TOWARDS EFFICIENT GRADIENT-BASED META-LEARNING IN HETEROGENOUS ENVIRONMENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

A challenging problem for machine learning is few-shot learning, as its models usually require many training samples. Since meta-learning models have strong fine-tuning capabilities for the distribution of tasks, many of them have been applied to few-shot learning. Model-agnostic meta-learning (MAML) is one of the most popular ones. Recent studies showed that MAML-trained models tend to reuse learned features and do not perform strong adaption, especially in the earlier layers. This paper presents an in-detail analysis of this phenomenon by analyzing MAML's components of different variants. Our results show an interesting relationship between the importance of fine-tuning earlier layers and the difference in the distribution between training and testing. As a result, we determine a fundamental weakness of existing MAML variants when the task distribution is heterogeneous, e.g., the numbers of classes do not match during testing and training. We propose a novel nonparametric version of MAML that overcomes these issues while still being able to perform cross-domain adaption.

1 INTRODUCTION

Learning tasks from only a few observations is known as *few-shot learning* and of major interest in the machine learning community (Finn et al., 2017; Vinyals et al., 2016; Snell et al., 2017; Cai & Shen, 2020; Tseng et al., 2020). Usually, a problem is solved by minimizing the empirical risk with many training samples in many iterations. However, humans learn new tasks very quickly by using knowledge obtained in their lives a priori (Salakhutdinov et al., 2012). Meta-learning is motivated by how humans learn, where the goal is to learn to learn and a common approach to solve few-shot learning tasks.

Model-Agnostic Meta-Learning (MAML) has been one of the most successful meta-learning algorithms for few-shot learning in recent years (Finn et al., 2017). In MAML, the network is meta-optimized to be fine-tuned on an unseen task quickly by gradient descent. Its formulation of the meta-learning objective inspired a plethora of research (Yoon et al., 2018; Li et al., 2017; Vuorio et al., 2019; Finn et al., 2018), to the extent that MAML exists both as a concrete meta-learning algorithm but also as a paradigm that influences meta-learning methods to this day.

Previous work has discussed whether MAML allows rapid fine-tuning or simply reuses features effectively. Raghu et al. (2020) found out that freezing the first part of a network during fine-tuning improves the performance and indicates that MAML rather reuses features. Oh et al. (2021) discovered that in the case of a cross-domain adaptation, a change in the earlier layers is beneficial and proposed to fix the last layer to enforce earlier weight change. However, its fixed final layer is impractical when the numbers of classes differ across tasks, which is a considerable limitation in real-world scenarios.

In this paper, we develop a novel technique called NP-MAML, which has a nonparametric head but is still trainable via gradients. Similar to BOIL, NP-MAML enforces changes in earlier layers to solve cross-domain tasks. In addition, it is flexible to the heterogeneous task distribution. We compare these methods' performance and representation change under different challenges: cross-domain adaption and task dimensionality. We further deliver an analysis of the different components of the network and want to specify the role of all different components.

2 META-LEARNING FOR FEW-SHOT LEARNING

In few-shot learning, a dataset consists of tasks coming from a task distribution. Each task contains of a support set S with labels samples and a query set Q where predictions have to be made on the samples (Vinyals et al., 2016). A typical support set consists of K examples for each of N classes, wherefore a problem is usually described as N-*way*-K-*shot* classification problem.

Next to non-episodic approaches (Gidaris & Komodakis, 2018; Qi et al., 2018; Chen et al., 2019), many meta-learning methods have been applied to few-shot learning problems (Vinyals et al., 2016; Li et al., 2017; Yoon et al., 2018; Snell et al., 2017). They are particularly useful as they can learn a configuration from which it is easy to solve various tasks. The configuration has seen similar features or recurring patterns shared across the tasks and allowed to be transferred to novel unseen tasks.

2.1 OPTIMIZATION-BASED META-LEARNING

Optimization-based meta-learning (Ravi & Larochelle, 2017) follows the idea of meta-learning a task-specific optimizer U, that transforms some initial set of parameters θ into task-parameters ϕ. Although U can be chosen arbitrarily, it is typically modelled as a m-step, gradient-based update scheme, denoted $U^{(m)}(\theta)$. While methods like Meta-LSTMs model $U^{(m)}(\theta)$ explicitly via an LSTM (Ravi & Larochelle, 2017), which iteratively transforms θ, given both loss and loss gradient, the popularity and success of MAML is due to its efficient and model-agnostic design, allowing $U^{(m)}$ to take any differentiable form and rather meta-optimizing θ, the initial parameters, leading to superior performance and more flexibility.

2.2 MODEL-AGNOSTIC META-LEARNING

MAML (Finn et al., 2017) introduces a couple of new interpretations of the meta-learning problem that allow its efficient and flexible training. Firstly, task-data is assumed to come from a *task distribution* $p(\tau)$, which in turn will allow us to form an expectation over task performance. Secondly, let $L_\tau(\mathbf{x}; \theta^*)$ denote the loss of a model f_θ on task-data $\mathbf{x} = \{(x_1, y_1), ..., (x_T, y_T)\}$, parameterized by θ^*. To stay consistent with the notation introduced hitherto, we will denote with S_τ the support set of task τ and with Q_τ the query set of task τ. Then, we can express the optimization objective of MAML as

$$\min_\theta \mathbb{E}_{\tau \sim p(\tau)} \Big[L_\tau(Q_\tau; U^{(m)}(S_\tau; \theta)) \Big], \tag{1}$$

where we write $U^{(m)}(S_\tau; \theta)$ to denote a m-step optimizer, transforming meta-parameters θ given the support set S_τ. Intuitively, MAML improves on-task performance on average by optimizing the initial parameters of model f_θ and subsequently fine-tuning those parameters on the support set S_τ, where on-task performance is measured by evaluating the fine-tuned model on the query set Q_τ.

As $p(\tau)$ is high-dimensional and typically unknown, computing the actual expectation integral is not feasible, which is why we define the meta-loss of MAML as

$$\mathcal{L}(\theta) = \frac{1}{|\mathcal{T}|} \sum_{\tau \in \mathcal{T}} L_\tau(Q_\tau; U^{(m)}(S_\tau; \theta)), \tag{2}$$

where $\mathcal{T}$ is a batch of tasks, sampled from $p(\tau)$, and where we replace the expectation with an empirical mean. This meta-loss is then optimized with standard gradient descent with step-size β, i.e.,

$$\theta^{(t)} = \theta^{(t-1)} - \beta \nabla_{\theta^{(t-1)}} \mathcal{L}(\theta^{(t-1)}). \tag{3}$$

A crucial detail of MAML is that the gradient descent update in Equation 3 involves second-order derivatives if task optimizer $U^{(m)}$ is gradient-based. The authors thus propose a first-order approximation to the update that performs just as well as the full-gradient version of MAML.

2.2.1 RAPID LEARNING OR FEATURE REUSE

After investigating the adaptation behavior of MAML, Raghu et al. (2020) claimed that MAML rather allows rapidly fine-tuning than simply reuses features effectively. They analyzed the representation similarity of the layers and observed hardly any changes in the early layers. Their proposed technique ANIL (almost no inner loop) slightly outperforms MAML on most of several tasks like MiniImageNet and is additionally faster. Oh et al. (2021) proposed BOIL (body only inner loop), a complementary approach by freezing the last layer, showing that in the case of (cross-)domain adaption, a change of weights in the earlier layers, especially the penultimate layer, is beneficial.

2.2.2 MODULARIZATION OF MAML

To accumulate the perspective of previous research and ours, we formulate a *modular interpretation* of MAML. Raghu et al. (2020) and Oh et al. (2021) have found out that fine-tuning different components of the architecture effect the runtime and the overall performance. They analyze the effects of MAML by doing experiments on image datasets such as MiniImageNet with an architecture consisting of four convolutional blocks and a linear layer, which we further call *conv4*. We will later discuss each component of the conv4 architecture in the context of the algorithmic and architectural challenges.

