

Semantics-driven Attentive Few-shot Learning over Clean and Noisy Samples

Orhun Bugra Baran^a, Ramazan Gokberk Cinbis^a

^a*Department of Computer Engineering, Middle East Technical University, 06800 Ankara, Turkey*

Abstract

Over the last couple of years few-shot learning (FSL) has attracted great attention towards minimizing the dependency on labeled training examples. An inherent difficulty in FSL is the handling of ambiguities resulting from having too few training samples per class. To tackle this fundamental challenge in FSL, we aim to train meta-learner models that can leverage prior semantic knowledge about novel classes to guide the classifier synthesis process. In particular, we propose semantically-conditioned feature attention and sample attention mechanisms that estimate the importance of representation dimensions and training instances. We also study the problem of sample noise in FSL, towards the utilization of meta-learners in more realistic and imperfect settings. Our experimental results demonstrate the effectiveness of the proposed semantic FSL model with and without sample noise.

Keywords: Few-shot learning, vision and language integration

1. Introduction

Contemporary supervised learning approaches combined with large training datasets yield excellent results on a variety of recognition problems. A major challenge, however, is learning to model concepts with limited samples. Few-shot learning (FSL) [1, 2, 3, 4] techniques aim to tackle this problem by learning to synthesize effective models based on a few examples.

A major source of motivation for studying FSL is the observation that humans, starting at young ages, can learn new concepts with limited examples [5], [6], [7]. In addition, in most real-world classification problems, such as object recognition [8], class distributions can be heavily

Email addresses: bugra@ceng.metu.edu.tr (Orhun Bugra Baran), gcinbis@ceng.metu.edu.tr (Ramazan Gokberk Cinbis)

Preprint submitted to Elsevier

January 11, 2022



Figure 1: (a) All dog instances appear in a misleadingly consistent beach background context. (b) Spuriously consistent background-foreground relation may cause FSL models to ignore other salient features of classes. (c) It is difficult to understand when a consistent foreground texture is informative or misleading, e.g., the Dalmatian texture is distinctive but the *Dalmatian dog* class but otherwise misleading for the generic *dog* class. (d) Some examples can be more *prototypical* than the others. A closely related problem is having *sample noise* in the training set.

long-tailed [9]. FSL research can be seen as the focused study of learning to recognize in the low-data regime, which can play a central role in building semantically comprehensive and rich models.

A variety of FSL approaches have been introduced in recent years. Most of the recent work can be summarized as follows: metric learning based FSL [10, 11, 12, 13], generative model based statistical data augmentation [14, 15, 16, 17, 18, 19, 20], non-generative data augmentation [21, 22], feed-forward classifier synthesis [23, 24, 25], model initialization for few-shot adaptation [26, 27, 28], learning-to-optimize for FSL [29] and memory-based FSL approaches [30, 31]. In addition, recent work has highlighted the importance of implementations details in improving and evaluating FSL models, including batch normalization details [32], feature extraction backbones [33] and pretraining strategies [34]. Variations of FSL, such as cross-domain [35, 36, 37, 38, 39] and variable-shot [40] learning have also been introduced.

An inherent difficulty in FSL, independent of the method being used, is the ambiguity result-

ing from having few training samples per class. Particularly, it is difficult to figure out whether a cue that appears consistently in the limited set of examples is truly indicative. For example, few-shot samples may contain misleadingly similar contextual information (Figure 1a), spurious foreground-background relationships (Figure 1b), or suspiciously consistent foreground features (Figure 1c). Similarly, some samples might actually be less *prototypical* [41] than the others or completely *noisy*, due to a number of factors, such as background clutter, viewpoint, occlusions, overall representativeness and sample noise (See Figure 1 for further discussion). Hence, FSL model may *overfit* to incorrect features, misleading or noisy samples, or under-utilize distinctive cues, due to the fundamental difficulty of disambiguating spurious cues from the informative ones purely based on few examples.

To tackle these fundamental challenges in FSL, we aim to guide the meta-learner via semantic priors, which we call *semantic few-shot learning*. To this end, we build meta-learning models that can benefit from text-based semantic representations of classes of interest when synthesizing target classifiers. For this purpose, we focus on one of the most popular metric based few-shot learners Prototypical Networks [10] (PNs). In the context of PN formulation, we introduce semantically-conditioned *feature attention* and *sample attention* mechanisms, towards reducing the risk of overfitting to misleading features or samples and improving the data efficiency in few-shot learning.

The use of semantic vector-space representations of classes, i.e., *class embeddings*, is prominent in *zero-shot learning* (ZSL). Most mainstream ZSL approaches learn to estimate the degree of relation between a given input (image) and a class embedding, so that previously unseen classes can be recognized purely based on class embeddings, e.g., [42, 43, 44, 45, 18, 46, 47, 48]. In ZSL, class embeddings can be interpreted as class summaries from which valuable discriminative or generative knowledge can be extracted. In our work, instead of relying purely on class embeddings for building classification models, we aim to benefit from them in improving sample efficiency in FSL.

The idea of jointly benefiting from class semantics in FSL has received attention only recently. In [35], semantics-based class prototype priors are estimated and adaptively combined with the data-driven class prototypes. [36] extends this model to multiple semantic information sources. [37] aims to obtain augmented feature representations based on embedding images into the semantic space and then decoding the sampled semantic features to obtain new samples for

training purposes. [39] proposes to reformulate the PN loss based on task dependent similarities measured from semantic priors to regulate the margin between classes in a task.

The prior work most related to ours is the AM3 model [35], which we use as our starting point. While AM3 makes an important step forward by defining class semantics based models priors, the approach does not leverage prior semantic knowledge in knowledge accumulation. To this end, we explore the uses of semantic class knowledge in a deeper way in the following two main ways. First, we aim to estimate the importance of provided few-shot samples for each class by evaluating the consistency across samples and semantics. Second, we estimate per-dimension feature importance factors for the data-driven class prototypes based on prior knowledge. Third, we define a *noisy FSL* problem, where some of the support samples can be incorrectly labeled, and evaluate the proposed approach in this setting. We believe that noise-tolerant FSL models can have applications in variety of real-world problems, e.g., the construction of large-vocabulary models over automatically retrieved web samples based on meta-data where top ranking items tend to be much more coherent yet not without noise; large-scale fine-grained FSL problems where human annotators are prone to making mistakes; and on-the-fly construction of models from few interactively provided samples in robotic systems.

Our contributions can, therefore, be summarized as follows: (i) we propose semantics-driven feature and sample attention mechanisms to improve FSL data efficiency in a principled way; (ii) we define an experimental setting for noisy FSL and investigate the applicability of the proposed sample attention approach to this problem; (iii) we present a detailed experimental analysis towards understanding the effects and dynamics of the proposed attention mechanisms. Our quantitative and qualitative experimental results demonstrate the effectiveness of the proposed semantic FSL model both in clean and noisy settings.

2. Related work

In this section, we first present an overview of mainstream few-shot learning approaches. We then present an overview of works on the recently emerging topic of semantic FSL. Finally, we briefly discuss zero-shot learning and its relation to our work.

Few-shot learning. The most related mainstream FSL approaches can be summarized within *initialization based*, *metric learning based* and *generative model based* groups. Initialization based FSL methods aim to learn the *ideal* initial model such that the model can perform well

even when fine-tuned using just a few examples. MAML [26] is arguably the most well-known example of this category. In MAML, the main idea is to learn the initial model that minimizes the loss of validation samples when the initial model is fine-tuned using one or few gradient-based updates. Several other related and follow-up works exist, such as [27, 28, 49, 29].

