

## Few-shot Learning with Noisy Labels

Kevin J Liang<sup>1</sup> Samrudhdi B. Rangrej<sup>2</sup> Vladan Petrovic<sup>1</sup> Tal Hassner<sup>1</sup>  
<sup>1</sup>Facebook AI Research <sup>2</sup>McGill University

kevinjliang@fb.com

### Abstract

*Few-shot learning (FSL) methods typically assume clean support sets with accurately labeled samples when training on novel classes. This assumption can often be unrealistic: support sets, no matter how small, can still include mislabeled samples. Robustness to label noise is therefore essential for FSL methods to be practical, but this problem surprisingly remains largely unexplored. To address mislabeled samples in FSL settings, we make several technical contributions. (1) We offer simple, yet effective, feature aggregation methods, improving the prototypes used by ProtoNet, a popular FSL technique. (2) We describe a novel Transformer model for Noisy Few-Shot Learning (TraNFS). TraNFS leverages a transformer’s attention mechanism to weigh mislabeled versus correct samples. (3) Finally, we extensively test these methods on noisy versions of MiniImageNet and TieredImageNet. Our results show that TraNFS is on-par with leading FSL methods on clean support sets, yet outperforms them, by far, in the presence of label noise.*

### 1. Introduction

Modern few-shot learning (FSL) methods aim to learn classifiers for novel classes from only a handful of examples. These methods, however, generally assume that the few *support set* samples used for training were carefully selected to represent their class. Unfortunately, real-world settings rarely offer such guarantees. In fact, even carefully annotated and curated datasets often contain mislabeled samples [9, 33, 45, 53], due to automated weakly supervised annotation, ambiguity, or even human error.

Whereas there are plenty of methods designed for learning with noise in many-shot supervised settings [1, 19, 20, 25, 31, 51], noise in few-shot settings remains largely unexplored. This dearth is surprising considering the utility of FSL methods in settings where human supervision cannot easily be provided: *e.g.* in fully automated systems which learn many novel classes [12, 43, 56, 57], making human curation of the labels of every support set, unrealistic.

Fig. 1 shows the challenge of learning from few, possibly



Figure 1. **Few-shot learning with mislabeled samples.** A 5-shot, 5-way support set of MiniImageNet [48] images. Rows show support set samples of each novel class. Two samples in each row were mislabeled by symmetric label flips (Sec. 6.1). Can you spot which ones? See Supp. for answers and more examples.

mislabeled, examples. It presents a sample 5-shot, 5-way support set from MiniImageNet [48]. Each row includes the support set training images of one of the five classes. Two of the samples in each row are mislabeled with symmetric label noise (Sec. 6.1). With so few examples, spotting mislabeled images can be difficult, even for humans with considerable prior knowledge, which FSL methods lack.

As we later demonstrate empirically, FSL methods are especially vulnerable to such label noise. When training from few samples, each sample represents a significant contribution to the final decision boundary. Thus, even a single noisy example can be destructive to the model’s accuracy. We illustrate this observation in Fig. 2, which reports the performance of ProtoNet [40], a popular FSL method, on MiniImageNet with noisy labels. ProtoNet averages the convolutional features of each class’s support set into class *prototypes*. Queries are then classified by the class of their nearest neighbor prototype. Fig. 2 shows the effect of increasing the number of mislabeled samples, compared with a model trained after mislabeled samples were removed

(*i.e.*, smaller, but cleaner, support sets). The widening gap between the two curves reflects the degradation of accuracy when mislabeled samples are not accounted for.

We address the vulnerability of FSL methods to label noise by making a number of technical innovations. We begin by exploring simple, yet effective alternatives to the design of ProtoNet [40]. Specifically, we replace the mean operator, used by ProtoNet for aggregating support set features, with more robust methods. We evaluate an unweighted option, the median, and options which weigh support set samples based on feature similarities. We show that these changes already improve robustness to label noise.

We then introduce our **Transformer model for Noisy Few-Shot Learning (TraNFS)**. Unlike previous methods, TraNFS *learns* to aggregate support samples into class representations. The transformer architecture offers a natural means for processing variable numbers of shots and ways with permutation invariance. Robustness to label noise is achieved by leveraging a modified version of the transformer’s self-attention mechanism [47]. This modified self-attention used by TraNFS compares support set samples and downweights samples considered likely to be mislabeled.

We test our proposed methods extensively on versions of MiniImageNet [48] and TieredImageNet [38] with three methods of adding label noise. Our results show that the proposed TraNFS (and even the simpler modifications of ProtoNet) surpass popular FSL methods by wide margins in the presence of label noise, while offering comparable performance in the absence of label noise.

To summarize, we make the following contributions.

- We propose median and similarity weighting as simple yet effective substitutes to ProtoNet’s mean prototypes.
- We present TraNFS, a novel transformer model adapted to FSL with noisy labels.
- We extensively benchmark many popular FSL methods on three types of support set noise pollution: symmetric, paired, and outlier.

## 2. Related work

**Few-shot learning.** The field of FSL methods is vast; we refer to surveys for comprehensive overviews [6, 7].

Metric-based methods classify query samples based on their similarity to each class’s support examples, learning a transferable embedding space for which such comparisons can be made. Metrics such as cosine similarity [48], Euclidean distance [40], Mahalanobis distance [6], and Earth Mover’s Distance (EMD) [59] have been shown to be effective. RelationNet [41] and Satorras *et al.* [39] used convolutional and graph neural networks, respectively, to learn a similarity metric. TADAM [34], FEAT [55], and TAFE-Net [50] proposed task-specific adaptation of embeddings. CrossTransformers [15] used attention for spatially-aware similarity between local features.

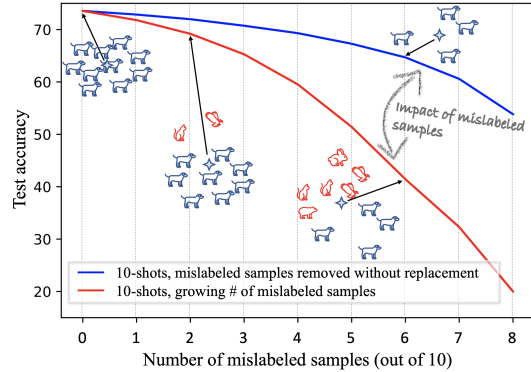


Figure 2. **Number of mislabeled samples versus few-shot learning accuracy.** Accuracy for 10-shot, 5-way classification on MiniImageNet [48] reported for a vanilla ProtoNet [40]. The animals represent support set embeddings, with the mean prototype (star) being pulled out of the clean class (dog) distribution with increasing mislabeled samples. **Blue:** Accuracy if mislabeled samples are known and ignored. **Red:** Accuracy when using full support sets without removing mislabeled samples. The gap between these two curves reflects the vulnerability of few-shot learning to label noise.

Optimization-based methods fine-tune model parameters on few support examples. MAML [2, 16] learned model parameter initializations that allow fast fine-tuning on few samples. REPTILE [32] simplified MAML with a first-order formulation. MetaNet [30] introduced fast and slow weights for fast parameterization and rapid generalization. Bertinetto *et al.* [8] and MetaOptNet [23] presented closed form solutions and differentiable solvers for task-dependent Ridge Regression, Logistic Regression (LR), and Support Vector Machines (SVMs). Tian *et al.* [42] showed learning of generalizable feature embeddings used to train linear classifiers on novel tasks.