Let f_θ be a model parameterized by θ. We modularize f_θ by partitioning the f_θ into sub-components, e.g., a single layer, a group of layers that form one modular unit, or an entire sub-network. We call $\mathbf{g} = \{g^{(1)}(\cdot, \theta_1), ..., g^{(n)}(\cdot, \theta_n)\}$ a modularization of size n of f_θ if and only if

$$f_\theta = g^{(n)}(\cdot, \theta_n) \circ ... \circ g^{(1)}(\cdot, \theta_1) \tag{4}$$

and

$$\theta = \theta_n \cup ... \cup \theta_1. \tag{5}$$

Furthermore, we denote with $\mathbf{g}^* \subseteq \mathbf{g}$ the components that receive task-specific fine-tuning and call $\mathbf{g}^*$ a *meta-learning configuration*. Then, the meta-update according to MAML for parameter θ_i becomes

$$\theta_i = \theta_i - \beta \cdot \nabla_{\theta_i} L(Q; \phi_i), \tag{6}$$

with

$$\phi_i = \begin{cases} U^{(m)}(S; \theta_i), & \text{if } g^{(i)} \in \mathbf{g}^* \\ \theta_i, & \text{otherwise,} \end{cases} \tag{7}$$

where S is the support set and $U^{(m)}$ is an optimizer with m steps. We can recover the original MAML by setting $\mathbf{g}^* = \mathbf{g}$.

We study the advantages and challenges of different configurations of $\mathbf{g}^*$. We partition the conv4 network into three components: $g_{\textbf{early}}$ denotes the first three convolutional blocks, $\theta_{\textbf{penult}}$ the last convolutional block (penultimate layer) and with $\theta_{\textbf{penult}}$ the last linear layer. Previously proposed methods can be obtained by only optimizing the weights of selected sub-components (see Table 5 in Appendix A.

3 CHALLENGES OF FEW-SHOT LEARNING

Standard few-shot learning benchmark datasets like Omniglot or MiniImageNet are single-modal, homogenous, and balanced, whereas real-world datasets are multi-modal, heterogeneous, and imbalanced (Vuorio et al., 2019; Triantafillou et al., 2020). The data can come from different domains (multi-modal) and have a flexible number of shots (imbalanced) and ways (heterogeneous). We divide the challenges into two groups and distinguish between *algorithmic* and *architectural* challenges[1] and discuss how the methods from Section 2.2.1 can deal with that.

[1]architectural challenges take the form of simple, physical limitations like mismatches of in- and output dimensions and algorithmic refers to problems with performance, lack of generalization or instability of the results (such as very high variance on the classification accuracy)

3.1 ALGORITHMIC CHALLENGE: DOMAIN SHIFT

Domain shift is as a straightforward, but still a challenging problem of cross-domain adaptation (Oh et al., 2021; Triantafillou et al., 2020; Cai & Shen, 2020; Tseng et al., 2020). Our model is meta-trained with a task from a domain A and is meta-tested with the task from domain B, which has not been available during training. Oh et al. (2021) showed that BOIL outperforms MAML and ANIL in this setup, as they learn the general features for domain A and are not able to adapt these features for the different domain B. Thus, our proposed technique in Section 4 aims at ensuring the cross-domain adaptation capabilities of BOIL. A more detailed discussion on domain adaption can be found in Appendix B.

3.2 ARCHITECTURAL CHALLENGE: TASK DIMENSIONALITY

Under *task dimensionality*, we summarize the dimensionality of both the samples and the labels within a task. Whereas it is always possible to scale images to a common $n \times m$ pixel grid[2], for the output dimensionality of a task (ways), there is no trivial modification besides re-initialization. This limitation is caused by the fully-connected linear layer whose weights need a fixed output size to be transferred. The head of the network turning a hidden representation into logits and probabilities is typically expressed as

$$p(y \mid x) = \sigma\Big(W^T \cdot g_{\textbf{body}}(x; \theta) + b\Big), \tag{8}$$

where x is an input sample, parameters $W \in \mathbb{R}^{T \times N}$ and $b \in \mathbb{R}^N$ form the linear layer, $g_{\textbf{body}} = g_{\textbf{penult}} \circ g_{\textbf{early}}$, T is the output dimension of the network body and σ is the softmax operation. Since T has to be fixed, a specific architecture is limited because it cannot adapt distributions with varying output dimensions. A parametric solution of this is proposed by Triantafillou et al. (2020) setting $W_k = 2c_k$ and $b_k = -||c_k||^2$ in Equation 8 which they call Proto-MAML where c_k are the prototypes of all samples having the specific class in the support set. The parameters can be fine-tuned in the inner loop.

4 NONPARAMETRIC MODEL-AGNOSTIC META-LEARNING

In this section, we propose a novel technique, nonparametric model-agnostic meta-learning (NP-MAML), which can deal with both algorithmic and architectural challenges of Section 3.

Raghu et al. (2020) proposed a nonparametric head for MAML and suggested no updates in the inner loop (NIL := No Inner Loop) during meta-testing. Utilizing metric-based approaches of Vinyals et al. (2016) and Snell et al. (2017) in NIL, a similarity metric d is applied to the representation $g_{\textbf{body}}(\hat{x})$ and the prototype $c_i := \frac{1}{K} \sum_{x \in S_i} g_{\textbf{body}}(x)$ for every class label of the support set S

$$p(\hat{y} \mid \hat{x}, S) = \sum_{i=1}^{N} l(-d(g_{\textbf{body}}(\hat{x}), c_i)) \cdot y_i \tag{9}$$

with $k(\cdot)$ as a kernel normalizing the distance measures into a probability. Unfortunately, this approach does not allow fine-tuning the weights of earlier layers in the network. The representation of the penultimate layer needs to be meaningful. Applying a NIL head to the BOIL-trained model is synthetically possible but leads to non-meaningful results, as it is trained to receive updates on the first part. Therefore, it is impossible to benefit from BOIL's cross-domain ability.

To overcome this problem, we propose a new method, nonparametric model-agnostic meta-learning (NP-MAML), which replaces the parametric head of MAML with a nonparametric head also for training. It allows updating the earlier layers during fine-tuning and can be incorporated during meta-optimization. Since the nonparametric head does not require a fixed output size, it can deal with architectural challenges. Yet, it addresses the problem of cross-domain adaptation, as we enforce fine-tuning the earlier part of the network similar to BOIL.

In NP-MAML, we use the predictive distribution of Equation 9 and extend it such that it allows gradients to flow through the fine-tuning stage:

[2]Although we will not study the effect of rescaling in this work, we encourage a further investigation into whether rescaling has an impact on the predictive performance of meta-learning models for image processing.

$$p(\hat{y} \mid \hat{x}, S) = \sum_{i=1}^{N} k(-d_{\textbf{query}}(g_{\textbf{body}}(\hat{x}; U(S;\theta)) - c_i)) \cdot y_i, \tag{10}$$

where U is a special one-step optimizer that is similar to a contrastive loss (Andonian et al., 2021)

$$U(S;\theta) = \theta - \frac{\alpha}{K} \nabla_\theta \Big(\sum_{i,j} -d_{\textbf{support}}(g_{\textbf{body}}(S_i;\theta) - c_j) \Big). \tag{11}$$

Features from different classes should be less similar than features from the same class.

The prototypes c_k also depend on θ, but we stop the gradient flow for the prototype when calculating the meta-gradient, achieving similar runtime performance as the other methods. Although this is not a first-order approximation of the meta-gradient, this simplification is based on previous evidence that for successful meta-learning, not the entire meta-gradient is required (Finn et al., 2017). A visualization of the fine-tuning effect on the network's features can be found in Appendix C.

5 Experiments

In this section, we show the results of several experiments, which empirically support that NP-MAML can deal with heterogeneous challenges while ensuring cross-domain adaption and outperforming previous adaptations of MAML. In addition, we analyze the different components of MAML and show how each part of the network contributes. All models are trained in a *homogeneous* environment and afterward tested in homogenous and heterogeneous environments. This allows us to have more fine-grained control of cause and effect relationships with regard to predictive performance than we would have if we applied it on Meta-Dataset Triantafillou et al. (2020). This setting is also most realistic for real-world applications, where pre-trained models are often used to solve novel problems. Nevertheless, we regard a validation of our findings on Meta-Dataset as the next step.