In metric learning based FSL, the goal is to learn a metric space where the similarity of feature representations of sample pairs can be used to classify pairs as same class/different class pairs. One of the most well-known examples of this category is Prototypical Networks [10], which we also use as the basis of our approach (see Section 3). Due to its simplicity and high FSL performance, many other metric learning FSL approaches have also been introduced, e.g., [11, 12, 23, 24]. Despite these explorations, recent works show that a carefully tuned Prototypical Network can yield state-of-the-art few-shot learning results [34, 50, 51]. We adapt the *modernized* PN implementation of [51] as the non-semantic FSL baseline in the construction of our models.

Generative modeling based FSL approaches aim to learn a sample-synthesizing model, which can be used for augmenting a few-shot training set. For example, [14] learns a mapping that can be used to transform existing train samples into new ones, [52, 16] propose GAN [53] based generative models towards synthesizing novel examples.

Another important research direction is learning generalizable feature representations [54, 50, 32]. The representation generalizability is a major problem in FSL as backbone networks are utilized on novel classes at test time. [55] uses self-attention over spatial locations to improve representations, similar to non-local networks [56], for a relation-network based FSL approach. [57] proposes a transductive FSL approach that aims to obtain query-specific feature representations via attention mechanisms. These approaches are orthogonal to ours, as we focus on leveraging semantic priors, in a query-independent way.

Sample noise in FSL is largely an overlooked problem. To the best of knowledge, the only directly relevant work is the few-shot text classification approach of [58], which looks into noisy annotations for few-shot relation classification. [58] defines query-to-support sentence and support samples driven feature-level attention mechanisms. Our work and focus fundamentally differs as we (i) leverage prior knowledge for building conditional attention mechanisms, (ii) estimate sample importance in a query-agnostic way, and (iii) define an experimental protocol for studying the sample noise problem on the mainstream FSL image classification benchmarks

MiniImageNet [12] and TieredImageNet [59].

Semantic few-shot learning. Semantic FSL refers to an FSL problem variant where a supplementary class-wise knowledge source is made available. Since such additional knowledge often comes from a new data modality, semantic FSL is also sometimes referred to as *multi-modal FSL*. There exists only a few and recent works on semantic FSL.

In a pioneering work, [35] proposes to define class prototypes as convex combinations of average visual features and transformed semantic priors. [36] extends the approach of [35], mainly by introducing multiple semantic priors (label, description or attribute) jointly to obtain richer semantic information. In [37], visual features are mapped into a semantic space via an encoder, where semantic features, belonging to the same class, can be sampled to augment visual features via a decoder model. The approach then performs classification by using both real and augmented features. Similarly, [18] aims to learn aligned auto-encoders with reconstruction losses across modalities. In a more recent work, [39] uses semantic prior information with a similarity function to obtain margin scores that are used as an additional term in a distance based loss such as [10]. [60] uses semantic priors in addition to visual features to obtain two different classification weights for a class. This work obtains visual feature based classification weights by l2-normalizing support image features, averaging them per-class and learning a semantic embedding vector that can be used as a classifier weight maximizing the inner product between the support image based classification weights. [61] proposes a model that re-weights support samples according to a generic weighting function and an auto-encoder that can use either class embeddings or Gaussian noise vectors for encoding regularization. Among these models, the most closely related one is AM3 [35], as we build our models on top of it. While nearly all others can be considered as complementary to ours, instead of being alternatives, we provide empirical comparisons to semantic FSL works in Section 4.

Zero Shot Learning. Zero shot learning aims to build recognition models that can handle classes with no training examples, purely based on prior class semantic knowledge. Mainstream ZSL approaches include learning mappings between the space of visual features and the semantic class representations [42, 43, 44], and semantics conditional generative models [45, 18, 46, 47, 48]. In our work, we use class semantics to build attention mechanisms in the presence of few training examples, instead of aiming to remove the need for training samples completely as in zero-shot learning.

3. Method

In this section, we first provide a formal definition of semantics-driven few-shot learning and summarize the *episodic training* framework, which we embrace in our approach. We then provide a summary of Prototypical Networks (PN) model [10] for few-shot learning and *adaptive cross-modal few-shot learning* (AM3) [35] model that extends PNs by utilizing semantic knowledge to construct model priors. Finally, we present the proposed feature-attention and sample-attention mechanisms and explain how we integrate them into the PN and AM3 models.

3.1. Problem definition and training framework

In our work, we focus on few-shot learning of image classification models. The goal is to estimate a new classification model, for a set of target classes ($\mathcal{C} = \{c_1, \dots, c_n\}$), based on a limited set of labeled training examples. In our discussion, an n -way classifier is expressed in terms of a scoring function $f(x; \theta)$ that maps the input $x \in \mathcal{X}$ to an n -dimensional vector, according to the model parameters θ . In a standard supervised training problem, a typical way to estimate θ is to find the model parameters minimizing some regularized empirical loss function on the train set.

In the case of few-shot learning, however, the main challenge is to estimate a successful classification model based on a few training samples per class. In most practical cases, it is fundamentally difficult to achieve generalization based on few examples, and, therefore, a model learned by minimizing a generic empirical loss function is unlikely to perform well on novel (test) data. To tackle this problem, the main interest in meta-learning based FSL is to learn a meta-learner $\xi(\mathcal{D}; \beta)$ that can take a new limited set \mathcal{D} of training examples and synthesize the corresponding classification model parameters. β represents the trainable parameters of the meta-learner model ξ .

Episodic training. A popular approach for training meta-learning models is *episodic training*. The main idea is to construct a set or series of few-shot learning tasks and update the meta-learning model based on the regularized empirical loss of meta-learned classification models. More specifically, at each iteration, a new task T is created by sampling a subset \mathcal{C}_T of training classes, and then sampling training \mathcal{D}_T^s and validation \mathcal{D}_T^q samples from the whole training set. Each \mathcal{D}_T^s consists of few-shot training samples of classes \mathcal{C}_T and is commonly referred to as the *support set*. Similarly, each \mathcal{D}_T^q consists of task-specific validation samples and is commonly

referred to as the *query set*. The meta-learning, then, is achieved by minimizing the expected empirical loss of meta-learned models over the pairs of support and query sets:

$$\min_{\beta} \mathbb{E}_T \left[\mathbb{E}_{(x,y) \sim \mathcal{D}_T^q} [l(f(x; \theta = \xi(\mathcal{D}_T^s; \beta)), y)] \right] \quad (1)$$

where $l(\cdot, y)$ is a classification loss function for label y . θ represents the task-specific model parameters inferred by the meta-learner ξ . No regularization term is shown for brevity.

Semantic FSL. Arguably, the main premise of meta-learning is to learn domain-specific inductive biases better than what can be provided by general-purpose supervised learning formulations. However, small-sized training sets can be inherently misleading and/or ambiguous due to spurious patterns, as previously illustrated in Figure 1. In this respect, prior knowledge about classes can be a crucial source of information towards overcoming the limitations of few-shot training sets. In our work, we presume that semantic knowledge about each class c is a d_e -dimensional vector, represented by ψ_c . The meta-learner ξ , therefore, can additionally access to semantic class embeddings during both training and testing. In the following parts, we define the baseline and proposed meta-learner ξ variants, in terms of formulations for inferring task-specific parameters θ from a set of support samples and semantic priors.

3.2. Semantics-driven attention mechanisms

We build our semantics-driven attention mechanisms on top of the semantic few-shot learning model AM3, which is based on the prototypical network (PN) model. Below we first summarize the PN and AM3 models and then present our approach in its context.