**Noisy labels and outliers.** Methods for learning noise transition matrices are common [26, 54, 60]. Estimating a noise transition matrix from a handful of potentially mislabeled samples, however, is an ill-posed problem. Other methods [19, 20, 58] leverage deep neural networks’ tendency to learn easier (and thus likely correctly labeled) samples first [4, 44] to select reliable samples to learn from, but such behavior cannot be relied upon when only a few samples are available. Deep out-of-distribution (OOD) detection is also extensively explored [10, 14, 37, 49], but these methods typically focus on identifying test-time outliers that are out-of-distribution relative to the training set. In FSL, disjoint base and novel sets mean that *all* meta-test samples are considered OOD, including correctly labeled support set samples. Finally, there are several works that take a meta-learning approach to learning noisy labels [24, 52, 61], but these methods typically assume known label spaces with abundant data (*i.e.* many-shots), rather than our few-shot setting. With only few training samples, these methods fail.

**Robust FSL.** With few previous works, noisy labels have

largely been ignored by FSL methods. RNNP [29] combined data augmentation with repeated applications of  $k$ -means to produce refined prototypes, but such unsupervised clustering implicitly assumes that noisy data is from one of the support set classes. RapNets [28] proposed a BiLSTM-based attentive module to overcome representation or label noise. Alternatively, RW-MAML [22] learned to weigh support samples by extending MAML to bi-bi-level optimization, but it considers the less realistic setting of mixing in OOD tasks during metatraining rather than noisy few-shot meta-test. Finally, robustness of meta-learners to adversarial attacks have also been considered [18].

### 3. Preliminaries

FSL classification tasks are often referred to as  $K$ -shot  $N$ -way, where  $N$  is the number of classes being learned and  $K$  is the number of labeled samples per class to learn from. These  $KN$  samples  $S = \{x_1^{(1)}, x_2^{(1)}, \dots, x_{K-1}^{(N)}, x_K^{(N)}\}$  are often referred to as the *support set*. After training, unlabeled queries are classified into one of these  $N$  classes. To produce an effective classifier of novel classes  $C^n$  from few samples, FSL models typically use knowledge transfer, leveraging a dataset of *base classes*  $C^b$  with ample labeled data. It is commonly assumed that the classes in  $C^n$  are unknown in advance and thus absent from  $C^b$  (i.e.  $C^b \cap C^n = \emptyset$ ). Recent FSL methods often adopt a meta-learning paradigm, simulating the desired inference-time behavior by meta-training the model with many episodes of  $K$ -shot,  $N$ -way tasks, optimizing for accuracy on  $Q$  query samples from each of the episode’s  $N$  classes.

**ProtoNets.** One such relevant FSL method is Prototypical Networks (ProtoNets) [40]. ProtoNets use a convolutional feature extractor  $\mathcal{F}$  to convert each sample in the support set to an embedding  $h_i^{(c)} = \mathcal{F}(x_i^{(c)}) \in \mathbb{R}^D$ . These embeddings are then aggregated into  $N$  class prototypes  $p^{(c)}$  using a simple mean of the embeddings for each class  $c$ :

$$p^{(c)} = \frac{1}{K} \sum_i \mathcal{F}(x_i^{(c)}). \quad (1)$$

A query sample,  $x_q$ , is then classified based on the nearest prototype in embedding space:

$$y = \underset{c}{\operatorname{argmin}} d(\mathcal{F}(x_q), p^{(c)}). \quad (2)$$

Despite its simplicity, ProtoNets remain a strong baseline, and its easy implementation makes it compelling for real-world use cases at scale. Using mean to aggregate embeddings, however, implies sensitivity to mislabeled samples, especially when only few samples are provided. Indeed, as we show in Fig. 2, incorrectly labeled samples can easily degrade accuracy of the resulting classifiers. This is a symptom of the prototypes being *pulled away* from the class’s true (unknown) mean by mislabeled samples.

## 4. Static alternatives to the mean

Using mean as proposed by ProtoNet [40] to aggregate features is not the only way to combine embeddings into prototypes: other aggregation methods may be better suited when mislabeled samples are expected. We begin by exploring simple alternatives to the mean, intended to make prototypes more robust to mislabeled samples while maintaining accuracy if all labels are correct.

### 4.1. Spatial median prototypes

The median is a natural alternative to the mean in noisy settings. While order statistics like the median are well defined for scalars, this is not the case for vectors. For scalars, there is a connection between various distribution statistics (e.g. mean, median, mode) and minimization of the appropriate loss functions [5]. For example, empirical mean minimizes total squared error between the mean and values in the set. Similarly, empirical median minimizes total absolute error between the median and the set, so finding a median is equivalent to minimizing the total absolute error.

This minimization generalizes well to higher dimensional spaces. We thus define a cost function to be the sum of distances to embedding vectors  $h_i^{(c)}, i \in \{1, 2, \dots, K\}$ , in the set for each class,  $c$ , and find the median vector  $p^{(c)}$  minimizing this cost. For brevity, we drop the class index  $c$  in derivations that follow. To make the loss differentiable at all points, we work with a smooth version of the loss, usually referred to as the *pseudo-Huber* loss:

$$\mathcal{L}(p) = \sum_{i=1}^K \left( \sqrt{\|p - h_i\|_2^2 + \epsilon^2} - \epsilon \right), \quad (3)$$

where  $K$  is the number of vectors in the set,  $\|\cdot\|_2$  is an  $L^2$  vector norm, and  $\epsilon$  is a small constant.

No closed-form solution for this minimization problem exists, so we use Newton’s method for an iterative solution:

$$p(t+1) = p(t) - \mathcal{H}^{-1}(p(t)) \cdot \nabla \mathcal{L}(p(t)). \quad (4)$$

We find the gradient,  $\nabla \mathcal{L}(p)$ , and the Hessian matrix,  $\mathcal{H}(p)$ , using matrix calculus with numerator layout as,

$$\nabla \mathcal{L}(p) = \sum_{i=1}^K \frac{p - h_i}{\sqrt{\|p - h_i\|_2^2 + \epsilon^2}}, \quad (5)$$

$$\mathcal{H}(p) = \left( \sum_{i=1}^K \frac{1}{\sqrt{\|p - h_i\|_2^2 + \epsilon^2}} \right) I_{D \times D} - UU^T, \quad (6)$$

where  $D$  is the dimension of the vector space,  $I_{D \times D}$  is the identity matrix, and  $U = [u_1, u_2, \dots, u_K]$  is a  $D \times K$  matrix formed by stacking vectors  $u_i = \frac{p - h_i}{(\|p - h_i\|_2^2 + \epsilon^2)^{\frac{3}{4}}}$ . As an approximation, we can neglect the second, non-diagonal

term in the Hessian, in which case the iteration becomes:

$$p(t+1) = p(t) - \frac{\sum_{i=1}^K \frac{p(t) - h_i}{\sqrt{\|p(t) - h_i\|_2^2 + \epsilon^2}}}{\sum_{i=1}^K \frac{1}{\sqrt{\|p(t) - h_i\|_2^2 + \epsilon^2}}}. \quad (7)$$

Note that the choice of pseudo-Huber loss with a small constant  $\epsilon$  avoids division by zero even when the median estimate falls exactly at one of the vectors in the support set.

