

Improving Few-Shot Learning through Multi-task Representation Learning Theory

Quentin Bouniot^{1,2}[0000-0002-0982-372X], Ievgen Redko²[0000-0002-3860-5502],
Romaric Audigier¹[0000-0002-4757-2052], Angélique Loesch¹[0000-0001-5427-3010],
Amaury Habrard^{2,3}[0000-0003-3038-9347]

¹ Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France
{firstname.lastname}@cea.fr

² Université de Lyon, UJM-Saint-Etienne, CNRS, IOGS,
Laboratoire Hubert Curien UMR 5516, F-42023, Saint-Etienne, France
{firstname.lastname}@univ-st-etienne.fr

³ Institut Universitaire de France (IUF)

Abstract. In this paper, we consider the framework of multi-task representation (MTR) learning where the goal is to use source tasks to learn a representation that reduces the sample complexity of solving a target task. We start by reviewing recent advances in MTR theory and show that they can provide novel insights for popular meta-learning algorithms when analyzed within this framework. In particular, we highlight a fundamental difference between gradient-based and metric-based algorithms in practice and put forward a theoretical analysis to explain it. Finally, we use the derived insights to improve the performance of meta-learning methods via a new spectral-based regularization term and confirm its efficiency through experimental studies on few-shot classification benchmarks. To the best of our knowledge, this is the first contribution that puts the most recent learning bounds of MTR theory into practice for the task of few-shot classification.

Keywords: Few-shot learning, meta-learning, multi-task learning.

1 Introduction

Even though many machine learning methods now enjoy a solid theoretical justification, some more recent advances in the field are still in their preliminary state which requires the hypotheses put forward by the theoretical studies to be implemented and verified in practice. One such notable example is the success of *meta-learning*, also called *learning to learn* (LTL), methods where the goal is to produce a model on data coming from a set of (meta-train) source tasks to use it as a starting point for learning successfully a new previously unseen (meta-test) target task. The success of many meta-learning approaches is directly related to their capacity of learning a good representation [37] from a set of tasks making it closely related to multi-task representation learning (MTR). For this latter, several theoretical studies [2, 5, 29, 34, 53] provided probabilistic learning bounds

that require the amount of data in the meta-train source task *and* the number of meta-train tasks to tend to infinity for it to be efficient. While capturing the underlying general intuition, these bounds do not suggest that all the source data is useful in such learning setup due to the additive relationship between the two terms mentioned above and thus, for instance, cannot explain the empirical success of MTR in few-shot classification (FSC) task. To tackle this drawback, two very recent studies [15, 45] aimed at finding deterministic assumptions that lead to faster learning rates allowing MTR algorithms to benefit from all the source data. Contrary to probabilistic bounds that have been used to derive novel learning strategies for meta-learning algorithms [2, 53], there has been no attempt to verify the validity of the assumptions leading to the fastest known learning rates in practice or to enforce them through an appropriate optimization procedure.

In this paper, we aim to use the recent advances in MTR theory [15, 45] to explore the inner workings of these popular meta-learning methods. Our rationale for such an approach stems from a recent work [47] proving that the optimization problem behind the majority of meta-learning algorithms can be written as an MTR problem. Thus, we believe that looking at meta-learning algorithms through the recent MTR theory lens, could provide us a better understanding for the capacity to work well in the few-shot regime. In particular, we take a closer look at two families of meta-learning algorithms, notably: gradient-based algorithms [6, 10, 26, 28, 31–33, 37, 38] including MAML [16] and metric-based algorithms [1, 22, 27, 41–43, 46] with its most prominent example given by PROTONET [42].

Our main contributions are then two-fold:

1. We empirically show that tracking the validity of assumptions on optimal predictors used in [15, 45] reveals a striking difference between the behavior of gradient-based and metric-based methods in how they learn their optimal feature representations. We provide elements of theoretical analysis that explain this behavior and explain the implications of it in practice. Our work is thus complementary to Wang et al. [47] and connects MTR, FSC and Meta-Learning from both theoretical and empirical points of view.
2. Following the most recent advances in the MTR field that leads to faster learning rates, we show that theoretical assumptions mentioned above can be forced through simple yet effective learning constraints which improve performance of the considered algorithms for FSC baselines: gradient- and metric-based methods using episodic training, as well as non-episodic algorithm such as Multi-Task Learning (MTL [47]).

The rest of the paper is organized as follows. We introduce the MTR problem and the considered meta-learning algorithms in Section 2. In Section 3, we investigate and explain how they behave in practice. We further show that one can force meta-learning algorithms to satisfy such assumptions through adding an appropriate spectral regularization term to their objective function. In Section 4, we provide an experimental evaluation of several state-of-the-art meta-learning and MTR methods. We conclude and outline the future research perspectives in Section 5.

2 Preliminary Knowledge

2.1 Multi-Task Representation Learning Setup

Given a set of T source tasks observed through finite size samples of size n_1 grouped into matrices $\mathbf{X}_t = (\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,n_1}) \in \mathbb{R}^{n_1 \times d}$ and vectors of outputs $y_t = (y_{t,1}, \dots, y_{t,n_1}) \in \mathbb{R}^{n_1}$, $\forall t \in [[T]] := \{1, \dots, T\}$ generated by their respective distributions μ_t , the goal of MTR is to learn a shared representation ϕ belonging to a certain class of functions $\Phi := \{\phi \mid \phi : \mathbb{X} \rightarrow \mathbb{Y}, \mathbb{X} \subseteq \mathbb{R}^d, \mathbb{Y} \subseteq \mathbb{R}^k\}$, generally represented as (deep) neural networks, and linear predictors $\mathbf{w}_t \in \mathbb{R}^k$, $\forall t \in [[T]]$ grouped in a matrix $\mathbf{W} \in \mathbb{R}^{T \times k}$. More formally, this is done by solving the following optimization problem:

$$\hat{\phi}, \hat{\mathbf{W}} = \arg \min_{\phi \in \Phi, \mathbf{W} \in \mathbb{R}^{T \times k}} \frac{1}{T n_1} \sum_{t=1}^T \sum_{i=1}^{n_1} \ell(y_{t,i}, \langle \mathbf{w}_t, \phi(\mathbf{x}_{t,i}) \rangle), \quad (1)$$

where $\ell : \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}_+$, with $\mathbb{Y} \subseteq \mathbb{R}$, is a loss function. Once such a representation is learned, we want to apply it to a new previously unseen target task observed through a pair $(\mathbf{X}_{T+1} \in \mathbb{R}^{n_2 \times d}, y_{T+1} \in \mathbb{R}^{n_2})$ containing n_2 samples generated by the distribution μ_{T+1} . We expect that a linear classifier \mathbf{w} learned on top of the obtained representation leads to a low true risk over the whole distribution μ_{T+1} . For this, we first use $\hat{\phi}$ to solve the following problem:

$$\hat{\mathbf{w}}_{T+1} = \arg \min_{\mathbf{w} \in \mathbb{R}^k} \frac{1}{n_2} \sum_{i=1}^{n_2} \ell(y_{T+1,i}, \langle \mathbf{w}, \hat{\phi}(\mathbf{x}_{T+1,i}) \rangle). \quad (2)$$