Prototypical networks and AM3. The core idea in PN is to estimate *class prototypes* based on train samples provided in a task and then perform classification based on query-to-prototype similarities. In PNs, a class prototype θ^{PN} for a class c is obtained by averaging the support sample representations:

$$\theta^{\text{PN}}(c) = \mathbb{E}_{\mathcal{D}^{s,c}} [\phi(x)] \quad (2)$$

where $\phi(x)$ is the feature embedding function parameterized by (a subset of) β that is being trained as part of the PN model. $\phi(x)$ is ResNet12 in our experiments (see Section 4). $\mathcal{D}^{s,c}$ refers to the subset of examples belonging to class c within a given support set. In the de facto standard PN formulation, the classification score f for class c is given by the negative Euclidean distance

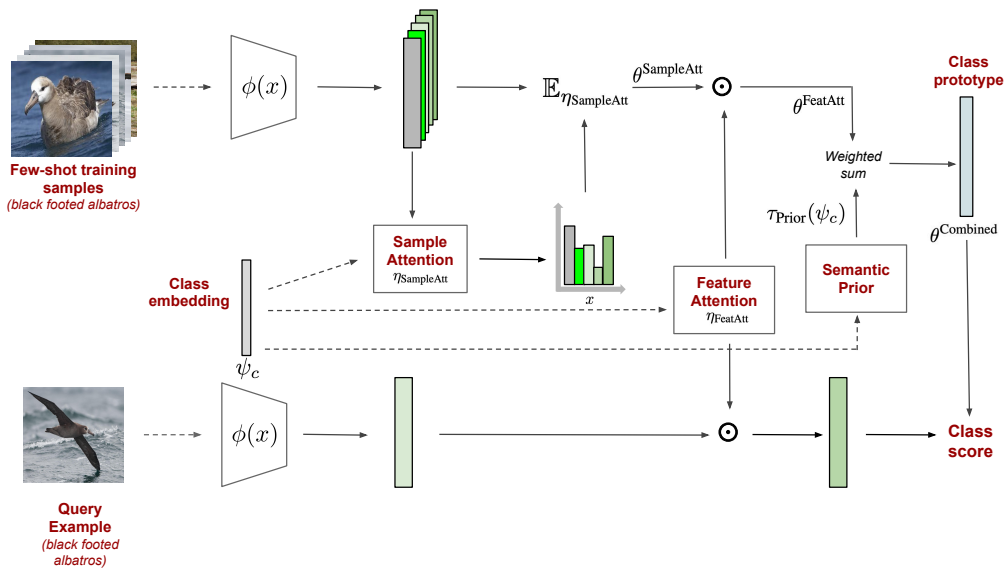


Figure 2: Summary of the proposed meta-learner (ξ), illustrating the process of a new class prototype estimation in the 5-shot setting. Arrows indicate the input dependencies across model components. Dashed lines are used to indicate externally provided data input. Sample attention model ($\theta^{\text{SampleAtt}}$) uses class embeddings (ψ_c) to assign relative importance weights to the provided support examples. Feature attention model (θ^{FeatAtt}) similarly uses class embeddings to estimate task-specific feature dimension scaling coefficients. The final model (θ^{Combined}) is obtained by combining the sample and feature attention driven task parameters with the semantics-only prior (θ^{Prior}). The resulting model is applied to the test input(s).

between the feature-space embedding of the input and the corresponding class prototype:

$$[f(x)]_c = -\|\phi(x) - \theta^{\text{PN}}(c)\|^2. \quad (3)$$

The parameters β are estimated by minimizing cross entropy loss of query samples over the episodes.

The AM3 model aims to improve the PN model by using semantic knowledge about classes as prototype priors. More specifically, AM3 redefines a class prototype θ^{Prior} as the weighted combination of transformed class semantic embeddings and feature averages:

$$\theta^{\text{Prior}}(c) = \alpha\theta^{\text{PN}}(c) + (1 - \alpha)\tau_{\text{Prior}}(\psi_c) \quad (4)$$

where τ_{Prior} is a trainable transformation that maps the semantic class embeddings to the space of class prototypes. α is the cross-validated hyper-parameter that controls the weight between the original PN prototypes and class semantics based prototypes. τ_{Prior} is parameterized by β .

The AM3 model, therefore, can be interpreted as a way to build classification model priors in the PN formulation. Consistent with the experimental results, such a prior is particularly valuable in the case of one-shot learning, where the training data size is at its extreme minimum. We propose and explore two novel ways to leverage semantic information in a more expressive way towards tackling the inherent difficulties in few-shot learning: (i) *sample attention*, (ii) *feature attention*. Below we provide the details of the proposed mechanisms, which act as the components of the whole model, a scheme of which can be found in Figure 2.

Sample attention. We observe that, in the original PN model, each class prototype is defined as a plain average of support examples. However, the information content of samples can vary greatly due to a number of factors, including background clutter, viewpoint and occlusion. Towards estimating the prototypicality of training samples, we introduce a sample attention mechanism into the final model. Ultimately, we aim to build a model that can estimate the importance of each sample based on the compatibility of samples and prior semantic information. Therefore, we define sample attention module $\eta_{\text{SampleAtt}}(x, c)$ as a function that computes the normalized attention scores of a sample $x \in \mathcal{D}^{s,c}$ in the context of support sample set $\mathcal{D}^{s,c}$, conditioned on the class semantic embedding vector ψ_c :

$$\eta_{\text{SampleAtt}}(x, c) = \frac{\exp(\gamma_{\text{vis}}(\phi(x))^\top \gamma_{\text{sem}}(\psi_c))}{\sum_{x' \in \mathcal{D}^{s,c}} \exp(\gamma_{\text{vis}}(\phi(x'))^\top \gamma_{\text{sem}}(\psi_c))} \quad (5)$$

where γ_{vis} and γ_{sem} are trainable models that are used to obtain visual and semantic feature embeddings. In our experiments, γ_{vis} and γ_{sem} are implemented as MLPs that take $\phi(x)$ and ψ_c respectively and return embeddings whose inner products yield the attention scores. We note that $\eta_{\text{SampleAtt}}$ defines a distribution over the support samples, which we use to re-define a class prototype:

$$\theta^{\text{SampleAtt}}(c) = \mathbb{E}_{\eta_{\text{SampleAtt}}(x,c)} [\phi(x)], \quad (6)$$

which amounts to computing the attention-weighted average of sample features. This sample attention mechanism also naturally handles potential sample noise in FSL, which we explore in Section 4.

Feature attention. The motivation in feature attention is to tackle the problems resulting from having too few examples. Overall, the goal is to enhance or attenuate certain prototype dimensions as a function of semantic class embeddings. For this purpose, we introduce the feature attention function η_{FeatAtt} and re-define the class prototype as follows:

$$\theta^{\text{FeatAtt}}(c) = \eta_{\text{FeatAtt}}(\psi_c) \odot \theta^0(c) \quad (7)$$

where $\theta^0(c)$ refers to averaging based class prototypes, which can either be the vanilla PN prototype (θ^{PN}) or the sample attention based prototype ($\theta^{\text{SampleAtt}}$). \odot represents the Hadamard product operator. Here, the trainable feature attention function η_{FeatAtt} predicts per-dimension scaling coefficients as a function of class embeddings. The feature attention output is applied to both the prototype and the query, resulting in the following scoring function:

$$[f(x)]_c^{\text{FeatAtt}} = -\|\eta_{\text{FeatAtt}}(\psi_c) \odot \phi(x) - \theta^{\text{FeatAtt}}(c)\|^2. \quad (8)$$

The final model. We build our final model by integrating our semantic-driven sample and feature attention mechanisms into the semantic FSL framework defined by the AM3 model. More specifically, in the final *combined attention* model, we first estimate sample attention based class prototypes $\theta^{\text{SampleAtt}}$ and then update them into θ^{FeatAtt} using the feature attention model. We obtain the final class prototypes, which we call θ^{Combined} , by computing the α -weighted combination of the attention-driven prototypes and the pure semantic embedding based prototypes given by $\tau_{\text{Prior}}(\psi_c)$:

$$\theta^{\text{Combined}}(c) = \alpha \eta_{\text{FeatAtt}}(\psi_c) \odot \theta^{\text{SampleAtt}}(c) + (1 - \alpha) \tau_{\text{Prior}}(\psi_c). \quad (9)$$

The combined attention model corresponds to the scheme presented in Figure 2.