## 4.2. Similarity weighted prototypes

ProtoNet-style mean aggregation uniformly weighs all shots of a class’s support set. A clear extension to this scheme is a non-uniform weighting which suppresses outliers and amplifies clean samples. Of course, if we knew which samples were mislabeled, we could remove them from the support set, but this information is typically unavailable. Instead, we can try to identify mislabeled samples based on how the support set is arranged in feature space.

Specifically, we assume that a well-trained feature extractor  $\mathcal{F}$  embeds correctly labeled samples close to one another [21], thus on average being closer in the induced metric space than any mislabeled samples. This intuition implies that the subset of correctly labeled samples is larger than any subset of mislabeled samples from a single, unrelated class. This assumption, however, is typical to many robust estimators, including, *e.g.*, the Random Sample Consensus (RANSAC) [17]. Building on this assumption, we offer the following similarity measures.

**Squared Euclidean distance.** This distance is the same one minimized by ProtoNets and thus a natural choice for measuring distances when attempting to identify mislabel samples. We compute the similarity score as:

$$a_i^{(c)} = -\frac{1}{K-1} \sum_{i \neq j} \|h_i^{(c)} - h_j^{(c)}\|_2^2. \quad (8)$$

Lower distance implies being closer to other support samples, so we negate the average distance for the final score.

**Absolute distance.** The  $L^2$  norm can heavily penalize large distances in few feature dimensions: a large difference in only few dimensions may result in a large distance between the features, even if they share similar values in all other dimensions. We thus also consider  $L^1$ :

$$a_i^{(c)} = -\frac{1}{K-1} \sum_{i \neq j} |h_i^{(c)} - h_j^{(c)}|. \quad (9)$$

As with the Euclidean distance, we use a factor of  $-1$  so that lower distances produce higher scores.

**Cosine similarity.** While not a proper distance metric, the cosine angle between two features is a common measure of feature similarity in the few-shot literature [11].

$$a_i^{(c)} = \frac{1}{K-1} \sum_{i \neq j} \frac{h_i^{(c)} \cdot h_j^{(c)}}{\|h_i^{(c)}\| \|h_j^{(c)}\|}. \quad (10)$$

As the inputs are normalized, cosine similarity is less sensitive to the magnitude of the embeddings.

**Aggregating features with weighted similarity.** Once we obtain the average distance of each feature to others in the same support set, using one of the scores above, we produce an aggregated prototype by weighting the support samples using these scores, normalizing the result with a softmax.

$$w_i^{(c)} = \frac{\exp(a_i^{(c)}/T)}{\sum_j \exp(a_j^{(c)}/T)}, \quad (11)$$

$$p^{(c)} = \sum_i w_i^{(c)} \mathcal{F}(x_i^{(c)}), \quad (12)$$

where  $T$  is a temperature term controlling the diffuseness of the softmax. As  $T \rightarrow 0$ , this method picks the support sample with minimum distance to the other samples as a class prototype, while as  $T \rightarrow \infty$ , this reduces to the mean (*i.e.* ProtoNets [40]). We choose soft-weighting of support samples, rather than top- $k$  selection or a hard threshold, as the latter two require either knowing the number of noisy samples or threshold tuning, which may vary depending on the class or support sample distribution.

## 5. Learning a prototype aggregator

The aggregation methods discussed in Sec. 4, weighted or otherwise, are hard coded: They do not adjust to differences in support set feature distributions. We hypothesize that a learned mechanism that compares support set embeddings for similarity *and then refines them*, can potentially outperform these static methods. Crucially, in typical FSL settings, the number and order of support samples and classes are arbitrary. Thus, any learned alternative to the methods of Sec. 4 must process arbitrary numbers of shots or ways while remaining permutationally invariant to both.

### 5.1. A transformer model for noisy FSL

Given these requirements, we propose a **Transformer** model for **Noisy Few-Shot Learning** (TraNFS) (Fig. 3). Transformers are designed to process sequences of arbitrary length while offering permutation invariance. Importantly, we note that a transformer’s self-attention mechanism [47] can be leveraged to compute similarities between support set samples and naturally weigh them when aggregating them into prototypes. To this end, we concatenate the convolutional features of a support set’s samples to form an input sequence  $\mathbf{h} = [h_1^{(1)}, h_2^{(1)}, \dots, h_{K-1}^{(N)}, h_K^{(N)}]$  to the transformer,  $\mathcal{T}$ . We then make the following adaptations to the transformer to enable it to process a typical FSL support set.

**Class token.** Partly inspired by BERT [13], we use a set of classification tokens  $\text{CLS}^{(c)}$ ,  $c \in \{1, \dots, N\}$  to denote the positions representing the prototype for each class and concatenate  $[\text{CLS}^{(1)}, \dots, \text{CLS}^{(N)}]$  to the support set embedding

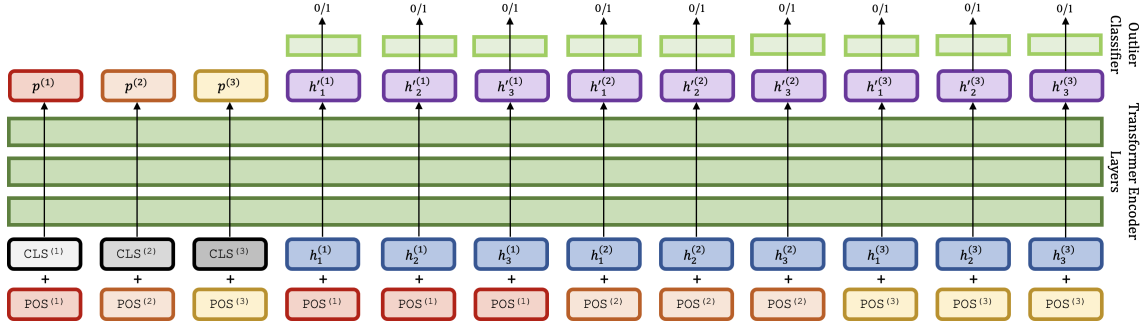


Figure 3. Visualization of our proposed TraNFS architecture, for a 3-shot 3-way support set example input / output sequence.

sequence  $h$ . By taking the output at the position of  $\text{CLS}^{(c)}$  to be the prototype  $p^{(c)}$  for class  $c$ , we motivate the transformer to learn to aggregate the information in all support set samples into this position. There are multiple choices for instantiating  $\text{CLS}^{(c)}$ , including as a random constant, the mean of the support set embeddings for class  $c$  (*i.e.* a mean prototype), or a learnable embedding. We report comparisons of these variations in the supplemental material.

**Positional encoding.** Shot and class order are both typically arbitrary in FSL and should therefore not be encoded. Still, we require some means of informing the transformer the class identity of each support sample. Vaswani *et al.* [47] utilized a sinusoidal positional encoding added to the input sequence to indicate word order. We repurpose this mechanism and use it to encode the class  $c$  associated with each position in the input sequence. Specifically, we create  $N$   $D$ -dimensional embeddings corresponding to special tokens  $\text{POS}^{(c)}$ ,  $c \in \{1, \dots, N\}$  and add each  $\text{POS}^{(c)}$  to all support sample embeddings  $h_i^{(c)}$  and the class tokens  $\text{CLS}^{(c)}$ , as seen in Fig. 3. With the positional embeddings added to the input sequence, the transformer can learn to attend to the positional encoding to associate support set embeddings and the prototype position of each class together.

