



A survey of deep meta-learning

Mike Huisman¹ · Jan N. van Rijn¹ · Aske Plaat¹

Published online: 19 April 2021
© The Author(s) 2021

Abstract

Deep neural networks can achieve great successes when presented with large data sets and sufficient computational resources. However, their ability to learn new concepts *quickly* is limited. Meta-learning is one approach to address this issue, by enabling the network to learn how to learn. The field of *Deep Meta-Learning* advances at great speed, but lacks a unified, in-depth overview of current techniques. With this work, we aim to bridge this gap. After providing the reader with a theoretical foundation, we investigate and summarize key methods, which are categorized into (i) metric-, (ii) model-, and (iii) optimization-based techniques. In addition, we identify the main open challenges, such as performance evaluations on heterogeneous benchmarks, and reduction of the computational costs of meta-learning.

Keywords Meta-learning · Learning to learn · Few-shot learning · Transfer learning · Deep learning

1 Introduction

In recent years, deep learning techniques have achieved remarkable successes on various tasks, including game-playing (Mnih et al. 2013; Silver et al. 2016), image recognition (Krizhevsky et al. 2012; He et al. 2015), machine translation (Wu et al. 2016), and automatic classification in biomedical domains (Goceri 2019a; Goceri and Karakas 2020; Iqbal et al. 2019a, b, 2020). Despite these advances and recent solutions (Goceri 2019b, 2020), ample challenges remain to be solved, such as the large amounts of data and training that are needed to achieve good performance. These requirements severely constrain the ability of deep neural networks to learn new concepts quickly, one of the defining aspects of human intelligence (Jankowski et al. 2011; Lake et al. 2017).

Meta-learning has been suggested as one strategy to overcome this challenge (Naik and Mammon 1992; Schmidhuber 1987; Thrun 1998). The key idea is that meta-learning agents improve their learning ability over time, or equivalently, learn to learn. The learning process is primarily concerned with tasks (set of observations) and takes place at two

✉ Mike Huisman
m.huisman@liacs.leidenuniv.nl

¹ Leiden Institute of Advanced Computer Science, Niels Bohrweg 1, 2333CA Leiden, The Netherlands

different levels: an inner- and an outer-level. At the *inner-level*, a new task is presented, and the agent tries to quickly learn the associated concepts from the training observations. This quick adaptation is facilitated by knowledge that it has accumulated across earlier tasks at the *outer-level*. Thus, whereas the inner-level concerns a single task, the outer-level concerns a multitude of tasks.

Historically, the term meta-learning has been used with various scopes. In its broadest sense, it encapsulates all systems that leverage prior learning experience in order to learn new tasks more quickly (Vanschoren 2018). This broad notion includes more traditional algorithm selection and hyperparameter optimization techniques for Machine Learning (Brazdil et al. 2008). In this work, however, we focus on a subset of the meta-learning field which develops meta-learning procedures to learn a good *inductive bias* for (deep) neural networks.¹ Henceforth, we use the term *Deep Meta-Learning* to refer to this subfield of meta-learning.

The field of Deep Meta-Learning is advancing at a quick pace, while it lacks a coherent, unifying overview, providing detailed insights into the key techniques. Vanschoren (2018) has surveyed meta-learning techniques, where meta-learning was used in the broad sense, limiting its account of Deep Meta-Learning techniques. Also, many exciting developments in deep meta-learning have happened after the survey was published. A more recent survey by Hospedales et al. (2020) adopts the same notion of deep meta-learning as we do, but aims for a broad overview, omitting technical details of the various techniques.

We attempt to fill this gap by providing detailed explications of contemporary Deep Meta-Learning techniques, using a unified notation. More specifically, we cover modern techniques in the field for supervised and reinforcement learning, that have achieved state-of-the-art performance, obtained popularity in the field, and presented novel ideas. Extra attention is paid to MAML (Finn et al. 2017), and related techniques, because of their impact on the field. We show how the techniques relate to each other, detail their strengths and weaknesses, identify current challenges, and provide an overview of promising future research directions. One of the observations that we make is that the network complexity is highly related to the few-shot classification performance (see Fig. 1). One might expect that in a few-shot setting, where only a few examples are available to learn from, the number of network parameters should be kept small to prevent overfitting. Clearly, the figure shows that this does not hold, as techniques that use larger backbones tend to achieve better performance. One important factor might be that due to the high amount of tasks that have been seen by the network, we are in a setting where similarly large amounts of observations have been evaluated. This result suggests that the size of the network should be taken into account when comparing algorithms.

This work can serve as an educational introduction to the field of Deep Meta-Learning, and as reference material for experienced researchers in the field. Throughout, we will adopt the taxonomy used by Vinyals (2017), which identifies three categories of Deep Meta-Learning approaches: (i) metric-, (ii) model-, and (iii) optimization-based meta-learning techniques.

The remainder of this work is structured as follows. Sect. 2 builds a common foundation on which we will base our overview of Deep Meta-Learning techniques. Sections 3, 4, and 5 cover the main metric-, model-, and optimization-based meta-learning techniques, respectively. Section 6 provides a helicopter view of the field and

¹ Here, inductive bias refers to the assumptions of a model which guide predictions on unseen data (Mitchell 1980).

