

# ADVERSARIALLY ROBUST FEW-SHOT LEARNING: A META-LEARNING APPROACH

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Previous work on adversarially robust neural networks for image classification requires large training sets and computationally expensive training procedures. On the other hand, few-shot learning methods are highly vulnerable to adversarial examples. The goal of our work is to produce networks which both perform well at few-shot classification tasks and are simultaneously robust to adversarial examples. We develop an algorithm for producing adversarially robust meta-learners, and we thoroughly investigate factors which contribute to adversarial vulnerability. Moreover, our method achieves far superior robust performance on few-shot image classification tasks, such as Mini-ImageNet and CIFAR-FS, than robust transfer learning.

## 1 INTRODUCTION

For safety-critical applications like facial recognition, traffic sign detection, and copyright control, adversarial attacks pose an actionable threat (Zhao et al., 2018; Eykholt et al., 2017; Saadatpanah et al., 2019). Conventional adversarial training and pre-processing defenses aim to produce networks that resist attack (Madry et al., 2017; Zhang et al., 2019; Samangouei et al., 2018), but such defenses rely heavily on the availability of large training data sets. In applications that require *few-shot learning*, such as face recognition from few images, recognition of a video source from a single clip, or recognition of a new object from few example photos, the conventional robust training pipeline breaks down.

When data is scarce or new classes arise frequently, neural networks must adapt quickly (Duan et al., 2017; Kaiser et al., 2017; Pfister et al., 2014; Vartak et al., 2017). In these situations, *meta-learning* methods conduct few-shot learning by creating networks that learn quickly from little data and with computationally cheap fine-tuning. While state-of-the-art meta-learning methods perform well on benchmark few-shot classification tasks, these naturally trained neural networks are highly vulnerable to adversarial examples. In fact, even adversarially trained feature extractors fail to resist attacks in the few-shot setting (see Section 4.1).

We propose a new approach, called *adversarial querying*, in which the network is exposed to adversarial attacks during the query step of meta-learning. This algorithm-agnostic method produces a feature extractor that is robust, even without adversarial training during fine-tuning. In the few-shot setting, we show that adversarial querying outperforms other robustness techniques by a wide margin in terms of both clean accuracy and adversarial robustness (see Table 1). We solve the following minimax problem:

$$\min_{\theta} \mathbb{E}_{S, (\mathbf{x}, y)} \left[ \max_{\|\delta\|_p < \epsilon} \mathcal{L}(F_{A(\theta, S)}, \mathbf{x} + \delta, y) \right], \quad (1)$$

where  $S$  and  $(\mathbf{x}, y)$  are data sampled from the training distribution,  $A$  is a fine-tuning algorithm for the model parameters,  $\theta$ , and  $\epsilon$  is a  $p$ -norm bound for the attacker. In Section 4, we further motivate adversarial querying and exhibit a wide range of experiments. To motivate the necessity for adversarial querying, we test methods, such as adversarial fine-tuning and pre-processing defenses, which if successful, would eliminate the need for expensive adversarial training routines. We find that these methods are far less effective than adversarial querying.

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$
AT transfer learning (R2-D2 backbone)	39.13%	25.33%
Naturally Trained R2-D2	73.01%	0.00%
ADML	47.75%	18.49%
AQ R2-D2 (ours)	57.87%	<b>31.52%</b>

Table 1: Adversarially trained transfer learning, R2-D2 (Bertinetto et al., 2018), ADML (Yin et al., 2018), and our adversarially queried (AQ) R2-D2 model on 5-shot Mini-ImageNet. Natural accuracy is denoted  $\mathcal{A}_{nat}$ , and robust accuracy,  $\mathcal{A}_{adv}$ , is computed with a 20-step PGD attack as in Madry et al. (2017) with  $\epsilon = 8/255$ . A description of our training regime can be found in Appendix A.5.

## 2 RELATED WORK

### 2.1 LEARNING WITH LESS DATA

Before the emergence of meta-learning, a number of approaches existed to cope with few-shot problems. One simple approach is *transfer learning*, in which pre-trained feature extractors are trained on large data sets and fine-tuned on new tasks (Bengio, 2012). Metric learning methods avoid overfitting to the small number of training examples in new classes by instead performing classification using nearest-neighbors in feature space with a feature extractor that is trained on a large corpus of data and not re-trained when classes are added (Snell et al., 2017; Gidaris & Komodakis, 2018; Mensink et al., 2012). Metric learning methods are computationally efficient when adding new classes at inference, since the feature extractor is not re-trained.

Meta-learning algorithms create a “base” model that quickly adapts to new tasks by fine-tuning. This model is created using a set of training tasks  $\{\mathcal{T}_i\}$  that can be sampled from a task distribution. Each task comes with *support* data,  $\mathcal{T}_i^s$ , and *query* data,  $\mathcal{T}_i^q$ . Support data is used for fine-tuning, and query data is used to measure the performance of the resulting network. In practice, each task is taken to be a classification problem involving only a small subset of classes in a large many-class data set. The number of examples per class in the support set is called the *shot*, so that fine-tuning on five support examples per class is 5-shot learning.

An iteration of training begins by sampling tasks  $\{\mathcal{T}_i\}$  from the task distribution. In the “inner loop” of training, the base model is fine-tuned on the support data from the sampled tasks. In the “outer loop” of training, the fine-tuned network is used to make predictions on the query data, and the base model parameters are updated to improve the accuracy of the resulting fine-tuned model. The outer loop requires backpropagation through the fine-tuning steps. A formal treatment of the prototypical meta-learning routine can be found in Algorithm 1.

---

#### Algorithm 1: The meta-learning framework

---

**Require:** Base model,  $F_\theta$ , fine-tuning algorithm,  $A$ , learning rate,  $\gamma$ , and distribution over tasks,  $p(\mathcal{T})$ .