Then, we define the true target risk of the learned linear classifier $\hat{\mathbf{w}}_{T+1}$ as: $\mathcal{L}(\hat{\phi}, \hat{\mathbf{w}}_{T+1}) = \mathbb{E}_{(\mathbf{x}, y) \sim \mu_{T+1}} [\ell(y, \langle \hat{\mathbf{w}}_{T+1}, \hat{\phi}(\mathbf{x}) \rangle)]$ and want it to be as close as possible to the ideal true risk $\mathcal{L}(\phi^*, \mathbf{w}_{T+1}^*)$ where \mathbf{w}_{T+1}^* and ϕ^* satisfy:

$$\forall t \in [[T+1]] \text{ and } (\mathbf{x}, y) \sim \mu_t, \quad y = \langle \mathbf{w}_t^*, \phi^*(\mathbf{x}) \rangle + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2). \quad (3)$$

Equivalently, most of the works found in the literature seek to upper-bound the *excess risk* defined as $\text{ER}(\hat{\phi}, \hat{\mathbf{w}}_{T+1}) := \mathcal{L}(\hat{\phi}, \hat{\mathbf{w}}_{T+1}) - \mathcal{L}(\phi^*, \mathbf{w}_{T+1}^*)$.

2.2 Learning Bounds and Assumptions

First studies in the context of MTR relied on the probabilistic assumption [2, 5, 29, 34, 53] stating that meta-train and meta-test tasks distributions are all sampled i.i.d. from the same random distribution. A natural improvement to this bound was then proposed by [15] and [45] that obtained the bounds on the excess risk behaving as

$$\text{ER}(\hat{\phi}, \hat{\mathbf{w}}_{T+1}) \leq O \left(\frac{C(\Phi)}{n_1 T} + \frac{k}{n_2} \right), \quad (4)$$

where $C(\Phi)$ is a measure of the complexity of Φ . Both these results show that all the source and target samples are useful in minimizing the excess risk. Thus, in

the FSC regime where target data is scarce, all source data helps to learn well. From a set of assumptions made by the authors in both of these works, we note the following two:

Assumption 1: Diversity of the source tasks The matrix of optimal predictors \mathbf{W}^* should cover all the directions in \mathbb{R}^k evenly. More formally, this can be stated as $\kappa(\mathbf{W}^*) := \frac{\sigma_1(\mathbf{W}^*)}{\sigma_k(\mathbf{W}^*)} = O(1)$,

where $\sigma_i(\cdot)$ denotes the i^{th} singular value of \mathbf{W}^* . As pointed out by the authors, such an assumption can be seen as a measure of diversity between the source tasks that are expected to be complementary to each other to provide a useful representation for a previously unseen target task. In the following, we will refer to $\kappa(\mathbf{W})$ as the *condition number* for matrix \mathbf{W} .

Assumption 2: Consistency of the classification margin The norm of the optimal predictors \mathbf{w}^* should not increase with the number of tasks seen during meta-training⁴. This assumption says that the classification margin of linear predictors should remain constant thus avoiding over- or under-specialization to the seen tasks.

An intuition behind these assumptions and a detailed review can be found in the Appendix. While being highly insightful, the authors did not provide any experimental evidence suggesting that verifying these assumptions in practice helps to learn more efficiently in the considered learning setting.

2.3 Meta-Learning Algorithms

Meta-learning algorithms considered below learn an optimal representation sequentially via the so-called episodic training strategy introduced by [46], instead of jointly minimizing the training error on a set of source tasks as done in MTR. Episodic training mimics the training process at the task scale with each task data being decomposed into a training set (*support set* S) and a testing set (*query set* Q). Recently, [8] showed that the episodic training setup used in meta-learning leads to a generalization bounds of $O(\frac{1}{\sqrt{T}})$. This bound is independent of the task sample size n_1 , which could explain the success of this training strategy for FSC in the asymptotic limit. However, unlike the results obtained by [15] studied in this paper, the lack of dependence on n_1 makes such a result un-insightful in practice as we are in a finite-sample size setting. This bound does not give information on other parameters to leverage when the task number cannot increase. We now present two major families of meta-learning approaches below.

Metric-based methods These methods learn an embedding space in which feature vectors can be compared using a similarity function (usually a L_2 distance or cosine similarity) [1, 22, 27, 41–43, 46]. They typically use a form of contrastive loss as their objective function, similarly to Neighborhood Component Analysis (NCA) [18] or Triplet Loss [20]. In this paper, we focus our analysis on the popular Prototypical Networks [42] (PROTONET) that computes prototypes as the mean vector of support points belonging to the same class: $\mathbf{c}_i = \frac{1}{|S_i|} \sum_{\mathbf{s} \in S_i} \phi(\mathbf{s})$,

⁴ While not stated separately, this assumption is used in [15] to derive the final result on p.5 after the discussion of Assumption 4.3.

with S_i the subset of support points belonging to class i . PROTONET minimizes the negative log-probability of the true class i computed as the softmax over distances to prototypes \mathbf{c}_i :

$$\mathcal{L}_{proto}(S, Q, \phi) := \mathbb{E}_{\mathbf{q} \sim Q} \left[-\log \frac{\exp(-d(\phi(\mathbf{q}), \mathbf{c}_i))}{\sum_j \exp(-d(\phi(\mathbf{q}), \mathbf{c}_j))} \right] \quad (5)$$

with d being a distance function used to measure similarity between points in the embedding space.