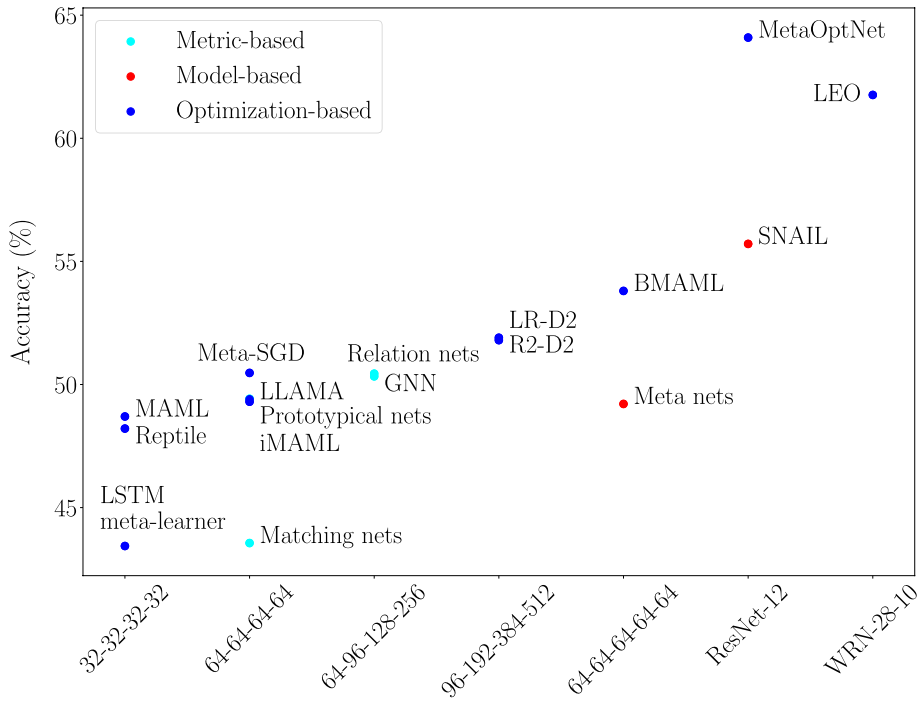


Fig. 1 The accuracy scores of the covered techniques on 1-shot miniImageNet classification. The used feature extraction backbone is displayed on the x-axis. As one can see, there is a strong relationship between the network complexity and the classification performance

summarizes the key challenges and open questions. Table 1 gives an overview of the notation that we will use throughout this paper.

2 Foundation

In this section, we build the necessary foundation for investigating Deep Meta-Learning techniques in a consistent manner. To begin with, we contrast regular learning and meta-learning. Afterwards, we briefly discuss how Deep Meta-Learning relates to different fields, what the usual training and evaluation procedure looks like, and which benchmarks are often used for this purpose. We finish this section by describing the context and some applications of the meta-learning field.

2.1 The meta abstraction

In this subsection, we contrast base-level (regular) learning and meta-learning for two different paradigms, i.e., supervised and reinforcement learning.

Table 1 Some notation and meaning, which we use throughout this paper

Expression	Meaning
Meta-learning	Learning to learn
$\mathcal{T}_j = (D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test})$	A task consisting of a labeled support and query set
Support set	The train set $D_{\mathcal{T}_j}^{tr}$ associated with a task \mathcal{T}_j
Query set	The test set $D_{\mathcal{T}_j}^{test}$ associated with a task \mathcal{T}_j
x_i	Example input vector i in the support set
y_i	(One-hot encoded) label of example input x_i from the support set
k	Number of examples per class in the support set
N	Number of classes in the support and query sets of a task
x	Input in the query set
y	A (one-hot encoded) label for input x
$(f/g/h)_\circ$	Neural network function with parameters \circ
Inner-level	At the level of a single task
Outer-level	At meta-level: across tasks
Fast weights	A term used in the literature to denote task-specific parameters
Base-learner	Learner that works at the inner-level
Meta-learner	Learner that operates at the outer-level
θ	The parameters of the base-learner network
\mathcal{L}_D	Loss function with respect to task/dataset D
Input embedding	Penultimate layer representation of the input
Task embedding	An internal representation of a task in a network/system
SL	Supervised Learning
RL	Reinforcement Learning

2.1.1 Regular supervised learning

In *supervised learning*, we wish to learn a function $f_\theta : X \rightarrow Y$ that learns to map inputs $x_i \in X$ to their corresponding outputs $y_i \in Y$. Here, θ are model parameters (e.g. weights in a neural network) that determine the function’s behavior. To learn these parameters, we are given a data set of m observations: $D = \{(x_i, y_i)\}_{i=1}^m$. Thus, given a data set \mathcal{D} , learning boils down to finding the correct setting for θ that minimizes an empirical loss function \mathcal{L}_D , which must capture how the model is performing, such that appropriate adjustments to its parameters can be made. In short, we wish to find

$$\theta_{SL} := \arg \min_{\theta} \mathcal{L}_D(\theta), \tag{1}$$

where *SL* stands for “supervised learning”. Note that this objective is specific to the data set \mathcal{D} , meaning that our model f_θ may not *generalize* to examples outside of \mathcal{D} . To measure generalization, one could evaluate the performance on a separate test data set, which contains unseen examples. A popular way to do this is through *cross-validation*, where one repeatedly creates train and test splits $D^{tr}, D^{test} \subset D$ and uses these to train and evaluate a model respectively (Hastie et al. 2009).

Finding globally optimal parameters θ_{SL} is often computationally infeasible. We can, however, approximate them, guided by *pre-defined* meta-knowledge ω (Hospedales et al.

2020), which includes, e.g., the initial model parameters θ , choice of optimizer, and learning rate schedule. As such, we approximate

$$\theta_{SL} \approx g_\omega(D, \mathcal{L}_D), \tag{2}$$

where g_ω is an optimization procedure that uses *pre-defined* meta-knowledge ω , data set D , and loss function \mathcal{L}_D , to produce updated weights $g_\omega(D, \mathcal{L}_D)$ that (presumably) perform well on D .

2.1.2 Supervised meta-learning

In contrast, *supervised meta-learning* does not assume that any meta-knowledge ω is given, or pre-defined. Instead, the goal of meta-learning is to find the best ω , such that our (regular) base-learner can learn new *tasks* (data sets) as quickly as possible. Thus, whereas supervised regular learning involves one data set, supervised meta-learning involves a group of data sets. The goal is to learn meta-knowledge ω such that our model can learn many different tasks well. Thus, our model is learning to learn.