We design the experiment in a way that the effects of individual parts of the environment (domain, number of ways, and shots) are investigated partially isolated. We train the conv4 architecture, the standard architecture in previous literature on few-shot learning with MAML (Finn et al., 2017; Raghu et al., 2020; Oh et al., 2021; Vuorio et al., 2019; Yoon et al., 2018; Finn et al., 2018) on MiniImageNet to solve a N-way-K-shot classification task. We model domain shift by testing the FC100 dataset (Oreshkin et al., 2018) and measure the test accuracy. We choose the hyperparameters of the original MAML-approaches (Finn et al., 2017; Raghu et al., 2020; Oh et al., 2021) summarized in Appendix D. For all performance results, we show classification accuracy on a hold-out set, where the $\pm$ indicates the 95% confidence interval over tasks, following Finn et al. (2017).

5.1 Domain Shift

In our first experiment, we train each model on 5-way-1-shot and 5-way-5-shot on MiniImageNet and test it on MiniImageNet (in-domain) and FC100 (cross-domain), where the numbers of shots K matches its training conditions.

In Table 1, we observe that our approach NP-MAML is competitive with BOIL on 1-shot MiniImageNet and better in all other cases. As expected, ANIL is especially bad in the cross-domain setup.

5.2 Meta-Feature Quality when Varying Number of Ways

In this experiment, we look at the quality of the features of the penultimate layer, which we call *meta-features*. The model is trained like in Section 5.1 on MiniImageNet on the homogeneous 5-way-K-shot problems. We test the model on both MiniImageNet and FC100 but vary the number of ways of each task ($N = 4, 7$ and 10). Since MAML, ANIL and BOIL can architecturally not fine-tune the last layer (see Section 4, we use the NIL-testing proposed by Raghu et al. (2020). NP-MAML can be fine-tuned, and we depict both results without fine-tuning (0 inner loop steps) and with fine-tuning NP-MAML (1 inner loop step).

	FC100			
	1-shot	5-shot	1-shot	5-shot
ANIL	47.69 ± 0.62	62.58 ± 0.54	32.83 ± 0.53	43.20 ± 0.53
BOIL	$\mathbf{50.16 \pm 0.64}$	65.31 ± 0.53	36.03 ± 0.57	47.83 ± 0.52
MAML	45.42 ± 0.61	61.84 ± 0.55	34.23 ± 0.55	44.50 ± 0.54
NP-MAML	49.82 ± 0.64	$\mathbf{67.39 \pm 0.53}$	$\mathbf{38.57 \pm 0.57}$	$\mathbf{50.86 \pm 0.54}$

Table 1: In this table, the **classification accuracy [%]** of all models on a test set of MiniImageNet and FC100 are depicted. All models have been trained on MiniImageNet.

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
ANIL	55.29 ± 0.77	49.64 ± 0.63	40.78 ± 0.49	32.96 ± 0.36	39.03 ± 0.67	33.61 ± 0.53	26.39 ± 0.42	20.62 ± 0.31
BOIL	30.98 ± 0.45	25.90 ± 0.38	19.17 ± 0.28	14.47 ± 0.21	29.68 ± 0.45	24.50 ± 0.36	18.31 ± 0.28	13.48 ± 0.21
MAML	52.60 ± 0.73	47.11 ± 0.63	38.61 ± 0.47	30.79 ± 0.34	40.84 ± 0.69	35.09 ± 0.57	27.22 ± 0.44	21.23 ± 0.30
NP-MAML (0)	49.18 ± 0.70	43.44 ± 0.60	35.20 ± 0.44	28.38 ± 0.34	40.24 ± 0.64	34.76 ± 0.52	27.38 ± 0.40	21.24 ± 0.31
NP-MAML (1)	$\mathbf{55.50 \pm 0.75}$	$\mathbf{49.82 \pm 0.64}$	$\mathbf{41.48 \pm 0.47}$	$\mathbf{33.59 \pm 0.36}$	$\mathbf{44.09 \pm 0.67}$	$\mathbf{38.57 \pm 0.57}$	$\mathbf{30.37 \pm 0.43}$	$\mathbf{23.83 \pm 0.30}$

Table 2: The **classification accuracy** [%] of the models on 1-shot in-domain for in-domain (MiniImageNet) and cross-domain (FC100) tasks, where the number of ways varies during test time. All models have been meta-trained on MiniImageNet.

In Table 2 the results for 1-shot in- and cross-domain can be found. In the case of an in-domain adaptation, NP-MAML (fine-tuned) performs the best; however, it is only slightly better than ANIL (the difference is smaller than 1%). This holds true across all varying numbers of ways, which does not seem to have an impact on the performance difference. NP-MAML is, without fine-tuning, significantly worse than ANIL or MAML, but with one fine-tuning step, its performance is improved by $\approx 5\%$. BOIL performs comparatively badly, which is not surprising because it is trained to enforce representation change during fine-tuning, which is not possible and is one of the motivations of our nonparametric approach.

In the case of a cross-domain adaptation, NP-MAML performs again the best, followed by MAML with around about 3.5% across ways. NP-MAML without fine-tuning is only slightly worse than MAML. We observe similar results in the case of 5-shots (Appendix E).

5.3 Performance when Varying Number of Shots

We repeat the experiments from Section 5.1 and 5.2 and additionally interchange the numbers of shots while testing the models (1-shot to 5-shot and vice versa).

In Table 4a, we see the results when K is constant. Especially in the case of a 1-to-5-shot problem, NP-MAML is good and outperforms when we are in a cross-domain setup. Interestingly, ANIL is very bad when we have five shots in training but only 1 in training.

When also the number of shots varies, BOIL is again not able to achieve good results. When increasing the number of shots during testing (see Table 4b), NP-MAML and ANIL are overall the best performers in the case of in-domain, and NP-MAML is the best in the case of cross-domain. The only difference is that fine-tuning in this heterogeneous setup is no longer beneficial for NP-MAML.

When decreasing the number of shots, NP-MAML suffers from the small numbers of samples in the case of in-domain (Table 4c). Interestingly, ANIL is very good in contrast to the fixed setup of Table 4a. We fine-tune the last layer in the fixed case, whereas we use NIL-testing in the mixed setup, which seems to influence the performance in case of varying the numbers of shots. In the case of cross-domain, NP-MAML outperforms the other approaches.

5.4 Effect of Training Early Layers

We want to investigate the role of the earlier layer $g_{\textbf{early}}$, which are in the standard setup of MAML and BOIL fine-tuned. We observe the significance of updating the early layers by comparing the

(a) Switching number of shots with fixed ways

	MiniImageNet		FC100	
	5-to-1-shot	1-to-5-shot	5-to-1-shot	1-to-5-shot
ANIL	27.23 ± 0.42	61.56 ± 0.56	25.05 ± 0.35	41.71 ± 0.56
BOIL	$\mathbf{49.20 \pm 0.63}$	$\mathbf{63.92 \pm 0.54}$	35.51 ± 0.53	46.83 ± 0.52
MAML	43.23 ± 0.60	60.42 ± 0.57	32.37 ± 0.52	45.44 ± 0.55
NP-MAML	45.68 ± 0.58	63.41 ± 0.51	$\mathbf{35.57 \pm 0.52}$	$\mathbf{50.28 \pm 0.53}$

(b) 1-to-5-shot with mixed ways

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
ANIL	68.74 ± 0.61	63.20 ± 0.55	54.96 ± 0.44	46.54 ± 0.34	50.38 ± 0.63	44.64 ± 0.54	36.15 ± 0.43	29.24 ± 0.31
BOIL	38.33 ± 0.50	32.68 ± 0.42	25.34 ± 0.32	19.45 ± 0.23	35.31 ± 0.53	29.84 ± 0.44	23.02 ± 0.33	17.59 ± 0.24
MAML	66.40 ± 0.64	60.81 ± 0.56	52.26 ± 0.43	44.08 ± 0.33	51.12 ± 0.66	45.20 ± 0.55	36.69 ± 0.43	29.53 ± 0.30
NP-MAML (0)	$\mathbf{69.23 \pm 0.61}$	$\mathbf{64.16 \pm 0.54}$	$\mathbf{56.10 \pm 0.45}$	$\mathbf{48.43 \pm 0.35}$	$\mathbf{56.42 \pm 0.64}$	$\mathbf{50.91 \pm 0.54}$	$\mathbf{42.26 \pm 0.44}$	$\mathbf{34.65 \pm 0.32}$
NP-MAML (1)	69.18 ± 0.59	63.41 ± 0.51	54.31 ± 0.42	45.82 ± 0.33	56.18 ± 0.65	50.28 ± 0.53	41.25 ± 0.42	33.40 ± 0.32

