim, S.: Fast AutoAugment. Tech. rep. (2019)
27. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: Lecture Notes in Computer Science. vol. 8693 LNCS, pp. 740–755 (2014). [https://doi.org/10.1007/978-3-319-10602-1\\_48](https://doi.org/10.1007/978-3-319-10602-1_48)
28. Liu, Y., Lee, J., Park, M., Kim, S., Yang, E., Hwang, S.J., Yang, Y.: Learning To Propagate Labels: Transductive Propagation Network For Few-Shot Learning (2019)
29. Lloyd, S.P., Lloyd, S.P.: Least squares quantization in pcm. IEEE TRANSACTIONS ON INFORMATION THEORY **28**, 129–137 (1982), <https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.131.1338>
30. Munkhdalai, T., Yu, H.: Meta Networks. arXiv:1703.00837 (2017). <https://doi.org/10.1093/mnrasl/slx008>, <http://arxiv.org/abs/1703.00837>
31. Nakamura, A., Harada, T.: REVISITING FINE-TUNING FOR FEW-SHOT LEARNING. Tech. rep.
32. Oreshkin, B.N., Rodriguez, P., Lacoste, A.: TADAM: Task dependent adaptive metric for improved few-shot learning. NeurIPS (5 2018), <http://arxiv.org/abs/1805.10123>
33. Park, D., Ramanan, D.: Articulated pose estimation with tiny synthetic videos. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) **2015-October**, 58–66 (2015). <https://doi.org/10.1109/CVPRW.2015.7301337>
34. Pearson, K.: On lines and planes of closest fit to systems of points in space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **2**(11), 559–572 (11 1901). <https://doi.org/10.1080/14786440109462720>
35. Qiao, L., Shi, Y., Li, J., Wang, Y., Huang, T., Tian, Y.: Transductive Episodic-Wise Adaptive Metric for Few-Shot Learning (2019), <http://arxiv.org/abs/1910.02224>
36. Ravi, S., Larochelle, H.: Optimization As a Model for Few-Shot Learning. ICLR pp. 1–11 (2017), <https://openreview.net/pdf?id=rJY0-Kcll>
37. Reed, S., Chen, Y., Paine, T., van den Oord, A., Eslami, S.M.A., Rezende, D., Vinyals, O., de Freitas, N.: Few-shot autoregressive density estimation: towards learning to learn distributions. arXiv:1710.10304 (2016), 1–11 (2018)
38. Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-Learning for Semi-Supervised Few-Shot Classification. ICLR (3 2018), <http://arxiv.org/abs/1803.00676> <http://bair.berkeley.edu/blog/2017/07/18/>



39. Rippel, O., Paluri, M., Dollar, P., Bourdev, L.: Metric Learning with Adaptive Density Discrimination. arXiv:1511.05939 pp. 1–15 (2015), <http://arxiv.org/abs/1511.05939>
40. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. IJCV (9 2015), <http://arxiv.org/abs/1409.0575>
41. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-Learning with Latent Embedding Optimization. In: ICLR (7 2018), <http://arxiv.org/abs/1807.05960>
42. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised Domain Adaptation via Minimax Entropy. In: ICCV (4 2019), <http://arxiv.org/abs/1904.06487>
43. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T.: Meta-Learning with Memory-Augmented Neural Networks. Journal of Machine Learning Research **48**(Proceedings of The 33rd International Conference on Machine Learning), 1842–1850 (2016). <https://doi.org/10.1002/2014GB005021>
44. Schwartz, E., Karlinsky, L., Feris, R., Giryes, R., Bronstein, A.M.: Baby steps towards few-shot learning with multiple semantics pp. 1–11 (2019), <http://arxiv.org/abs/1906.01905>
45. Schwartz, E., Karlinsky, L., Shtok, J., Harary, S., Marder, M., Kumar, A., Feris, R., Giryes, R., Bronstein, A.M.: Delta-Encoder: an Effective Sample Synthesis Method for Few-Shot Object Recognition. NeurIPS (2018), <https://arxiv.org/pdf/1806.04734.pdf>
46. Simon, C., Koniusz, P., Harandi, M.: Projective Sub-Space Networks For Few-Sot Learning. In: ICLR 2019 OpenReview. <https://openreview.net/pdf?id=rkzfuiA9F7>
47. Snell, J., Swersky, K., Zemel, R.: Prototypical Networks for Few-shot Learning. In: NIPS (2017), <http://arxiv.org/abs/1703.05175>
48. Su, H., Qi, C.R., Li, Y., Guibas, L.J.: Render for CNN Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views.pdf. IEEE International Conference on Computer Vision (ICCV) pp. 2686–2694 (2015)
49. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H.S., Hospedales, T.M.: Learning to Compare: Relation Network for Few-Shot Learning <https://arxiv.org/pdf/1711.06025.pdf>
50. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H.S., Hospedales, T.M.: Learning to Compare: Relation Network for Few-Shot Learning. arXiv:1711.06025 (11 2017), <http://arxiv.org/abs/1711.06025>
51. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching Networks for One Shot Learning. NIPS (2016). <https://doi.org/10.1109/CVPR.2016.95>, <http://arxiv.org/abs/1606.04080>
52. Wang, Y., Chao, W.L., Weinberger, K.Q., van der Maaten, L.: SimpleShot: Revisiting Nearest-Neighbor Classification for Few-Shot Learning (11 2019), <http://arxiv.org/abs/1911.04623>
53. Wang, Y.X., Girshick, R., Hebert, M., Hariharan, B.: Low-Shot Learning from Imaginary Data. arXiv:1801.05401 (2018), <http://arxiv.org/abs/1801.05401>
54. Weinberger, K.Q., Saul, L.K.: Distance Metric Learning for Large Margin Nearest Neighbor Classification. The Journal of Machine Learning Research **10**, 207–244 (2009). <https://doi.org/10.1126/science.277.5323.215>

55. Xing, C., Rostanzadeh, N., Oreshkin, B.N., Pinheiro, P.O.: Adaptive Cross-Modal Few-Shot Learning (2019), <https://arxiv.org/pdf/1902.07104.pdf>  
<http://arxiv.org/abs/1902.07104>
56. Yu, A., Grauman, K.: Semantic Jitter: Dense Supervision for Visual Comparisons via Synthetic Images. Proceedings of the IEEE International Conference on Computer Vision **2017-October**, 5571–5580 (2017). <https://doi.org/10.1109/ICCV.2017.594>
57. Zhang, J., Zhao, C., Ni, B., Xu, M., Yang, X.: Variational Few-Shot Learning. In: IEEE International Conference on Computer Vision (ICCV) (2019)
58. Zhou, F., Wu, B., Li, Z.: Deep Meta-Learning: Learning to Learn in the Concept Space. Tech. rep. (2 2018), <http://arxiv.org/abs/1802.03596>