More formally, we have a probability distribution of tasks $p(\mathcal{T})$, and wish to find optimal meta-knowledge

$$\omega^* := \arg \min_{\omega} \underbrace{\mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})}}_{\text{Outer-level}} \underbrace{[\mathcal{L}_{\mathcal{T}_j}(g_\omega(\mathcal{T}_j, \mathcal{L}_{\mathcal{T}_j}))]}_{\text{Inner-level}}. \tag{3}$$

Here, the inner-level concerns task-specific learning, while the outer-level concerns multiple tasks. One can now easily see why this is meta-learning: we learn ω , which allows for quick learning of tasks \mathcal{T}_j at the inner-level. Hence, we are learning to learn.

2.1.3 Regular reinforcement learning

In *reinforcement learning*, we have an agent that learns from experience. That is, it interacts with an environment, modeled by a Markov Decision Process (MDP) $M = (S, A, P, r, p_0, \gamma, T)$. Here, S is the set of states, A the set of actions, P the transition probability distribution defining $P(s_{t+1} | s_t, a_t)$, $r : S \times A \rightarrow \mathbb{R}$ the reward function, p_0 the probability distribution over initial states, $\gamma \in [0, 1]$ the discount factor, and T the time horizon (maximum number of time steps) (Sutton and Barto 2018; Duan et al. 2016).

At every time step t , the agent finds itself in the state s_t , in which the agent performs an action a_t , computed by a policy function π_θ (i.e., $a_t = \pi_\theta(s_t)$), which is parameterized by weights θ . In turn, it receives a reward $r_t = r(s_t, \pi_\theta(s_t)) \in \mathbb{R}$ and a new state s_{t+1} . This process of interactions continues until a termination criterion is met (e.g. fixed time horizon T reached). The goal of the agent is to learn how to act in order to maximize its expected reward. The reinforcement learning (RL) goal is to find

$$\theta_{RL} := \arg \min_{\theta} \mathbb{E}_{\text{traj}} \sum_{t=0}^T \gamma^t r(s_t, \pi_\theta(s_t)), \tag{4}$$

where we take the expectation over the possible *trajectories* $\text{traj} = (s_0, \pi_\theta(s_0), \dots, s_T, \pi_\theta(s_T))$ due to the random nature of MDPs (Duan et al. 2016). Note that γ is a hyperparameter that can prioritize short- or long-term rewards by decreasing or increasing it, respectively.

Also in the case of reinforcement learning, it is often infeasible to find the global optimum θ_{RL} , and thus we settle for approximations. In short, given a learning method ω , we approximate

$$\theta_{RL} \approx g_{\omega}(\mathcal{T}_j, \mathcal{L}_{\mathcal{T}_j}), \quad (5)$$

where again \mathcal{T}_j is the given MDP, and g_{ω} is the optimization algorithm, guided by pre-defined meta-knowledge ω .

Note that in a Markov Decision Process (MDP), the agent knows the state at any given time step t . When this is not the case, it becomes a Partially Observable Markov Decision Process (POMDP), where the agent receives only observations O and uses these to update its belief with regard to the state it is in Sutton and Barto (2018).

2.1.4 Meta reinforcement learning

The meta abstraction has as its object a group of tasks, or Markov Decision Processes (MDPs) in the case of reinforcement learning. Thus, instead of maximizing the expected reward on a single MDP, the meta reinforcement learning objective is to maximize the expected reward over various MDPs, by learning meta-knowledge ω . Here, the MDPs are sampled from some distribution $p(\mathcal{T})$. So, we wish to find a set of parameters

$$\omega^* := \arg \min_{\omega} \underbrace{\mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})}}_{\text{Outer-level}} \left[\underbrace{\mathbb{E}_{\text{traj}} \sum_{i=0}^T \gamma^i r(s_i, \pi_{g_{\omega}(\mathcal{T}_j, \mathcal{L}_{\mathcal{T}_j})}(s_i))}_{\text{Inner-level}} \right]. \quad (6)$$

2.1.5 Contrast with other fields

Now that we have provided a formal basis for our discussion for both supervised and reinforcement meta-learning, it is time to contrast meta-learning briefly with two related areas of machine learning that also have the goal to improve the speed of learning. We will start with transfer learning.

Transfer Learning In Transfer Learning, one tries to *transfer* knowledge of previous tasks to new, unseen tasks (Pan and Yang 2009; Taylor and Stone 2009), which can be challenging when a new task comes from a different distribution than the one used for training Iqbal et al. (2018). The distinction between Transfer Learning and Meta-Learning has become more opaque over time. A key property of meta-learning techniques, however, is their *meta-objective*, which explicitly aims to optimize performance across a distribution over tasks (as seen in previous sections by taking the expected loss over a distribution of tasks). This objective need not always be present in Transfer Learning techniques, e.g., when one *pre-trains* a model on a large data set and *fine-tunes* the learned weights on a smaller data set.

Multi-task learning Another, closely related field, is that of multi-task learning. In multi-task learning, a model is jointly trained to perform well on multiple fixed tasks

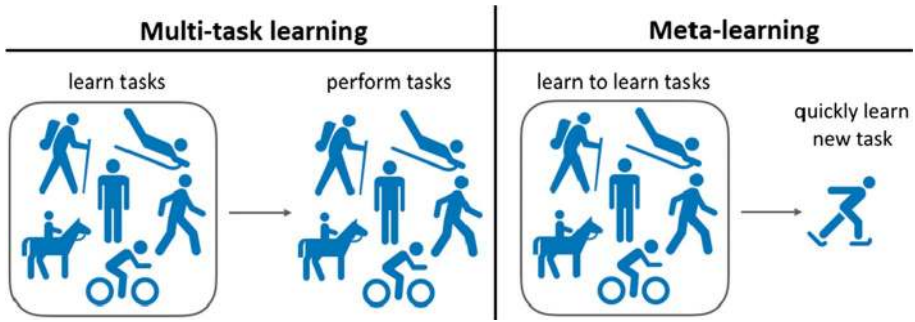


Fig. 2 The difference between multi-task learning and meta-learning. Adapted from <https://meta-world.github.io/>.