(c) 5-to-1-shot with mixed ways

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
ANIL	$\mathbf{55.84 \pm 0.73}$	$\mathbf{50.47 \pm 0.64}$	$\mathbf{42.25 \pm 0.48}$	$\mathbf{34.22 \pm 0.36}$	40.83 ± 0.63	35.34 ± 0.53	27.62 ± 0.40	21.53 ± 0.29
BOIL	31.15 ± 0.44	25.96 ± 0.37	19.44 ± 0.29	14.57 ± 0.21	29.94 ± 0.43	24.45 ± 0.34	18.40 ± 0.27	13.52 ± 0.19
MAML	53.86 ± 0.73	48.34 ± 0.63	39.93 ± 0.46	31.99 ± 0.35	40.49 ± 0.63	34.26 ± 0.53	27.05 ± 0.41	20.95 ± 0.29
NP-MAML (0)	46.52 ± 0.67	41.18 ± 0.59	33.15 ± 0.44	26.35 ± 0.34	39.75 ± 0.59	34.11 ± 0.53	26.98 ± 0.40	21.20 ± 0.31
NP-MAML (1)	50.91 ± 0.68	45.68 ± 0.58	37.31 ± 0.44	30.28 ± 0.35	$\mathbf{41.56 \pm 0.62}$	$\mathbf{35.57 \pm 0.52}$	$\mathbf{27.96 \pm 0.40}$	$\mathbf{21.88 \pm 0.29}$

Table 3: In these tables, the number of shots is switched between training and testing. We see the **classification accuracy** [%] of the models on a test of MiniImageNet and FC100, all trained on MiniImageNet. In (a), the number of ways is fixed, whereas in (b) and (c), we have a flexible number of ways. In (b), the number of shots increases during testing, and in (c) decreases.

predictive performance of BOIL and NP-MAML with only updating the penultimate layer (suffix -4) in the inner loop versus updating all layersIn Table 4, we present the results for 5-way-few-shot and mixed-way-1-shot.

In the case of fixed-way (Table 4a), the penultimate only variant performs similarly for in-domain tasks but is significantly worse for cross-domain tasks (2% for NP-MAML and 4% for BOIL). For the mixed-way setup (Table 4b), it is similar. In in-domain problems, the penultimate only variant is slightly better, but updating all weights in the inner loop is beneficial for the cross-domain setup. For the mixed-way in-domain case (b), NP-MAML-4 consistently performs best across all ways after fine-tuning.However, for the in-domain case, the fine-tuned NP-MAML-4 improves only slightly over the fine-tuned NP-MAML with a performance difference of about 0.5%. As for all mixed-way experiments, BOIL, and by extension BOIL-4, cannot compete with the other methods, and we also observe no notable performance difference between BOIL-4 and BOIL in neither the in- nor the cross-domain case. For the mixed-way cross-domain case (c), the fine-tuned NP-MAML-4 falls behind the fine-tuned NP-MAML by about 3%. In all mixed-way experiments, fine-tuning improves the performance of NP-MAML.

5.5 Representation Similarity Analysis

In our last experiment, we approach the question of Raghu et al. (2020), rapid learning or feature reuse, for NP-MAML. Similar to Raghu et al. (2020); Oh et al. (2021); Goerttler & Obermayer (2021), we measure the centered kernel alignment (CKA) similarity of representations of the query set before and after adaptation for the three components $g_{\textbf{early}}$, $g_{\textbf{penult}}$ and $g_{\textbf{head}}$ of conv4 of all models we compared.

In Figure 1, we see the CKA similarities for different scenarios of our experiments. We can generally observe that the representation of the head changes much less for NP-MAML than for the other ones, meaning a good performance does not require a huge change there. In the case of cross-domain tasks, the adaption for MAML and NP-MAML is larger than in the case of in-domain tasks.

(a) 5-way-few-shot

	MiniImageNet		FC100	
	1-shot	5-shot	1-shot	5-shot
BOIL	50.16 $\pm$ 0.64	65.31 $\pm$ 0.53	36.03 $\pm$ 0.57	47.83 $\pm$ 0.52
BOIL-4	49.61 $\pm$ 0.62	64.15 $\pm$ 0.53	32.97 $\pm$ 0.52	43.48 $\pm$ 0.51
NP-MAML	49.82 $\pm$ 0.64	**67.39 $\pm$ 0.53**	**38.57 $\pm$ 0.57**	**50.86 $\pm$ 0.54**
NP-MAML-4	**50.73 $\pm$ 0.62**	67.13 $\pm$ 0.53	36.12 $\pm$ 0.55	48.40 $\pm$ 0.54

(b) mixed-way-1-shot

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
BOIL	30.98 $\pm$ 0.45	25.90 $\pm$ 0.38	19.17 $\pm$ 0.28	14.47 $\pm$ 0.21	29.68 $\pm$ 0.45	24.50 $\pm$ 0.36	18.31 $\pm$ 0.28	13.48 $\pm$ 0.21
BOIL-4	29.67 $\pm$ 0.40	24.44 $\pm$ 0.33	17.98 $\pm$ 0.25	13.14 $\pm$ 0.18	29.08 $\pm$ 0.41	24.10 $\pm$ 0.35	17.90 $\pm$ 0.27	13.18 $\pm$ 0.20
NP-MAML (0)	49.18 $\pm$ 0.70	43.44 $\pm$ 0.60	35.20 $\pm$ 0.44	28.38 $\pm$ 0.34	40.24 $\pm$ 0.64	34.76 $\pm$ 0.52	27.38 $\pm$ 0.40	21.24 $\pm$ 0.31
NP-MAML (1)	55.50 $\pm$ 0.75	49.82 $\pm$ 0.64	41.48 $\pm$ 0.47	33.59 $\pm$ 0.36	**44.09 $\pm$ 0.67**	**38.57 $\pm$ 0.57**	**30.37 $\pm$ 0.43**	**23.83 $\pm$ 0.30**
NP-MAML-4 (0)	49.26 $\pm$ 0.70	43.11 $\pm$ 0.58	34.74 $\pm$ 0.45	27.69 $\pm$ 0.34	37.33 $\pm$ 0.57	31.80 $\pm$ 0.49	24.63 $\pm$ 0.40	19.22 $\pm$ 0.29
NP-MAML-4 (1)	**56.02 $\pm$ 0.75**	**50.73 $\pm$ 0.62**	**41.98 $\pm$ 0.48**	**34.06 $\pm$ 0.37**	41.74 $\pm$ 0.65	36.12 $\pm$ 0.55	28.44 $\pm$ 0.43	22.45 $\pm$ 0.30

Table 4: **Classification accuracy** [%] of BOIL, BOIL-4, NP-MAML-4 and NP-MAML, each trained on 5-way-1-shot MiniImageNet and evaluated on both fixed-way (a) and mixed-way (b) problems. Since only NP-MAML and NP-MAML-4 can be fine-tuned in the mixed-way setting, NP-MAML and NP-MAML-4 results are suffixed with the number of fine-tuning steps in parentheses.

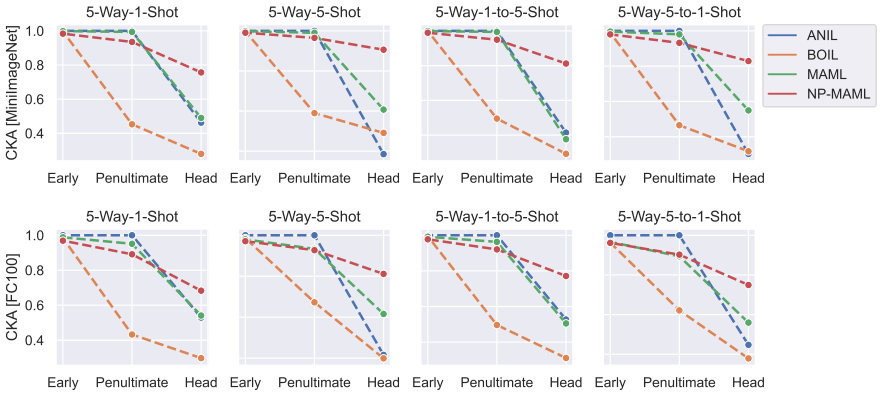

Figure 1: **CKA similarity** of network representations before and after adaptation for the models of different scenarios from previous experiments. Representation similarity was obtained for the three components $g_{\textbf{early}}$, $g_{\textbf{penult}}$ and $g_{\textbf{head}}$ of the conv4 architecture.