## 5.2. Optimization

We meta-train TraNFS to minimize the standard ProtoNet loss. Logits are computed as the negative distance  $d$  between prototypes predicted by the model at the  $\text{CLS}$  token positions and the embedded query sample  $\mathcal{F}(x_q)$ :

$$\mathcal{L}_{\text{xent}} = - \sum_{c=1}^N y_q \cdot \log \left( \frac{\exp(-d(p^{(c)}, \mathcal{F}(x_q)))}{\sum_{c'} \exp(-d(p^{(c')}, \mathcal{F}(x_q)))} \right) \quad (13)$$

where  $\cdot$  is the dot product and  $y_q$  is the one-hot ground truth label of query  $x_q$ .

When meta-training TraNFS, we found it essential to expose the model to support sets with noisy samples (Sec. 6.4). We do this by artificially introducing label noise to the support set, using label  $o_i^{(c)} \in \{0, 1\}$  to track the positions of the noisy samples. This step ensures that the transformer learns a noise rejection mechanism. Without noisy

samples, the transformer is not motivated to learn anything beyond recreating ProtoNet by averaging support samples.

**Clean prototype loss.** Besides optimizing the position of predicted prototypes relative to the meta-training query samples, we also encourage the predicted prototype for each class to be close to a *clean* prototype,  $\hat{p}^{(c)}$ , aggregated from correctly labeled samples in the support set:

$$\hat{p}^{(c)} = \frac{1}{K - \sum_i o_i^{(c)}} \sum_i \mathbf{1}[o_i^{(c)} = 0] \mathcal{F}(x_i^{(c)}), \quad (14)$$

$$\mathcal{L}_{\text{clean}} = \frac{1}{N} \sum_c \|p^{(c)} - \hat{p}^{(c)}\|_2^2. \quad (15)$$

We choose mean squared error here, but other alternatives such as negative cosine similarity are also viable.

**Binary outlier classification loss.** The ProtoNet and clean prototype losses described above both implicitly encourage identification of noisy samples. We found it helpful to also explicitly train the model to classify support set samples as either mislabeled or not.

We instantiate the binary classifier as a fully connected layer  $\mathcal{B}$  applied to the transformer’s output at positions corresponding to the support set samples. We share weights for  $\mathcal{B}$  across all such positions, with loss term:

$$\mathcal{L}_{\text{bin}} = - \frac{1}{KN} \sum_{i,c} o_i^{(c)} \log \sigma(\mathcal{B}(h_i^{(c)})) + (1 - o_i^{(c)}) \log (1 - \sigma(\mathcal{B}(h_i^{(c)}))), \quad (16)$$

where  $\sigma$  is the sigmoid function and  $h_i^{(c)}$  is the transformer output corresponding to  $h_i^{(c)}$ .

Our final optimization objective combines the three losses described above:

$$\mathcal{L} = \mathcal{L}_{\text{xent}} + \lambda_c \mathcal{L}_{\text{clean}} + \lambda_b \mathcal{L}_{\text{bin}}, \quad (17)$$

where  $\lambda_c$  and  $\lambda_b$  are weighting terms for the clean prototype and binary outlier classification losses, respectively.

Table 1. **Few-shot with symmetric label swap noise.** 5-way 5-shot Acc.  $\pm$  95% CI on [MiniImageNet](#) [48], [TieredImageNet](#) [38]. Our TraNFS is comparable to existing methods at 0% noise, with a growing gap in its favor as noise levels increase. Best viewed in color.

Model \ Noise Proportion		0%		20%		40%		60%	
Baselines	Oracle	68.18 $\pm$ 0.16	71.42 $\pm$ 0.18	66.08 $\pm$ 0.17	69.19 $\pm$ 0.19	62.60 $\pm$ 0.17	66.14 $\pm$ 0.20	56.89 $\pm$ 0.18	60.39 $\pm$ 0.21
	Nearest $k = 1$	55.91 $\pm$ 0.17	58.81 $\pm$ 0.20	47.27 $\pm$ 0.18	49.48 $\pm$ 0.19	38.68 $\pm$ 0.18	40.25 $\pm$ 0.19	29.20 $\pm$ 0.16	29.84 $\pm$ 0.17
	Nearest $k = 3$	55.29 $\pm$ 0.18	58.44 $\pm$ 0.20	48.43 $\pm$ 0.17	51.11 $\pm$ 0.19	39.14 $\pm$ 0.17	41.09 $\pm$ 0.18	29.66 $\pm$ 0.15	30.69 $\pm$ 0.15
	Nearest $k = 5$	56.15 $\pm$ 0.18	59.22 $\pm$ 0.20	50.92 $\pm$ 0.17	53.75 $\pm$ 0.19	42.12 $\pm$ 0.17	44.14 $\pm$ 0.19	32.62 $\pm$ 0.16	33.99 $\pm$ 0.17
	Linear Classifier	66.65 $\pm$ 0.16	69.89 $\pm$ 0.18	58.41 $\pm$ 0.17	61.96 $\pm$ 0.19	47.23 $\pm$ 0.17	50.08 $\pm$ 0.19	34.04 $\pm$ 0.16	35.75 $\pm$ 0.17
	Matching Networks [48]	62.16 $\pm$ 0.17	64.92 $\pm$ 0.19	56.21 $\pm$ 0.18	59.20 $\pm$ 0.20	46.18 $\pm$ 0.18	49.12 $\pm$ 0.20	34.66 $\pm$ 0.18	36.80 $\pm$ 0.19
	MAML [16]	63.25 $\pm$ 0.18	63.96 $\pm$ 0.19	53.28 $\pm$ 0.18	54.62 $\pm$ 0.19	42.58 $\pm$ 0.18	43.71 $\pm$ 0.19	31.01 $\pm$ 0.17	31.74 $\pm$ 0.17
	Vanilla ProtoNet [40]	68.27 $\pm$ 0.16	71.36 $\pm$ 0.18	62.43 $\pm$ 0.17	66.15 $\pm$ 0.19	51.41 $\pm$ 0.19	55.05 $\pm$ 0.21	38.33 $\pm$ 0.19	40.61 $\pm$ 0.21
	Baseline++ [11]	67.91 $\pm$ 0.16	71.24 $\pm$ 0.18	61.87 $\pm$ 0.17	65.58 $\pm$ 0.19	51.87 $\pm$ 0.18	55.00 $\pm$ 0.20	38.36 $\pm$ 0.19	40.19 $\pm$ 0.20
	RNNP [29]	68.38 $\pm$ 0.16	71.36 $\pm$ 0.18	62.43 $\pm$ 0.17	65.95 $\pm$ 0.19	51.62 $\pm$ 0.19	54.86 $\pm$ 0.21	38.45 $\pm$ 0.19	40.63 $\pm$ 0.21
Ours	Median	68.45 $\pm$ 0.16	71.28 $\pm$ 0.18	63.19 $\pm$ 0.17	66.65 $\pm$ 0.20	51.86 $\pm$ 0.19	55.09 $\pm$ 0.21	39.32 $\pm$ 0.19	41.94 $\pm$ 0.21
	Absolute	68.24 $\pm$ 0.16	71.27 $\pm$ 0.18	63.46 $\pm$ 0.17	66.87 $\pm$ 0.20	52.06 $\pm$ 0.20	55.26 $\pm$ 0.22	39.78 $\pm$ 0.20	42.54 $\pm$ 0.22
	Euclidean	68.32 $\pm$ 0.16	<b>71.48 <math>\pm</math> 0.18</b>	63.02 $\pm$ 0.17	66.69 $\pm$ 0.19	52.09 $\pm$ 0.19	55.62 $\pm$ 0.21	39.33 $\pm$ 0.20	41.75 $\pm$ 0.21
	Cosine	68.20 $\pm$ 0.16	70.59 $\pm$ 0.18	63.46 $\pm$ 0.17	66.62 $\pm$ 0.20	52.42 $\pm$ 0.20	55.78 $\pm$ 0.22	39.90 $\pm$ 0.20	42.56 $\pm$ 0.22
	TraNFS-2	68.29 $\pm$ 0.17	70.92 $\pm$ 0.19	64.74 $\pm$ 0.18	67.33 $\pm$ 0.21	56.14 $\pm$ 0.21	58.76 $\pm$ 0.23	42.24 $\pm$ 0.23	44.17 $\pm$ 0.25
	TraNFS-3	<b>68.53 <math>\pm</math> 0.17</b>	71.17 $\pm$ 0.19	<b>65.08 <math>\pm</math> 0.18</b>	<b>67.67 <math>\pm</math> 0.20</b>	<b>56.65 <math>\pm</math> 0.21</b>	<b>58.88 <math>\pm</math> 0.23</b>	<b>42.60 <math>\pm</math> 0.24</b>	<b>44.21 <math>\pm</math> 0.25</b>