(Hospedales et al. 2020). Meta-learning, in contrast, aims to find a model that can learn new (previously unseen) tasks quickly. This difference is illustrated in Fig. 2.

2.2 The meta-setup

In the previous section, we have described the learning objectives for (meta) supervised and reinforcement learning. We will now describe the general setting that can be used to achieve these objectives. In general, one optimizes a meta-objective by using various tasks, which are data sets in the context of supervised learning, and (Partially Observable) Markov Decision Processes in the case of reinforcement learning. This is done in three stages: the (i) *meta-train* stage, (ii) *meta-validation* stage, and (iii) *meta-test* stage, each of which is associated with a set of tasks.

First, in the meta-train stage, the meta-learning algorithm is applied to the meta-train tasks. Second, the meta-validation tasks can then be used to evaluate the performance on unseen tasks, which were not used for training. Effectively, this measures the *meta-generalization* ability of the trained network, which serves as feedback to tune, e.g., hyperparameters of the meta-learning algorithm. Third, the meta-test tasks are used to give a final performance estimate of the meta-learning technique.

2.2.1 N-way, k-shot Learning

A frequently used instantiation of this general meta-setup is called *N*-way, *k*-shot classification (see Fig. 3). This setup is also divided into the three stages—meta-train, meta-validation, and meta-test—which are used for meta-learning, meta-learner hyperparameter optimization, and evaluation, respectively. Each stage has a corresponding set of disjoint labels, i.e., $L^{tr}, L^{val}, L^{test} \subset Y$, such that $L^{tr} \cap L^{val} = \emptyset, L^{tr} \cap L^{test} = \emptyset$, and $L^{val} \cap L^{test} = \emptyset$. In a given stage *s*, tasks/episodes $\mathcal{T}_j = (D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test})$ are obtained by sampling examples (x_i, y_i) from the full data set \mathcal{D} , such that every $y_i \in L^s$. Note that this requires access to a data set \mathcal{D} . The sampling process is guided by the *N*-way, *k*-shot principle, which states that every training data set $D_{\mathcal{T}_j}^{tr}$ should contain exactly *N* classes

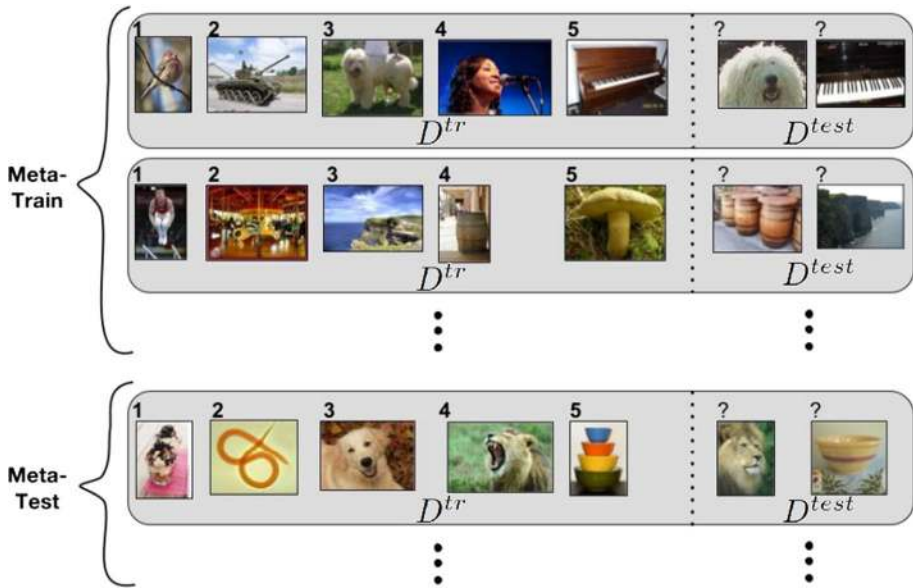


Fig. 3 Illustration of N -way, k -shot classification, where $N = 5$, and $k = 1$. Meta-validation tasks are not displayed. Adapted from Ravi and Larochelle (2017)

and k examples per class, implying that $|D_{T_j}^{tr}| = N \cdot k$. Furthermore, the true labels of examples in the test set $D_{T_j}^{test}$ must be present in the train set $D_{T_j}^{tr}$ of a given task T_j . $D_{T_j}^{tr}$ acts as a *support set*, literally supporting classification decisions on the *query set* $D_{T_j}^{test}$. Importantly, note that with this terminology, the query set (or test set) of a task is used during the meta-training phase. Furthermore, the fact that the labels across stages are disjoint ensures that we test the ability of a model to learn *new* concepts.

The meta-learning objective in the training phase is to minimize the loss function of the model predictions on the query sets, conditioned on the support sets. As such, for a given task T_j , the model ‘sees’ the support set, and extracts information from the support set to guide its predictions on the query set. By applying this procedure to different episodes/tasks T_j , the model will slowly accumulate meta-knowledge ω , which can ultimately speed up learning on new tasks.

The easiest way to achieve this is by doing this with regular neural networks, but as was pointed out by various authors (see, e.g., Finn et al. 2017) more sophisticated architectures will vastly outperform such networks. In the remainder of this work, we will review such architectures.

