

Tensor feature hallucination for few-shot learning

Michalis Lazarou¹ Tania Stathaki¹ Yannis Avrithis²

¹Imperial College London

²Athena RC

Abstract

Few-shot learning addresses the challenge of learning how to address novel tasks given not just limited supervision but limited data as well. An attractive solution is synthetic data generation. However, most such methods are overly sophisticated, focusing on high-quality, realistic data in the input space. It is unclear whether adapting them to the few-shot regime and using them for the downstream task of classification is the right approach. Previous works on synthetic data generation for few-shot classification focus on exploiting complex models, e.g. a Wasserstein GAN with multiple regularizers or a network that transfers latent diversities from known to novel classes.

We follow a different approach and investigate how a simple and straightforward synthetic data generation method can be used effectively. We make two contributions, namely we show that: (1) using a simple loss function is more than enough for training a feature generator in the few-shot setting; and (2) learning to generate tensor features instead of vector features is superior. Extensive experiments on miniImagenet, CUB and CIFAR-FS datasets show that our method sets a new state of the art, outperforming more sophisticated few-shot data augmentation methods. The source code can be found at https://github.com/MichalisLazarou/TFH_fewshot.

1. Introduction

Deep learning continuously improves the state of the art in different fields, such as *natural language understanding* [42] and *computer vision* [30]. However, a fundamental limitation of representation learning from raw data is the dependence on large amounts of task-specific or domain-specific data, labeled or not. This limitation inhibits the application of deep learning to real-world problems, such as rare species classification, where the cost of obtaining and annotating data from a new domain is high.

To address this limitation, *few-shot learning* [65, 58, 10] has attracted significant interest in recent years. Few-shot learning is concerned with learning not only under limited




Figure 1. CUB original images (row 1) followed by images generated from separately trained reconstructors using as input tensor features (row 2) or vector features (row 3). More results and implementation details are given in the supplementary material.

supervision, but also from limited data. This constraint excludes representation learning from scratch and inhibits adapting the representation, which is otherwise common in *transfer learning* [9, 28] *domain/task adaptation* [11, 51] and *continual learning* [52].

Data augmentation, commonly based on simple input transformations, is a universal way of regularizing and improving the generalization ability of a model [30], as well as exploiting unlabeled data [5, 59]. In few-shot learning, recent methods go beyond input transformations towards *synthetic data generation* and *hallucination*, either in the *image space* [7, 73] or in the *feature space* [8, 34, 40]. Hence, they address the data deficiency by augmenting real data with synthetic, achieving a greater extent of diversity.

The vast majority of generative models focuses on high-quality, high-resolution images, assuming a large amount of data. Also, the metrics used to evaluate generative models, focus on whether the generated data is realistic [18, 54]. Generating high quality, realistic data may not be necessary in “downstream tasks” such as classification. It is unclear whether and how state of the art generative models in the image space can succeed in the few-shot setting.

Most of the recent few-shot feature hallucination methods focus on generating vectors in the feature space [8, 34, 40]. These vectors are most commonly obtained by *global average pooling* (GAP) on the output feature maps. This

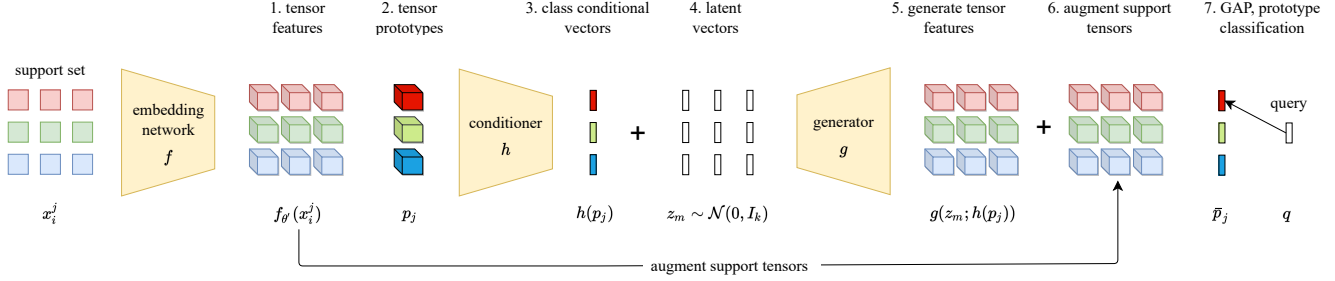


Figure 2. *Overview of our method.* At inference: 1) Map the support examples x_i^j (each color indicates a different class j) into tensor features $f_{\theta'}(x_i^j)$ through the pre-trained embedding network $f_{\theta'}$. 2) Average x_i^j into a tensor prototype p_j per class j (3). 3) Map each p_j to a class conditional vector $h(p_j)$ through the *conditioner network* h . Draw M samples z_m per class from a k -dimensional normal distribution $\mathcal{N}(\mathbf{0}, I_k)$. 5) Generate M class-conditional tensor features $g(z_m; h(p_j))$ per class j using *generator network* g . 6) Augment the support tensors with the generated tensors. 7) Perform *global average pooling* (GAP) and average the augmented features into vector prototype \bar{p}_j per class j (5) and classify queries q to nearest prototype. At training (not shown): a) Train f_{θ} using cross-entropy (1). b) Fine-tune f_{θ} to $f_{\theta'}$ using self-distillation (2). c) Train *tensor feature hallucinator* (TFH) $\{h, g\}$ using reconstruction loss (4).

discards spatial details that might be necessary to model the underlying data distribution. We hypothesize that working with the feature map tensors directly may be more effective. To investigate this, we train two image reconstructors separately: one for tensor features and the other for vector features obtained by GAP. The latter has the same architecture as the former, except for one additional upsampling layer. As shown in Figure 1, the feature map tensors preserve more information indeed.