Initialize  $\theta$ , the weights of  $F$ ;

**while** not done **do**

Sample batch of tasks,  $\{\mathcal{T}_i\}_{i=1}^n$ , where  $\mathcal{T}_i \sim p(\mathcal{T})$  and  $\mathcal{T}_i = (\mathcal{T}_i^s, \mathcal{T}_i^q)$ .

**for**  $i = 1, \dots, n$  **do**

Fine-tune model on  $\mathcal{T}_i$  (inner loop). New network parameters are written  $\theta_i = A(\theta, \mathcal{T}_i^s)$ .

Compute gradient  $g_i = \nabla_{\theta} \mathcal{L}(F_{\theta_i}, \mathcal{T}_i^q)$ .

Update base model parameters (outer loop):  $\theta \leftarrow \theta - \frac{\gamma}{n} \sum_i g_i$

---

Note that the fine-tuned parameters,  $\theta_i = A(\theta, \mathcal{T}_i^s)$ , in Algorithm 1, are a function of the base model’s parameters so that the gradient computation in the outer loop may backpropagate through  $A$ . For validation after training, the base model is fine-tuned on the support set of hold-out tasks, and accuracy on the query set is reported. In this work, we report performance on Omniglot, Mini-ImageNet, and CIFAR-FS (Lake et al., 2015; Vinyals et al., 2016; Bertinetto et al., 2018).

We focus on four meta-learning algorithms: MAML, R2-D2, MetaOptNet, and ProtoNet (Finn et al., 2017; Bertinetto et al., 2018; Lee et al., 2019; Snell et al., 2017). During fine-tuning, MAML uses SGD to update all parameters, minimizing cross-entropy loss. Since unrolling SGD steps into a deep computation graph is expensive, first-order variants have been developed to avoid computing second-order derivatives. We use the original MAML formulation. R2-D2 and MetaOptNet, on the other hand, only update the final linear layer during fine-tuning, leaving the “backbone network” that extracts these features frozen at test time. R2-D2 replaces SGD with a closed-form differentiable solver for regularized ridge regression, while MetaOptNet achieves its best performance when replacing SGD with a solver for SVM. Because the objective of these linear problems is convex, differentiable convex optimizers can be efficiently deployed to find optima, and differentiate these optima with respect to the backbone parameters at train time. ProtoNet takes an approach inspired by metric learning. It constructs class prototypes as centroids in feature space for each task. These centroids are then used to classify the query set in the outer loop of training. Because each class prototype is a simple geometric average of feature representations, it is easy to differentiate through the fine-tuning step.

## 2.2 ADVERSARIAL ATTACKS AND DEFENSES

Following standard practices, we assess the robustness of models by attacking them with  $\ell_\infty$ -bounded perturbations. We craft adversarial perturbations using the projected gradient descent attack (PGD) since it has proven to be one of the most effective algorithms both for attacking as well as for adversarial training (Madry et al., 2017). A detailed description of the PGD attack can be found in Algorithm 2. We consider perturbations with  $\ell_\infty$  bound  $8/255$  and a step size of  $2/255$  as described by Madry et al. (2017).

*Adversarial training* is the industry standard for creating robust models that maintain good clean-label performance (Madry et al., 2017). This method involves replacing clean examples with adversarial examples during the training routine. A simple way to harden models to attack is adversarial training, which solves the minimax problem

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[ \max_{\|\delta\|_p < \epsilon} \mathcal{L}_{\theta}(\mathbf{x} + \delta, y) \right], \quad (2)$$

where  $\mathcal{L}_{\theta}(\mathbf{x} + \delta, y)$  is the loss function of a network with parameters  $\theta$ ,  $\mathbf{x}$  is an input image with label  $y$ , and  $\delta$  is an adversarial perturbation. Adversarial training finds network parameters which maintain low loss (and correct class labels) even when adversarial perturbations are added.

---

### Algorithm 2: PGD Attack

---

**Require:** network,  $F_{\theta}$ , input data,  $(\mathbf{x}, y)$ , number of steps,  $n$ , step size,  $\gamma$ , and attack bound,  $\epsilon$ .

Initialize  $\delta \in \mathcal{B}_{\epsilon}(\mathbf{x})$  randomly;

**for**  $i = 1, \dots, n$  **do**

Compute  $g = \text{sign}(\nabla_{\delta} \mathcal{L}_{\theta}(\mathbf{x} + \delta, y))$ .  
 Update  $\delta = \delta + \gamma g$ .  
 If  $\|\delta\|_p > \epsilon$ , then project  $\delta$  onto the surface of  $\mathcal{B}_{\epsilon}(\mathbf{x})$ .

**return** perturbed image  $\mathbf{x} + \delta$

---

## 2.3 ROBUST LEARNING WITH LESS DATA

Several authors have tried to learn robust models in the data scarce regime. The authors of Shafahi et al. (2019) study robustness properties of transfer learning. They find that retraining earlier layers of the network during fine-tuning impairs the robustness of the network, while only retraining later layers can largely preserve robustness. ADML is the first attempt at achieving robustness through meta-learning. ADML is a MAML variant, specifically designed for robustness, which employs adversarial training (Yin et al., 2018). However, this method for robustness is only compatible with MAML, an outdated meta-learning algorithm. Moreover, ADML is computationally expensive, and the authors only test their method against a weak attacker. We implement ADML and test it against a strong attacker. We show that our method simultaneously achieves higher robustness and higher natural accuracy.

### 3 NATURALLY TRAINED META-LEARNING METHODS ARE NOT ROBUST

In this section, we benchmark the robustness of existing meta-learning methods. Similarly to classically trained classifiers, we expect that few-shot learners are highly vulnerable to attack when adversarial defenses are not employed. We test prominent meta-learning algorithms against a 20-step PGD attack as in Madry et al. (2017). Table 2 contains natural and robust accuracy on the Mini-ImageNet and CIFAR-FS 5-shot tasks (Vinyals et al., 2016; Bertinetto et al., 2018). Experiments in the 1-shot setting can be found in Appendix A.1.