At the meta-validation and meta-test stages, or evaluation phases, the learned meta-information in ω is fixed. The model is, however, still allowed to make task-specific updates to its parameters θ (which implies that it is learning). After task-specific updates, we can evaluate the performance on the test sets. In this way, we test how well a technique performs at meta-learning.

N -way, k -shot classification is often performed for small values of k (since we want our models to learn new concepts quickly, i.e., from few examples). In that case, one can refer to it as *few-shot learning*.

2.2.2 Common benchmarks

Here, we briefly describe some benchmarks that can be used to evaluate meta-learning algorithms.

- *Omniglot* (Lake et al. 2011): This data set presents an image recognition task. Each image corresponds to one out of 1 623 characters from 50 different alphabets. Every character was drawn by 20 people. Note that in this case, the characters are the classes/labels.
- *ImageNet* (Deng et al. 2009): This is the largest image classification data set, containing more than 20K classes and over 14 million colored images. *miniImageNet* is a mini variant of the large ImageNet data set (Deng et al. 2009) for image classification, proposed by Vinyals et al. (2016) to reduce the engineering efforts to run experiments. The mini data set contains 60 000 colored images of size 84×84 . There are a total of 100 classes present, each accorded by 600 examples. *tieredImageNet* (Ren et al. 2018) is another variation of the large ImageNet data set. It is similar to miniImageNet, but contains a hierarchical structure. That is, there are 34 classes, each with its own sub-classes.
- *CIFAR-10 and CIFAR-100* (Krizhevsky 2009): Two other image recognition data sets. Each one contains 60K RGB images of size 32×32 . CIFAR-10 and CIFAR-100 contain 10 and 100 classes respectively, with a uniform number of examples per class (6 000 and 600 respectively). Every class in CIFAR-100 also has a super-class, of which there are 20 in the full data set. Many variants of the CIFAR data sets can be sampled, giving rise to e.g. *CIFAR-FS* (Bertinetto et al. 2019) and *FC-100* (Oreshkin et al. 2018).
- *CUB-200-2011* (Wah et al. 2011): The CUB-200-2011 data set contains roughly 12K RGB images of birds from 200 species. Every image has some labeled attributes (e.g. crown color, tail shape).
- *MNIST* (LeCun et al. 2010): MNIST presents a hand-written digit recognition task, containing ten classes (for digits 0 through 9). In total, the data set is split into a 60K train and 10K test gray scale images of hand-written digits.
- *Meta-Dataset* (Triantafillou et al. 2020): This data set comprises several other data sets such as Omniglot (Lake et al. 2011), CUB-200 (Wah et al. 2011), ImageNet (Deng et al. 2009), and more (Triantafillou et al. 2020). An episode is then constructed by sampling a data set (e.g. Omniglot), selecting a subset of labels to create train and test splits as before. In this way, broader generalization is enforced since the tasks are more distant from each other.
- *Meta-world* (Yu et al. 2019): A meta reinforcement learning data set, containing 50 robotic manipulation tasks (control a robot arm to achieve some pre-defined goal, e.g. unlocking a door, or playing soccer). It was specifically designed to cover a broad range of tasks, such that meaningful generalization can be measured (Yu et al. 2019).

2.2.3 Some applications of meta-learning

Deep neural networks have achieved remarkable results on various tasks from image recognition, text processing, game playing to robotics (Silver et al. 2016; Mnih et al. 2013; Wu et al. 2016), but their success depends on the amount of available data (Sun et al. 2017) and computing resources. Deep meta-learning reduces this dependency by allowing deep

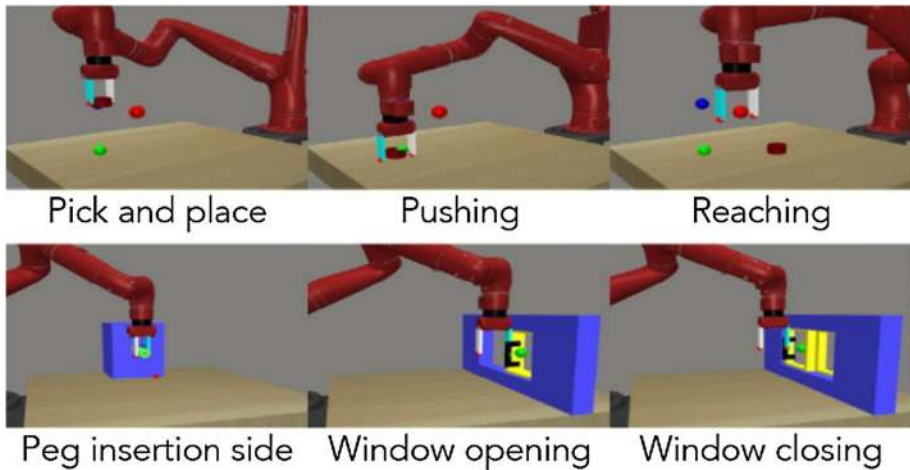


Fig. 4 Learning continuous robotic control tasks is an important application of Deep Meta-Learning techniques. Image is taken from (Yu et al. 2019)

neural networks to learn new concepts quickly. As a result, meta-learning widens the applicability of deep learning techniques to many application domains. Such areas include few-shot image classification (Finn et al. 2017; Snell et al. 2017; Ravi and Larochelle 2017), robotic control policy learning (Gupta et al. 2018; Nagabandi et al. 2019) (see Fig. 4), hyperparameter optimization (Antoniou et al. 2019; Schmidhuber et al. 1997), meta-learning learning rules (Bengio et al. 1991, 1997; Miconi et al. 2018, 2019), abstract reasoning (Barrett et al. 2018), and many more. For a larger overview of applications, we refer interested readers to Hospedales et al. (2020).