4. Experiments

In this section, we first explain our experimental setup and our implementation details. We then present our main experimental results and analyses. Finally, we present and discuss our results.

4.1. Experimental setup

Prior work [33, 32, 34] shows that implementation details, such as (batch) normalization schemes, backbones, hyper-parameter tuning strategies and data augmentation schemes, can make a great impact on few-shot learning results. Therefore, we systematically tune the hyper-parameters, including learning rates, number of iterations, and dropout rates on the validation sets for all results that we report based on our own implementation. Below we provide additional details regarding our experimental setup and our efforts to make fair comparisons.

Datasets. We evaluate our model on the MiniImageNet [12] and TieredImageNet [59] datasets. The MiniImageNet dataset is a subset of ImageNet dataset [62] with 100 classes and 600 images per class. We use the split of [29] for MiniImageNet where the 64, 16 and 20 classes are used as the train, validation and test subsets, respectively. TieredImageNet is a separate benchmark based on [62], with 351, 97, 160 classes for training, validation and testing, respectively. In all our experiments, we use Glove [63] vectors of class names to extract semantic embeddings. Following [35], we use the Common Crawl version trained on 840B tokens with 300 dimensional embeddings.

Evaluation. We report our test set results over 10000 random tasks and report the average accuracy and 95% confidence interval scores. In all k-shot experiments, we use 15 query examples for each of the n classes in an episode. We execute all of our experiments using the same batch structure for consistency across the experiments.

Backbone architecture. In all models we use the same ResNet-12 architecture as the feature extraction backbone. We use Batch Normalization [64] only in backbone layers, using the *eval* mode [65] for meta-testing to avoid *accidental* transductive setting, which is known to potentially result in misleadingly better few-shot learning results [32].

Backbone pretraining. Following [51] we use supervised pretraining for ResNet-12 backbone and the exact details can be found in [51]. After the pretraining we employ a two-staged training

approach. In the first stage we train everything except the backbone with learning rate 0.1, momentum 0.9 and weight decay of 0.0005 for 200 epochs for MiniImageNet and 50 epochs for TieredImageNet where an epoch contains 100 episodes. After every epoch we validate our results by using 600 episodes for both MiniImageNet and TieredImageNet. After we complete the 200 epochs and 50 epochs for MiniImageNet and TieredImageNet respectively we train the whole model end to end for another 200 epochs. In the second stage the learning rate for backbone is selected to be 0.002 and 0.02 for the rest of the network. After every 40 epoch we halve the learning rates. We use SGD as optimizer for every model and experiment.

PN implementation. Our code base is built upon the PN implementation of [51]. For a fair comparison, therefore, we report the results that we obtain for PN and [51] using the publicly available official source codes for [51]. Similar to [54], we use distance scaling and found that it is better to divide the distances by 32 for MiniImageNet and 16 for TieredImageNet based on the validation results.

AM3 implementation. In our re-implementation of AM3 [35], we have obtained the highest validation accuracy scores using 0.4 as the dropout rate for the fully-connected (FC) layers, on both MiniImageNet and TieredImageNet datasets. In the proposed model, we use 2 FC layers to encode the visual and semantic information for the sample attention module. These FC layers have the structure of FC-Dropout-ReLU-FC and the dimensions are reduced to 32. The dropout probabilities are selected to be 0.2 and 0.6 for visual and semantic branches respectively by using the validation sets. For feature attention module we use 2-layer FC module with the structure of FC-Softmax-FC and the dimensionality is first reduced to 32 and then increased to the feature dimensionality, which is 640 for ResNet-12. Finally, we also add the semantic prototype branch of AM3 [35] exactly as it is with the only difference of keeping the α a fixed hyper-parameter based on the validation set instead of predicting it from the semantic prior. We utilize the same Euclidean scaling in AM3 as in PN to obtain fair results. Overall, our implementation significantly improves the performance of the baseline AM3 model, compared to the scores originally reported in [35].

Source code. Training of our final model takes approximately 4 hours on a single V100 GPU. We are committed to making the PyTorch implementation of our approach publicly available upon publication.

4.2. Main results

In this section, we present our main results, ablative studies and experimental analysis of the proposed model on few-shot learning benchmarks. We then present a comparison to state-of-the-art approaches. Finally, we evaluate the approach on the more challenging 10-way and 15-way few-shot learning setups.

Baselines. Table 1 presents our main results for 1-shot or 5-shot, 5-way classification, including baselines that are carefully tuned in the same way to make fair comparisons. First of all, we validate our PN and AM3 baselines. For this purpose, we compare our PN and AM3 results (lower part) to the results reported in the original AM3 work [35] (upper part of Table 1). Overall, our results for both baselines are significantly better than the originally reported ones. Noticeably, our PN results on MiniImageNet are higher for nearly 7 and 4 points in 1-shot and 5-shot cases, respectively. Similarly, our AM3 results are higher for nearly 4 and 5 points in 1-shot and 5-shot cases, respectively. We observe even larger improvements for both baselines on TieredImageNet. These results re-highlight the importance of implementation details in FSL and validate the strength of our main baselines. A major factor in obtaining strong baselines is *backbone pretraining*, for which we present additional results in the context of our work later in this section.

Main results and ablative experiments. The very last row of Table 1 contains our main results using the final combined attention models. The preceding lines present the results for the ablated versions of the model with only sample attention or only feature attention based AM3 extensions. We note that in the case of 1-shot learning, sample attention has no difference by definition. From the results, first, we observe that sample attention in the case of 5-shot learning slightly degrades by 0.3 points on MiniImageNet but improves on TieredImageNet by nearly 0.5 points. Second, feature attention improves the AM3 model by approximately 2 points for 1-shot and slightly degrades by 0.5 points for 5-shot on MiniImageNet. We observe similar patterns on TieredImageNet.

Looking into the final combined attention model results in Table 1, we observe consistent improvements in all cases. The proposed combined attention improves 1-shot learning from 63.62 (PN) and 67.55 (AM3) to 69.76 on MiniImageNet. Similarly, 5-shot results improve from 78.37 (PN) and 80.22 (AM3) to 81.19. We observe similar improvements on TieredImageNet: 1-shot results improve from 67.58 (PN) and 72.60 (AM3) to 72.69, and 5-shot results improve

Table 1: Evaluation of our sample-attention, feature-attention and combined-attention models with comparisons to our implementations of PN and AM3, in the 5-way classification setting.