Motivated by this finding, we explore the potential of using *tensor features* instead of vector features in a simple generative model to improve few-shot classification. We employ a simple *conditioner-generator* architecture and we introduce a simple *reconstruction loss* between the generated tensor and the corresponding class prototype tensor. This allows the generation of a diverse set of synthetic data, not necessarily realistic, from a limited amount of real data from a previously unseen task. An overview is shown in Figure 2. We demonstrate empirically that our model provides state of the art results, outperforming more sophisticated generative models on a number of benchmarks. Our contributions are summarized as follows:

1. We are the first to generate *tensor features* instead of vector features in the few-shot setting and to leverage their structural properties (subsection 3.3).
2. We introduce a novel loss function that is simpler than alternatives in state of the art few-shot synthetic data generation methods [34, 40, 4] (subsection 3.4).
3. Our *tensor feature hallucinator* (TFH) sets new state of the art on three common few-shot classification benchmarks: *miniImagenet*, CUB and CIFAR-FS.
4. We demonstrate the robustness of our hallucinator against using different backbone networks and classifiers, as well as its applicability to the challenging

setting of *cross-domain* few-shot learning.

2. Related work

2.1. Few-shot learning

In few-shot learning, the objective is to learn from an abundant labeled set of *base* classes how to solve tasks from a limited *support* set over a distinct set of *novel* classes. We briefly discuss different approaches to this objective, followed by a more detailed account of *synthetic data generation*, where our contribution lies.

Meta-learning The objective of few-shot classification fits naturally within *meta-learning* [56, 64], referring to learning at two levels, where generic knowledge is acquired before adapting to more specific tasks. There are different instantiations of this idea, all sharing the fact that the loss is computed on a *set* of examples, called an *episode*.

Optimization meta-learning aims to learn how to quickly update the model parameters to novel tasks without overfitting. This includes updates in closed form [2], iterative updates according to the gradient descent [10, 74, 44] and learnable iterative updates, such as LSTM [50].

Model-based meta-learning aims to learn how to update specific model architectures to novel tasks. This includes memory-augmented neural networks [55] and meta-networks, designed explicitly according to the two-level learning paradigm [43].

Metric learning is a standalone field that overlaps meta-learning and aims to learn how to compare examples of unseen classes [45]. Most often, it amounts to learning an *embedding space* where distances or similarities are taken between a query and individual examples [27], all examples of a class [65], or the class centroid [58]. A metric or similarity function may also be learned directly [60].

Representation learning Instead of learning in episodes, it is simpler to compute the loss on one example at a time, like standard cross-entropy. Learning a classifier on the base classes then amounts to *representation learning*. This simplified approach has been popularized in few-shot classification with the *cosine-based classifier* [49, 13, 36, 6]. Any method that helps in learning a better representation is applicable in this sense, including pretext tasks [14], self-distillation [61] and manifold mixup [41].

Task adaptation In few-shot learning, the challenge is to adapt the representation to novel tasks on limited data without overfitting. This is possible using metric scaling [47], attention mechanisms [13], task-adaptive projections [68], set-to-set functions [67], identifying task-relevant features [33] or growing the architecture [36].

Unlabeled data The constraint of few-shot learning may be relaxed by accessing more novel-class unlabeled data (*semi-supervised learning*) or multiple queries at a time (*transductive inference*). This allows exploiting the *manifold structure* of the additional data, for instance using graph neural networks [12, 24], label propagation [39], embedding propagation [53] or label cleaning [31].

2.2. Synthetic data generation

Generative models aim to model the underlying data distribution of the training set in a *latent space*. *Generative adversarial networks* (GAN) [15] are by far the most popular approach. The idea is a zero-sum game between a generator and a discriminator, such that generated data are realistic enough to be indistinguishable from real. Several improvements concern the architecture, as well as training and regularization methods [23, 3, 21].

Other approaches include *variational autoencoders* (VAE) [26], imposing a prior distribution in the latent space, *autoregressive* (AR) *models* [63, 46], iteratively generating samples conditioned on previous steps, and *flow-based models* [25, 20], modeling the data distribution explicitly via invertible transformations.

Most state of the art generative models do not focus on improving the performance of downstream tasks such as classification, but rather on image quality metrics such as *Fréchet inception distance* (FID) [18] and *inception score* (IS) [54]. They also assume access to abundant training data, which is in direct contrast to the few-shot setting. Even though recent methods address small datasets [22, 37], they are limited to unconditional image generation, which inhibits their use in novel tasks.

2.3. Synthetic data generation for few-shot learning

In few-shot learning, a generative model can be learned on base-class data to augment real novel-class data with synthetic. One of the first ideas in this direction is *fea-*

ture hallucination [16]. There are several approaches based on GANs, including MetaGAN [73], which integrates MAML [10] with a conditional GAN to generate examples in the input space; AFHN [34], a feature hallucinator using wGAN [1]; and FUNIT [38], a GAN-based method for few-shot image-to-image translation.

There are also alternative approaches to GANs, including VI-Net [40], which uses a class-conditional VAE as a feature hallucinator; *diversity transfer network* (DTN) [4], which learns how to transfer latent diversities from base to novel classes; SalNet [71], which hallucinates features by combining foregrounds and backgrounds from different images; and IDeMe-Net [7], which combines novel-class support with similar base-class images.

All the aforementioned methods, including the current state of the art, use overly complex training regimes, adapting them to the few-shot setting. For instance, AFHN [34] adapts the Wasserstein GAN [1] and VI-Net [40] adapts a variational autoencoder [26]. It is not clear whether such sophisticated generative models and loss functions are necessary for a “downstream task” like few-shot classification. At the same time, generating vector features incurs information loss as demonstrated in Figure 1.

2.4. On our contribution

Our work falls within generating synthetic data in the *feature space*. However, we are the first to train a model to generate *tensor features* instead of vector features in few-shot learning, exploiting the spatial and structural properties of tensors (subsection 3.3). This allows us to use a *simple reconstruction loss* between the generated tensors and their class prototype (subsection 3.4), while still outperforming methods using overly complex generative models.

Our model bears similarities to a VAE [26], also used by VI-Net [40]. Our *conditioner* and *hallucinator* networks play a similar role to encoder and decoder, respectively. However, rather than predicting the variance and imposing a prior distribution in the latent space, we condition the model on *class prototypes*, also represented by tensor features, and we use the same prototypes in the reconstruction loss.