Model	$\mathcal{A}_{nat}$ MI	$\mathcal{A}_{adv}$ MI	$\mathcal{A}_{nat}$ FS	$\mathcal{A}_{adv}$ FS
ProtoNet	70.23%	0.00%	79.66%	0.00%
R2-D2	73.02%	0.00%	82.81%	0.00%
MetaOptNet	78.12%	0.00%	84.11%	0.00%

Table 2: 5-shot MiniImageNet (MI) and CIFAR-FS (FS) results comparing naturally trained meta-learners.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack.

We find that these algorithms are completely unable to resist the attack. Interestingly, MetaOptNet uses SVM for fine-tuning, which is endowed with a wide margins property. The failure of even SVM to express robustness during testing suggests that using robust fine-tuning methods (at test time) on naturally trained meta-learners is insufficient for robust performance. To further examine this, we consider MAML, which updates the entire network during fine-tuning. We use a naturally trained MAML model and perform adversarial training during fine-tuning (see Table 3). Adversarial training is performed with 7-PGD as in (Madry et al., 2017). If adversarial fine-tuning yielded robust classification, then we could avoid expensive adversarial training variants during meta-learning.

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$	$\mathcal{A}_{nat(AT)}$	$\mathcal{A}_{adv(AT)}$
1-shot Mini-ImageNet	45.04%	0.03%	33.18%	0.20%
5-shot Mini-ImageNet	60.25%	0.03%	32.45%	1.55%
1-shot Omniglot	91.50%	68.46%	91.60%	74.66%
5-shot Omniglot	97.12%	82.28%	97.71%	87.94%

Table 3: MAML models on Mini-ImageNet and Omniglot.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack.  $\mathcal{A}_{nat(AT)}$  and  $\mathcal{A}_{adv(AT)}$  are natural and robust test accuracy with 7-PGD fine-tuning.

While clean trained MAML models with adversarial fine-tuning are slightly more robust than their naturally fine-tuned counterparts, they achieve almost no robustness on Mini-ImageNet. Omniglot is an easier data set, and the performance of adversarially fine-tuned MAML on the 5-shot version is below a reasonable tolerance for robustness. We conclude from these experiments that naturally trained meta-learners are vulnerable to adversarial examples, and robustness techniques specifically for few-shot learning are required.

### 4 ADVERSARIAL QUERYING: A META-LEARNING ALGORITHM FOR ADVERSARIAL ROBUSTNESS

We now introduce adversarial querying (AQ), an adversarial training algorithm for meta-learning. Let  $A(\theta, S)$  denote a fine-tuning algorithm. Then,  $A$  is a map from support data set,  $S$ , and network parameters,  $\theta$ , to parameters for the fine-tuned network. Then, we seek to solve the following minimax problem (Equation 1 revisited):

$$\min_{\theta} \mathbb{E}_{S, (\mathbf{x}, y)} \left[ \max_{\|\delta\|_p < \epsilon} \mathcal{L}(F_{A(\theta, S)}, \mathbf{x} + \delta, y) \right],$$

where  $S$  and  $(\mathbf{x}, y)$  are support and query data, respectively, sampled from the training distribution, and  $\epsilon$  is a  $p$ -norm bound for the attacker. Thus, the objective is to find a central parameter vector

which, when fine-tuned on support data, minimizes the expected query loss against an attacker. We approach this minimax problem with an alternating algorithm consisting of the following steps:

1. Sample support and query data
2. Fine-tune on the support data (inner loop)
3. Perturb query data to maximize loss
4. Minimize query loss, backpropagating through the fine-tuning algorithm (outer loop)

A formal treatment of this method is presented in Algorithm 3. We test adversarial querying across multiple data sets and meta-learning protocols. It is important to note that adversarial querying is algorithm-agnostic. We test this method on the ProtoNet, R2-D2, and MetaOptNet algorithms on CIFAR-FS and Mini-ImageNet (see Table 4 and Table 5).

---

**Algorithm 3:** Adversarial Querying
 

---

**Require:** Base model,  $F_\theta$ , fine-tuning algorithm,  $A$ , learning rate,  $\gamma$ , and distribution over tasks,  $p(\mathcal{T})$ .

Initialize  $\theta$ , the weights of  $F$ ;

**while** not done **do**

Sample batch of tasks,  $\{\mathcal{T}_i\}_{i=1}^n$ , where  $\mathcal{T}_i \sim p(\mathcal{T})$  and  $\mathcal{T}_i = (\mathcal{T}_i^s, \mathcal{T}_i^q)$ .

**for**  $i = 1, \dots, n$  **do**

Fine-tune model on  $\mathcal{T}_i$ . New network parameters are written  $\theta_i = A(\theta, \mathcal{T}_i^s)$ .

Construct adversarial query data,  $\widehat{\mathcal{T}}_i^q$ , by maximizing  $\mathcal{L}(F_{\theta_i}, \widehat{\mathcal{T}}_i^q)$  constrained to  $\|\widehat{\mathbf{x}}_j^q - \mathbf{x}_j^q\|_p < \epsilon$  for query examples,  $\mathbf{x}_j^q$ , and their associated adversaries,  $\widehat{\mathbf{x}}_j^q$ .

Compute gradient  $g_i = \nabla_{\theta} \mathcal{L}(F_{\theta_i}, \widehat{\mathcal{T}}_i^q)$ .