## 6. Experiments

### 6.1. Experimental setup

**Datasets.** We experiment on two common FSL datasets: MiniImageNet [48] and TieredImageNet [38]. Both include  $84 \times 84$  pixel images. MiniImageNet contains 64, 16, and 20 classes for train, validation, and test, with 60K images in total. TieredImageNet consists of 351, 97, and 160 classes for train, validation, and test, with  $\sim 0.78$ M images in total.

**Label noise types.** We explore three forms of label noise:

- *Symmetric label swap* noise [46] draws mislabeled samples, uniformly at random, from the other  $N - 1$  classes of the episode, with the restriction that a noisy class does not tie or outnumber the original clean class.
- *Paired label swap* noise [19] is more challenging: we always draw mislabeled samples from the same class by assigning each class with a *noisy class* counterpart, simulating real-world tendencies to confuse certain classes with others during labeling. We randomly generate these assignments in each episode as a derangement, to prevent models from learning these pairings across episodes.
- *Outlier* noise is sampled from classes *outside* the  $N$ -way episode. We use 600 images from each of the 350 ImageNet classes unincorporated from MiniImageNet and TieredImageNet. We split these classes in half for meta-training and meta-test, to ensure that meta-test episode outliers represent previously unseen classes.

The amount of noise in a support set is specified as the percent of the total number of shots. We only consider settings where the clean class can reasonably be identified. Thus, for example, we only consider paired label swap noise under 50%, as at 50% noise and above, the clean class is ambiguous or a minority. We also exclude paired label swap settings that are identical to the corresponding symmetric label swap setting (e.g. 20% for 5-shot 5-way).

Table 2. **Few-shot with paired label swap noise.** 5-way 5-shot Acc.  $\pm$  95% CI on [MiniImageNet](#) [48], [TieredImageNet](#) [38].

Model \ Noise Proportion		40%	
Baselines	Oracle	62.60 $\pm$ 0.17	66.14 $\pm$ 0.20
	Nearest $k = 1$	37.97 $\pm$ 0.17	39.40 $\pm$ 0.18
	Nearest $k = 3$	37.84 $\pm$ 0.16	39.70 $\pm$ 0.18
	Nearest $k = 5$	40.39 $\pm$ 0.17	42.17 $\pm$ 0.18
	Linear Classifier	44.49 $\pm$ 0.17	46.70 $\pm$ 0.18
	Matching Networks [48]	43.53 $\pm$ 0.17	46.13 $\pm$ 0.19
	MAML [16]	40.67 $\pm$ 0.18	41.66 $\pm$ 0.18
	Vanilla ProtoNet [40]	47.77 $\pm$ 0.19	50.85 $\pm$ 0.21
	Baseline++ [11]	47.82 $\pm$ 0.18	50.69 $\pm$ 0.20
	RNNP [29]	47.88 $\pm$ 0.19	50.91 $\pm$ 0.20
Ours	Median	48.81 $\pm$ 0.19	51.91 $\pm$ 0.21
	Absolute	49.38 $\pm$ 0.20	52.40 $\pm$ 0.22
	Euclidean	48.67 $\pm$ 0.19	51.90 $\pm$ 0.21
	Cosine	49.40 $\pm$ 0.19	52.72 $\pm$ 0.22
	TraNFS-2	50.63 $\pm$ 0.22	54.82 $\pm$ 0.24
	TraNFS-3	<b>53.96 <math>\pm</math> 0.23</b>	<b>55.12 <math>\pm</math> 0.24</b>

**Model.** Our models are implemented in PyTorch [35], using learn2learn [3] as a starting point. We set  $T = 25$  for similarity weighted prototypes with squared euclidean and absolute distances and  $T = 0.2$  for cosine similarity. For TraNFS, we instantiate the transformer  $\mathcal{T}$  using 2 or 3-layers with eight heads, learnable positional embeddings  $\text{POS}^{(c)}$ , and random constant class tokens  $\text{CLS}^{(c)}$ . We apply an orthogonally initialized pair of down-projection and up-projection weight matrices before and after the transformer, reducing the transformer’s dimensionality to 128. We found these orthogonal projections stabilize training [36], while also greatly reducing the number of transformer parameters. Finally, we set the hyperparameters of Eq. (17) as  $\lambda_b = 0.5$  and  $\lambda_c = 5$ . See Supp. for hyperparameter sweeps. Note that while the transformer must be meta-trained and used to generate robust prototypes during meta-test, it is not used during inference on individual query samples. Hence, the number of parameters and computation cost during inference is similar to methods like ProtoNet [40].

**Training and testing.** To isolate the effect of the method

Table 3. **Few-shot with outlier noise.** 5-way 5-shot Acc.  $\pm$  95% CI on [MiniImageNet](#) [48], [TieredImageNet](#) [38]. Our TraNSF is comparable to existing methods at 0% or low noise, with a growing gap in its favor as noise levels increase. Best viewed in color.