To improve representation learning, we use *self-distillation* as an auxiliary loss term, following [61] (subsection 3.2). We also perform task adaptation by *fine-tuning* the hallucinator to novel class data for few iterations.

3. Method

3.1. Problem formulation

We are given a labeled dataset $D_{\text{base}} := \{(x_i, y_i)\}_{i=1}^I$ of I examples, with each example x_i having a label y_i in one of the classes in C_{base} . This dataset is used to learn the parameters θ of a mapping $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{d \times h \times w}$ from an input image space \mathcal{X} to a *feature* or *embedding* space, where *fea-*

ture tensors have d dimensions (channels) and spatial resolution $h \times w$ (height \times width).

The knowledge acquired at representation learning is used to solve *novel tasks*, assuming access to a dataset D_{novel} , with each example being associated with one of the classes in C_{novel} , where C_{novel} is disjoint from C_{base} . In *few-shot classification* [65], a novel task is defined by sampling a *support set* S from D_{novel} , consisting of N classes with K labeled examples per class, for a total of $L := NK$ examples. Given the mapping f_θ and the support set S , the problem is to learn an N -way classifier that makes predictions on unlabeled *queries*, also sampled from novel classes. Queries are treated independently of each other. This is referred to as *inductive inference*.

3.2. Representation learning

The goal of *representation learning* is to learn the embedding function f_θ that can be applied to D_{novel} to extract embeddings and solve novel tasks. We use f_θ followed by *global average pooling* (GAP) and a parametric *base classifier* c_ϕ to learn the representation. We denote by $\bar{f}_\theta : \mathcal{X} \rightarrow \mathbb{R}^d$ the composition of f_θ and GAP. We follow the two-stage regime by [61] to train our embedding model. In the *first stage*, we train f_θ on D_{base} using standard cross-entropy loss L_{CE} :

$$J(D_{\text{base}}; \theta, \phi) := \sum_{i=1}^I \ell_{\text{CE}}(c_\phi(\bar{f}_\theta(x_i)), y_i) + R(\phi), \quad (1)$$

where R is a regularization term. In the *second stage*, we adopt a *self-distillation* process: The embedding model f_θ and classifier c_ϕ from the first stage serve as the teacher and we distill their knowledge to a new student model $f_{\theta'}$ and classifier $c_{\phi'}$, with identical architecture. The student is trained using a linear combination of the standard cross-entropy loss, as in stage one, and the Kullback-Leibler (KL) divergence between the student and teacher predictions:

$$J_{\text{KD}}(D_{\text{base}}; \theta', \phi') := \alpha J(D_{\text{base}}; \theta', \phi') + \beta \text{KL}(c_{\phi'}(\bar{f}_{\theta'}(x_i)), c_\phi(\bar{f}_\theta(x_i))), \quad (2)$$

where α and β are scalar weights and θ, ϕ are fixed.

3.3. Feature tensor hallucinator

Existing feature hallucination methods [34, 4, 40, 72, 16] are trained using *vector features*, losing significant spatial and structural information. By contrast, our hallucinator is trained on tensor features before GAP and generates *tensor features* as well. In particular, we use the student model $f_{\theta'} : \mathcal{X} \rightarrow \mathbb{R}^{d \times h \times w}$, pre-trained using (2), as our embedding network to train our tensor feature hallucinator.

The hallucinator consists of two networks: a *conditioner* network h and a *generator* network g . The conditioner

aids the generator in generating class-conditional examples. Given a set $X_j := \{x_i^j\}_{i=1}^K$ of examples associated with each class $j = 1, \dots, N$, conditioning is based on the *prototype tensor* $p_j := p(X_j) \in \mathbb{R}^{d \times h \times w}$ of class j ,

$$p(X_j) := \frac{1}{K} \sum_{i=1}^K f_{\theta'}(x_i^j). \quad (3)$$

The conditioner $h : \mathbb{R}^{d \times h \times w} \rightarrow \mathbb{R}^{d'}$ maps the prototype tensor to the *class-conditional vector* $s_j := h(p_j) \in \mathbb{R}^{d'}$. The generator $g : \mathbb{R}^{k+d'} \rightarrow \mathbb{R}^{d \times h \times w}$ takes as input this vector as well as a *latent vector* $z \sim \mathcal{N}(\mathbf{0}, I_k)$ drawn from a k -dimensional standard normal distribution and generates a *class-conditional tensor feature* $g(z; s_j) \in \mathbb{R}^{d \times h \times w}$ for each class j .

Algorithm 1: Meta-training of tensor hallucinator

input : training set D_{base}
input : pre-trained embedding f_θ
output: trained tensor hallucinator $\{h, g\}$

- 1 **while** *not done* **do**
- 2 Sample an N -way K -shot episode $E := \{E_j\}_{j=1}^N$ from D_{base}
- 3 **for** class $j = 1, \dots, N$ **do**
- 4 Obtain class prototype tensor $p_j := p(E_j)$ by (3)
- 5 Map p_j to class-conditional vector $s_j := h(p_j)$
- 6 Draw M samples $\{z_m\}_{m=1}^M$ from $\mathcal{N}(\mathbf{0}, I_k)$
- 7 Generate M class-conditional tensor features $\{g(z_m; s_j)\}_{m=1}^M$
- 8 Update parameters of hallucinator $\{h, g\}$ by (4)

3.4. Training the hallucinator

We train our hallucinator using a meta-training regime, similar to [34, 8, 57]. At every iteration, we sample a new episode by randomly sampling N classes and K examples $E_j := \{x_i^j\}_{i=1}^K$ for each class j from D_{base} . For each class $j = 1, \dots, N$, we obtain the prototype tensor $p_j := p(E_j)$ using (3) and the class-conditional vector $s_j := h(p_j)$ by the conditioner h . We then draw M samples $\{z_m\}_{m=1}^M$ from the standard normal distribution $\mathcal{N}(\mathbf{0}, I_k)$ and generate M class-conditional tensor features $\{g(z_m; s_j)\}_{m=1}^M$ using the generator g . We train our hallucinator $\{h, g\}$ on the episode data $E := \{E_j\}_{j=1}^N$ by minimizing the *mean squared error* (MSE) of generated class-conditional tensor features of class j to the corresponding class prototype p_j :

$$J_{\text{H}}(E; h, g) = \frac{1}{MN} \sum_{j=1}^N \sum_{m=1}^M \|g(z_m; h(p_j)) - p_j\|^2. \quad (4)$$

Algorithm 1 summarizes the overall training process of the hallucinator.