2.3 The meta-learning field

As mentioned in the introduction, meta-learning is a broad area of research, as it encapsulates all techniques that leverage prior learning experience to learn new tasks more quickly (Vanschoren 2018). We can classify two distinct communities in the field with a different focus: (i) algorithm selection and hyperparameter optimization for machine learning techniques, and (ii) search for inductive bias in deep neural networks. We will refer to these communities as group (i) and group (ii) respectively. Now, we will give a brief description of the first field, and a historical overview of the second.

Group (i) uses a more traditional approach, to select a suitable machine learning algorithm and hyperparameters for a new data set \mathcal{D} (Peng et al. 2002). This selection can for example be made by leveraging prior model evaluations on various data sets \mathcal{D}' , and by using the model which achieved the best performance on the most similar data set (Vanschoren 2018). Such traditional approaches require (large) databases of prior model evaluations, for many different algorithms. This has led to initiatives such as OpenML (Vanschoren et al. 2014), where researchers can share such information. The usage of these systems would limit the freedom in picking the neural network architecture as they would be constrained to using architectures that have been evaluated beforehand.

Driven by advances in neural networks another approach, taken by group (ii), is to adopt the view of a self-improving agent, which improves its learning ability over time by finding

Table 2 A high-level overview of the three Deep Meta-Learning categories, i.e., (i) metric-, (ii) model-, and (iii) optimization-based techniques, and their main strengths and weaknesses

	Metric	Model	Optimization
Key idea	Input similarity	Internal task representation	Optimize for fast adaptation
Strength	Simple and effective	Flexible	More robust generalizability
$p_\theta(Y \mathbf{x}, D_{\mathcal{T}_j}^{tr})$	$\sum_{(x_i, y_i) \in D_{\mathcal{T}_j}^{tr}} k_\theta(\mathbf{x}, \mathbf{x}_i) y_i$	$f_\theta(\mathbf{x}, D_{\mathcal{T}_j}^{tr})$	$f_{g_\theta(\theta, D_{\mathcal{T}_j}^{tr}, \mathcal{C}_{D_{\mathcal{T}_j}^{tr}})}(\mathbf{x})$

Recall that \mathcal{T}_j is a task, $D_{\mathcal{T}_j}^{tr}$ the corresponding support set, $k_\theta(\mathbf{x}, \mathbf{x}_i)$ a kernel function returning the similarity between the two inputs \mathbf{x} and \mathbf{x}_i , y_i are true labels for known inputs \mathbf{x}_i , θ are base-learner parameters, and g_θ is a (learned) optimizer with parameters φ

a good inductive bias (a set of assumptions that guide predictions). We now present a brief historical overview of developments in this field of Deep Meta-Learning, based on Hospedales et al. (2020).

Pioneering work was done by Schmidhuber (1987) and Hinton and Plaut (1987). Schmidhuber developed a theory of *self-referential* learning, where the weights of a neural network can serve as input to the model itself, which then predicts updates (Schmidhuber 1987, 1993). In that same year, Hinton and Plaut (1987) proposed to use two weights per neural network connection, i.e., *slow* and *fast* weights, which serve as long- and short-term memory respectively. Later came the idea of meta-learning learning rules (Bengio et al. 1991, 1997). Meta-learning techniques that use gradient-descent and backpropagation were proposed by Hochreiter et al. (2001) and Younger et al. (2001). These two works have been pivotal to the current field of Deep Meta-Learning, as the majority of techniques rely on backpropagation, as we will see on our journey of contemporary Deep Meta-Learning techniques. We will now cover the three categories metric-, model-, and optimization-based techniques, respectively.

2.4 Overview of the rest of this work

In the remainder of this work, we will look in more detail at individual meta-learning methods. As indicated before, the techniques can be grouped into three main categories (Vinyals 2017), namely (i) metric-, (ii) model-, and (iii) optimization-based methods. We will discuss them in sequence.

To help give an overview of the methods, we draw your attention to the following figure and tables. Table 2 summarizes the three categories and provides key ideas, strengths, and weaknesses of the approaches. The terms and technical details are explained more fully in the remainder of this paper. Table 3 contains an overview of all techniques that are discussed further on.

3 Metric-based meta-learning

At a high level, the goal of metric-based techniques is to acquire—among others—meta-knowledge ω in the form of a good feature space that can be used for various new tasks. In the context of neural networks, this feature space coincides with the weights θ of the

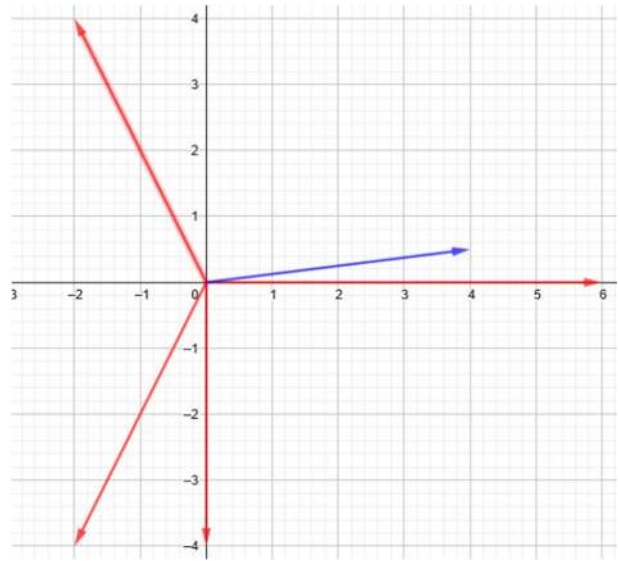
Table 3 Overview of the discussed Deep Meta-Learning techniques