Gradient-based methods These methods learn through end-to-end or two-step optimization [6, 10, 26, 28, 31–33, 37, 38] where given a new task, the goal is to learn a model from the task’s training data specifically adapted for this task. MAML [16] updates its parameters θ using an end-to-end optimization process to find the best initialization such that a new task can be learned quickly, *i.e.* with few examples. More formally, given the loss ℓ_t for each task $t \in [[T]]$, MAML minimizes the expected task loss after an *inner loop* or *adaptation* phase, computed by a few steps of gradient descent initialized at the model’s current parameters:

$$\mathcal{L}_{MAML}(\theta) := \mathbb{E}_{t \sim \eta} [\ell_t(\theta - \alpha \nabla \ell_t(\theta))], \quad (6)$$

with η the distribution of the meta-training tasks and α the learning rate for the adaptation phase. For simplicity, we take a single step of gradient update in this equation.

In what follows, we establish our theoretical analysis for the popular methods PROTONET and MAML. We add their improved variations respectively called Infinite Mixture Prototypes [1] (IMP) and Meta-Curvature [33] (MC) in the experiments to validate our findings.

3 Understanding Meta-learning Algorithms through MTR Theory

3.1 Link between MTR and Meta-learning

Recently, [47] has shown that meta-learning algorithms that only optimize the last layer in the inner-loop, solve the same underlying optimization procedure as multi-task learning. In particular, their contributions have the following implications:

1. For *Metric-based algorithms*, the majority of methods can be seen as MTR problems. This is true, in particular, for PROTONET and IMP algorithms considered in this work.
2. In the case of *Gradient-based algorithms*, such methods as ANIL [37] and MetaOptNet [26] that do not update the embeddings during the inner-loop, can be also seen as multi-task learning. However, MAML and MC in principal do update the embeddings even though there exists strong evidence suggesting that the changes in the weights during their inner-loop are mainly affecting the last layer [37]. Consequently, we follow [47] and use this assumption to analyze MAML and MC in MTR framework as well.

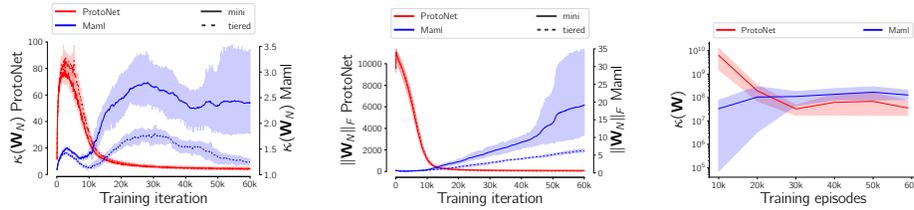


Fig. 1: Evolution of $\kappa(\mathbf{W}_N)$, $\|\mathbf{W}_N\|_F$ and $\kappa(\mathbf{W})$ (in log scale) during the training of PROTONET (red, left axes) and MAML (blue, right axes) on miniImageNet (mini, solid lines) and tieredImageNet (tiered, dashed lines) with 5-way 1-shot episodes.

3. In practice, [47] showed that the mismatch between the multi-task and the actual episodic training setup leads to a negligible difference.

In the following section, we start by empirically verifying that the behavior of meta-learning methods reveals very distinct features when looked at through the prism of the considered MTR theoretical assumptions.

3.2 What happens in practice?

To verify whether theoretical results from MTR setting are also insightful for episodic training used by popular meta-learning algorithms, we first investigate the natural behavior of MAML and PROTONET when solving FSC tasks on the popular *miniImageNet* [38] and *tieredImageNet* [39] datasets. The full experimental setup is detailed in Section 4.1 and in the Appendix. Additional experiments for *Omniglot* [25] benchmark dataset portraying the same behavior are also postponed to the Appendix.

To verify Assumption 1 from MTR theory, we want to compute singular values of \mathbf{W} during the meta-training stage and to follow their evolution. In practice, as T is typically quite large, we propose a more computationally efficient solution that is to calculate the condition number only for the last batch of N predictors (with $N \ll T$) grouped in the matrix $\mathbf{W}_N \in \mathbb{R}^{N \times k}$ that capture the latest dynamics in the learning process. We further note that $\sigma_i(\mathbf{W}_N \mathbf{W}_N^T) = \sigma_i^2(\mathbf{W}_N)$, $\forall i \in [[N]]$ implying that we can calculate the SVD of $\mathbf{W}_N \mathbf{W}_N^T$ (or $\mathbf{W}_N^T \mathbf{W}_N$ for $k \leq N$) and retrieve the singular values from it afterwards. We now want to verify whether \mathbf{w}_t cover all directions in the embedding space and track the evolution of the ratio of singular values $\kappa(\mathbf{W}_N)$ during training. For the first assumption to be satisfied, we expect $\kappa(\mathbf{W}_N)$ to decrease gradually during the training thus improving the generalization capacity of the learned predictors and preparing them for the target task. To verify the second assumption, the norm of the linear predictors should not increase with the number of tasks seen during training, *i.e.*, $\|\mathbf{w}\|_2 = O(1)$ or, equivalently, $\|\mathbf{W}\|_F^2 = O(T)$ and $\|\mathbf{W}_N\|_F = O(1)$.

For gradient-based methods, linear predictors are directly the weights of the last layer of the model. Indeed, for each task, the model learns a batch of linear

predictors and we can directly take the weights for \mathbf{W}_N . Meanwhile, metric-based methods do not use linear predictors but compute a similarity between features. In the case of PROTONET, the similarity is computed with respect to class prototypes that are mean features of the instances of each class. For the Euclidean distance, *this is equivalent to a linear model* with the prototype of a class acting as the linear predictor of this class [42]. This means that we can apply our analysis directly to the *prototypes* computed by PROTONET. In this case, the matrix \mathbf{W}^* will be the matrix of the optimal prototypes and we can then take the *prototypes* computed for each task as our matrix \mathbf{W}_N .

From Fig. 1, we can see that for MAML (*blue*), both $\|\mathbf{W}_N\|_F$ (*left*) and $\kappa(\mathbf{W}_N)$ (*middle*) increase with the number of tasks seen during training, whereas PROTONET (*red*) naturally learns prototypes with a good coverage of the embedding space, and minimizes their norm. Since we compute the singular values of the last N predictors in $\kappa(\mathbf{W}_N)$, we can only compare the *overall behavior* throughout training between methods. For the sake of completeness, we also compute $\kappa(\mathbf{W})$ (*right*) at different stages in the training. To do so, we fix the encoder ϕ_T learned after T episodes and recalculate the linear predictors of the T past training episodes with this fixed encoder. We can see that $\kappa(\mathbf{W})$ of PROTONET also decreases during training and reach a lower final value than $\kappa(\mathbf{W})$ of MAML. This confirms that the dynamics of $\kappa(\mathbf{W}_N)$ and $\kappa(\mathbf{W})$ are similar whereas the values $\kappa(\mathbf{W}_N)$ between methods should not be directly compared. The behavior of $\kappa(\mathbf{W})$ also validate our finding that PROTONET learns to cover the embedding space with prototypes. This behavior is rather peculiar as neither of the two methods explicitly controls the theoretical quantities of interest, and still, PROTONET manages to do it implicitly.