Algorithm 2: Using hallucinator at inference

input : support set $S := \{S_j\}_{j=1}^N$
input : pre-trained embedding f_θ
input : pre-trained hallucinator $\{h, g\}$
output: predicted label for query $q \in \mathcal{X}$

- 1 **for** class $j = 1, \dots, N$ **do**
 - 2 Obtain class prototype tensor $p_j := p(S_j)$ by (3)
 - 3 Draw M samples $\{z_m\}_{m=1}^M$ from $\mathcal{N}(\mathbf{0}, I_k)$
 - 4 Generate M class-conditional tensor features
 $G_j := \{g(z_m; h(p_j))\}_{m=1}^M$
 - 5 Augment the support S_j with generated features G_j
 - 6 Apply GAP and obtain vector class prototypes \bar{p}_j by (5)
 - 7 Assign query q to class of nearest prototype to feature $\bar{f}_{\theta'}(q)$
-

3.5. Inference

At inference, we are given a few-shot task with a support set $S := \{S_j\}_{j=1}^N$, containing N novel classes with K examples $S_j := \{x_i^j\}_{i=1}^K$ for each class j . For each class $j = 1, \dots, N$, we use our trained backbone network $f_{\theta'}$ to compute the tensor feature $f_{\theta'}(x_i^j) \in \mathbb{R}^{d \times h \times w}$ of each example in S_j and we obtain the prototype $p_j := p(S_j)$ by (3). Then, using our trained tensor feature hallucinator $\{h, g\}$, we generate M class-conditional tensor features $G_j := \{g(z_m; h(p_j))\}_{m=1}^M$, also in $\mathbb{R}^{d \times h \times w}$, where z_m are drawn from $\mathcal{N}(\mathbf{0}, I_k)$. We augment the support features $f_{\theta'}(S_j)$ with the generated features G_j , resulting in $K + M$ labeled tensor features per class in total. We now apply GAP to those tensor features and obtain new, *vector class prototypes* in \mathbb{R}^d :

$$\bar{p}_j := \frac{1}{K + M} \left(\sum_{i=1}^K \bar{f}_{\theta'}(x_i^j) + \sum_{m=1}^M \bar{g}(z_m; h(p_j)) \right), \quad (5)$$

where \bar{g} denotes the composition of g and GAP. Finally, given a query $q \in \mathcal{X}$, we apply GAP to the tensor feature $f_{\theta'}(q)$ and assign it to the class of the nearest vector prototype. Algorithm 2 summarizes the inference process.

We refer to the above approach as *prototypical classifier*. In subsection 4.6 we experiment with alternative classifiers such as logistic regression and support vector machine on the same augmented (support + generated) features.

4. Experiments

4.1. Datasets

We use three common few-shot classification datasets: *miniImageNet* [65, 50], CUB [6, 19] and CIFAR-FS [6, 29]. More details are given in the supplementary material.

4.2. Networks

Our *tensor feature hallucinator* (TFH) consists of a conditioner network and a generator network. Their architecture depends on the backbone network.

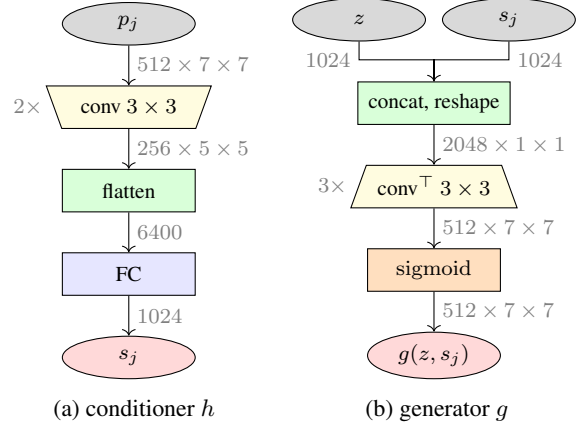


Figure 3. Architecture of our (a) conditioner network h , (b) generator network g for ResNet-18. See text for more details and for ResNet-12. conv^\top : transpose convolution.

Backbone Many recent data augmentation methods [8, 7, 34, 40] use ResNet-18 [17] as a backbone embedding model. To perform as fair comparison as possible, we use this backbone by default. To investigate the transferability and robustness of our tensor hallucinator, we also use a pre-trained ResNet-12 backbone from the publicly available code of DeepEMD [70].

For ResNet-18, the embedding dimension is $d = 512$ and the resolution $h \times w = 7 \times 7$. For ResNet-12, it is $d = 640$ and $h \times w = 5 \times 5$.

Conditioner As shown in Figure 3(a), our conditioner $h : \mathbb{R}^{d \times h \times w} \rightarrow \mathbb{R}^{d'}$ consists of two convolutional layers with a ReLU activation in-between, followed by flattening and a fully-connected layer. The convolutional layers use kernels of size 3×3 , and stride 1. In the first convolutional layer, we also use padding 1. The output channels are d and $d/2$ in the first and second layer, respectively. The dimension of the *class-conditional vector* is $d' = 1024$.

For ResNet-18, the tensor dimensions of all conditioner layers are $[512 \times 7 \times 7]$, $[512 \times 7 \times 7]$, $[256 \times 5 \times 5]$, $[6400]$ (flattening) and $[1024]$. For ResNet-12, they are $[640 \times 5 \times 5]$, $[640 \times 5 \times 5]$, $[320 \times 3 \times 3]$, $[2880]$ (flattening) and $[1024]$.

Generator As shown in Figure 3(b), for ResNet-18, our generator $g : \mathbb{R}^{k+d'} \rightarrow \mathbb{R}^{d \times h \times w}$ consists of concatenation of z and s_j into $(z; s_j) \in \mathbb{R}^{k+d'}$, followed by reshaping to $(k + d') \times 1 \times 1$, three transpose-convolutional layers with ReLU activations in-between and a sigmoid function at the end. For ResNet-12, the generator architecture is the same, except that it has only *two* transpose-convolutional layers. The dimension of the latent vector z is $k = 1024$. All transpose-convolutional layers use kernels of size 3×3 , stride 1 and d output channels.