Model \ Noise Proportion		0%		20%		40%		60%	
Baselines	Oracle	68.18 $\pm$ 0.16	71.42 $\pm$ 0.18	66.08 $\pm$ 0.17	69.19 $\pm$ 0.19	62.60 $\pm$ 0.17	66.14 $\pm$ 0.20	56.89 $\pm$ 0.18	60.39 $\pm$ 0.21
	Nearest $k = 1$	55.87 $\pm$ 0.18	58.89 $\pm$ 0.20	50.90 $\pm$ 0.18	54.57 $\pm$ 0.20	45.28 $\pm$ 0.18	49.45 $\pm$ 0.20	38.75 $\pm$ 0.18	43.20 $\pm$ 0.19
	Nearest $k = 3$	55.28 $\pm$ 0.18	58.38 $\pm$ 0.20	50.53 $\pm$ 0.17	53.98 $\pm$ 0.20	44.40 $\pm$ 0.17	48.06 $\pm$ 0.19	37.03 $\pm$ 0.16	40.11 $\pm$ 0.18
	Nearest $k = 5$	56.34 $\pm$ 0.17	59.25 $\pm$ 0.19	52.32 $\pm$ 0.17	55.30 $\pm$ 0.19	46.49 $\pm$ 0.17	49.34 $\pm$ 0.19	38.44 $\pm$ 0.16	40.56 $\pm$ 0.17
	Linear Classifier	66.70 $\pm$ 0.16	69.60 $\pm$ 0.18	61.13 $\pm$ 0.17	64.58 $\pm$ 0.19	53.86 $\pm$ 0.18	57.57 $\pm$ 0.20	44.05 $\pm$ 0.18	47.90 $\pm$ 0.20
	Matching Networks [48]	62.05 $\pm$ 0.17	64.99 $\pm$ 0.19	57.69 $\pm$ 0.18	60.74 $\pm$ 0.20	51.32 $\pm$ 0.19	54.28 $\pm$ 0.21	42.39 $\pm$ 0.19	44.93 $\pm$ 0.20
	MAML [16]	63.21 $\pm$ 0.18	63.90 $\pm$ 0.19	57.35 $\pm$ 0.19	58.14 $\pm$ 0.19	50.00 $\pm$ 0.19	51.11 $\pm$ 0.20	40.90 $\pm$ 0.17	42.01 $\pm$ 0.20
	Vanilla ProtoNet [40]	68.18 $\pm$ 0.16	<b>71.42 <math>\pm</math> 0.18</b>	63.92 $\pm$ 0.17	67.58 $\pm$ 0.19	57.07 $\pm$ 0.18	60.97 $\pm$ 0.20	46.99 $\pm$ 0.20	50.29 $\pm$ 0.21
	Baseline++ [11]	67.85 $\pm$ 0.16	71.29 $\pm$ 0.18	63.49 $\pm$ 0.17	67.07 $\pm$ 0.19	56.84 $\pm$ 0.18	60.64 $\pm$ 0.20	46.96 $\pm$ 0.19	50.07 $\pm$ 0.21
	RNNP [29]	68.17 $\pm$ 0.16	71.28 $\pm$ 0.18	63.80 $\pm$ 0.17	67.29 $\pm$ 0.19	56.97 $\pm$ 0.18	60.83 $\pm$ 0.20	46.92 $\pm$ 0.20	50.09 $\pm$ 0.21
Ours	Median	68.37 $\pm$ 0.16	71.28 $\pm$ 0.18	64.46 $\pm$ 0.17	67.79 $\pm$ 0.19	57.85 $\pm$ 0.18	61.63 $\pm$ 0.21	47.19 $\pm$ 0.20	50.63 $\pm$ 0.21
	Absolute	68.13 $\pm$ 0.16	71.17 $\pm$ 0.18	64.69 $\pm$ 0.17	<b>68.00 <math>\pm</math> 0.19</b>	58.30 $\pm$ 0.18	61.98 $\pm$ 0.21	47.39 $\pm$ 0.20	50.59 $\pm$ 0.22
	Euclidean	<b>68.51 <math>\pm</math> 0.16</b>	71.28 $\pm$ 0.18	64.57 $\pm$ 0.17	67.89 $\pm$ 0.19	58.01 $\pm$ 0.18	61.61 $\pm$ 0.20	47.25 $\pm$ 0.20	50.49 $\pm$ 0.21
	Cosine	68.20 $\pm$ 0.16	70.79 $\pm$ 0.18	64.78 $\pm$ 0.17	67.94 $\pm$ 0.19	58.36 $\pm$ 0.18	62.37 $\pm$ 0.21	47.34 $\pm$ 0.20	51.12 $\pm$ 0.22
	TraNFS-2	67.76 $\pm$ 0.17	70.83 $\pm$ 0.19	64.47 $\pm$ 0.19	67.52 $\pm$ 0.21	58.29 $\pm$ 0.20	61.76 $\pm$ 0.22	47.37 $\pm$ 0.23	51.40 $\pm$ 0.23
	TraNFS-3	68.11 $\pm$ 0.17	71.13 $\pm$ 0.19	<b>64.96 <math>\pm</math> 0.18</b>	67.93 $\pm$ 0.20	<b>59.03 <math>\pm</math> 0.20</b>	<b>62.39 <math>\pm</math> 0.22</b>	<b>47.69 <math>\pm</math> 0.22</b>	<b>51.82 <math>\pm</math> 0.23</b>

from the learned features, we use the same, frozen, 4-layer convolutional backbone [48], trained with the ProtoNet objective, for all models except MAML [16]. We chose this simple backbone to emphasize the method, not the feature extractor. The backbone is trained with AdamW [27], with weight decay of 0.01, initial learning rate  $1 \times 10^{-3}$ , and learning rate decay of  $\times 0.7$  every 10K episodes for 100K episodes for MiniImageNet, and every 25K episodes for 250K episodes for TieredImageNet. Meta-validation is used for accuracy model selection. Our TraNFS is similarly optimized, with initial learning rate  $5 \times 10^{-4}$ , decayed after every 25K episodes, for 200K episodes. We use random horizontal flips, resized crops, and color jitters as data augmentations for all models. Finally, each meta-train and meta-test episode has 15 queries. We report mean accuracy and 95% confidence interval for 10K meta-test episodes. All experiments are run on a single Nvidia V100 GPU.

## 6.2. Noisy few-shot results

We compare all our proposed methods for noisy few-shot learning—Median, Absolute, Euclidean, Cosine, and TraNFS—with several baselines (see Supp. for details on these baselines). We report results for 5-way and 5-shot<sup>1</sup> on MiniImageNet and TieredImageNet using symmetric noise (Table 1), paired noise (Table 2), and outlier noise (Table 3). We further report results for an *Oracle*: a ProtoNet [40] that knows which samples are mislabeled and ignores them by removing them from each support set, thereby representing perfect noise rejection (the blue line in Fig. 2).

Unsurprisingly, noisy labels negatively affect all methods. Our proposed median and similarity weighting alternatives to ProtoNet’s mean suffer less than the baselines, on all three noise types. This is expected, as their aggregation methods are less sensitive to outliers. Furthermore, our transformer-based TraNFS is clearly superior to its baselines. For example, consider the challenging 5-way

<sup>1</sup>See Supp. for 3-shot and 10-shot experiments.

Table 4. **Various amounts of injected meta-training artificial noise.** TraNFS-3 5-way 5-shot Acc.  $\pm$  95% CI on MiniImageNet.

0%		20%		40%		60%	
✓				<b>69.10 <math>\pm</math> 0.16</b>	63.56 $\pm$ 0.18	52.85 $\pm$ 0.19	39.19 $\pm$ 0.21
	✓			68.67 $\pm$ 0.17	64.85 $\pm$ 0.18	55.76 $\pm$ 0.21	41.73 $\pm$ 0.23
		✓		67.37 $\pm$ 0.17	63.97 $\pm$ 0.19	55.65 $\pm$ 0.21	41.63 $\pm$ 0.24
			✓	50.40 $\pm$ 0.19	48.26 $\pm$ 0.19	43.11 $\pm$ 0.21	35.44 $\pm$ 0.23
	✓	✓		68.53 $\pm$ 0.17	<b>65.08 <math>\pm</math> 0.18</b>	56.65 $\pm$ 0.21	42.60 $\pm$ 0.24
✓	✓	✓		68.90 $\pm$ 0.17	<b>65.08 <math>\pm</math> 0.18</b>	<b>56.73 <math>\pm</math> 0.21</b>	<b>42.69 <math>\pm</math> 0.24</b>
	✓	✓	✓	66.92 $\pm$ 0.17	63.52 $\pm$ 0.19	54.98 $\pm$ 0.22	42.01 $\pm$ 0.24
✓	✓	✓	✓	67.64 $\pm$ 0.17	63.83 $\pm$ 0.18	54.81 $\pm$ 0.21	41.33 $\pm$ 0.24

Table 5. **Meta-training mean/median prototype models with noise.** 5-way 5-shot Acc.  $\pm$  95% CI on MiniImageNet.