Name	RL	Key idea	Bench.
<i>Metric-based</i>		<i>Input similarity</i>	–
Siamese networks	✗	Two-input, shared-weight, class identity network	1, 8
Matching networks	✗	Learn input embeddings for cosine-similarity weighted predictions	1, 2
Prototypical networks	✗	Input embeddings for class prototype clustering	1, 2, 7
Relation networks	✗	Learn input embeddings and similarity metric	1, 2, 7
ARC	✗	LSTM-based input fusion through interleaved glimpses	1, 2
GNN	✗	Propagate label information to unlabeled inputs in a graph	1, 2
<i>Model-based</i>		<i>Internal and stateful latent task representations</i>	–
Recurrent ml.	✓	Deploy Recurrent networks on RL problems	–
MANNs	✗	External short-term memory module for fast learning	1
Meta networks	✓	Fast reparameterization of base-learner by distinct meta-learner	1, 2
SNAIL	✓	Attention mechanism coupled with temporal convolutions	1, 2
CNP	✗	Condition predictive model on embedded contextual task data	1, 8
Neural stat.	✗	Similarity between latent task embeddings	1, 8
<i>Opt.-based</i>		<i>Optimize for fast task-specific adaptation</i>	–
LSTM optimizer	✗	RNN proposing weight updates for base-learner	6, 8
LSTM ml.	✓	Embed base-learner parameters in cell state of LSTM	2
RL optimizer	✗	View optimization as RL problem	4, 6
MAML	✓	Learn initialization weights θ for fast adaptation	1, 2
iMAML	✓	Approx. higher-order gradients, independent of optimization path	1, 2
Meta-SGD	✓	Learn both the initialization and updates	1, 2
Reptile	✓	Move initialization towards task-specific updated weights	1, 2
LEO	✗	Optimize in lower-dimensional latent parameter space	2, 3
Online MAML	✗	Accumulate task data for MAML-like training	4, 8
LLAMA	✗	Maintain probability distribution over post-update parameters θ'_j	2
PLATIPUS	✗	Learn a probability distribution over weight initialization θ	–
BMAML	✓	Learn multiple initializations Θ , jointly optimized by SVGD	2
Diff. solvers	✗	Learn input embeddings for simple base-learners	1, 2, 3, 4, 5

The table is partitioned into three sections, i.e., metric-, model-, and optimization-based. All methods in one section adhere to the key idea of its corresponding category, which is mentioned in bold font. The columns RL and Bench show whether the techniques are applicable to reinforcement learning settings and the used benchmarks for testing the performance of the techniques. Note that all techniques are applicable to supervised learning, with the exception of recurrent meta-learners. The benchmark column displays which benchmarks from Sect. 2.2.2 were used in the paper proposing the technique. The used coding scheme for this column is the following. 1: Omniglot, 2: miniImageNet, 3: tieredImageNet, 4: CIFAR-100, 5: CIFAR-FS, 6: CIFAR-10, 7: CUB, 8: MNIST, “–”: used other evaluation methods that are non-standard in Deep Meta-Learning and thus not covered in Sect. 2.2.2. Used abbreviations: “opt.”: optimization, “diff.”: differentiable, “bench.”: benchmarks

networks. Then, new tasks can be learned by comparing new inputs to example inputs (of which we know the labels) in the meta-learned feature space. The higher the similarity between a new input and an example, the more likely it is that the new input will have the same label as the example input.

Fig. 5 Illustration of our metric-based example. The blue vector represents the new input from the query set, whereas the red vectors are inputs from the support set which can be used to guide our prediction for the new input



Metric-based techniques are a form of meta-learning as they leverage their prior learning experience (meta-learned feature space) to ‘learn’ new tasks more quickly. Here, ‘learn’ is used in a non-standard way since metric-based techniques do not make any network changes when presented with new tasks, as they rely solely on input comparisons in the already meta-learned feature space. These input comparisons are a form of *non-parametric learning*, i.e., new task information is not absorbed into the network parameters.

More formally, metric-based learning techniques aim to learn a similarity kernel, or equivalently, *attention mechanism* k_θ (parameterized by θ), that takes two inputs \mathbf{x}_1 and \mathbf{x}_2 , and outputs their similarity score. Larger scores indicate larger similarities. Class predictions for new inputs \mathbf{x} can then be made by comparing \mathbf{x} to example inputs \mathbf{x}_i , of which we know the true labels y_i . The underlying idea is that the larger the similarity between \mathbf{x} and \mathbf{x}_i , the more likely it becomes that \mathbf{x} also has label y_i .

Given a task $\mathcal{T}_j = (D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test})$ and an unseen input vector $\mathbf{x} \in D_{\mathcal{T}_j}^{test}$, a probability distribution over classes Y is computed/predicted as a weighted combination of labels from the support set $D_{\mathcal{T}_j}^{tr}$, using similarity kernel k_θ , i.e.,

$$p_\theta(Y|\mathbf{x}, D_{\mathcal{T}_j}^{tr}) = \sum_{(\mathbf{x}_i, y_i) \in D_{\mathcal{T}_j}^{tr}} k_\theta(\mathbf{x}, \mathbf{x}_i) y_i. \tag{7}$$

Importantly, the labels y_i are assumed to be *one-hot encoded*, meaning that they are represented by zero vectors with a ‘1’ on the position of the true class. For example, suppose there are five classes in total and our example \mathbf{x}_1 has true class 4. Then, the one-hot encoded label is $y_1 = [0, 0, 0, 1, 0]$. Note that the probability distribution $p_\theta(Y|\mathbf{x}, D_{\mathcal{T}_j}^{tr})$ over classes is a vector of size $|Y|$, in which the i -th entry corresponds to the probability that input \mathbf{x} has class Y_i (given the support set). The predicted class is thus $\hat{y} = \arg \max_{i=1,2,\dots,|Y|} p_\theta(Y|\mathbf{x}, S)_i$, where $p_\theta(Y|\mathbf{x}, S)_i$ is the computed probability that input \mathbf{x} has class Y_i .