Model	MiniImageNet		TieredImageNet	
	1-Shot	5-Shot	1-Shot	5-Shot
Results from [35] (ResNet-12 backbone)				
PN [10]	56.52 \pm 0.45	74.28 \pm 0.20	58.47 \pm 0.64	78.41 \pm 0.41
AM3 [35]	65.21 \pm 0.30	75.20 \pm 0.27	67.23 \pm 0.34	78.95 \pm 0.22
Our implementation (ResNet-12 backbone)				
PN [10]	63.62 \pm 0.23	78.37 \pm 0.21	67.58 \pm 0.22	84.71 \pm 0.17
AM3 [35]	67.55 \pm 0.24	80.22 \pm 0.18	72.60 \pm 0.21	84.59 \pm 0.18
Sample attention (ours)	67.55 \pm 0.24	79.93 \pm 0.17	72.60 \pm 0.21	85.02 \pm 0.16
Feature attention (ours)	69.76 \pm 0.21	79.84 \pm 0.15	72.69 \pm 0.20	84.24 \pm 0.16
Combined (ours)	69.76 \pm 0.21	81.19 \pm 0.18	72.69 \pm 0.20	85.29 \pm 0.17

Table 2: Comparison of PN and Our approach with and without pretraining of the ResNet-12 backbone on MiniImageNet.

Model	MiniImageNet			
	1-Shot Val	1-Shot Test	5-Shot Val	5-Shot Test
PN w/o Pretraining [10]	43.43 \pm 0.67	43.70 \pm 0.24	71.90 \pm 0.66	69.67 \pm 0.22
Ours w/o Pretraining	64.55 \pm 0.70	61.57 \pm 0.23	72.46 \pm 0.68	68.78 \pm 0.19
PN with Pretraining [10]	68.73 \pm 0.68	63.62 \pm 0.20	80.88 \pm 0.67	78.37 \pm 0.22
Ours with Pretraining	76.44 \pm 0.65	69.76 \pm 0.21	85.02 \pm 0.65	81.19 \pm 0.18

from 84.71 (PN) and 84.59 (AM3) to 85.29.

Noticeably for 5-shot learning on TieredImageNet, the performance gaps between PN, AM3 and our semantic FSL models are relatively smaller, compared to those on MiniImageNet. This can be due to the fact that the backbone is pretrained with a more diverse training set, which is likely to yield better feature representations, make few-shot inference tasks less challenging and reduce the need for leveraging semantic priors. Consistently, we also observe that FSL models typically yield higher results on the test set of TieredImageNet, in comparison to MiniImageNet.

Importance of backbone pretraining. In Table 2 we inspect how pretraining and backbone quality affects the few-shot learning performance. Here we use both PN and our approach to see the effect in both uni-modal and multi-modal approaches. Models without pretraining are trained in an end-to-end fashion. In order to create a fair comparison, we train the models by using the same number of epochs in both cases. All of the models are trained for 400 epochs and the results reported based on best validation score. As can be seen from Table 2, model

Table 3: Few-shot classification accuracy on the test set of MiniImageNet for uni-modal (non-semantic) and multi-modal FSL approaches. * indicates our own implementation.

Model	Backbone	MiniImageNet	
		1-Shot	5-Shot
Uni-modal few-shot learning baselines			
Prototypical Networks* [10]	ResNet-12	63.62 \pm 0.23	78.37 \pm 0.21
Matching Networks [12]	ResNet-18	52.91 \pm 0.91	68.88 \pm 0.69
Relation Net [66]	ResNet-18	52.48 \pm 0.86	69.83 \pm 0.68
MAML [26]	ResNet-18	49.61 \pm 0.92	65.72 \pm 0.77
LogReg [33]	ResNet-18	51.75 \pm 0.80	74.27 \pm 0.63
LogReg Cosine [33]	ResNet-18	51.87 \pm 0.77	75.68 \pm 0.63
SARN [55]	Conv-4	51.62 \pm 0.31	66.16 \pm 0.51
TADAM [54]	ResNet-12	58.56 \pm 0.39	76.65 \pm 0.35
STANet [57]	ResNet-12	58.35 \pm 0.57	71.07 \pm 0.39
SimpleShot [67]	ResNet-18	62.85 \pm 0.20	80.02 \pm 0.14
MetaOptNet [68]	ResNet-12	62.64 \pm 0.61	78.63 \pm 0.46
LEO [28]	WRN-28-10	61.76 \pm 0.08	77.59 \pm 0.12
FEAT* [51]	ResNet-12	65.38 \pm 0.20	77.79 \pm 0.15
Multi-modal few-shot learning baselines			
KTN [60]	Conv-128	64.42 \pm 0.72	74.16 \pm 0.56
RIN [61]	ResNet-12	56.92 \pm 0.81	75.62 \pm 0.62
TriNet [37]	ResNet-18	58.12 \pm 1.37	76.92 \pm 0.69
ACAM [69]	ResNet-12	66.43 \pm 0.57	75.74 \pm 0.48
Baby steps [36]	ResNet-10	67.2	74.8
AM3 + TRAML [39]	ResNet-12	67.10 \pm 0.52	79.54 \pm 0.60
AM3* [35]	ResNet-12	67.55 \pm 0.24	80.22 \pm 0.18
Ours	ResNet-12	69.76 \pm 0.21	81.19 \pm 0.18

pretraining improves both approaches very significantly, and is crucial for achieving state-of-the-art performance. An interesting result is the comparison of PN and our approach for the 1-shot setting without pretraining since it highlights the significance of semantic information when the backbone quality and/or visual feature quality is low. As the backbone becomes better with pretraining the gap becomes smaller.

Comparison to the state-of-the-art. Although comparing classification results across models with different implementation details is particularly problematic in few-shot learning as previously stated, it is still of interest to show how well our results are in general, compared to the relevant state-of-the-art methods. Table 3 and Table 4 present the results of FSL methods (upper

Table 4: Few-shot classification accuracy on the test set of TieredImageNet for uni-modal (non-semantic) and multi-modal FSL approaches. * indicates our own implementation.

Model	Backbone	TieredImageNet	
		1-Shot	5-Shot
Uni-modal few-shot learning baselines			
Prototypical Networks* [10]	ResNet-12	67.58 \pm 0.22	84.71 \pm 0.19
Relation Net [66]	ResNet-12	54.48 \pm 0.93	71.32 \pm 0.78
MAML [26]	ResNet-12	51.67 \pm 1.81	70.30 \pm 0.08
MetaOptNet [68]	ResNet-12	65.99 \pm 0.72	81.56 \pm 0.63
SimpleShot [67]	ResNet-18	69.09 \pm 0.22	84.58 \pm 0.16
LEO [28]	ResNet-12	66.33 \pm 0.05	81.44 \pm 0.09
FEAT* [51]	ResNet-12	70.53 \pm 0.22	84.71 \pm 0.15
Multi-modal few-shot learning baselines			
ACAM [69]	ResNet-12	67.89 \pm 0.69	79.23 \pm 0.52
AM3* [35]	ResNet-12	72.60 \pm 0.21	84.59 \pm 0.15
Ours	ResNet-12	72.69 \pm 0.21	85.24 \pm 0.16

half) and semantic FSL methods (lower half), on MiniImageNet and TieredImageNet datasets respectively. First, we observe that our modernized PN baselines are strong baselines, outperforming many other more recently proposed approaches. Second, our final model outperforms all methods in FSL and semantic FSL categories on both MiniImageNet and TieredImageNet datasets. Not surprisingly, the most significant improvements are observed in the 1-shot settings on both datasets in comparison to the uni-modal (non-semantic) approaches since a single training sample is often insufficient and semantic prior becomes most valuable in this setting. Overall, the results highlight the value of semantic priming for few-shot learning.