Update base model parameters:  $\theta \leftarrow \theta - \frac{\gamma}{n} \sum_i g_i$

---

Model	$\mathcal{A}_{nat}$ 1-Shot	$\mathcal{A}_{adv}$ 1-Shot	$\mathcal{A}_{nat}$ 5-Shot	$\mathcal{A}_{adv}$ 5-Shot
ProtoNet AQ	42.33%	26.48%	63.53%	40.11%
R2-D2 AQ	52.38%	<b>32.33%</b>	69.25%	<b>44.80%</b>
MetaOptNet AQ	53.27%	30.74%	71.07%	43.79%

Table 4: Comparison of adversarially queried (AQ) meta-learners on 1-shot and 5-shot CIFAR-FS.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top 1-shot and 5-shot robust accuracy is bolded.

Model	$\mathcal{A}_{nat}$ 1-Shot	$\mathcal{A}_{adv}$ 1-Shot	$\mathcal{A}_{nat}$ 5-Shot	$\mathcal{A}_{adv}$ 5-Shot
ProtoNet AQ	33.31%	17.69%	52.04%	27.99%
R2-D2 AQ	37.91%	<b>20.59%</b>	57.87%	<b>31.52%</b>
MetaOptNet AQ	43.74%	18.37%	60.71%	28.08%

Table 5: Comparison of adversarially queried (AQ) meta-learners on 1-shot and 5-shot Mini-ImageNet.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top 1-shot and 5-shot robust accuracy is bolded.

In our tests, R2-D2 outperforms MetaOptNet in robust accuracy despite having a less powerful backbone architecture. In Section 4.5, we dissect the effects of backbone architecture and classification head on robustness of meta-learned models. In Section 4.7, we verify that adversarial querying generates networks robust to a wide array of strong attacks.

#### 4.1 ADVERSARIAL QUERYING IS MORE ROBUST THAN TRANSFER LEARNING FROM ADVERSARIALLY TRAINED MODELS

We observe above that few-shot learning methods with a non-robust feature extractor break under attack. But what if we use a robust feature extractor? In the following section, we consider both transfer learning and meta-learning with a robust feature extractor.

In order to compare robust transfer learning and meta-learning, we train the backbone networks from meta-learning algorithms on all training data simultaneously in the fashion of standard adversarial training using 7-PGD (not meta-learning). We then fine-tune using the head from a meta-learning algorithm on top of the transferred feature extractor. We compare the performance of these feature extractors to that of those trained using adversarially queried meta-learning algorithms with the same backbones and heads. This experiment provides a direct comparison of feature extractors produced by robust transfer learning and robust meta-learning (see Table 6). Meta-learning exhibits far superior robustness than transfer learning for all algorithms we test. Additional experiments on CIFAR-FS and on 1-shot Mini-ImageNet can be found in Appendix A.2.

Model	$\mathcal{A}_{nat}$ Transfer	$\mathcal{A}_{adv}$ Transfer	$\mathcal{A}_{nat}$ Meta	$\mathcal{A}_{adv}$ Meta
MAML	32.79%	18.03%	<b>33.45%</b>	<b>23.07%</b>
ProtoNet	31.14%	22.31%	<b>52.04%</b>	<b>27.99%</b>
R2-D2	39.13%	25.33%	<b>57.87%</b>	<b>31.52%</b>
MetaOptNet	50.23%	22.45%	<b>60.71%</b>	<b>28.08%</b>

Table 6: Adversarially trained transfer learning and adversarially queried meta-learning on 5-shot Mini-ImageNet.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top natural and robust accuracy for each architecture is bolded.

#### 4.2 WHY ATTACK ONLY QUERY DATA?

In the adversarial querying procedure detailed in Algorithm 3, we only attack query data. Consider that the loss value on query data represents performance on testing data after fine-tuning on the support data. Thus, low loss on perturbed query data represents robust accuracy on testing data after fine-tuning. Simply put, minimizing loss on adversarial query data moves the parameter vector towards a network with high robust test accuracy. It follows that attacking only support data is not an option for achieving robust meta-learners. Attacking support data but not query data can be seen as maximizing clean test accuracy when fine-tuned in a robust manner, but since we want to maximize *robust* test accuracy, this would be inappropriate. One question remains: should we attack both query and support data?

One reason to perturb only query data is computational efficiency. The bulk of computation in adversarial training is spent computing adversarial attacks since perturbing each batch requires an iterative algorithm. Perturbing support data doubles the number of adversarial attacks computed during training. Thus, only attacking query data significantly accelerates training.

Another reason to avoid attacking support data is that some meta-learning fine-tuning algorithms based on metric learning, such as ProtoNet, do not involve loss minimization during fine-tuning, so it is not clear what loss the attacker would maximize during the inner loop. But if attacking support data during the inner loop of training were to significantly improve robust performance, we would like to know.

We now compare adversarial querying to a variant in which support data is also perturbed during training. We use MAML to conduct this comparison on the Omniglot and Mini-ImageNet data sets. Additional experiments on 1-shot tasks can be found in Appendix A.3.

In these experiments, we find that adversarially attacking the support data during the inner loop of meta-learning does not improve performance over adversarial querying. Furthermore, networks trained in this fashion require adversarial fine-tuning during test time or else they suffer a massive loss in robust test accuracy. Following these results and the significant reasons to avoid attacking support data, we subsequently only attack query data.

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$	$\mathcal{A}_{nat(AT)}$	$\mathcal{A}_{adv(AT)}$
MAML (naturally trained)	97.12%	82.28%	97.71%	87.94%
MAML adv. query	97.27%	<b>95.85%</b>	97.51%	<b>96.14%</b>
MAML adv. query and support	95.61%	77.73%	97.46%	95.65%
ADML	97.31%	94.19%	97.56%	94.82%

Table 7: Performance on 5-shot Omniglot. Robust accuracy,  $\mathcal{A}_{adv}$ , is computed with respect to a 20-step PGD attack.  $\mathcal{A}_{nat(AT)}$  and  $\mathcal{A}_{adv(AT)}$  are natural and robust test accuracy with 7-PGD training during fine-tuning. Top robust accuracy with and without adversarial fine-tuning is bolded.