For ResNet-18, the tensor dimensions of all generator layers are $[2048 \times 1 \times 1]$, $[512 \times 3 \times 3]$, $[512 \times 5 \times 5]$,

BACKBONE	1-SHOT		5-SHOT	
	k	γ	k	γ
ResNet-18	15	10^{-7}	10	10^{-4}
ResNet-12	10	10^{-7}	10	10^{-4}

Table 1. Number of steps t and learning rate η for TFH-ft, chosen on *miniImageNet* validation set and used in all experiments.

and $[512 \times 7 \times 7]$. For ResNet-12, they are $[2048 \times 1 \times 1]$, $[640 \times 3 \times 3]$, and $[640 \times 5 \times 5]$.

4.3. Training

Embedding model Similarly to [61], we use SGD optimizer with learning rate 0.05, momentum 0.9 and weight decay 0.0005. For data augmentation, as in [32], we adopt random crop, color jittering, and horizontal flip.

Tensor feature hallucinator (TFH) Our TFH is trained in episodes of $N = 5$ classes, $K = 20$ examples per class and generation of $M = 50$ class-conditional examples. We train for 50 epochs, where each epoch consists of 600 episodes. We use Adam optimizer with initial learning rate 10^{-4} , decaying by half at every 10 epochs.

Novel-task fine-tuning (TFH-ft) Given a novel task, we also provide an improved solution, TFH-ft, where our hallucinator is *fine-tuned* on novel-class support examples. We use exactly the same loss function as in hallucinator training (4) and we fine-tune for t steps using Adam optimizer and learning rate η . Table 1 shows the values t and η used in all experiments.

4.4. Setup

Tasks We consider N -way, K -shot classification tasks with $N = 5$ randomly sampled novel classes and $K \in \{1, 5\}$ examples drawn at random per class as support set S , that is, $L = 5K$ examples in total. For the query set Q , we draw 15 additional examples per class, that is, 75 examples in total, which is the most common choice [39, 35, 69]. We measure the *classification accuracy* as the percentage of correctly classified queries per task. To reduce the variance, we report mean accuracy and 95% confidence interval over 600 tasks per experiment, similarly to [34].

Hyperparameters For ResNet-18, we generate $M = 100$ and $M = 2$ features per class in 1-shot and 5-shot tasks respectively using TFH and $M = 100$ using TFH-ft. For ResNet-12, we generate $M = 1000$ and $M = 1$ features per class in 1-shot and 5-shot tasks respectively using TFH and $M = 5$ using TFH-ft.

Baselines: no augmentation We define baselines consisting only of the embedding network f_θ (1) or $f_{\theta'}$ (2) at representation learning and a prototypical classifier at inference, without feature hallucination. We refer to them as

Baseline (1) and Baseline-KD (2), respectively.

Baseline: vector feature hallucinator (VHF) To validate the benefit of generating tensor features, we also generate *vector* features by using $f_{\theta'} : \mathcal{X} \rightarrow \mathbb{R}^d$ including GAP (2) as embedding model. In this case, the *conditioner* $h : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ consists of two fully-connected layers with a ReLU activation in-between. The *generator* $g : \mathbb{R}^{k+d'} \rightarrow \mathbb{R}^d$ also consists of two fully-connected layers with a ReLU activation in-between and a sigmoid function at the end. The dimension d' of the class-conditional vector as well as the dimensions of the hidden layers of both the conditioner and the generator are all set to 512.

Competitors We compare our method with state-of-the-art data augmentation methods for few-shot learning, including MetaGAN [73], Δ -encoder [57], *salient network* (SalNet) [72], *diversity transfer network* (DTN) [4], dual TriNet [8], *image deformation meta-network* (IDeMe-Net) [7], *adversarial feature hallucination network* (AFHN) [34] and *variational inference network* (VI-Net) [40].

4.5. Comparison with the state of the art

Standard few-shot classification Table 2 compares our method with baselines and the state of the art. Most important are the comparisons with [8, 7, 34, 40], which use the same backbone. Our TFH provides new state of the art in all datasets in both 1-shot and 5-shot tasks, outperforming all competing few-shot data augmentation methods.

TFH is superior to VFH, especially in *miniImageNet* 1-shot, providing almost 3% performance improvement. The importance of tensor features is clear from the fact that VHF is worse than AFHN [34] while TFH is better than AFHN by more than 2% in *miniImageNet* 1-shot. VFH still outperforms the state of the art in all other experiments, highlighting the effectiveness of our new loss function. Novel-task *fine-tuning* (TFH-ft) is mostly beneficial, impressively even in 1-shot tasks. This shows that our tensor hallucinator is robust and avoids overfitting. *Self-distillation* provides a significant gain in all experiments.

Cross-domain few-shot classification We investigate the ability of our tensor hallucinator to address domain shift, carrying out experiments in the *cross-domain* few-shot classification setting proposed by [6]. We are not aware of any other synthetic data generation method that has been applied to this setting. We train the ResNet-18 backbone and our tensor hallucinator on *miniImageNet* as D_{base} and solve novel tasks on CUB and CIFAR-FS as D_{novel} . As shown in Table 3, our tensor hallucinator can address the domain shift effectively, especially in 1-shot tasks, even without fine-tuning. Even though the results in 5-shot tasks are less impressive, our hallucinator still outperforms the Baseline-KD on *miniImageNet*→CIFAR-FS, while being on par on