10-way and 15-way few-shot learning. In Table 5 and Table 6, we compare PN, AM3 and the proposed approach for the more challenging 10-way and 15-way settings, respectively. As expected, the performances of all models decrease as the way count increases, especially in comparison to the commonly studied 5-way setting. Our model outperforms both PN and AM3 in the 1-shot setting, in both 10-way and 15-way experiments and the performance margin between those models increases. In the 5-shot setting, although our approach obtains nearly 1% better accuracy on validation sets, it falls 0.5 points behind AM3 on test sets. This points to a less than ideal correlation between the validation and test set performances. This is likely to degrade as the way count increases with a fixed sized validation (or test) set, since the overlap across

Table 5: Evaluation of 10-way few-shot classification on MiniImageNet.

MiniImageNet (10-way)				
Model	1-Shot Val	1-Shot Test	5-Shot Val	5-Shot Test
PN [10]	53.61 \pm 0.60	46.04 \pm 0.13	70.30 \pm 0.39	66.03 \pm 0.10
AM3 [35]	61.46 \pm 0.55	51.37 \pm 0.12	72.92 \pm 0.43	67.51 \pm 0.10
Ours	62.63 \pm 0.50	52.22 \pm 0.12	73.76 \pm 0.41	67.36 \pm 0.10

Table 6: Evaluation of 15-way few-shot classification on MiniImageNet.

MiniImageNet (15-way)				
Model	1-Shot Val	1-Shot Test	5-Shot Val	5-Shot Test
PN [10]	45.64 \pm 0.43	37.42 \pm 0.10	63.47 \pm 0.28	58.40 \pm 0.08
AM3 [35]	49.23 \pm 0.44	39.10 \pm 0.10	66.06 \pm 0.30	60.27 \pm 0.08
Ours	55.24 \pm 0.35	43.54 \pm 0.09	66.91 \pm 0.28	59.72 \pm 0.08

tasks increase and the variance across tasks reduce with larger way settings, reducing the overall richness and reliability of the evaluation.

4.3. Few-shot learning with noisy samples

While the problem of having support sample noise is not a commonly studied problem in FSL, possibly due to the presumption that few samples are likely to be correct, there can be various real-world scenarios where few-shot learner needs to operate over noisy support examples, e.g., learning in robotic systems through visual demonstration or through on-the-fly crawled samples of a desired class. In this sense, FSL techniques can be interpreted in a broader context as fundamental learning mechanisms as parts of approaches tackling noisy supervision problems. In addition, the support sample noise setting also provides a challenging experimental setup for understanding the effectiveness of the proposed sample attention mechanism.

In our sample noise experiments, we focus on the 5-shot 5-way setting on MiniImageNet. We artificially introduce noise into the support batches during training and testing using the following procedure: for each class in a task, we guarantee the existence of at least three correctly labeled support examples. With 50% probability, we apply label noise to the remaining samples by randomly permuting their class labels, resulting in up to two noisy samples per class, per task. We note that such within-task label noise is more challenging and practically more relevant than

Table 7: Evaluation of sample attention for handling noisy support samples.

Method	5-way 5-Shot Acc.
PN [10]	62.61 +- 0.22
Only Sample Attention	70.48 +- 0.20

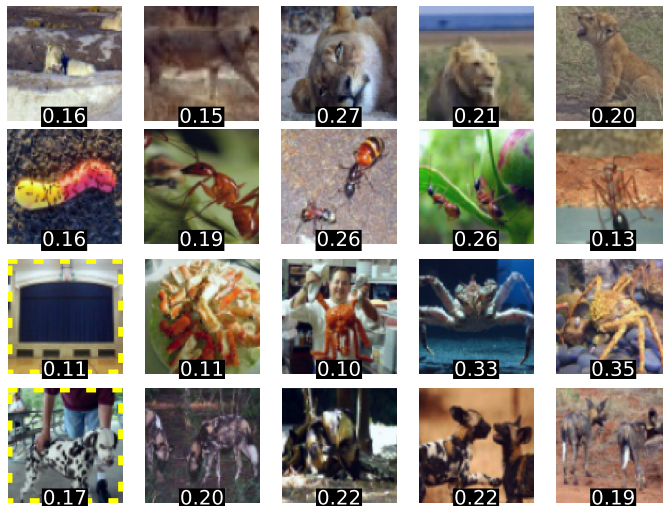


Figure 3: Sample attention examples for clean (top two rows) and noisy (lower two rows) support sample settings. Yellow dashed lines represent the noisy samples. Numbers indicate attention scores.

having completely irrelevant support samples, as irrelevant samples are less likely to cause cross-class confusion in the few-shot classification results. To evaluate sample attention in an isolated manner, we use sample attention directly on top of PN in these experiments.

We report the results in Table 7. We observe a large performance margin of approximately 7.8 points between the baseline (62.61) and noise-aware sample attention (70.48). This strongly suggests the ability of sample attention to utilize semantic priors for selecting the most informative support samples. To better understand how sample attention improves the results here, we also provide qualitative results both with and without noise in support samples, in Figure 3. Each image 5-tuple corresponds to a 5-shot support set of a class. For comparison, we present the sample attention results from both the clean support settings (the upper two rows) and the noisy support settings (the lower two rows). The dashed borders indicate noisy samples and the numbers indicate the resulting attention values. In the the top-most example, we observe that the




Figure 4: Distribution of sample attention weights for noisy and clean support samples.

model puts higher attention scores to support examples where the target *lion* is most recognizable, and estimates significantly lower attention to the images where the lion is small (the first example in the row) or pretty much unrecognizable (the second one). Similarly in the second row, we observe that the model attends more strongly to the support samples where the target *ant* instances are clearly recognizable. The example on the third row shows that the model estimates a low attention score to the noisy (the first one in the row) and cluttered (the second and third examples for the class) inputs for the target *king crab* classes. The fourth row shows an example where the attention to the noisy sample is comparable to the correct ones, due to the similarity across the noisy *dalmatian dog* sample and the target *African wild dog* class.

Sample attention distributions. The quantitative and qualitative examples highlight that sample the proposed attention mechanisms is able to leverage semantics to a large extent in prioritizing the training samples, but it can also unsurprisingly make (relative) mistakes due to variety of reasons. To understand the *correctness* of the attention estimates in a quantitative manner, beyond the accuracy values, we present the distribution of the attention score estimates for clean and noisy samples in Figure 4. Each histogram bin corresponds to an attention score range, and the shown frequencies are obtained over 100 randomly sampled 5-way, 5-shot tasks with noisy support batches. In these two distributions, we make the following observations: (i) attention scores given to noisy samples tend to be lower, especially where lower than 0.10 attention values are assigned almost exclusively to noisy samples, (ii) noise *detection* has room for improvement, as the assignment of relatively higher attention scores to noise samples is not rare, and (iii) the learned importance estimator works in a non-degenerative way where attention scores vary from

almost-zero values to 0.35 and above values.

5. Conclusions

In this paper, we propose an approach for leveraging semantic prior knowledge in few-shot learning of classifiers. The main motivation of our approach is to utilize semantic prior knowledge about classes to estimate importance of support samples and representation dimensions.

Our method performs well against not only the approaches that use only visual data but also the ones that use additional semantic information. The performance gains are more prominent in lower shot settings where the need for auxiliary semantic information is typically higher. We also study the case of sample noise in support sets. We also study the approach in a noisy FSL setting, where some of the few-shot support samples are stochastically corrupted. While we focus on a single semantic modality throughout our experiments, we believe that incorporating multiple modalities, as in [36], and integration into a generative FSL model based approach, as in [37], are promising future work directions.