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$	$\mathcal{A}_{nat(AT)}$	$\mathcal{A}_{adv(AT)}$
MAML (naturally trained)	60.25%	0.03%	32.45%	1.55%
MAML adv. query	33.45%	<b>23.07%</b>	33.03%	<b>23.29%</b>
MAML adv. query and support	29.98%	22.55%	30.44%	23.03%
ADML	47.75%	18.49%	47.27%	20.23%

Table 8: Performance on 5-shot Mini-ImageNet. Robust accuracy,  $\mathcal{A}_{adv}$ , is computed with respect to a 20-step PGD attack.  $\mathcal{A}_{nat(AT)}$  and  $\mathcal{A}_{adv(AT)}$  are natural and robust test accuracy with 7-PGD training during fine-tuning. Top robust accuracy with and without adversarial fine-tuning is bolded.

#### 4.3 ADVERSARIAL QUERYING MAY BE USED TO ADAPT OTHER ROBUSTNESS TECHNIQUES TO META-LEARNING

Adversarial querying can also be used to construct meta-learning analogues for other variants of adversarial training. We explore this by using the TRADES loss function in the querying step of AQ (Zhang et al., 2019). We refer to this method as meta-TRADES. While meta-TRADES can marginally outperform our initial adversarial querying method in robust accuracy with a careful hyperparameter choice,  $\lambda$ , we find that networks trained with meta-TRADES severely sacrifice natural accuracy (see Table 9). Additional experiments on 1-shot tasks can be found in Appendix A.4.

Model	$\mathcal{A}_{nat}$ MI	$\mathcal{A}_{adv}$ MI	$\mathcal{A}_{nat}$ FS	$\mathcal{A}_{adv}$ FS
R2-D2 adversarial queried	<b>57.87%</b>	31.52%	<b>69.25%</b>	44.80%
R2-D2 TRADES ( $1/\lambda = 1$ )	56.02%	30.96%	66.29%	45.59%
R2-D2 TRADES ( $1/\lambda = 3$ )	51.51%	<b>32.30%</b>	61.41%	<b>46.54%</b>
R2-D2 TRADES ( $1/\lambda = 6$ )	34.29%	22.04%	58.32%	45.89%

Table 9: 5-shot Mini-ImageNet (MI) and CIFAR-FS (FS) results comparing meta-TRADES to adversarial querying.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack.

#### 4.4 FOR BETTER NATURAL AND ROBUST ACCURACY, ONLY FINE-TUNE THE LAST LAYER.

High performing meta-learning models, like MetaOptNet and R2-D2, fix their feature extractor and only update their last linear layer during fine-tuning. In the setting of transfer learning, robustness is a feature of early convolutional layers, and re-training these early layers leads to a significant drop in robust test accuracy (Shafahi et al., 2019). We verify that re-training only the last layer leads to improved natural and robust accuracy in adversarially queried meta-learners by training a MAML model but only updating the final fully-connected layer during fine-tuning including during the inner loop of meta-learning. We find that the AQ model trained by fine-tuning only the last layer during the inner loop decisively outperforms the standard AQ MAML algorithm in both natural and robust accuracy (see Table 10).

Layers updated	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$	$\mathcal{A}_{nat(AT)}$	$\mathcal{A}_{adv(AT)}$
All layers	33.45%	23.07%	33.03%	23.29%
FC only	<b>40.06%</b>	<b>25.15%</b>	<b>39.94%</b>	<b>25.32%</b>

Table 10: Adversarially queried MAML compared with a MAML variant where only the last layer is re-trained during fine-tuning on 5-shot Mini-ImageNet.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack.  $\mathcal{A}_{nat(AT)}$  and  $\mathcal{A}_{adv(AT)}$  are natural and robust test accuracy, respectively with 7-PGD training during fine-tuning. Layers are fine-tuned for 10 steps with a learning rate of 0.01.

#### 4.5 THE R2-D2 HEAD, NOT EMBEDDING, IS RESPONSIBLE FOR SUPERIOR ROBUST PERFORMANCE.

The naturally trained R2-D2 algorithm performs worse than MetaOptNet in natural accuracy, but previous research has found that performance discrepancies between meta-learning algorithms might be an artifact of different backbone networks (Chen et al., 2019). We confirm that MetaOptNet with the R2-D2 backbone performs similarly to R2-D2 in the natural meta-learning setting (see Table 11). However, we find that the performance discrepancy in the adversarial setting is *not* explained by differences in backbone architecture. In our adversarial querying experiments, we see that MetaOptNet is less robust than R2-D2. This discrepancy remains when we train MetaOptNet with the R2-D2 backbone (see Table 12). We conclude that MetaOptNet’s backbone is not responsible for its inferior robustness. These experiments suggest that ridge regression may be a more effective fine-tuning technique than SVM for robust performance. ProtoNet with R2-D2 backbone also performs worse than the other two adversarially queried models with the same backbone architecture.

Model	1-shot MI	5-shot MI	1-shot FS	5-shot FS
R2-D2	55.22%	73.02%	68.36%	82.81%
MetaOptNet	<b>60.65%</b>	<b>78.12%</b>	<b>70.99%</b>	<b>84.11%</b>
MetaOptNet (R2-D2 backbone)	55.78%	73.15%	68.37%	82.71%

Table 11: Natural test accuracy of naturally trained R2-D2, MetaOptNet, and the MetaOptNet head with R2-D2 backbone on the Mini-ImageNet (MI) and CIFAR-FS (FS) data sets.