METHOD	BACKBONE	<i>mini</i> IMAGENET		CUB		CIFAR-FS	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
MetaGAN [73]	ConvNet-4	52.71 \pm 0.64	68.63 \pm 0.67	–	–	–	–
Δ -Encoder \dagger [57]	VGG-16	59.90	69.70	69.80 \pm 0.46	82.60 \pm 0.35	66.70	79.80
SalNet [72]	ResNet-101	62.22 \pm 0.87	77.95 \pm 0.65	–	–	–	–
DTN [4]	Resnet-12	63.45 \pm 0.86	77.91 \pm 0.62	72.00	85.10	71.50	82.80
Dual TriNet [8]	ResNet-18	58.80 \pm 1.37	76.71 \pm 0.69	69.61	84.10	63.41 \pm 0.64	78.43 \pm 0.64
IDeMe-Net [7]	ResNet-18	59.14 \pm 0.86	74.63 \pm 0.74	–	–	–	–
AFHN [34]	ResNet-18	62.38 \pm 0.72	78.16 \pm 0.56	70.53 \pm 1.01	83.95 \pm 0.63	68.32 \pm 0.93	81.45 \pm 0.87
VI-Net [40]	ResNet-18	61.05	78.60	74.76	86.84	–	–
Baseline (1)	ResNet-18	56.81 \pm 0.81	78.31 \pm 0.59	67.14 \pm 0.89	86.22 \pm 0.50	65.71 \pm 0.95	84.68 \pm 0.61
Baseline-KD (2)	ResNet-18	59.62 \pm 0.85	79.31 \pm 0.62	70.85 \pm 0.90	87.64 \pm 0.48	69.15 \pm 0.94	85.89 \pm 0.59
VFH (ours)	ResNet-18	61.88 \pm 0.85	79.63 \pm 0.61	75.44 \pm 0.85	87.82 \pm 0.47	72.31 \pm 0.91	85.64 \pm 0.64
TFH (ours)	ResNet-18	64.49 \pm 0.84	79.94 \pm 0.60	75.66 \pm 0.85	88.39 \pm 0.49	73.77 \pm 0.85	86.68 \pm 0.63
TFH-ft (ours)	ResNet-18	65.07 \pm 0.82	80.81 \pm 0.61	75.76 \pm 0.83	88.60 \pm 0.47	74.77 \pm 0.90	86.88 \pm 0.59

Table 2. *Standard few-shot classification.* Comparison of our TFH, variants and baselines to SOTA few-shot data augmentation methods. Baseline (1), Baseline-KD (2): prototypical classifier at inference, no feature generation. VFH: our vector feature hallucinator; TFH: our tensor feature hallucinator; TFH-ft: our tensor feature hallucinator followed by novel-task fine-tuning. \dagger : Delta-encoder uses VGG-16 backbone for *miniImageNet* and CIFAR-FS and *ResNet-18* for CUB.

METHOD	<i>mIN</i> →CUB		<i>mIN</i> →CIFAR-FS	
	1-shot	5-shot	1-shot	5-shot
Baseline (1)	43.14 \pm 0.78	62.20 \pm 0.70	50.25 \pm 0.86	69.43 \pm 0.74
Baseline-KD (2)	44.40 \pm 0.82	63.83 \pm 0.73	51.54 \pm 0.89	70.40 \pm 0.72
VFH (ours)	44.77 \pm 0.79	62.61 \pm 0.73	50.36 \pm 0.87	69.31 \pm 0.74
TFH (ours)	45.67 \pm 0.80	63.08 \pm 0.73	51.82 \pm 0.89	69.77 \pm 0.76
TFH-ft (ours)	45.96 \pm 0.80	63.64 \pm 0.74	53.07 \pm 0.86	71.29 \pm 0.75

Table 3. *Few-shot cross-domain classification.* Comparison of our TFH to variants and baselines as defined in Table 2, using ResNet-18 trained on *miniImageNet* (*mIN*).

miniImageNet→CUB.

4.6. Ablations

Alternative classifiers We investigate the effect of replacing the prototypical classifier [58] by alternative classifiers at inference, applied to the same augmented support set according the features generated by our tensor hallucinator. We consider *logistic regression* and *support vector machine* (SVM) classifiers, both using the scikit-learn framework [48]. As shown in Table 4, our tensor hallucinator clearly provides the best accuracy results irrespective of the classifier, while the performance of all three classifiers is on par. This indicates that individual generated features are also useful, not just the centroid per class.

Alternative backbone network To investigate the effect of using alternative backbone networks, we replace ResNet-18 by ResNet-12 and use the pre-trained networks

METHOD	<i>mini</i> IMAGENET		CUB	
	1-shot	5-shot	1-shot	5-shot
LOGISTIC REGRESSION				
Baseline (1)	59.20 \pm 0.82	77.71 \pm 0.61	69.44 \pm 0.86	86.19 \pm 0.49
Baseline-KD (2)	61.83 \pm 0.82	79.27 \pm 0.61	72.74 \pm 0.88	87.71 \pm 0.49
VFH (ours)	62.37 \pm 0.83	79.70 \pm 0.59	75.06 \pm 0.88	87.95 \pm 0.46
TFH (ours)	64.03 \pm 0.84	79.93 \pm 0.60	75.08 \pm 0.85	88.82 \pm 0.46
TFH-ft (ours)	64.83 \pm 0.82	80.49 \pm 0.61	75.43 \pm 0.85	88.52 \pm 0.48
SUPPORT VECTOR MACHINE				
Baseline (1)	57.12 \pm 0.84	76.45 \pm 0.62	67.24 \pm 0.87	84.72 \pm 0.52
Baseline-KD (2)	60.21 \pm 0.84	78.28 \pm 0.61	71.23 \pm 0.89	86.34 \pm 0.52
VFH (ours)	61.95 \pm 0.85	79.15 \pm 0.60	75.43 \pm 0.85	87.46 \pm 0.47
TFH (ours)	64.20 \pm 0.83	79.62 \pm 0.60	75.64 \pm 0.85	88.74 \pm 0.45
TFH-ft (ours)	65.06 \pm 0.82	80.33 \pm 0.60	75.77 \pm 0.83	88.22 \pm 0.46

Table 4. *Alternative classifiers.* Our TFH, variants and baselines as defined in Table 2, using ResNet-18, where at inference, the prototypical classifier is replaced by logistic regression or SVM.

on *miniImageNet* and CUB provided by DeepEMD [70]¹. We use all three classifiers: prototypical, logistic regression and SVM. As shown in Table 5, our tensor hallucinator provides the best results in all settings with a significant performance gain of around 3-5% in 1-shot tasks in prototypical and support vector machines classifiers in both datasets.