3.3 The case of Meta-Learning algorithms

The differences observed above for the two methods call for a deeper analysis of their behavior.

PROTONET We start by first explaining why PROTONET learns prototypes that cover the embedding space efficiently. This result is given by the following theorem (cf. Appendix for the proof):

Theorem 1. (Normalized PROTONET)

If $\forall i \|\mathbf{c}_i\| = 1$, then $\forall \hat{\phi} \in \arg \min_{\phi} \mathcal{L}_{proto}(S, Q, \phi)$, the matrix of the optimal prototypes \mathbf{W}^ is well-conditioned, i.e. $\kappa(\mathbf{W}^*) = O(1)$.*

This theorem explains the empirical behavior of PROTONET in FSC task: the minimization of its objective function naturally minimizes the condition number when the norm of the prototypes is low.

In particular, it implies that norm minimization seems to initiate the minimization of the condition number seen afterwards due to the contrastive nature of the loss function minimized by PROTONET. We confirm this latter implication through experiments in Section 4 showing that norm minimization is enough for considered metric-based methods to obtain the well-behaved condition number and that minimizing both seems redundant.

MAML Unfortunately, the analysis of MAML in the most general case is notoriously harder, as even expressing its loss function and gradients in the case of an overparametrized linear regression model with only 2 parameters requires using a symbolic toolbox for derivations [3]. To this end, we resort to the linear regression model considered in this latter paper and defined as follows. We assume for all $t \in [[T]]$ that the task parameters θ_t are normally distributed with $\theta_t \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, the inputs $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ and the output $y_t \sim \mathcal{N}(\langle \theta_t, \mathbf{x}_t \rangle, 1)$. For each t , we consider the following learning model and its associated square loss:

$$\hat{y}_t = \langle \mathbf{w}_t, \mathbf{x}_t \rangle, \quad \ell_t = \mathbb{E}_{p(\mathbf{x}_t, y_t | \theta_t)} (y_t - \langle \mathbf{w}_t, \mathbf{x}_t \rangle)^2. \quad (7)$$

We can now state the following result.

Proposition 1. *Let $\forall t \in [[T]]$, $\theta_t \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ and $y_t \sim \mathcal{N}(\langle \theta_t, \mathbf{x}_t \rangle, 1)$. Consider the learning model from Eq. (7), let $\Theta_i := [\theta_i, \theta_{i+1}]^T$, and denote by $\widehat{\mathbf{W}}_2^i$ the matrix of last two predictors learned by MAML at iteration i starting from $\widehat{\mathbf{w}}_0 = \mathbf{0}_d$. Then, we have that:*

$$\forall i, \quad \kappa(\widehat{\mathbf{W}}_2^{i+1}) \geq \kappa(\widehat{\mathbf{W}}_2^i), \quad \text{if } \sigma_{\min}(\Theta_i) = 0. \quad (8)$$

This proposition provides an explanation of why MAML may tend to increase the ratio of singular values during the iterations. Indeed, the condition when this happens indicates that the optimal predictors forming matrix Θ_i are linearly dependent implying that its smallest singular values becomes equal to 0. While this is not expected to be the case for all iterations, we note, however, that in FSC task the draws from the dataset are in general not i.i.d. and thus may correspond to co-linear optimal predictors. In every such case, the condition number is expected to remain non-decreasing, as illustrated in Fig. 1 (*left*) where for MAML, contrary to PROTONET, $\kappa(\mathbf{W}_N)$ exhibits plateaus but also intervals where it is increasing.

This highlights a major difference between the two approaches: MAML does not specifically seek to diversify the learned predictors, while PROTONET does. Proposition 1 gives a hint to the essential difference between the methods studied. On the one hand, PROTONET constructs its classifiers directly from the data of the current task and they are independent of the other tasks. On the other hand, for MAML, the weights of the classifiers are reused between tasks and only slightly adapted to be specific to each task. This limits the generalization capabilities of the linear predictors learned by MAML since they are based on predictors from previous tasks.

3.4 Enforcing the assumptions

Why should we force the assumptions ? From the results obtain by [15], and with the same assumptions, we can easily make appear $\kappa(\mathbf{W}^*)$ to obtain a more explicit bound:

Proposition 2. *If $\forall t \in [[T]], \|\mathbf{w}_t^*\| = O(1)$ and $\kappa(\mathbf{W}^*) = O(1)$, and \mathbf{w}_{T+1} follows a distribution ν such that $\|\mathbb{E}_{\mathbf{w} \sim \nu}[\mathbf{w}\mathbf{w}^\top]\| \leq O(\frac{1}{k})$, then*

$$ER(\hat{\phi}, \hat{\mathbf{w}}_{T+1}) \leq O\left(\frac{C(\Phi)}{n_1 T} \cdot \kappa(\mathbf{W}^*) + \frac{k}{n_2}\right). \quad (9)$$

Proposition 2 suggests that the terms $\|\mathbf{w}_t^*\|$ and $\kappa(\mathbf{W}^*)$ underlying the assumptions directly impact the tightness of the established bound on the excess risk. The full proof can be found in the Appendix.

Can we force the assumptions ? According to the previous result, satisfying the assumptions from MTR theory is expected to come in hand with better performance. However all the terms involved refer to optimal predictors, that we cannot act upon. Thus, we aim to answer the following question:

Given \mathbf{W}^ such that $\kappa(\mathbf{W}^*) \gg 1$, can we learn $\widehat{\mathbf{W}}$ with $\kappa(\widehat{\mathbf{W}}) \approx 1$ while solving the underlying classification problems equally well?*

While obtaining such a result for any distribution seems to be very hard in the considered learning setup, we provide a constructive proof for the existence of a distribution for which the answer to the above-mentioned question is positive in the case of two tasks. The latter restriction comes out of the necessity to analytically calculate the singular values of \mathbf{W} but we expect our example to generalize to more general setups and a larger number of tasks as well.