For the penultimate layer, the similarity of BOIL is much less than for NP-MAML. Interestingly, all observations do not change significantly when the number of shots changes during training and testing.

6 Discussion

In this section, we discuss all major components of MAML during training and testing under domain adaptation: training and testing (data processing, importance of low-level features and high-level features, network's predictor.

6.1 Task Distribution

The results of Section 5 show that ANIL and MAML perform very well on in-domain tasks, even if the number of shots and ways is different in testing. However, in the case of cross-domain adaption, we see that these algorithms become poor, as their features are not updated to the shift. BOIL,

which is good in cross-domain adaptation, fails in the case of flexible numbers of ways, as it needs a change of an earlier layer in the inner loop, which is not possible architecturally. NP-MAML solves this issue by keeping the strength of BOIL (enforcing representation change) but still can deal with flexible numbers of ways and shots.

6.2 Fine-Tuning

As discussed in Section 3, fine-tuning on heterogeneous test environments is challenging for existing architecture, whereas our approaches NP-MAML and NP-MAML are architecturally able to fine-tune. From the results of Table 5.3b, we see that fine-tuning is not improving the results when increasing the number of shots, which we believe is partially attributed to the class prototypes that average across five samples; however, the model is trained to update on one.

In addition, we observed that NP-MAML reaches their cross-domain performance with just one fine-tuning step (similar to BOIL), whereas ANIL and MAML use 5 and 10 for training and testing, respectively. We think that NP-MAML has many similarities to BOIL, e.g., strong cross-domain adaption or body-only inner learning.

Unless in BOIL, where weights of the head are frozen, the reason to not fine-tune the head in NP-MAML is that there is no real head. But when viewing NP-MAML as a classic linear layer whose dimensionality is implicitly formed by the output (Snell et al., 2017), one could say that our pseudo-head is fine-tuned because the parameters are formed by the network's features, which do change.

6.3 Early Layers

The results of Section 5.4 show that penultimate-only models perform at least as good as full models in in-domain tasks but are significantly worse for cross-domain adaption. The effect of only updating the penultimate layer harms BOIL more than NP-MAML, which indicates a higher resilience of our approach to this effect. Since for NP-MAML, the penultimate version is always better for in-domain tasks, it can be summarized that fine-tuning them is beneficial for cross-domain tasks, but for in-domain tasks, we better leave them not-updated. Especially in the case of larger models (e.g., ResNet), this is important to note, as it speeds up the fine-tuning.

6.4 Penultimate Layer

Oh et al. (2021) have already found out that the penultimate layer is very important for few-shot learning, especially for cross-domain adaptation. Our results from Section 5.1 and 5.3 approve this and show that also in our heterogeneous test environments, this holds true, wherefore there is a demand of our novel approach NP-MAML does not perform well there.

6.5 Predictor (Head)

The choice of the predictor (head) influences the generalization performance the most. In our experiments, we could show the superiority of nonparametric heads. Especially for heterogeneous tasks, it allows for a good adaptation. In addition, NP-MAML is the only model where training and test conditions match precisely. That this is desirable has been discussed by Vinyals et al. (2016).

7 Conclusion

In this paper, we proposed an NP-MAML nonparametric version of MAML, which can deal with architectural and algorithmic challenges. In addition, we explored the key components of MAML. Our experiments' results show that NP-MAML outperforms established, parameterized methods. When there is a distribution shift between meta-training and meta-testing, it beats ANIL and MAML due to representation change also in the earlier layer. In the case of heterogenous task distribution, NP-MAML can still update during fine-tuning, which BOIL is not able to. In our depth analysis of MAML key components, we revealed the hitherto unknown importance of fine-tuning early layers for the fine-tuning task. Our representation similarity analysis showed that NP-MAML adapts less in the network body during the inner loop than BOIL; however, enough to outperform it. Further analysis on that will be an interesting topic for upcoming research.

REFERENCES

Milad Abdollahzadeh, Touba Malekzadeh, and Ngai-Man Cheung. Revisit multimodal meta-learning through the lens of multi-task learning. *CoRR*, abs/2110.14202, 2021. URL https://arxiv.org/abs/2110.14202.

Alex Andonian, Taesung Park, Bryan Russell, Phillip Isola, Jun-Yan Zhu, and Richard Zhang. Contrastive feature loss for image prediction. In *IEEE/CVF International Conference on Computer Vision Workshops, ICCVW 2021, Montreal, BC, Canada, October 11-17, 2021*, pp. 1934–1943. IEEE, 2021. doi: 10.1109/ICCVW54120.2021.00220. URL https://doi.org/10.1109/ICCVW54120.2021.00220.

Sébastien M R Arnold, Praateek Mahajan, Debajyoti Datta, Ian Bunner, and Konstantinos Saitas Zarkias. learn2learn: A library for Meta-Learning research. August 2020. URL http://arxiv.org/abs/2008.12284.

John Cai and Sheng Mei Shen. Cross-domain few-shot learning with meta fine-tuning. *CoRR*, abs/2005.10544, 2020. URL https://arxiv.org/abs/2005.10544.

Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=HkxLXnAcFQ.

Tristan Deleu, Tobias Würfl, Mandana Samiei, Joseph Paul Cohen, and Yoshua Bengio. Torchmeta: A Meta-Learning library for PyTorch, 2019. URL https://arxiv.org/abs/1909.06576. Available at: https://github.com/tristandeleu/pytorch-meta.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135. PMLR, 2017. URL http://proceedings.mlr.press/v70/finn17a.html.

Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 9537–9548, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/8e2c381d4dd04f1c55093f22c59c3a08-Abstract.html.

Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 4367–4375. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00459. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Gidaris_Dynamic_Few-Shot_Visual_CVPR_2018_paper.html.

Thomas Goerttler and Klaus Obermayer. Exploring the similarity of representations in model-agnostic meta-learning. *CoRR*, abs/2105.05757, 2021. URL https://arxiv.org/abs/2105.05757.

Hal Daumé III and Daniel Marcu. Domain adaptation for statistical classifiers. *J. Artif. Intell. Res.*, 26:101–126, 2006. doi: 10.1613/jair.1872. URL https://doi.org/10.1613/jair.1872.

Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few shot learning. *CoRR*, abs/1707.09835, 2017. URL http://arxiv.org/abs/1707.09835.

Vivek Myers and Nikhil Sardana. Bayesian meta-learning through variational gaussian processes. *CoRR*, abs/2110.11044, 2021. URL https://arxiv.org/abs/2110.11044.

Jaehoon Oh, Hyungjun Yoo, ChangHwan Kim, and Se-Young Yun. BOIL: towards representation change for few-shot learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=umIdUL8rMH.

Boris N. Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. TADAM: task dependent adaptive metric for improved few-shot learning. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 719–729, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/66808e327dc79d135ba18e051673d906-Abstract.html.

Hang Qi, Matthew Brown, and David G. Lowe. Low-shot learning with imprinted weights. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 5822–5830. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00610. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Qi_Low-Shot_Learning_With_CVPR_2018_paper.html.

Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of MAML. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=rkgMkCEtPB.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006. ISBN 026218253X. URL https://www.worldcat.org/oclc/61285753.

Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=rJY0-Kcll.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3): 211–252, 2015. doi: 10.1007/s11263-015-0816-y. URL https://doi.org/10.1007/s11263-015-0816-y.

Ruslan Salakhutdinov, Joshua B. Tenenbaum, and Antonio Torralba. One-shot learning with a hierarchical nonparametric bayesian model. In Isabelle Guyon, Gideon Dror, Vincent Lemaire, Graham W. Taylor, and Daniel L. Silver (eds.), *Unsupervised and Transfer Learning - Workshop held at ICML 2011, Bellevue, Washington, USA, July 2, 2011*, volume 27 of *JMLR Proceedings*, pp. 195–206. JMLR.org, 2012. URL http://proceedings.mlr.press/v27/salakhutdinov12a/salakhutdinov12a.pdf.

Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 4077–4087, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/cb8da6767461f2812ae4290eac7cbc42-Abstract.html.

Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=rkgAGAVKPr.

Hung-Yu Tseng, Hsin-Ying Lee, Jia-Bin Huang, and Ming-Hsuan Yang. Cross-domain few-shot classification via learned feature-wise transformation. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=SJl5Np4tPr.

Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3630–3638, 2016. URL https://proceedings.neurips.cc/paper/2016/hash/90e1357833654983612fb05e3ec9148c-Abstract.html.

Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 1–12, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/e4da3b7fbbce2345d7772b0674a318d5-Abstract.html.

Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. Bayesian model-agnostic meta-learning. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 7343–7353, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/e1021d43911ca2c1845910d84f40aeae-Abstract.html.

Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I*, volume 8689 of *Lecture Notes in Computer Science*, pp. 818–833. Springer, 2014. doi: 10.1007/978-3-319-10590-1_53. URL https://doi.org/10.1007/978-3-319-10590-1_53.

Model	Configuration $\mathbf{g}^*$
MAML (Finn et al., 2017)	$\{g_{\textbf{early}}, g_{\textbf{penult}}, g_{\textbf{head}}\}$
ANIL (Raghu et al., 2020)	$\{g_{\textbf{head}}\}$
BOIL (Oh et al., 2021)	$\{g_{\textbf{early}}, g_{\textbf{penult}}\}$
BOIL-4 (Oh et al., 2021)	$\{g_{\textbf{penult}}\}$

Table 5: Meta-learning configurations corresponding to methods from (Finn et al., 2017; Raghu et al., 2020; Oh et al., 2021). BOIL-4 indicates a variant of BOIL where only the penultimate layer of conv4 is meta-learned.

A Instantiation of previous methods of our modularization

Our modularization decomposes the network into different parts Previously proposed methods (Finn et al., 2017; Raghu et al., 2020; Oh et al., 2021) can be obtained by only optimizing the weights of selected sub-components (see Table 5).

B (Cross-)Domain Adaptation and Generalization

The problem of domain adaptation refers to the situation, where (labelled) data in one domain A exists plentiful, while (labelled) data in a different domain B (typically called the *target domain*) is scarce, requiring models to transfer knowledge gained in domain A to make predictions in domain B (Tseng et al., 2020). As stated by III & Marcu (2006), the idea is not to ignore data from domain B entirely but rather to reduce the acquisition effort by transferring from domain A. This is opposed to *domain generalization* where models have no access to any data from B during training. The latter is sometimes also referred to as *cross-domain transfer* (Triantafillou et al., 2020) or *cross-domain adaptation* (Oh et al., 2021).

In this work, we adopt the use of *cross-domain adaptation* as referring to both having little or no access to data from domain B, to stay consistent with the current literature in meta-learning (Triantafillou et al., 2020; Cai & Shen, 2020; Tseng et al., 2020).

B.1 Multi-Modal Meta-Learning

To adapt the above distinction for another knowledge transfer problem setting, in *multi-modal learning*, one presents a model with data from both A and B, hoping that predictive performance improves via knowledge transfer between the domains, that is, for example, reuse of general image features that can be found in both domains (Vuorio et al., 2019; Abdollahzadeh et al., 2021). Yet, knowledge transfer is not limited to such rather obvious commonalities. Vuorio et al. (2019) explicitly model such multi-modal settings with MAML and show performance improvement over a MAML model that treats each task coming from the same domain. Research in this direction has direct implications on mixtures-of-domain datasets, such as Meta-Dataset (Triantafillou et al., 2020).

B.2 Coarse- vs. Fine-Grained Domains

Another important notion in cross-domain adaptation is the notion of *coarse-* and *fine-grained domains* (Triantafillou et al., 2020; Oh et al., 2021), referring to the level of abstractions of the "topics" contained within the data. For example, in the context of image classification, a domain containing classes like dogs, cats, elephants, giraffes is considered coarser than a domain with classes like German Sheppard, Corgi, Husky, Bulldog. As "level of abstraction" is not clearly defined, this notion is typically used only in the distinction of two domains, for example, when performing domain adaptation from a coarse- to a fine-grained domain.

C Investigating the Effect of the Nonparametric Predictor

In Figure 2 we investigate the effect of the nonparametric predictor on the features of the network body visually.

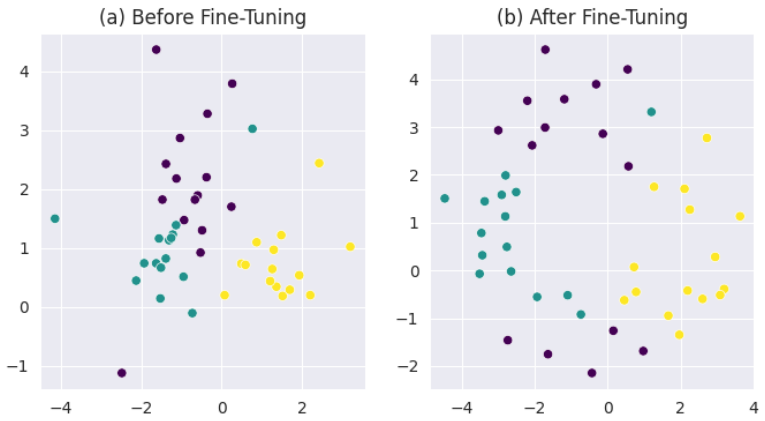

Figure 2: T-SNE plot of the features of a 3-way-1-shot task, produced by a network body trained with an NP-predictor. Features are computed from test samples of a MiniImageNet task, (a) before and (b) after fine-tuning and colored according to their ground-truth label. We observe that the NP-based fine-tuning, as introduced in Equation 11, resolves the cluttered features in the middle of the left plot, pulling features that do not belong to the same tasks apart. The model that produced the above example was trained for 30.000 iterations on 5-way-1-shot MiniImageNet, with the same hyperparameters as for our experiments in Section 5.

D Implementation Details

In this section, we outline the exact implementation details of each of the models we used throughout the experiments in Section 5.

All models were trained and tested on the conv4 architecture (Finn et al., 2017; Raghu et al., 2020; Oh et al., 2021), comprised of four convolutional blocks, each containing a 3×3 convolutional layer with one stride, a 2D Batch-Normalization layer, a ReLU non-linearity and max-pooling layer with two strides. For models with parametric heads (ANIL, MAML, and BOIL), we formed predictions from the convolutional features with a linear, fully-connected layer.

MiniImageNet and FC100 data has been preprocessed with the data-loaders of the torchmeta library (Deleu et al., 2019). Constructing the conv4 architecture, as well as fine-tuning the models, has been aided by the learn2learn framework (Arnold et al., 2020).

D.1 Model Configurations and Hyperparameters

We carefully selected model configurations and their corresponding hyperparameters based on best practices in previous work (Finn et al., 2017; Raghu et al., 2020; Oh et al., 2021). In addition, to run MAML training at a similar time than the other models, we used first-order MAML throughout all of our experiments. In addition, as outlined in Section 2.2.2, we used an approximation to the meta-gradient of NP-MAML and NP-MAML-4 models. We trained and evaluated each model on three different seeds and reported averages across seeds.

Gradient-based meta-learning typically requires the following hyperparameters: A number of training iterations $T_{\textbf{train}}$, a number of test iterations $T_{\textbf{test}}$, a number of fine-tuning steps during training $s_{\textbf{train}}$, a number of fine-tuning steps during test time $s_{\textbf{test}}$, a number of test shots (in the query set) $K_{\textbf{test}}$, an outer learning rate β, an inner learning rate α, a batch-size of the number of tasks sampled each iteration and over which the meta-gradient is formed, a flag indicating whether to approximate the meta-gradient and a method for forming predictions if nonparametric testing is required (e.g., NIL testing (Raghu et al., 2020)).

Following previous work (Raghu et al., 2020; Oh et al., 2021), we train each model on 30.000 iterations on MiniImageNet, and following Finn et al. (2017), we evaluate test performance on a total of 500 randomly sampled tasks from the test set of the target domain (either MiniImageNet

Method	s_{train}	s_{test}	β	α	Approx. MetaGrad	Non-Param. Pred.
ANIL	5	10	0.001	0.01	No	NIL
MAML	5	10	0.001	0.01	Yes	NIL
BOIL	1	1	0.001	0.5	No	NIL
NP-MAML	1	1	0.001	0.5	Yes	Pairw. distance

Table 6: **Hyperparameter settings** for ANIL, MAML, BOIL and NP-MAML, used for the experiments in Section 5.