Acknowledgements

This work was supported in part by the TUBITAK Grants 116E445 and 119E597. The numerical calculations reported in this paper were partially performed at TUBITAK ULAKBIM, High Performance and Grid Computing Center (TRUBA resources).

References

- [1] E. Bart, S. Ullman, Cross-generalization: learning novel classes from a single example by feature replacement, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., Vol. 1, 2005, pp. 672–679 vol. 1. [1](#)
- [2] M. Fink, Object classification from a single example utilizing class relevance metrics, in: Proc. Adv. Neural Inf. Process. Syst., 2004, pp. 449–456. [1](#)
- [3] Li Fei-Fei, R. Fergus, P. Perona, One-shot learning of object categories, IEEE Trans. Pattern Anal. Mach. Intell. 28 (4) (2006) 594–611. [1](#)
- [4] B. Lake, R. Salakhutdinov, J. Gross, J. Tenenbaum, One shot learning of simple visual concepts, Cognitive Science 33. [1](#)
- [5] B. Landau, L. B. Smith, S. S. Jones, The importance of shape in early lexical learning, Cognitive Development 3 (3) (1988) 299 – 321. [1](#)
- [6] W. E. Merriman, Categorization and naming in children: Problems of induction. ellen markman. cambridge, ma: Mit press, 1989. pp. 250., Applied Psycholinguistics 12 (3) (1991) 385–392. [1](#)

- [7] L. Smith, L. Slone, A developmental approach to machine learning?, *Frontiers in Psychology* 8. [1](#)
- [8] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, S. X. Yu, Large-Scale Long-Tailed Recognition in an Open World, in: *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Long Beach, CA, USA, 2019, pp. 2532–2541. [1](#)
- [9] W. J. Reed, The pareto, zipf and other power laws, *Economics letters* 74 (1) (2001) 15–19. [2](#)
- [10] J. Snell, K. Swersky, R. Zemel, Prototypical networks for few-shot learning, in: *Advances in Neural Information Processing Systems*, Vol. 30, 2017. [2](#), [3](#), [5](#), [6](#), [7](#), [15](#), [16](#), [17](#), [18](#), [19](#)
- [11] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, T. M. Hospedales, Learning to compare: Relation network for few-shot learning, in: *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 1199–1208. [2](#), [5](#)
- [12] O. Vinyals, C. Blundell, T. Lillicrap, k. kavukcuoglu, D. Wierstra, Matching networks for one shot learning, in: *Advances in Neural Information Processing Systems*, Vol. 29, 2016. [2](#), [5](#), [6](#), [12](#), [16](#)
- [13] G. R. Koch, Siamese neural networks for one-shot image recognition, 2015. [2](#)
- [14] B. Hariharan, R. Girshick, Low-shot Visual Recognition by Shrinking and Hallucinating Features, in: *ICCV*, 2017. [2](#), [5](#)
- [15] Y.-X. Wang, R. Girshick, M. Herbert, B. Hariharan, Low-shot learning from imaginary data, in: *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. [2](#)
- [16] H. Gao, Z. Shou, A. Zareian, H. Zhang, S.-F. Chang, Low-shot Learning via Covariance-Preserving Adversarial Augmentation Networks, in: *Advances in Neural Information Processing Systems* 31, 2018, pp. 983–993. [2](#), [5](#)
- [17] E. Schwartz, L. Karlinsky, J. Shtok, S. Harary, M. Marder, A. Kumar, R. Feris, R. Giryes, A. Bronstein, Delta-encoder: an effective sample synthesis method for few-shot object recognition, in: *Advances in Neural Information Processing Systems* 31, 2018, pp. 2845–2855. [2](#)
- [18] E. Schonfeld, S. Ebrahimi, S. Sinha, T. Darrell, Z. Akata, Generalized Zero- and Few-Shot Learning via Aligned Variational Autoencoders, in: *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Long Beach, CA, USA, 2019, pp. 8239–8247. [2](#), [3](#), [6](#)
- [19] A. Antoniou, A. Storkey, H. Edwards, Data Augmentation Generative Adversarial Networks, arXiv:1711.04340 [cs, stat]. [2](#)
- [20] K. Ridgeway, M. C. Mozer, Open-Ended Content-Style Recombination Via Leakage Filtering, arXiv e-prints (2018) arXiv:1810.00110 [arXiv:1810.00110](#). [2](#)
- [21] Z. Chen, Y. Fu, K. Chen, Y.-G. Jiang, Image Block Augmentation for One-Shot Learning, in: *AAAI*, 2019, p. 8. [2](#)
- [22] T. DeVries, G. W. Taylor, Dataset Augmentation in Feature Space, arXiv:1702.05538 [cs, stat]. [2](#)
- [23] S. Gidaris, N. Komodakis, Dynamic Few-Shot Visual Learning Without Forgetting, in: *CVPR*, 2018, p. 9. [2](#), [5](#)
- [24] H. Qi, M. Brown, D. G. Lowe, Low-shot learning with imprinted weights, in: *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. [2](#), [5](#)
- [25] S. Qiao, C. Liu, W. Shen, A. L. Yuille, Few-shot image recognition by predicting parameters from activations, in: *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. [2](#)
- [26] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in: *Proc. Int. Conf. Mach. Learn.*, Vol. 70 of *Proceedings of Machine Learning Research*, 2017, pp. 1126–1135. [2](#), [5](#), [16](#), [17](#)
- [27] A. Nichol, J. Achiam, J. Schulman, On First-Order Meta-Learning Algorithms, arXiv e-prints (2018) arXiv:1803.02999 [arXiv:1803.02999](#). [2](#), [5](#)
- [28] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, R. Hadsell, Meta-learning with latent