Proposition 3. *Let $T = 2$, $\mathbb{X} \subseteq \mathbb{R}^d$ be the input space and $\mathbb{Y} = \{-1, 1\}$ be the output space. Then, there exist distributions μ_1 and μ_2 over $\mathbb{X} \times \mathbb{Y}$, representations $\hat{\phi} \neq \phi^*$ and matrices of predictors $\widehat{\mathbf{W}} \neq \mathbf{W}^*$ that satisfy the data generating model (Eq. (3)) with $\kappa(\widehat{\mathbf{W}}) \approx 1$ and $\kappa(\mathbf{W}^*) \gg 1$.*

See Appendix for full proof and illustration. The established results show that even when \mathbf{W}^* does not satisfy Assumptions 1-2 in the ϕ^* space, it may still be possible to learn $\hat{\phi}$ such that the optimal predictors do satisfy them.

How to force the assumptions ? This can be done either by considering the constrained problem (Eq. (10)) or by using a more common strategy that consists in adding $\kappa(\mathbf{W})$ and $\|\mathbf{W}\|_F^2$ as regularization terms (Eq. (11)):

$$\hat{\phi}, \widehat{\mathbf{W}} = \arg \min_{\phi, \mathbf{W}} \frac{1}{Tn_1} \sum_{t=1}^T \sum_{i=1}^{n_1} \ell(y_{t,i}, \langle \mathbf{w}_t, \phi(\mathbf{x}_{t,i}) \rangle) \text{ s.t. } \kappa(\mathbf{W}) = O(1), \|\mathbf{w}_t\| = 1, \quad (10)$$

$$\hat{\phi}, \widehat{\mathbf{W}} = \arg \min_{\phi, \mathbf{W}} \frac{1}{Tn_1} \sum_{t=1}^T \sum_{i=1}^{n_1} \ell(y_{t,i}, \langle \mathbf{w}_t, \phi(\mathbf{x}_{t,i}) \rangle) + \lambda_1 \kappa(\mathbf{W}) + \lambda_2 \|\mathbf{W}\|_F^2. \quad (11)$$

To the best of our knowledge, such regularization terms based on insights from the advances in MTR theory have never been used in the literature before. We refer the reader to the Appendix for more details about them.

3.5 Positioning with respect to Previous Work

Understanding meta-learning While a complete theory for meta-learning is still lacking, several recent works aim to shed light on phenomena commonly

observed in meta-learning by evaluating different intuitive heuristics. For instance, [37] investigate whether MAML works well due to rapid learning with significant changes in the representations when deployed on target task, or due to feature reuse where the learned representation remains almost intact. In [19], the authors explain the success of meta-learning approaches by their capability to either cluster classes more tightly in feature space (task-specific adaptation approach), or to search for meta-parameters that lie close in weight space to many task-specific minima (full fine-tuning approach). Finally, the effect of the number of shots on the FSC accuracy was studied in [7] for PROTONET. More recently, [51] studied the impact of the permutation of labels when training MAML. Our paper investigates a new aspect of meta-learning that has never been studied before and, unlike [37], [7], [19] and [51], provides a more complete experimental evaluation with the two different approaches of meta-learning.

Normalization Multiple methods in the literature introduce a normalization of their features either to measure cosine similarity instead of Euclidean distance [9, 17, 35] or because of the noticed improvement in their performance [44, 47, 49]. In this work, we proved in Section 3.3 above that for PROTONET prototypes normalization is enough to achieve a good coverage of the embedding space, and we empirically show in Section 4.2 below that it leads to better performance. Since we only normalize the prototypes and not all the features, we do not measure cosine similarity. Moreover, with our Theorem 1, we give explanations through MTR theory regarding the link between the coverage of the representation space and performance.

Common regularization strategies In general, we note that regularization in meta-learning (i) is applied to either the weights of the whole neural network [4, 53], or (ii) the predictions [19, 21] or (iii) is introduced via a prior hypothesis biased regularized empirical risk minimization [12–14, 24, 34]. Contrary to the first group of methods and weight decay approach [23], we do not regularize the whole weight matrix learned by the neural network but the linear predictors of its last layer. While weight decay is used to avoid overfitting by penalizing large magnitudes of weights, our goal is to keep the classification margin unchanged during the training to avoid over-/under-specialization to source tasks. Similarly, spectral normalization proposed by [30] does not affect the condition number κ . Second, we regularize the singular values of the matrix of linear predictors obtained in the last batch of tasks instead of the predictions used by the methods of the second group (*e.g.*, using the theoretic-information quantities in [21]). Finally, the works of the last group are related to the online setting with convex loss functions only and do not specifically target the spectral properties of the learned predictors.

4 Impact of enforcing theoretical assumptions

4.1 Experimental Setup

We consider on three benchmark datasets for FSC, namely: 1) **Omniglot** [25] consisting of 1,623 classes with 20 images/class of size 28×28 ; 2) **miniIma-**

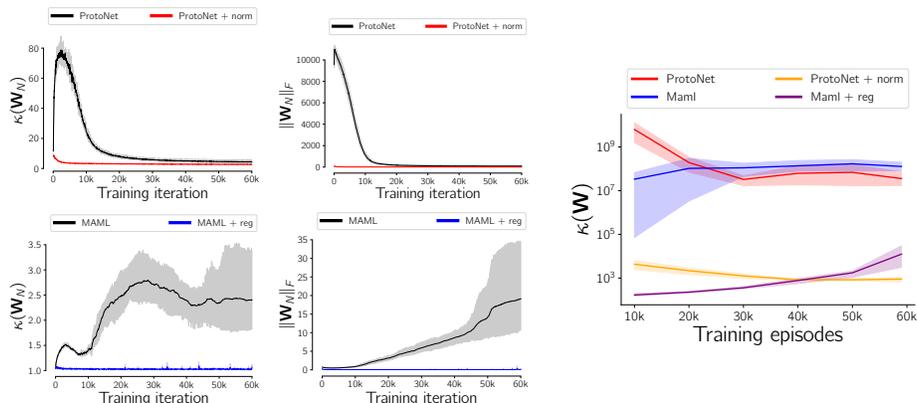


Fig. 2: Evolution of $\kappa(\mathbf{W}_N)$ (left), $\|\mathbf{W}_N\|_F$ (middle) and $\kappa(\mathbf{W})$ (right) on 5-way 1-shot episodes from *miniImageNet*, for PROTONET (top), MAML (bottom) and their regularized or normalized counterparts.

geNet [38] consisting of 100 classes with 600 images of size 84×84 /class; 3) **tieredImageNet** [39] consisting of 779,165 images divided into 608 classes.