Baseline \ Noise Proportion	0%	20%	40%	60%
Mean + Sym (0%, 20%, 40%)	67.89 $\pm$ 0.16	62.44 $\pm$ 0.18	51.66 $\pm$ 0.19	38.53 $\pm$ 0.20
Mean + Pair (0%, 20%, 40%)	-	-	48.05 $\pm$ 0.19	-
Mean + Out (0%, 20%, 40%)	66.88 $\pm$ 0.16	62.69 $\pm$ 0.17	56.00 $\pm$ 0.18	45.90 $\pm$ 0.19
Median + Sym (0%, 20%, 40%)	67.11 $\pm$ 0.16	62.39 $\pm$ 0.17	51.70 $\pm$ 0.19	39.57 $\pm$ 0.20
Median + Pair (0%, 20%, 40%)	-	-	48.64 $\pm$ 0.19	-
Median + Out (0%, 20%, 40%)	67.17 $\pm$ 0.16	63.46 $\pm$ 0.17	57.01 $\pm$ 0.18	46.44 $\pm$ 0.20

5-shot setting on MiniImageNet with 40% paired label swap noise. Our TraNFS provides a 6.19% absolute improvement in accuracy over ProtoNet, representing *a significant relative drop of 41.7% in error*, compared with the Oracle. As explained in Sec. 6.3, this gain is due to the transformer’s self-attention learning to compare support set examples and suppress samples suspected of being mislabeled.

Comparing noise types, we find that, as expected, FSL methods are more vulnerable to paired label swap than symmetric noise. Outlier noise has the least impact on model performance, with three outlier samples in a 5-way 5-shot test reducing accuracy similarly to two label swap samples. We reason that this is due to the direction in which these noisy samples push the model’s decision boundary: label swapped samples pull the decision boundary closer to the features of other classes in the  $N$ -way classification task; for paired label swaps, this effort is coordinated across noisy samples, amplifying the effect. In contrast, outlier samples have lower probability of being arranged in regions of feature space that interfere with the  $N$ -way classification.

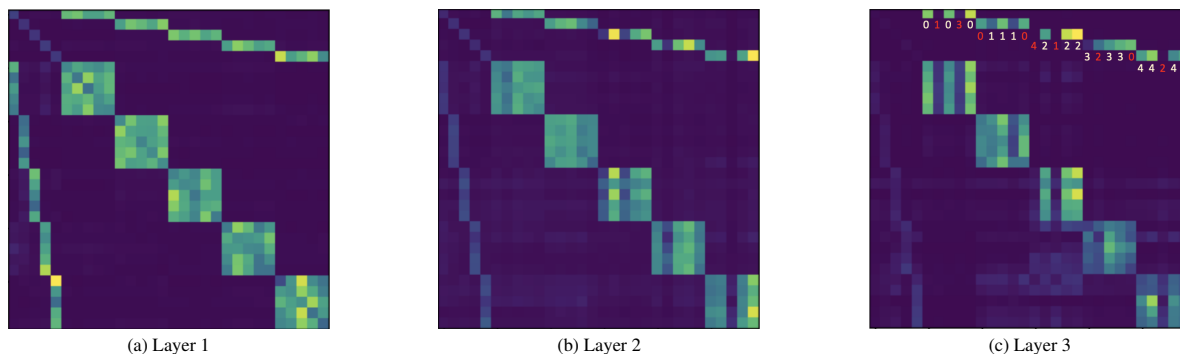


Figure 4. **Attention maps from selected TraNFS-3 self-attention heads.** 5-way 5-shot MiniImageNet, 40% symmetric noise. The first five rows of each map correspond to positions of CLS tokens, which produce the output prototypes. True class label of each support sample overlaid on the final layer’s attention maps. Evidently, later layers assign lower attention to noisy samples, effectively filtering them.

### 6.3. Visualizing transformer attention to noise

To understand how our proposed transformer model suppresses noisy samples, we visualize the self-attention from selected attention heads at each of its layers. The attention maps for a 5-way 5-shot test episode of MiniImageNet with 40% symmetric label swap noise are shown in Fig. 4. From Fig. 4a, attention at the CLS token positions suggests that the first layer of the transformer mainly uses positional encodings to focus on per-class examples, resulting in a representation reminiscent of ProtoNet’s mean.

Subsequent layers are visualized in Fig. 4b-4c. Evidently, the transformer is able to refine its class representations by decreasing attention to noisy samples, suppressing their influence on the aggregated representations. For example, our self-attention mechanism correctly learned to ignore the 2<sup>nd</sup> and 4<sup>th</sup> samples of the first class, which are indeed mislabeled. While this filtering ability is not perfect, we emphasize that learning class concepts from so few samples, without any prior concept of the class, is a challenging task (see Fig 1); ImageNet contains enough intra-class variation and label ambiguity that identifying mislabeled samples can be challenging even for humans, who have the advantage of conceptual priors of the ImageNet classes.

### 6.4. Ablations: Meta-training noise proportion

To test the influence of adding training noise,<sup>2</sup> we meta-train a 3-layer TraNFS model on 5-shot, 5-way MiniImageNet with various amounts of symmetric label swap noise, synthetically added to meta-training. Table 4 reports these results, clearly showing a few patterns.

First, training with a single noise percentage boosts performance on that noise level during meta-test. Those models, however, do not generalize well to other noise levels. Instead, training on varying noise levels seems to offer the best results across a range of meta-test support set noise levels. This is important, as the stochastic nature of label noise means we expect real world support sets to have varying levels of noise, and it is desirable to handle multiple noise

<sup>2</sup>See Supp. for more ablation studies on other design choices.

levels with a single model. In particular, training on support sets with  $\{0, 20, 40\}\%$  appears to achieve the best overall performance. Finally, we observe that training on extremely noisy support sets (*e.g.*, 60%) appears counter-productive. We believe that this is due to a mixture of having a more challenging task to learn while also diluting learnable information of the clean class of each support set.

Synthetically adding noise during meta-training proved essential for TraNFS. A similar strategy could conceivably also be applied other methods. As Table 5 shows, however, this approach was unhelpful. We believe that the absence of learnable mechanisms for rejecting noise encourages these baselines to learn strong feature extractors on the highest quality (*i.e.*, cleanest) data available. Thus, we do not add artificial noise to our baselines during meta-training.

## 7. Conclusion

We focus on a key vulnerability of modern FSL methods: noisy, mislabeled support sets samples. We propose several technical novelties to mitigate this vulnerability: replacing the mean aggregator used by ProtoNets with a median or similarity weighted aggregation. We then present a novel, transformer-based model designed to learn a dynamic noise rejection mechanism, leveraging the transformer’s attention mechanism. Experiments on MiniImageNet and TieredImageNet under varying types and levels of noise clearly show the effectiveness of our techniques.