- embedding optimization, in: Proc. Int. Conf. Learn. Represent., 2019. [2](#), [5](#), [16](#), [17](#)
- [29] S. Ravi, H. Larochelle, Optimization as a model for few-shot learning, in: Proc. Int. Conf. Learn. Represent., 2017. [2](#), [5](#), [12](#)
- [30] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, T. Lillicrap, Meta-learning with memory-augmented neural networks, in: Proc. Int. Conf. Mach. Learn., Vol. 48 of Proceedings of Machine Learning Research, New York, New York, USA, 2016, pp. 1842–1850. [2](#)
- [31] S. Sukhbaatar, J. Weston, R. Fergus, et al., End-to-end memory networks, in: Advances in neural information processing systems, 2015, pp. 2440–2448. [2](#)
- [32] J. Bronskill, J. Gordon, J. Requeima, S. Nowozin, R. E. Turner, TaskNorm: Rethinking Batch Normalization for Meta-Learning, in: Machine Learning and Systems, 2020. [2](#), [5](#), [12](#)
- [33] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. Wang, J.-B. Huang, A closer look at few-shot classification, in: Proc. Int. Conf. Learn. Represent., 2019. [2](#), [12](#), [16](#)
- [34] Y. Chen, X. Wang, Z. Liu, H. Xu, T. Darrell, A new meta-baseline for few-shot learning, ArXiv abs/2003.04390. [2](#), [5](#), [12](#)
- [35] C. Xing, N. Rostamzadeh, B. N. Oreshkin, P. H. O. Pinheiro, Adaptive cross-modal few-shot learning, in: NeurIPS, 2019. [2](#), [3](#), [4](#), [6](#), [7](#), [12](#), [13](#), [14](#), [15](#), [16](#), [17](#), [18](#)
- [36] E. Schwartz, L. Karlinsky, R. Feris, R. Giryes, A. M. Bronstein, Baby steps towards few-shot learning with multiple semantics, arXiv e-prints (2019) arXiv:1906.01905 [arXiv:1906.01905](#). [2](#), [3](#), [6](#), [16](#), [21](#)
- [37] Z. Chen, Y. Fu, Y. Zhang, Y.-G. Jiang, X. Xue, L. Sigal, Multi-Level Semantic Feature Augmentation for One-Shot Learning, IEEE Transactions on Image Processing 28 (9) (2019) 4594–4605. [arXiv:1804.05298](#). [2](#), [3](#), [6](#), [16](#), [21](#)
- [38] J. Mu, P. Liang, N. Goodman, Shaping Visual Representations with Language for Few-shot Classification, arXiv e-prints (2019) arXiv:1911.02683 [arXiv:1911.02683](#). [2](#)
- [39] A. Li, W. Huang, X. Lan, J. Feng, Z. Li, L. Wang, Boosting few-shot learning with adaptive margin loss, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2020. [2](#), [4](#), [6](#), [16](#)
- [40] E. Triantafillou, T. Zhu, V. Dumoulin, P. Lamblin, U. Evci, K. Xu, R. Goroshin, C. Gelada, K. Swersky, P.-A. Manzagol, H. Larochelle, Meta-dataset: A dataset of datasets for learning to learn from few examples, in: Proc. Int. Conf. Learn. Represent., 2020. [2](#)
- [41] E. H. ROSCH, On the internal structure of perceptual and semantic categories, in: Cognitive Development and Acquisition of Language, San Diego, 1973, pp. 111 – 144. [3](#)
- [42] A. Frome, G. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, T. Mikolov, Devise: A deep visual-semantic embedding model, in: Proc. Adv. Neural Inf. Process. Syst., 2013. [3](#), [6](#)
- [43] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. Corrado, J. Dean, Zero-shot learning by convex combination of semantic embeddings, in: Proc. Int. Conf. Learn. Represent., 2014. [3](#), [6](#)
- [44] L. Zhang, T. Xiang, S. Gong, Learning a deep embedding model for zero-shot learning, Proc. IEEE Conf. Comput. Vis. Pattern Recog. (2017) 3010–3019. [3](#), [6](#)
- [45] L. Chen, H. Zhang, J. Xiao, W. Liu, S. Chang, Zero-shot visual recognition using semantics-preserving adversarial embedding networks, Proc. IEEE Conf. Comput. Vis. Pattern Recog. (2018) 1043–1052. [3](#), [6](#)
- [46] Y. Xian, S. Sharma, B. Schiele, Z. Akata, f-vaegan-d2: A feature generating framework for any-shot learning, in:

- Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2019, pp. 10275–10284. [3](#), [6](#)
- [47] A. Mishra, S. Krishna Reddy, A. Mittal, H. A. Murthy, A generative model for zero shot learning using conditional variational autoencoders, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2018, pp. 2188–2196. [3](#), [6](#)
- [48] Y. Zhu, M. Elhoseiny, B. Liu, X. Peng, A. Elgammal, A generative adversarial approach for zero-shot learning from noisy texts, Proc. IEEE Conf. Comput. Vis. Pattern Recog. (2018) 1004–1013. [3](#), [6](#)
- [49] Z. Li, F. Zhou, F. Chen, H. Li, Meta-SGD: Learning to Learn Quickly for Few-Shot Learning, arXiv e-prints (2017) arXiv:1707.09835 [arXiv:1707.09835](#). [5](#)
- [50] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, P. Isola, Rethinking Few-Shot Image Classification: a Good Embedding Is All You Need?, in: Proc. European Conf. on Computer Vision, 2020. [5](#)
- [51] H.-J. Ye, H. Hu, D.-C. Zhan, F. Sha, Few-shot learning via embedding adaptation with set-to-set functions, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2020, pp. 8808–8817. [5](#), [12](#), [13](#), [16](#), [17](#)
- [52] R. ZHANG, T. Che, Z. Ghahramani, Y. Bengio, Y. Song, Metagan: An adversarial approach to few-shot learning, in: Advances in Neural Information Processing Systems, Vol. 31, 2018, pp. 2365–2374. [5](#)
- [53] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative Adversarial Networks, in: Advances in Neural Information Processing Systems, 2014. [5](#)
- [54] B. N. Oreshkin, P. R. Lopez, A. Lacoste, Tadam: Task dependent adaptive metric for improved few-shot learning, in: NeurIPS, 2018. [5](#), [13](#), [16](#)
- [55] B. Hui, P. Zhu, Q. Hu, Q. Wang, Self-attention relation network for few-shot learning, in: 2019 IEEE International Conference on Multimedia Expo Workshops (ICMEW), 2019, pp. 198–203. [doi:10.1109/ICMEW.2019.00041](#). [5](#), [16](#)
- [56] X. Wang, R. Girshick, A. Gupta, K. He, Non-local neural networks, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2018, pp. 7794–7803. [5](#)
- [57] S. Yan, S. Zhang, X. He, A dual attention network with semantic embedding for few-shot learning, in: AAAI Conference on Artificial Intelligence, 2019. [5](#), [16](#)
- [58] T. Gao, X. Han, Z. Liu, M. Sun, Hybrid Attention-Based Prototypical Networks for Noisy Few-Shot Relation Classification, in: AAAI, 2019. [5](#)
- [59] M. Ren, S. Ravi, E. Triantafillou, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, R. S. Zemel, Meta-learning for semi-supervised few-shot classification, in: Proc. Int. Conf. Learn. Represent., 2018. [6](#), [12](#)
- [60] Z. Peng, Z. Li, J. Zhang, Y. Li, G. Qi, J. Tang, Few-shot image recognition with knowledge transfer, Proc. IEEE Int. Conf. on Computer Vision (2019) 441–449. [6](#), [16](#)
- [61] Z. Ji, X. Chai, Y. Yu, Z. Zhang, Reweighting and information-guidance networks for Few-Shot Learning, Neurocomputing 423 (2021) 13–23. [6](#), [16](#)
- [62] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, Li Fei-Fei, Imagenet: A large-scale hierarchical image database, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2009, pp. 248–255. [12](#)
- [63] J. Pennington, R. Socher, C. D. Manning, Glove: Global vectors for word representation, in: Empirical Methods in Natural Language Processing (EMNLP), 2014, pp. 1532–1543. [12](#)
- [64] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: Proc. Int. Conf. Mach. Learn., Vol. 37 of Proceedings of Machine Learning Research, Lille, France, 2015, pp. 448–456. [12](#)

- [65] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, A. Lerer, Automatic differentiation in pytorch. [12](#)
- [66] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. S. Torr, T. M. Hospedales, Learning to Compare: Relation Network for Few-Shot Learning, in: CVPR, 2018. [16](#), [17](#)
- [67] Y. Wang, W.-L. Chao, K. Q. Weinberger, L. van der Maaten, SimpleShot: Revisiting Nearest-Neighbor Classification for Few-Shot Learning, arXiv e-prints (2019) arXiv:1911.04623[arXiv:1911.04623](#). [16](#), [17](#)
- [68] K. Lee, S. Maji, A. Ravichandran, S. Soatto, Meta-learning with differentiable convex optimization, in: Proc. IEEE Conf. Comput. Vis. Pattern Recog., 2019. [16](#), [17](#)
- [69] J. Lian, H. Wang, S. Xiong, Learning class prototypes via anisotropic combination of aligned modalities for few-shot learning, in: 2020 IEEE International Conference on Multimedia and Expo (ICME), 2020, pp. 1–6. [16](#), [17](#)