Model	1-shot MI	5-shot MI	1-shot FS	5-shot FS
R2-D2	<b>20.59%</b>	<b>31.52%</b>	<b>32.33%</b>	<b>44.80%</b>
MetaOptNet	18.37%	28.08%	30.74%	43.79%
MetaOptNet (R2-D2 backbone)	18.81%	24.68%	29.57%	41.90%
ProtoNet (R2-D2 backbone)	18.24%	28.39%	26.48%	40.59%

Table 12: Robust test accuracy of adversarially queried R2-D2, MetaOptNet, and the MetaOptNet and heads with R2-D2 backbone on Mini-ImageNet (MI) and CIFAR-FS (FS) data sets. Robust accuracy is computed with respect to a 20-step PGD attack.

#### 4.6 ENHANCING ROBUSTNESS WITH ROBUST ARCHITECTURAL FEATURES

In addition to adversarial training, architectural features have been used to enhance robustness (Xie et al., 2019). Feature denoising blocks pair classical denoising operations with learned  $1 \times 1$  convolutions to reduce the feature noise in feature maps at various stages of a network, and thus reduce the success of adversarial attacks. Massive architectures with these blocks have achieved state-of-the-art robustness against targeted adversarial attacks on ImageNet. However, when deployed on small networks for meta-learning, we find that denoising blocks do not improve robustness. We deploy denoising blocks identical to those in Xie et al. (2019) after various layers of the R2-D2 network. The best results for the denoising experiments are achieved by adding a denoising block after the fourth layer in the R2-D2 embedding network (see Table 13).

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$
R2-D2	73.02%	0.00%
R2-D2 AQ	57.87%	<b>31.52%</b>
R2-D2 AQ Denoising	57.68%	31.14%

Table 13: 5-shot MiniImageNet results for our highest performing R2-D2 with feature denoising blocks.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top robust accuracy is bolded.

#### 4.7 RESISTANCE TO OTHER ATTACKS

We test our method by exposing our adversarially queried R2-D2 model to a variety of powerful adversarial attacks. We implement the momentum iterated fast gradient sign method (MI-FGSM), DeepFool, and 20-step PGD with 20 random restarts (Dong et al., 2018; Moosavi-Dezfooli et al., 2016; Madry et al., 2017). Our adversarially queried model is indeed nearly as robust against the strongest  $\ell_\infty$  bounded attacker as it is against the 20-step PGD attack with a single random start we tested against previously. Note that DeepFool is not  $\ell_\infty$  bounded and thus the perturbed images are outside of the robustness radius enforced during adversarial querying. Additional experiments on CIFAR-FS can be found in Appendix A.6.

Model	$\mathcal{A}_{DF}$	$\mathcal{A}_{MI}$	$\mathcal{A}_{20-PGD}$
R2-D2	7.91%	0.01%	0.0%
R2-D2 AQ (ours)	<b>14.45%</b>	<b>31.87%</b>	<b>30.31%</b>
R2-D2 Transfer	0.42%	24.01%	19.75%

Table 14: 5-shot MiniImageNet results against DeepFool (DF) (2 iteration)  $\ell_\infty$  attack, MI-FGSM (MI) ( $\epsilon = 8/255$ ) attack, and PGD attack with 20 random restarts (20-PGD). We compare R2-D2 trained with adversarial-querying (AQ) to the adversarially trained transfer learning R2-D2 as in section 4.1.

## 5 PRE-PROCESSING DEFENSES AS AN ALTERNATIVE TO ADVERSARIAL TRAINING

Recent works have proposed pre-processing defenses for sanitizing adversarial examples before feeding them into a naturally trained classifier. If successful, these methods would avoid the expensive adversarial querying procedure during training. While this approach has found success in the mainstream literature, we find that it is ineffective in the few-shot regime.

In DefenseGAN, a GAN trained on natural images is used to sanitize an adversarial example by replacing (possible corrupted) test images with the nearest image in the output range of the GAN (Samangouei et al., 2018). Unfortunately, GANs are not expressive enough to preserve the integrity of testing images on complex data sets involving high-resolution natural images, and recent attacks have critically compromised the performance of this defense (Ilyas et al., 2017; Athalye et al., 2018). We found the expressiveness of the generator architecture used in the original DefenseGAN setup to be insufficient for even CIFAR-FS, so we substitute a stronger ProGAN generator to model the CIFAR-100 classes (Karras et al., 2017).

Another pre-processing method, the superresolution defense, first denoises data with sparse wavelet filters and then performs superresolution (Mustafa et al., 2019). This defense is also motivated by the principle of projecting adversarial examples onto the natural image manifold. We test the superresolution defense using the same wavelet filtering and superresolution network (SRResNet) used by Mustafa et al. (2019) and first introduced by Ledig et al. (2017). We train the SRResNet, in a similar fashion to the generator for DefenseGAN, on the entire CIFAR-100 data set before applying the superresolution defense.

We find that these methods are not well suited to the few-shot domain, in which the generative model or superresolution network may not be able to train on the little data available. Moreover, even

after training the generator on all CIFAR-100 classes, we find that DefenseGAN with a naturally trained R2-D2 meta-learner performs significantly worse in both natural and robust accuracy than an adversarially queried meta-learner of the same architecture. Similarly, the superresolution defense achieves little robustness. The results of these experiments can be found in Table 15.

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$
R2-D2	82.81%	0.00%
R2-D2 AQ (ours)	69.25%	<b>44.80%</b>
R2-D2 with SR defense	35.15%	23.00%
R2-D2 with DefenseGAN	35.15%	28.05%

Table 15: 5-shot CIFAR-FS results comparing the superresolution defense (SR defense) and DefenseGAN.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Both methods perform worse than their adversarially queried counterpart. Top robust accuracy is bolded.