Effect of number of generated features We investigate the effect of the number of generated features per class, M , in Table 6. The results are similar for both backbones, high-

¹<https://github.com/icoz69/DeepEMD>

METHOD	<i>mini</i> IMAGENET		CUB	
	1-shot	5-shot	1-shot	5-shot
PROTOTYPICAL CLASSIFIER				
Baseline	59.94 \pm 0.84	78.65 \pm 0.61	66.72 \pm 0.90	84.27 \pm 0.59
TFH (ours)	64.69 \pm 0.83	79.58 \pm 0.59	68.71 \pm 0.90	84.83 \pm 0.53
TFH-ft (ours)	65.16 \pm 0.86	79.83 \pm 0.65	70.52 \pm 0.87	85.02 \pm 0.55
LOGISTIC REGRESSION				
Baseline	61.83 \pm 0.85	79.60 \pm 0.60	68.49 \pm 0.49	84.75 \pm 0.59
TFH (ours)	64.20 \pm 0.84	80.21 \pm 0.58	68.37 \pm 0.92	85.12 \pm 0.55
TFH-ft (ours)	64.53 \pm 0.53	80.25 \pm 0.62	69.59 \pm 0.88	85.08 \pm 0.55
SUPPORT VECTOR MACHINE				
Baseline	59.94 \pm 0.84	78.43 \pm 0.60	66.72 \pm 0.90	83.29 \pm 0.60
TFH(ours)	64.62 \pm 0.83	79.56 \pm 0.58	68.75 \pm 0.90	84.05 \pm 0.56
TFH-ft(ours)	64.64 \pm 0.86	79.88 \pm 0.62	69.42 \pm 0.88	84.41 \pm 0.55

Table 5. *Alternative backbone network.* Our TFH, variants and baselines as defined in Table 2, using the publicly available pre-trained ResNet-12 backbones provided by [70].

# FEAT. M	RESNET-18		RESNET-12	
	1-shot	5-shot	1-shot	5-shot
0	59.62 \pm 0.85	79.31 \pm 0.62	59.94 \pm 0.84	78.65 \pm 0.61
1	62.43 \pm 0.83	79.33 \pm 0.62	63.88 \pm 0.85	79.58 \pm 0.59
2	63.57 \pm 0.87	79.94 \pm 0.60	64.01 \pm 0.83	79.44 \pm 0.61
5	64.42 \pm 0.83	79.77 \pm 0.61	64.12 \pm 0.83	79.49 \pm 0.61
10	63.90 \pm 0.81	79.64 \pm 0.60	64.01 \pm 0.85	79.56 \pm 0.63
50	63.44 \pm 0.86	79.43 \pm 0.63	63.67 \pm 0.84	79.18 \pm 0.65
100	64.49 \pm 0.84	79.27 \pm 0.60	63.66 \pm 0.86	78.89 \pm 0.65
500	64.25 \pm 0.85	79.71 \pm 0.62	64.11 \pm 0.86	79.38 \pm 0.65
1000	63.76 \pm 0.87	79.52 \pm 0.62	64.69 \pm 0.83	79.48 \pm 0.60

Table 6. *Effect of number of generated features, M .* Using our tensor feature hallucinator (TFH) with the default prototypical classifier and without fine-tuning on *mini*ImageNet.

lighting the generality of our TFH. In all settings, regardless of number of generated tensor features, our hallucinator provides better performance. Remarkably, even 1 or 2 generated features per class provide a significant gain of around 4% in 1-shot tasks for both networks when compared to 0 (Baseline-KD for ResNet-18 and Baseline for ResNet-12). Interestingly, the performance of our hallucinator does not degrade even when 1000 tensor features are generated, still outperforming the baseline models and providing the best 1-shot accuracy for ResNet-12. This is significant, since when generating 1000 features, the vast majority of the support set consists only of synthetic data.

Visualization of feature embeddings To investigate the generated feature distribution, we sample a 5-way, 1-shot novel task, use the support set to generate 500 novel features per class and visualize the augmented support set in 2D us-

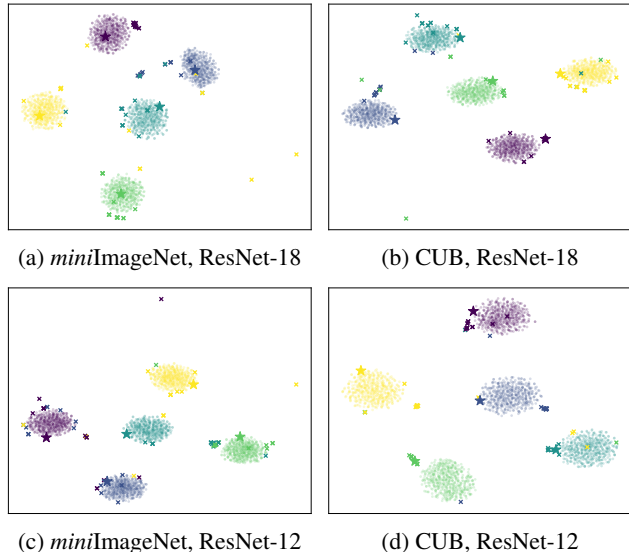


Figure 4. t-SNE visualization of the augmented support feature set of an 1-shot task using both ResNet-18 and ResNet-12 backbones on *mini*ImageNet and CUB. Colors indicate different classes. \star : support features; \bullet : generated features; \times : query features.

ing t-SNE [62]. As shown in Figure 4, the generated features of each class are clustered together, with distinct class boundaries. This synthetic class-conditional distribution is improving the classifier performance in the absence of true data. It can be seen that the query examples are much more scattered, showing that the variance of each novel class can be large, highlighting the inherent difficulty of the few-shot classification problem. If only the support examples were used, making predictions would be much harder because of how scattered the queries are.

5. Conclusion

In this work, we have introduced a conceptually simple tensor feature hallucinator that improves the state of the art on synthetic data generation for few-shot learning classification. We have provided evidence showing that the structural properties of tensors provide a significant performance gain, allowing for a simplification of the loss function and training regime. We have also shown the importance of complementing with improved representation learning, as well as task adaptation by fine-tuning on the augmented support set, which reduces the risk of overfitting.