For each dataset, we follow a common experimental protocol used in [9, 16] and use a four-layer convolution backbone with 64 filters (C64) as done by [9]. We also provide experiments with the ResNet-12 architecture (R12) [26] and we follow the recent practice to initialize the models with the weights pretrained on the entire meta-training set [36, 40, 52]. For PROTONET, we use the released code of [52] with a temperature of 1 instead of tuning it for all settings to be fair with other settings and methods. For MAML, we follow the exact implementation details from [51] since the code was not available. We measure the performance using the top-1 accuracy with 95% confidence intervals, reproduce the experiments with 4 different random seeds and average the results over 2400 test tasks. For all FSC experiments, unless explicitly stated, we use the regularization parameters $\lambda_1 = \lambda_2 = 1$ in the regularized problem (Eq. (11)). We refer the reader to the Appendix for all the hyperparameters used.⁵

4.2 Metric-based Methods

Theorem 1 tells us that with normalized class prototypes that act as linear predictors, PROTONET naturally decreases the condition number of their matrix. Furthermore, since the prototypes are directly the image features, adding a regularization term on the norm of the prototypes makes the model collapse to the trivial solution which maps all images to 0. To this end, we choose to ensure the theoretical assumptions for metric-based methods (PROTONET and IMP) only with prototype normalization, by using the normalized prototypes $\tilde{\mathbf{w}} = \frac{\mathbf{w}}{\|\mathbf{w}\|}$. Ac-

⁵ Code for the experiments is available at <https://github.com/CEA-LIST/MetaMTReg>.

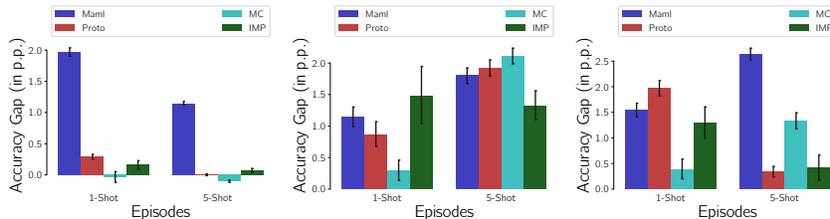


Fig. 3: Accuracy gap (in p.p.) when adding the normalization of prototypes for PROTONET (red) and IMP (green), and both spectral and norm regularization for MAML (blue) and MC (purple) enforcing the theoretical assumptions on Omniglot (left), miniImagenet (middle) and tieredImagenet (right) datasets.

According to Theorem 1, the normalization of the prototypes makes the problem similar to the constrained problem given in Eq. (10).

As can be seen in Fig. 2, the normalization of the prototypes has the intended effect on the condition number of the matrix of predictors. Indeed, $\kappa(\mathbf{W}_N)$ (left) stay constant and low during training, and we achieve a much lower $\kappa(\mathbf{W})$ (right) than without normalization. From Fig. 3, we note that normalizing the prototypes from the very beginning of the training process has an overall positive effect on the obtained performance, and this gain is statistically significant in most of the cases according to the *Wilcoxon signed-rank* test ($p < 0.05$) [11, 50].

In Table 1a, we compare the performance obtained against state-of-the-art algorithms behaving similarly to Instance Embedding algorithms [52] such as PROTONET, depending on the architecture used. Even with a ResNet-12 architecture, the proposed normalization still improves the performance to reach competitive results with the state-of-the-art. On the miniImageNet 5-way 5-shot benchmark, our normalized PROTONET achieves 80.95%, better than DSN (78.83%), CTM (78.63%) and SimpleShot (80.02%). We refer the reader to the Appendix for more detailed training curves.

4.3 Gradient-based Methods

Gradient-based methods learn a batch of linear predictors for each task, and we can directly take them as \mathbf{W}_N to compute its SVD. In the following experiments, we consider the regularized problem of Eq. (11) for MAML as well as Meta-Curvature (MC). As expected, the dynamics of $\|\mathbf{W}_N\|_F$ and $\kappa(\mathbf{W}_N)$ during the training of the regularized methods remain bounded and the effect of the regularization is confirmed with the lower value of $\kappa(\mathbf{W})$ achieved (cf. Fig. 2).

The impact of our regularization on the results is quantified in Fig. 3 where a statistically significant accuracy gain is achieved in most cases, according to the *Wilcoxon signed-rank* test ($p < 0.05$) [11, 50]. In Table 1a, we compare the performance obtained to state-of-the-art gradient-based algorithms. We can see that our proposed regularization is globally improving the results, even with a bigger architecture such as ResNet-12 and with an additional pretraining. On the