## 6 DISCUSSION & CONCLUSION

Naturally trained networks for few-shot image classification are vulnerable to adversarial attacks, and existing robust transfer learning methods do not perform well on few-shot tasks. Even, when adversarially fine-tuned, naturally trained networks suffer from adversarial vulnerability. We thus identify the need for few-shot methods for adversarial robustness. In particular, we study robustness in the context of meta-learning. We develop an algorithm-agnostic method, called adversarial querying, for hardening meta-learning models. We find that meta-learning models are most robust when the feature extractor is fixed and only the last layer is retrained during fine-tuning. We further identify that choice of classification head significantly impacts robustness. We believe that this paper is a starting point for developing adversarially robust methods for few-shot applications.

## REFERENCES

- Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018.
- Yoshua Bengio. Deep learning of representations for unsupervised and transfer learning. In *Proceedings of ICML workshop on unsupervised and transfer learning*, pp. 17–36, 2012.
- Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International conference on learning representations*, 2019.
- Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018. doi: 10.1109/cvpr.2018.00957. URL <http://dx.doi.org/10.1109/cvpr.2018.00957>.
- Yan Duan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. In *Advances in neural information processing systems*, pp. 1087–1098, 2017.
- Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning models. *arXiv preprint arXiv:1707.08945*, 2017.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1126–1135. JMLR. org, 2017.

- Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4367–4375, 2018.
- Andrew Ilyas, Ajil Jalal, Eirini Asteri, Constantinos Daskalakis, and Alexandros G Dimakis. The robust manifold defense: Adversarial training using generative models. *arXiv preprint arXiv:1712.09196*, 2017.
- Łukasz Kaiser, Ofir Nachum, Aurko Roy, and Samy Bengio. Learning to remember rare events. *arXiv preprint arXiv:1703.03129*, 2017.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation, 2017.
- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and et al. Photo-realistic single image super-resolution using a generative adversarial network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017. doi: 10.1109/cvpr.2017.19. URL <http://dx.doi.org/10.1109/CVPR.2017.19>.
- Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10657–10665, 2019.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Metric learning for large scale image classification: Generalizing to new classes at near-zero cost. In *European Conference on Computer Vision*, pp. 488–501. Springer, 2012.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2574–2582, 2016.
- Aamir Mustafa, Salman H. Khan, Munawar Hayat, Jianbing Shen, and Ling Shao. Image super-resolution as a defense against adversarial attacks, 2019.
- Tomas Pfister, James Charles, and Andrew Zisserman. Domain-adaptive discriminative one-shot learning of gestures. In *European Conference on Computer Vision*, pp. 814–829. Springer, 2014.
- Parsa Saadatpanah, Ali Shafahi, and Tom Goldstein. Adversarial attacks on copyright detection systems. *arXiv preprint arXiv:1906.07153*, 2019.
- Pouya Samangouei, Maya Kabkab, and Rama Chellappa. Defense-gan: Protecting classifiers against adversarial attacks using generative models. *arXiv preprint arXiv:1805.06605*, 2018.
- Ali Shafahi, Parsa Saadatpanah, Chen Zhu, Amin Ghiasi, Christoph Studer, David Jacobs, and Tom Goldstein. Adversarially robust transfer learning. *arXiv preprint arXiv:1905.08232*, 2019.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 4077–4087, 2017.
- Manasi Vartak, Arvind Thiagarajan, Conrado Miranda, Jeshua Bratman, and Hugo Larochelle. A meta-learning perspective on cold-start recommendations for items. In *Advances in neural information processing systems*, pp. 6904–6914, 2017.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Advances in neural information processing systems*, pp. 3630–3638, 2016.

Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 501–509, 2019.

Chengxiang Yin, Jian Tang, Zhiyuan Xu, and Yanzhi Wang. Adversarial meta-learning. *arXiv preprint arXiv:1806.03316*, 2018.

Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019.

Yue Zhao, Hong Zhu, Qintao Shen, Ruigang Liang, Kai Chen, and Shengzhi Zhang. Practical adversarial attack against object detector. *arXiv preprint arXiv:1812.10217*, 2018.

## A APPENDIX

### A.1 ADDITIONAL EXPERIMENTS TESTING THE ROBUSTNESS OF NATURALLY TRAINED META-LEARNING MODELS

Model	$\mathcal{A}_{nat}$ MI	$\mathcal{A}_{adv}$ MI	$\mathcal{A}_{nat}$ FS	$\mathcal{A}_{adv}$ FS
ProtoNet	43.26%	0.00%	59.56%	0.00%
R2-D2	55.22%	0.00%	68.36%	0.00%
MetaOptNet	60.65%	0.00%	70.99%	0.00%

Table 16: 1-shot MiniImageNet (MI) and CIFAR-FS (FS) results comparing naturally trained meta-learners.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack.

### A.2 ADDITIONAL EXPERIMENTS COMPARING ROBUST META-LEARNING AND ROBUST TRANSFER LEARNING

Model	$\mathcal{A}_{nat}$ Transfer	$\mathcal{A}_{adv}$ Transfer	$\mathcal{A}_{nat}$ Meta	$\mathcal{A}_{adv}$ Meta
ProtoNet	45.98%	35.63%	<b>63.53%</b>	<b>40.11%</b>
R2-D2	53.26%	33.33%	<b>69.25%</b>	<b>44.80%</b>
MetaOptNet	60.72%	35.11%	<b>71.07%</b>	<b>43.79%</b>

Table 17: Adversarially trained transfer learning and adversarially queried meta-learning on 5-shot CIFAR-FS.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top natural and robust accuracy for each architecture is bolded.