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
ANIL	70.89 ± 0.60	66.17 ± 0.54	58.24 ± 0.43	50.16 ± 0.34	52.43 ± 0.63	46.14 ± 0.54	37.84 ± 0.44	30.75 ± 0.29
BOIL	39.52 ± 0.50	33.47 ± 0.42	26.23 ± 0.31	20.02 ± 0.25	35.81 ± 0.49	30.30 ± 0.41	23.61 ± 0.32	17.96 ± 0.24
MAML	68.48 ± 0.62	63.15 ± 0.55	55.03 ± 0.43	46.62 ± 0.33	51.04 ± 0.64	44.75 ± 0.54	36.41 ± 0.43	29.60 ± 0.30
NP-MAML (0)	70.70 ± 0.60	65.49 ± 0.54	57.95 ± 0.42	50.47 ± 0.36	$\mathbf{57.89 \pm 0.63}$	$\mathbf{51.89 \pm 0.56}$	$\mathbf{43.66 \pm 0.45}$	$\mathbf{36.21 \pm 0.32}$
NP-MAML (1)	$\mathbf{72.27 \pm 0.60}$	$\mathbf{67.39 \pm 0.53}$	$\mathbf{59.59 \pm 0.43}$	$\mathbf{51.69 \pm 0.35}$	57.37 ± 0.60	50.86 ± 0.54	42.35 ± 0.44	34.58 ± 0.31

Table 7: **Classification accuracy** [%] of ANIL, BOIL, NP-MAML and MAML on 5-shot in-domain (MiniImageNet) and cross-domain (FC100) tasks with a varying number of ways during test time. Since only NP-MAML can be fine-tuned in the mixed-way setting, NP-MAML results are suffixed with the number of fine-tuning steps in parentheses. While NP-MAML (1) is the best performer for the in-domain case, for cross-domain tasks, fine-tuning NP-MAML hurts performance slightly. Varying the number of ways during test time has no or only small effects on performance differences between the different models. Since BOIL features require fine-tuning before being meaningful, BOIL's results fall behind consistently and substantially.

or FC100). Additionally, across all experiments and models we set $K_{\text{test}} = 15$. Further, for 5-shot training or testing, we set the batch-size to 2, and for 1-shot training or testing, we set the batch-size to 4. All other hyperparameters are method-specific and are presented for each of ANIL, BOIL, NP-MAML, and MAML in Table 6. Penultimate-only methods BOIL-4 and NP-MAML-4 follow the hyperparameter settings of BOIL and NP-MAML, respectively.

All experiments have been conducted on a single NVIDIA Titan Xp GPU.

E Additional Results for Varying the Ways

The results for 5-shot in- and cross-domain can be found in Table 7.

For both in- and cross-domain tasks, we observe generally smaller performance differences between methods than in the 1-shot case in Table 2, across ways. Most notably, for FC100 tasks, performance decreases after fine-tuning NP-MAML, despite the fact that both the non-fine-tuned and the fine-tuned BOIL are about 5% more accurate than third-place ANIL. Unlike for the 1-shot case, the non-fine-tuned NP-MAML performs roughly as well as ANIL on in-domain tasks on MiniImageNet, and subsequent fine-tuning improves NP-MAML further by around 1.5%. Further, performance differences between all methods are invariant to the number of ways during test time.

Again, unsurprisingly, non-fine-tuned BOIL features cannot compete in NIL-testing with any of the other methods, falling short of the best performers by about 30% on MiniImageNet and 20% on FC100.

F Additional Results of Experiments of the Penultimate Layer

In this section, we present the remaining results of Section 5.4. To this extent, we present comparisons of BOIL and NP-MAML with their penultimate-only counterparts BOIL-4 and NP-MAML-4 in the mixed-way-5-shot scenario in Table 9, as well as show comparisons for varying the number of shots between training and test time as in Section Tables 8, 10 and 11.

	MiniImageNet		FC100	
	5-to-1-shot	1-to-5-shot	5-to-1-shot	1-to-5-shot
BOIL	**49.20** $\pm$ **0.63**	63.92 $\pm$ 0.54	35.51 $\pm$ 0.53	46.83 $\pm$ 0.52
BOIL-4	48.66 $\pm$ 0.63	62.65 $\pm$ 0.55	33.08 $\pm$ 0.50	41.87 $\pm$ 0.49
NP-MAML (1)	45.68 $\pm$ 0.58	63.41 $\pm$ 0.51	**35.57** $\pm$ **0.52**	**50.28** $\pm$ **0.53**
NP-MAML-4 (1)	44.38 $\pm$ 0.60	**64.10** $\pm$ **0.51**	32.60 $\pm$ 0.48	47.39 $\pm$ 0.53

Table 8: **Classification accuracy** [%] of BOIL and NP-MAML, as well as their penultimate-only counterparts BOIL-4 and NP-MAML-4. Evaluation results for 1-shot are obtained from corresponding models trained on 5-way-5-shot MiniImageNet, and evaluation results for 5-shot are obtained from corresponding models trained on 5-way-1-shot MiniImageNet. The (1) after the NP-MAML and NP-MAML-4 results indicate that one fine-tuning step was performed before test performance was measured. Improvements of NP-MAML over its penultimate-only counterpart under domain shift become more evident when either in- or decreasing shots during test time. When switching from 5 to 1 shots in-domain (on MiniImageNet), BOIL and BOIL-4 are much stronger than the -NP variants. This effect is mirrored in favor of the -NP variants when switching from 1 to 5 shots cross-domain (on FC100).

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
BOIL	39.52 $\pm$ 0.50	33.47 $\pm$ 0.42	26.23 $\pm$ 0.31	20.02 $\pm$ 0.25	35.81 $\pm$ 0.49	30.30 $\pm$ 0.41	23.61 $\pm$ 0.32	17.96 $\pm$ 0.24
BOIL-4	35.97 $\pm$ 0.44	30.20 $\pm$ 0.37	23.08 $\pm$ 0.29	17.54 $\pm$ 0.22	35.10 $\pm$ 0.49	29.73 $\pm$ 0.42	23.01 $\pm$ 0.32	17.49 $\pm$ 0.23
NP-MAML (0)	70.70 $\pm$ 0.60	65.49 $\pm$ 0.54	57.95 $\pm$ 0.42	50.47 $\pm$ 0.36	**57.89** $\pm$ **0.63**	**51.89** $\pm$ **0.56**	**43.66** $\pm$ **0.45**	**36.21** $\pm$ **0.32**
NP-MAML (1)	**72.27** $\pm$ **0.60**	**67.39** $\pm$ **0.53**	59.59 $\pm$ 0.43	51.69 $\pm$ 0.35	57.37 $\pm$ 0.60	50.86 $\pm$ 0.54	42.35 $\pm$ 0.44	34.58 $\pm$ 0.31
NP-MAML-4 (0)	70.07 $\pm$ 0.60	65.09 $\pm$ 0.54	57.51 $\pm$ 0.43	49.78 $\pm$ 0.35	53.23 $\pm$ 0.60	47.08 $\pm$ 0.53	38.80 $\pm$ 0.42	32.15 $\pm$ 0.30
NP-MAML-4 (1)	72.11 $\pm$ 0.58	67.13 $\pm$ 0.53	**59.74** $\pm$ **0.43**	**51.88** $\pm$ **0.35**	54.54 $\pm$ 0.60	48.40 $\pm$ 0.54	40.06 $\pm$ 0.42	33.24 $\pm$ 0.30

Table 9: **Classification accuracy** [%] of BOIL and NP-MAML, as well as their penultimate-only counterparts BOIL-4 and NP-MAML-4. All models have been trained on 5-way-5-shot MiniImageNet. Since only NP-MAML and NP-MAML-4 can be fine-tuned in the mixed-way setting, NP-MAML and NP-MAML-4 results are suffixed with the number of fine-tuning steps in parentheses. Previously observed effects, such as that for five shots, fine-tuning the -NP variants amounts to little improvement in performance and that NP-MAML-4 and NP-MAML are roughly equal in performance for in-domain tasks, but NP-MAML being superior to NP-MAML-4 on cross-domain tasks, are confirmed in this experiment.

The results confirm the conclusions we draw in Section 6, such that fine-tuning early layers are beneficial for cross-domain adaptation and that 5-shot NP-MAML and NP-MAML-4 cannot substantially improve their performance through fine-tuning.