**Limitations.** As with other FSL methods, we assume a cleanly labeled meta-train dataset with noisy labels mostly affecting meta-testing. We believe this is reasonable: collecting meta-train data and meta-training are performed offline, before model deployment, with reasonable control over both. Sometimes, however, meta-training datasets can also be noisy. While we introduce noise during meta-training, our methods assume that the query set is correctly labeled. Queries with noisy labels could cause misleading gradients. In such cases, ideas from the noisy label literature [19,20] could offer promising direction for future work.



## References

- [1] Dana Angluin and Philip Laird. Learning from noisy examples. *Machine Learning*, 1988.
- [2] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. *arXiv preprint arXiv:1810.09502*, 2018.
- [3] Sébastien MR Arnold, Praateek Mahajan, Debajyoti Datta, Ian Bunner, and Konstantinos Saitas Zarkias. learn2learn: A library for meta-learning research. *arXiv preprint arXiv:2008.12284*, 2020.
- [4] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. *Int. Conf. Mach. Learning.*, 2017.
- [5] Jonathan T. Barron. A general and adaptive robust loss function. *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [6] Peyman Bateni, Raghav Goyal, Vaden Masrani, Frank Wood, and Leonid Sigal. Improved few-shot visual classification. *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [7] Nihar Bendre, Hugo Terashima Marín, and Peyman Najafirad. Learning from few samples: A survey. *arXiv preprint arXiv:2007.15484*, 2020.
- [8] Luca Bertinetto, Joao F Henriques, Philip Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *Int. Conf. Learn. Represent.*, 2018.
- [9] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xi-aohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020.
- [10] Koby Bibas, Meir Feder, and Tal Hassner. Single layer predictive normalized maximum likelihood for out-of-distribution detection. *Adv. Neural Inform. Process. Syst.*, 2021.
- [11] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019.
- [12] James Davidson, Benjamin Liebald, Junning Liu, Palash Nandy, Taylor Van Vleet, Ullas Gargi, Sujoy Gupta, Yu He, Mike Lambert, Blake Livingston, et al. The youtube video recommendation system. *ACM Conference on Recommender Systems*, 2010.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [14] Terrance DeVries and Graham W Taylor. Learning confidence for out-of-distribution detection in neural networks. *arXiv preprint arXiv:1802.04865*, 2018.
- [15] Carl Doersch, Ankush Gupta, and Andrew Zisserman. Crosstransformers: spatially-aware few-shot transfer. *arXiv preprint arXiv:2007.11498*, 2020.
- [16] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *Int. Conf. Mach. Learning.*, 2017.
- [17] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981.
- [18] Micah Goldblum, Liam Fowl, and Tom Goldstein. Adversarially robust few-shot learning: A meta-learning approach. *Adv. Neural Inform. Process. Syst.*, 2020.
- [19] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *arXiv preprint arXiv:1804.06872*, 2018.
- [20] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. *Int. Conf. Mach. Learning.*, 2018.
- [21] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 2020.
- [22] Krishnateja Killamsetty, Changbin Li, Chen Zhao, Rishabh Iyer, and Feng Chen. A reweighted meta learning framework for robust few shot learning. *arXiv preprint arXiv:2011.06782*, 2020.
- [23] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [24] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [25] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. *Int. Conf. Comput. Vis.*, 2017.
- [26] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2015.
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *Int. Conf. Learn. Represent.*, 2018.
- [28] Jiang Lu, Sheng Jin, Jian Liang, and Changshui Zhang. Robust few-shot learning for user-provided data. *IEEE trans. on neural networks and learning systems*, 2020.
- [29] Pratik Mazumder, Pravendra Singh, and Vinay P Namboodiri. Rnnp: A robust few-shot learning approach. *Proc. Winter Conf. on Applications of Comput. Vis.*, 2021.
- [30] Tsendsuren Munkhdalai and Hong Yu. Meta networks. *Int. Conf. Mach. Learning.*, 2017.
- [31] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with noisy labels. *Adv. Neural Inform. Process. Syst.*, 2013.
- [32] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [33] Curtis G Northcutt, Anish Athalye, and Jonas Mueller. Pervasive label errors in test sets destabilize machine learning benchmarks. *arXiv preprint arXiv:2103.14749*, 2021.
- [34] Boris N Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. *Adv. Neural Inform. Process. Syst.*, 2018.

- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inform. Process. Syst.*, 2019.
- [36] Samrudhdi Bharatkumar Rangrej, Kevin J Liang, Xi Yin, Guan Pang, Theofanis Karaletsos, Lior Wolf, and Tal Hassner. Revisiting linear decision boundaries for few-shot learning with transformer hypernetworks. 2021.
- [37] Jie Ren, Peter J Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark A DePristo, Joshua V Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. *arXiv preprint arXiv:1906.02845*, 2019.
- [38] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *Int. Conf. Learn. Represent.*, 2018.
- [39] Victor Garcia Satorras and Joan Bruna Estrach. Few-shot learning with graph neural networks. *Int. Conf. Learn. Represent.*, 2018.
- [40] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Adv. Neural Inform. Process. Syst.*, 2017.
- [41] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. *Conf. Comput. Vis. Pattern Recog.*, 2018.
- [42] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? *Eur. Conf. Comput. Vis.*, 2020.
- [43] Raciél Yera Toledo, Yailé Caballero Mota, and Luis Martínez. Correcting noisy ratings in collaborative recommender systems. *Knowledge-Based Systems*, 2015.
- [44] Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. *Int. Conf. Learn. Represent.*, 2019.
- [45] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. From imagenet to image classification: Contextualizing progress on benchmarks. *Int. Conf. Mach. Learning.*, 2020.
- [46] Brendan Van Rooyen, Aditya Krishna Menon, and Robert C Williamson. Learning with symmetric label noise: The importance of being unhinged. *arXiv preprint arXiv:1505.07634*, 2015.
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Adv. Neural Inform. Process. Syst.*, 2017.
- [48] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Adv. Neural Inform. Process. Syst.*, 2016.
- [49] Apoorv Vyas, Nataraj Jammalamadaka, Xia Zhu, Dipankar Das, Bharat Kaul, and Theodore L. Willke. Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. *Eur. Conf. Comput. Vis.*, 2018.
- [50] Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E Gonzalez. Tafe-net: Task-aware feature embeddings for low shot learning. *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [51] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. *Conf. Comput. Vis. Pattern Recog.*, 2018.
- [52] Zhen Wang, Guosheng Hu, and Qinghua Hu. Training noise-robust deep neural networks via meta-learning. *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [53] Yuewei Yang, Kevin J Liang, and Lawrence Carin. Object detection as a positive-unlabeled problem. *Brit. Mach. Vis. Conf.*, 2020.
- [54] Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual t: Reducing estimation error for transition matrix in label-noise learning. *arXiv preprint arXiv:2006.07805*, 2020.
- [55] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [56] Mang Ye and Pong C Yuen. Purifynet: A robust person re-identification model with noisy labels. *IEEE Transactions on Information Forensics and Security*, 2020.
- [57] Li Yin, Juan-Manuel Perez-Rua, and Kevin J Liang. Sylph: A hypernet-framework for incremental few-shot object detection. *Conf. Comput. Vis. Pattern Recog.*, 2022.
- [58] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? *Int. Conf. Mach. Learning.*, 2019.
- [59] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [60] Yivan Zhang, Gang Niu, and Masashi Sugiyama. Learning noise transition matrix from only noisy labels via total variation regularization. *arXiv preprint arXiv:2102.02414*, 2021.
- [61] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. Meta label correction for noisy label learning. *Proc. of the AAAI Conf. on Artificial Intelligence*, 2021.