Model	Arch.	miniImageNet 5-way		tieredImageNet 5-way	
		1-shot	5-shot	1-shot	5-shot
Gradient-based algorithms					
MAML [16]	C64	48.70 ± 1.84	63.11 ± 0.92	-	-
ANIL [37]	C64	48.0 ± 0.7	62.2 ± 0.5	-	-
Meta-SGD [28]	C64	50.47 ± 1.87	64.03 ± 0.94	-	-
TADAM [32]	R12	58.50 ± 0.30	76.70 ± 0.30	-	-
MC [33]	R12	61.22 ± 0.10	75.92 ± 0.17	66.20 ± 0.10	82.21 ± 0.08
MetaOptNet [26]	R12	62.64 ± 0.61	78.63 ± 0.46	<u>65.99 ± 0.72</u>	<u>81.56 ± 0.53</u>
MATE [10]	R12	62.08 ± 0.64	78.64 ± 0.46	-	-
MAML (Ours)	C64	47.93 ± 0.83	64.47 ± 0.69	50.08 ± 0.91	67.5 ± 0.79
MAML + reg. (Ours)	C64	49.15 ± 0.85	66.43 ± 0.69	51.5 ± 0.9	70.16 ± 0.76
MC (Ours)	C64	49.28 ± 0.83	63.74 ± 0.69	<u>55.16 ± 0.94</u>	71.95 ± 0.77
MC + reg. (Ours)	C64	<u>49.64 ± 0.83</u>	<u>65.67 ± 0.70</u>	55.85 ± 0.94	73.34 ± 0.76
MAML (Ours)	R12	63.52 ± 0.20	81.24 ± 0.14	63.96 ± 0.23	<u>81.79 ± 0.16</u>
MAML + reg. (Ours)	R12	64.04 ± 0.22	82.45 ± 0.14	64.32 ± 0.23	<u>81.28 ± 0.11</u>
Metric-based algorithms					
PROTONET [42]	C64	46.61 ± 0.78	65.77 ± 0.70	-	-
IMP [1]	C64	49.6 ± 0.8	68.1 ± 0.8	-	-
SimpleShot [49]	C64	<u>49.69 ± 0.19</u>	66.92 ± 0.17	51.02 ± 0.20	68.98 ± 0.18
Relation Nets [43]	C64	<u>50.44 ± 0.82</u>	65.32 ± 0.70	-	-
SimpleShot [49]	R18	62.85 ± 0.20	80.02 ± 0.14	69.09 ± 0.22	84.58 ± 0.16
CTM [27]	R18	<u>62.05 ± 0.55</u>	78.63 ± 0.06	64.78 ± 0.11	81.05 ± 0.52
DSN [41]	R12	<u>62.64 ± 0.66</u>	78.83 ± 0.45	66.22 ± 0.75	82.79 ± 0.48
PROTONET (Ours)	C64	49.53 ± 0.41	65.1 ± 0.35	51.95 ± 0.45	71.61 ± 0.38
PROTONET+ norm. (Ours)	C64	<u>50.29 ± 0.41</u>	<u>67.13 ± 0.34</u>	54.05 ± 0.45	71.84 ± 0.38
IMP (Ours)	C64	48.85 ± 0.81	66.43 ± 0.71	52.16 ± 0.89	71.79 ± 0.75
IMP + norm. (Ours)	C64	50.69 ± 0.8	<u>67.29 ± 0.68</u>	<u>53.46 ± 0.89</u>	72.38 ± 0.75
PROTONET (Ours)	R12	59.25 ± 0.20	77.92 ± 0.14	41.39 ± 0.21	83.06 ± 0.16
PROTONET+ norm. (Ours)	R12	62.69 ± 0.20	80.95 ± 0.14	68.44 ± 0.23	84.20 ± 0.16

(a)

Model	Arch.	miniImageNet 5-way		tieredImageNet 5-way	
		1-shot	5-shot	1-shot	5-shot
MTL	R12	55.73 ± 0.18	76.27 ± 0.13	62.49 ± 0.21	81.31 ± 0.15
MTL + norm.	R12	59.49 ± 0.18	77.3 ± 0.13	66.66 ± 0.21	83.59 ± 0.14
MTL + reg. (Ours)	R12	61.12 ± 0.19	<u>76.62 ± 0.13</u>	<u>66.28 ± 0.22</u>	81.68 ± 0.15

(b)

Table 1: Performance comparison on FSC benchmarks. (a) FSC models. (b) MTL models [47]. For a given architecture, **bold** values are the highest accuracy and underlined values are near-highest accuracies (less than 1-point lower).

miniImageNet 5-way 5-shot benchmark, with our regularization MAML achieves 82.45%, better than TADAM (76.70%), MetaOptNet (78.63%) and MATE with MetaOptNet (78.64%). We include in the Appendix ablative studies on the effect

of each term in our regularization scheme for gradient-based methods, and more detailed training curves.

4.4 Multi-Task learning Methods

We implement our regularization on a recent Multi-Task Learning (MTL) method [47], following the same experimental protocol. The objective is to empirically validate our analysis on a method using the MTR framework. As mentioned in Sec. 3.5, the authors introduce feature normalization in their method, speculating that it improves coverage of the representation space [48]. Using their code, we reproduce their experiments on three different settings compared in Table 1b: the vanilla MTL, the MTL with feature normalization, and MTL with our proposed regularization on the condition number and the norm of the linear predictors. We use $\lambda_1 = 1$ in all the settings, and $\lambda_2 = 1$ in the 1-shot setting and $\lambda_2 = 0.01$ in the 5-shot settings. We include in the Appendix an ablative study on the effect of each term of the regularization. Our regularization, as well as the normalization, globally improve the performance over the non-normalized models. Notably, our regularization is the most effective when there is the less data which is well-aligned with the MTR theory in few-shot setting. We can also note that in most of the cases, the normalized models and the regularized ones achieve similar results, hinting that they may have a similar effect. All of these results show that our analysis and our proposed regularization are also valid in the MTL framework.

5 Conclusion

In this paper, we studied the validity of the theoretical assumptions made in recent papers of Multi-Task Representation Learning theory when applied to popular metric- and gradient-based meta-learning algorithms. We found a striking difference in their behavior and provided both theoretical and experimental arguments explaining that metric-based methods satisfy the considered assumptions, while gradient-based don't. We further used this as a starting point to implement a regularization strategy ensuring these assumptions and observed that it leads to faster learning and better generalization.

While this paper proposes an initial approach to bridging the gap between theory and practice for Meta-Learning, some questions remain open on the inner workings of these algorithms. In particular, being able to take better advantage of the particularities of the training tasks during meta-training could help improve the effectiveness of these approaches. The similarity between the source and test tasks was not taken into account in this work, which is an additional assumption in the theory of [15]. We provide a preliminary study using different datasets between the meta-training and meta-testing in the Appendix to foster future work on this topic. Self-supervised meta-learning and multiple target tasks prediction are also important future perspectives for the application of meta-learning.

Acknowledgements This work was made possible by the use of the Factory-AI supercomputer, financially supported by the Ile-de-France Regional Council.