Model	$\mathcal{A}_{nat}$ Transfer	$\mathcal{A}_{adv}$ Transfer	$\mathcal{A}_{nat}$ Meta	$\mathcal{A}_{adv}$ Meta
MAML	<b>25.84%</b>	15.52%	21.42%	<b>17.9%</b>
ProtoNet	25.58%	18.06%	<b>33.31%</b>	<b>17.69%</b>
R2-D2	27.88%	17.72%	<b>37.91%</b>	<b>20.59%</b>
MetaOptNet	34.71%	16.01%	<b>43.74%</b>	<b>18.37%</b>

Table 18: Adversarially trained transfer learning and adversarially queried meta-learning on 1-shot Mini-ImageNet.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top natural and robust accuracy for each architecture is bolded.

Model	$\mathcal{A}_{nat}$ Transfer	$\mathcal{A}_{adv}$ Transfer	$\mathcal{A}_{nat}$ Meta	$\mathcal{A}_{adv}$ Meta
ProtoNet	35.15%	26.25%	<b>42.33%</b>	<b>26.48%</b>
R2-D2	38.15%	26.78%	<b>52.38%</b>	<b>32.33%</b>
MetaOptNet	42.98%	25.37%	<b>53.27%</b>	<b>30.74%</b>

Table 19: Adversarially trained transfer learning and adversarially queried meta-learning on 1-shot CIFAR-FS.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack. Top natural and robust accuracy for each architecture is bolded.

### A.3 ADDITIONAL EXPERIMENTS COMPARING ADVERSARIAL QUERYING TO BOTH ADVERSARIAL SUPPORT AND ADVERSARIAL QUERYING DURING TRAINING

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$	$\mathcal{A}_{nat(AT)}$	$\mathcal{A}_{adv(AT)}$
MAML (naturally trained)	91.50%	68.46%	91.60%	74.66%
MAML adv. query	91.11%	88.72%	91.31%	89.01%
MAML adv. query and support	90.58%	82.23%	91.36%	88.97%
ADML	91.99%	86.87%	92.24%	87.35%

Table 20: Performance on 1-shot Omniglot. Robust accuracy,  $\mathcal{A}_{adv}$ , is computed with respect to a 20-step PGD attack.  $\mathcal{A}_{nat(AT)}$  and  $\mathcal{A}_{adv(AT)}$  are natural and robust test accuracy with 7-PGD training during fine-tuning.

Model	$\mathcal{A}_{nat}$	$\mathcal{A}_{adv}$	$\mathcal{A}_{nat(AT)}$	$\mathcal{A}_{adv(AT)}$
MAML (naturally trained)	45.04%	0.03%	33.18%	0.20%
MAML adv. query	21.42%	17.9%	21.23%	17.87%
MAML adv. query and support	22.39%	19.07%	22.06%	19.14%
ADML	26.68%	16.63%	27.34%	17.78%

Table 21: Performance on 1-shot Mini-ImageNet. Robust accuracy,  $\mathcal{A}_{adv}$ , is computed with respect to a 20-step PGD attack.  $\mathcal{A}_{nat(AT)}$  and  $\mathcal{A}_{adv(AT)}$  are natural and robust test accuracy with 7-PGD training during fine-tuning.

### A.4 ADDITIONAL META-TRADES EXPERIMENTS

Model	$\mathcal{A}_{nat}$ MI	$\mathcal{A}_{adv}$ MI	$\mathcal{A}_{nat}$ FS	$\mathcal{A}_{adv}$ FS
R2-D2 adversarial queried	37.91%	20.59%	52.38%	32.33%
R2-D2 TRADES ( $1/\lambda = 1$ )	39.11%	20.25%	48.77%	31.99%
R2-D2 TRADES ( $1/\lambda = 6$ )	34.27%	22.00%	44.37%	33.55%

Table 22: 1-shot Mini-ImageNet (MI) and CIFAR-FS (FS) results comparing meta-TRADES to adversarial querying.  $\mathcal{A}_{nat}$  and  $\mathcal{A}_{adv}$  are natural and robust test accuracy, respectively, where robust accuracy is computed with respect to a 20-step PGD attack.

### A.5 TRAINING HYPERPARAMETERS

We train ProtoNet, R2-D2, and MetaOptNet models for 60 epochs with SGD. We use a learning rate of 0.1, momentum (Nesterov) of 0.9, and a weight decay term of  $5(10^{-4})$  for the parameters of both the head and the embedding. We decrease the learning rate to 0.06 after epoch 20, 0.012 after epoch 40, and 0.0024 after epoch 50. MAML is trained for 60000 epochs with meta learning rate of 0.001 and fine-tuning learning rate of 0.01. Fine-tuning is performed for 10 steps per task.

## A.6 MORE RESULTS AGAINST OTHER ATTACKS

Model	$\mathcal{A}_{DF}$	$\mathcal{A}_{MI}$	$\mathcal{A}_{20-PGD}$
R2-D2	0.00%	0.39%	0.01%
R2-D2 AQ (ours)	<b>14.45%</b>	<b>53.46%</b>	<b>46.57%</b>
R2-D2 AT (Transfer Learning)	1.41%	38.28%	33.17%

Table 23: 5-shot CIFAR-FS results against DeepFool (DF) (2 iteration)  $\ell_\infty$  attack, MI-FGSM (MI) ( $\epsilon = 8/255$ ) attack, and PGD attack with 20 random restarts (20-PGD). We compare R2-D2 trained with adversarial-querying (AQ) to the transfer learning R2-D2 as in section 4.1.

Model	$\mathcal{A}_{ResNet}$
R2-D2	0.00%
R2-D2 AQ (ours)	<b>59.68%</b>
R2-D2 AT (Transfer Learning)	42.02%

Table 24: 5-shot CIFAR-FS results against black-box transfer attacks crafted on an adversarially trained (transfer learning) ResNet-12 model using 7-PGD. We then test R2-D2 trained with adversarial-querying (AQ) and the transfer learning R2-D2 model on these crafted perturbations.