Potential future directions include: improving our hallucinator architecture; experimenting with different loss functions to train our hallucinator; and investigating the use of our hallucinator in different settings, such as long-tailed recognition or incremental learning.

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- [2] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [4] Mengting Chen, Yuxin Fang, Xinggang Wang, Heng Luo, Yifeng Geng, Xinyu Zhang, Chang Huang, Wenyu Liu, and Bo Wang. Diversity transfer network for few-shot learning. In *AAAI*, 2020.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019.
- [7] Zitian Chen, Yanwei Fu, Yu-Xiong Wang, Lin Ma, Wei Liu, and Martial Hebert. Image deformation meta-networks for one-shot learning. In *CVPR*, 2019.
- [8] Z. Chen, Y. Fu, Y. Zhang, Y. Jiang, X. Xue, and L. Sigal. Multi-level semantic feature augmentation for one-shot learning. *IEEE Transactions on Image Processing*, 2019.
- [9] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014.
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- [12] Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043*, 2017.
- [13] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018.
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018.
- [15] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.
- [16] Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *CVPR*, 2017.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [19] Nathan Hilliard, Lawrence Phillips, Scott Howland, Artëm Yankov, Courtney D Corley, and Nathan O Hodas. Few-shot learning with metric-agnostic conditional embeddings. *arXiv preprint arXiv:1802.04376*, 2018.
- [20] Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *ICML*, 2019.
- [21] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [22] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *arXiv preprint arXiv:2006.06676*, 2020.
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [24] Jongmin Kim, Taesup Kim, Sungwoong Kim, and Chang D Yoo. Edge-labeling graph neural network for few-shot learning. In *CVPR*, 2019.
- [25] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [27] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML workshop*, 2015.
- [28] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (BIT): General visual representation learning. In *ECCV*, 2020.
- [29] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, 2009.
- [30] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [31] Michalis Lazarou, Yannis Avrithis, and Tania Stathaki. Iterative label cleaning for transductive and semi-supervised few-shot learning. *arXiv preprint arXiv:2012.07962*, 2020.
- [32] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, 2019.
- [33] Hongyang Li, David Eigen, Samuel Dodge, Matthew Zeiler, and Xiaogang Wang. Finding task-relevant features for few-shot learning by category traversal. In *CVPR*, 2019.
- [34] Kai Li, Yulun Zhang, Kunpeng Li, and Yun Fu. Adversarial feature hallucination networks for few-shot learning. In *CVPR*, 2020.
- [35] Xinzhe Li, Qianru Sun, Yaoyao Liu, Qin Zhou, Shibao Zheng, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. In *NeurIPS*, 2019.
- [36] Yann Lifchitz, Yannis Avrithis, Sylvaine Picard, and Andrei Bursuc. Dense classification and implanting for few-shot learning. In *CVPR*, 2019.

- [37] Bingchen Liu, Yizhe Zhu, Kunpeng Song, and Ahmed El-gammal. Towards faster and stabilized GAN training for high-fidelity few-shot image synthesis. *arXiv e-prints*, pages arXiv-2101, 2021.
- [38] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *CVPR*, 2019.
- [39] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. *arXiv preprint arXiv:1805.10002*, 2018.
- [40] Qinxuan Luo, Lingfeng Wang, Jingguo Lv, Shiming Xiang, and Chunhong Pan. Few-shot learning via feature hallucination with variational inference. In *WACV*, 2021.
- [41] 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 *WACV*, 2020.
- [42] Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Eleventh annual conference of the international speech communication association*, 2010.
- [43] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *ICML*, 2017.
- [44] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [45] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *CVPR*, 2016.
- [46] Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelcnn decoders. *arXiv preprint arXiv:1606.05328*, 2016.
- [47] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. TADAM: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018.
- [48] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *JMLR*, 2011.
- [49] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *CVPR*, 2018.
- [50] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- [51] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *arXiv preprint arXiv:1705.08045*, 2017.
- [52] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, 2017.
- [53] Pau Rodríguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. *ECCV*, 2020.
- [54] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training GANs. In *NIPS*, 2016.
- [55] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, 2016.
- [56] Jürgen Schmidhuber. Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook. Master's thesis, Technische Universität München, 1987.
- [57] Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Abhishek Kumar, Rogerio Feris, Raja Giryes, and Alex Bronstein. Delta-encoder: an effective sample synthesis method for few-shot object recognition. In *NeurIPS*, 2018.
- [58] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [59] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.
- [60] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018.
- [61] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? *arXiv preprint arXiv:2003.11539*, 2020.
- [62] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *JMLR*, 9(11), 2008.
- [63] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016.
- [64] Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 18(2):77–95, 2002.
- [65] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NIPS*, 2016.
- [66] Davis Wertheimer, Luming Tang, and Bharath Hariharan. Few-shot classification with feature map reconstruction networks. In *CVPR*, 2021.
- [67] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *CVPR*, 2020.
- [68] Sung Whan Yoon, Jun Seo, and Jaekyun Moon. Tapnet: Neural network augmented with task-adaptive projection for few-shot learning. In *ICML*, 2019.
- [69] Zhongjie Yu, L. Chen, Zhongwei Cheng, and Jiebo Luo. TransMatch: A transfer-learning scheme for semi-supervised few-shot learning. *CVPR*, 2020.
- [70] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. DeepEMD: Few-shot image classification with differentiable earth mover's distance and structured classifiers. In *CVPR*, 2020.
- [71] Hongguang Zhang, Jing Zhang, and Piotr Koniusz. Few-shot learning via saliency-guided hallucination of samples. In *CVPR*, 2019.

- [72] Hongguang Zhang, Jing Zhang, and Piotr Koniusz. Few-shot learning via saliency-guided hallucination of samples. In *CVPR*, 2019.
- [73] Ruixiang Zhang, Tong Che, Zoubin Ghahramani, Yoshua Bengio, and Yangqiu Song. MetaGAN: An adversarial approach to few-shot learning. *NeurIPS*, 2018.
- [74] Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *ICML*, 2019.