References

1. Allen, K., Shelhamer, E., Shin, H., Tenenbaum, J.: Infinite mixture prototypes for few-shot learning. In: Proceedings of the 36th International Conference on Machine Learning (2019)
2. Amit, R., Meir, R.: Meta-learning by adjusting priors based on extended PAC-Bayes theory. In: International Conference on Machine Learning (2018)
3. Arnold, S., Iqbal, S., Sha, F.: When MAML can adapt fast and how to assist when it cannot. In: AISTATS (2021)
4. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: MetaReg: Towards domain generalization using meta-regularization. In: Advances in Neural Information Processing Systems. pp. 998–1008 (2018)
5. Baxter, J.: A model of inductive bias learning. *Journal of Artificial Intelligence Research* (2000)
6. Bertinetto, L., Henriques, J.F., Torr, P., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. In: International Conference on Learning Representations (2018)
7. Cao, T., Law, M.T., Fidler, S.: A theoretical analysis of the number of shots in few-shot learning. In: ICLR (2020)
8. Chen, J., Wu, X.M., Li, Y., Li, Q., Zhan, L.M., Chung, F.I.: A closer look at the training strategy for modern meta-learning. *Advances in Neural Information Processing Systems* (2020)
9. Chen, W.Y., Wang, Y.C.F., Liu, Y.C., Kira, Z., Huang, J.B.: A closer look at few-shot classification. In: ICLR (2019)
10. Chen, X., Wang, Z., Tang, S., Muandet, K.: MATE: Plugging in model awareness to task embedding for meta learning. In: Advances in neural information processing systems (2020)
11. David, F.N., Siegel, S.: *Nonparametric statistics for the behavioral sciences*. *Biometrika* (1956)
12. Denevi, G., Ciliberto, C., Grazi, R., Pontil, M.: Learning-to-learn stochastic gradient descent with biased regularization. In: International Conference on Machine Learning (2019)
13. Denevi, G., Ciliberto, C., Stamos, D., Pontil, M.: Incremental learning-to-learn with statistical guarantees. In: Conference on Uncertainty in Artificial Intelligence. pp. 457–466 (2018)
14. Denevi, G., Ciliberto, C., Stamos, D., Pontil, M.: Learning to learn around a common mean. In: Advances in Neural Information Processing Systems. pp. 10169–10179 (2018)
15. Du, S.S., Hu, W., Kakade, S.M., Lee, J.D., Lei, Q.: Few-shot learning via learning the representation, provably. In: International Conference on Learning Representations (2020)
16. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning (2017)
17. Gidaris, S., Komodakis, N.: Dynamic Few-Shot Visual Learning Without Forgetting. In: CVPR. pp. 4367–4375 (2018)
18. Goldberger, J., Hinton, G.E., Roweis, S., Salakhutdinov, R.R.: Neighbourhood components analysis. In: Saul, L., Weiss, Y., Bottou, L. (eds.) *NeurIPS* (2005)
19. Goldblum, M., Reich, S., Fowl, L., Ni, R., Cherepanova, V., Goldstein, T.: Unraveling meta-learning: Understanding feature representations for few-shot tasks. In: International Conference on Machine Learning (2020)

20. Hoffer, E., Ailon, N.: Deep metric learning using triplet network. In: Feragen, A., Pelillo, M., Loog, M. (eds.) *Similarity-Based Pattern Recognition*. pp. 84–92. *Lecture Notes in Computer Science*, Springer International Publishing, Cham (2015)
21. Jamal, M.A., Qi, G.J.: Task agnostic meta-learning for few-shot learning. In: *CVPR* (2019)
22. Koch, G., Zemel, R., Salakhutdinov, R.: Siamese neural networks for one-shot image recognition. In: *International Conference on Machine Learning deep learning workshop* (2015)
23. Krogh, A., Hertz, J.A.: A simple weight decay can improve generalization. In: *Advances in Neural Information Processing Systems* (1992)
24. Kuzborskij, I., Orabona, F.: Fast rates by transferring from auxiliary hypotheses. *Machine Learning* (2017)
25. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. *Science* (2015)
26. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: *CVPR* (2019)
27. Li, H., Eigen, D., Dodge, S., Zeiler, M., Wang, X.: Finding Task-Relevant Features for Few-Shot Learning by Category Traversal. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
28. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-SGD: Learning to Learn Quickly for Few-Shot Learning. [arXiv:1707.09835](https://arxiv.org/abs/1707.09835) (2017)
29. Maurer, A., Pontil, M., Romera-Paredes, B.: The benefit of multitask representation learning. *Journal of Machine Learning Research* **17**, 81:1–81:32 (2016)
30. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: *ICLR* (2018)
31. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. [arXiv:1803.02999 \[cs\]](https://arxiv.org/abs/1803.02999) (2018)
32. Oreshkin, B., Rodríguez López, P., Lacoste, A.: TADAM: Task dependent adaptive metric for improved few-shot learning. In: *Advances in neural information processing systems* (2018)
33. Park, E., Oliva, J.B.: Meta-curvature. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *NeurIPS*. pp. 3309–3319 (2019)
34. Pentina, A., Lampert, C.H.: A PAC-Bayesian bound for lifelong learning. In: *International Conference on Machine Learning* (2014)
35. Qi, H., Brown, M., Lowe, D.G.: Low-Shot Learning with Imprinted Weights. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Salt Lake City, UT (2018)
36. Qiao, S., Liu, C., Shen, W., Yuille, A.L.: Few-shot image recognition by predicting parameters from activations. In: *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7229–7238 (2018)
37. Raghu, A., Raghu, M., Bengio, S., Vinyals, O.: Rapid learning or feature reuse? Towards understanding the effectiveness of MAML. In: *ICLR* (2020)
38. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: *ICLR* (2017)
39. Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: *ICLR* (2018)

40. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: International Conference on Learning Representations (2018)
41. Simon, C., Koniusz, P., Nock, R., Harandi, M.: Adaptive Subspaces for Few-Shot Learning. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
42. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. In: Advances in Neural Information Processing Systems (2017)
43. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR. pp. 1199–1208 (2018)
44. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J.B., Isola, P.: Rethinking few-shot image classification: a good embedding is all you need? In: European Conference on Computer Vision. Springer (2020)
45. Tripuraneni, N., Jin, C., Jordan, M.: Provable meta-learning of linear representations. In: International Conference on Machine Learning. PMLR (2021)
46. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: Advances in Neural Information Processing Systems. pp. 3630–3638 (2016)
47. Wang, H., Zhao, H., Li, B.: Bridging multi-task learning and meta-learning: Towards efficient training and effective adaptation. In: Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, PMLR (2021)
48. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: ICML. pp. 9929–9939. Proceedings of machine learning research (2020)
49. Wang, Y., Chao, W.L., Weinberger, K.Q., van der Maaten, L.: SimpleShot: Revisiting Nearest-Neighbor Classification for Few-Shot Learning. arXiv:1911.04623 (2019)
50. Wilcoxon, F.: Individual comparisons by ranking methods. Biometrics Bulletin (1945)
51. Ye, H.J., Chao, W.L.: How to train your MAML to excel in few-shot classification. In: International Conference on Learning Representations (2022)
52. Ye, H.J., Hu, H., Zhan, D.C., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: Computer vision and pattern recognition (CVPR) (2020)
53. Yin, M., Tucker, G., Zhou, M., Levine, S., Finn, C.: Meta-learning without memorization. In: ICLR (2020)