G Extended Discussion of MAML's modules

In the following two paragraphs, we give further insight into some observations from Section 5.

G.1 Fine-Tuning Early Layers for Cross-Domain Adaptation

In this section, we discuss the interesting finding that fine-tuning earlier layers in the conv4 substantially improves classification accuracy under domain shift. In particular, the domain shift we studied shifts from MiniImageNet to FC100, both containing fairly general image features (Russakovsky et al., 2015). Previous work on understanding convolutional networks in the context of image processing has established the view that early convolutional layers typically learn low-level features like edges and corners, whereas high-level convolutional layers learn high-level features like faces or outlines of various objects (Zeiler & Fergus, 2014). We hypothesize that, thus, fine-tuning earlier layers becomes more important the more we can expect low-level features to vary. In the case of MiniImageNet and FC100, we can observe such a variation due to the fact that FC100 images are substantially smaller (32×32 pixels) than MiniImageNet images (84×84 pixels). As it has become customary in cross-domain few-shot-learning, images from all domains are typically resized to a

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
BOIL	38.33 $\pm$ 0.50	32.68 $\pm$ 0.42	25.34 $\pm$ 0.32	19.45 $\pm$ 0.23	35.31 $\pm$ 0.53	29.84 $\pm$ 0.44	23.02 $\pm$ 0.33	17.59 $\pm$ 0.24
BOIL-4	35.70 $\pm$ 0.44	29.87 $\pm$ 0.38	22.53 $\pm$ 0.28	16.91 $\pm$ 0.20	34.15 $\pm$ 0.49	28.86 $\pm$ 0.42	21.92 $\pm$ 0.31	16.66 $\pm$ 0.23
NP-MAML (0)	69.23 $\pm$ 0.61	64.16 $\pm$ 0.54	56.10 $\pm$ 0.45	**48.43 $\pm$ 0.35**	**56.42 $\pm$ 0.64**	**50.91 $\pm$ 0.54**	**42.26 $\pm$ 0.44**	**34.65 $\pm$ 0.32**
NP-MAML (1)	69.18 $\pm$ 0.59	63.41 $\pm$ 0.51	54.31 $\pm$ 0.42	45.82 $\pm$ 0.33	56.18 $\pm$ 0.65	50.28 $\pm$ 0.53	41.25 $\pm$ 0.42	33.40 $\pm$ 0.32
NP-MAML-4 (0)	**69.53 $\pm$ 0.60**	**64.64 $\pm$ 0.53**	**56.58 $\pm$ 0.43**	48.39 $\pm$ 0.34	53.29 $\pm$ 0.64	47.56 $\pm$ 0.53	39.14 $\pm$ 0.44	31.92 $\pm$ 0.31
NP-MAML-4 (1)	69.37 $\pm$ 0.59	64.10 $\pm$ 0.51	55.60 $\pm$ 0.41	47.05 $\pm$ 0.32	53.05 $\pm$ 0.63	47.39 $\pm$ 0.53	38.89 $\pm$ 0.43	31.70 $\pm$ 0.31

Table 10: **Classification accuracy** [%] of BOIL and NP-MAML, as well as their penultimate-only counterparts BOIL-4 and NP-MAML-4, each trained on 5-way-1-shot MiniImageNet tasks and evaluated on a varying number of ways and five shots on in-domain (MiniImageNet) and cross-domain (FC100) tasks. Since only NP-MAML and NP-MAML-4 can be fine-tuned in the mixed-way setting, NP-MAML and NP-MAML-4 results are suffixed with the number of fine-tuning steps in parentheses. In-domain results for all -NP variants differ only slightly from each other. Cross-domain results show improvements in favor of NP-MAML, highlighting the importance of early layer fine-tuning in cross-domain settings, as discussed in Section 6. BOIL-4 performance is generally worse compared to BOIL, independent of in- or cross-domain.

	MiniImageNet				FC100			
	4-way	5-way	7-way	10-way	4-way	5-way	7-way	10-way
BOIL	31.15 $\pm$ 0.44	25.96 $\pm$ 0.37	19.44 $\pm$ 0.29	14.57 $\pm$ 0.21	29.94 $\pm$ 0.43	24.45 $\pm$ 0.34	18.40 $\pm$ 0.27	13.52 $\pm$ 0.19
BOIL-4	29.64 $\pm$ 0.39	24.23 $\pm$ 0.33	17.93 $\pm$ 0.25	13.31 $\pm$ 0.18	29.52 $\pm$ 0.42	24.19 $\pm$ 0.35	18.07 $\pm$ 0.27	13.37 $\pm$ 0.20
NP-MAML (0)	46.52 $\pm$ 0.67	41.18 $\pm$ 0.59	33.15 $\pm$ 0.44	26.35 $\pm$ 0.34	39.75 $\pm$ 0.59	34.11 $\pm$ 0.53	26.98 $\pm$ 0.40	21.20 $\pm$ 0.31
NP-MAML (1)	**50.91 $\pm$ 0.68**	**45.68 $\pm$ 0.58**	**37.31 $\pm$ 0.44**	**30.28 $\pm$ 0.35**	**41.56 $\pm$ 0.62**	**35.57 $\pm$ 0.52**	**27.96 $\pm$ 0.40**	**21.88 $\pm$ 0.29**
NP-MAML-4 (0)	43.95 $\pm$ 0.69	38.04 $\pm$ 0.61	30.07 $\pm$ 0.46	23.17 $\pm$ 0.37	35.02 $\pm$ 0.52	29.44 $\pm$ 0.45	22.95 $\pm$ 0.35	17.72 $\pm$ 0.27
NP-MAML-4 (1)	49.91 $\pm$ 0.69	44.38 $\pm$ 0.60	36.30 $\pm$ 0.44	29.04 $\pm$ 0.35	38.28 $\pm$ 0.56	32.60 $\pm$ 0.48	25.81 $\pm$ 0.38	20.19 $\pm$ 0.28

Table 11: **Classification accuracy** [%] of BOIL and NP-MAML, as well as their penultimate-only counterparts BOIL-4 and NP-MAML-4, each trained on 5-way-5-shot MiniImageNet tasks and evaluated on a varying number of ways and 1 shot on in-domain (MiniImageNet) and cross-domain (FC100) tasks. Since only NP-MAML and NP-MAML-4 can be fine-tuned in the mixed-way setting, NP-MAML and NP-MAML-4 results are suffixed with the number of fine-tuning steps in parentheses. When switching from 5 shots during training to 1 shot during test time, NP-MAML (1) generally performs best with higher performance differences under domain shift.

Figure 3: Image samples from FC100 (top) vs. image samples from MiniImageNet (bottom). Images have been resized to a common 512×512 pixels. We observe that FC100 samples are much blurrier when resized, requiring different features for edges, corners, etc.

common shape to fit the model architectures (Vuorio et al., 2019; Triantafillou et al., 2020; Oh et al., 2021). The effects of such a resizing on low-level image features can be observed in Figure 3.

We further highlight that conv4 with a nonparametric predictor is flexible enough to deal with different image sizes, albeit not with a different number of channels. A study on the performance differences with and without resizing on NP-MAML and especially NP-MAML-4 is suggested as future work.

G.2 Potential of Nonparametric Predictors in Meta-Learning Research

Our nonparametric addition to the BOIL algorithm performs favorably compared to parametric solutions in cross-domain few-shot classification. Further, as established in Section 2.2.2, our method is flexible w.r.t. the in- and output dimensionality of a given task, making it straightforward to apply in heterogeneous environments. As a result, we encourage the use of nonparametric components, especially in meta-learning, as it enables architectures to keep up with increasing task diversity (Triantafillou et al., 2020). However, we believe that the potential of rapid learning in MAML goes far beyond the scope of image classification. Consequently, future meta-learning architectures are required to become increasingly flexible without losing the ability to rapidly learn new tasks from only a few samples. We leave as future work to challenge the meta-learning system even further, posing task distributions as mixtures of, e.g., predictions on time series, audio, and speech. We believe that favorable network initialization can be learned that bridges all of those domains.

Further, we recognize the use of Gaussian Processes (Rasmussen & Williams, 2006) as another way to form nonparametric predictive distributions over few-shot tasks, which can be employed for both regression and classification. Applying Gaussian Processes to meta-learning has already received attention (Myers & Sardana, 2021).