Table 4 Example showing pair-wise input comparisons

x_i	y_i	$\ x_i\ $	$\frac{x_i}{\ x_i\ }$	$\frac{x_i}{\ x_i\ } \cdot \frac{x}{\ x\ }$
[0, -4]	[1, 0, 0, 0]	4	[0, -1]	-0.12
[-2, -4]	[0, 1, 0, 0]	4.47	[-0.48, -0.89]	-0.58
[-2, 4]	[0, 0, 1, 0]	4.47	[-0.48, 0.89]	-0.37
[6, 0]	[0, 0, 0, 1]	6	[1, 0]	0.99

Numbers were rounded to two decimals

3.1 Example

Suppose that we are given a task $\mathcal{T}_j = (D_{\mathcal{T}_j}^r, D_{\mathcal{T}_j}^{test})$. Furthermore, suppose that $D_{\mathcal{T}_j}^r = \{([0, -4], 1), ([-2, -4], 2), ([-2, 4], 3), ([6, 0], 4)\}$, where a tuple denotes a pair (x_i, y_i) . For simplicity, the example will not use an embedding function, which maps example inputs onto an (more informative) embedding space. Our query set only contains one example $D_{\mathcal{T}_j}^{test} = \{([4, 0.5], y)\}$. Then, the goal is to predict the correct label for new input [4, 0.5] using only examples in $D_{\mathcal{T}_j}^r$. The problem is visualized in Fig. 5, where red vectors correspond to example inputs from our support set. The blue vector is the new input that needs to be classified. Intuitively, this new input is most similar to the vector [6, 0], which means that we expect the label for the new input to be the same as that for [6, 0], i.e., 4.

Suppose we use a fixed similarity kernel, namely the cosine similarity, i.e., $k(x, x_i) = \frac{x \cdot x_i^T}{\|x\| \cdot \|x_i\|}$, where $\|v\|$ denotes the length of vector v , i.e., $\|v\| = \sqrt{(\sum_n v_n^2)}$. Here, v_n denotes the n -th element of placeholder vector v (substitute v by x or x_i). We can now compute the cosine similarity between the new input [4, 0.5] and every example input x_i , as done in Table 4, where we used the facts that $\|x\| = \|[4, 0.5]\| = \sqrt{4^2 + 0.5^2} \approx 4.03$, and $\frac{x}{\|x\|} \approx \frac{[4, 0.5]}{4.03} = [0.99, 0.12]$.

From this table and Eq. 7, it follows that the predicted probability distribution $p_\theta(Y|x, D_{\mathcal{T}_j}^r) = -0.12y_1 - 0.58y_2 - 0.37y_3 + 0.99y_4 = -0.58[0, 1, 0, 0] - 0.37[0, 0, 1, 0] + 0.99[0, 0, 0, 1] = [-0.12, -0.58, -0.37, 0.99]$. Note that this is not really a probability distribution. That would require normalization such that every element is at least 0 and the sum of all elements is 1. For the sake of this example, we do not perform this normalization, as it is clear that class 4 (the class of the most similar example input [6, 0]) will be predicted.

One may wonder why such techniques are meta-learners, for we could take any single data set \mathcal{D} and use pair-wise comparisons to compute predictions. At the outer-level, metric-based meta-learners are trained on a distribution of different tasks, in order to learn (among others) a good input embedding function. This embedding function facilitates inner-level learning, which is achieved through pair-wise comparisons. As such, one learns an embedding function across tasks to facilitate task-specific learning, which is equivalent to “learning to learn”, or meta-learning.

After this introduction to metric-based methods, we will now cover some key metric-based techniques.

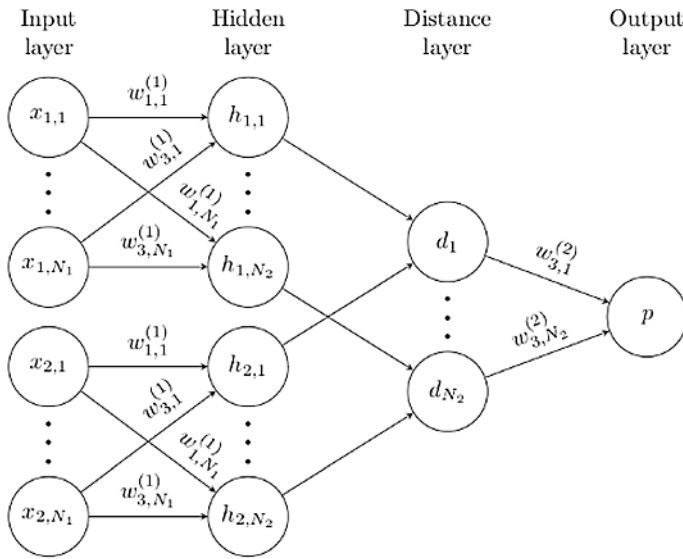


Fig. 6 Example of a Siamese neural network. Source: Koch et al. (2015)

3.2 Siamese neural networks

A Siamese neural network (Koch et al. 2015) consists of two neural networks f_θ that share the same weights θ . Siamese neural networks take two inputs $\mathbf{x}_1, \mathbf{x}_2$, and compute two hidden states $f_\theta(\mathbf{x}_1), f_\theta(\mathbf{x}_2)$, corresponding to the activation patterns in the final hidden layers. These hidden states are fed into a distance layer, which computes a distance vector $\mathbf{d} = |f_\theta(\mathbf{x}_1) - f_\theta(\mathbf{x}_2)|$, where d_i is the absolute distance between the i -th elements of $f_\theta(\mathbf{x}_1)$ and $f_\theta(\mathbf{x}_2)$. From this distance vector, the similarity between $\mathbf{x}_1, \mathbf{x}_2$ is computed as $\sigma(\boldsymbol{\alpha}^T \mathbf{d})$, where σ is the sigmoid function (with output range $[0, 1]$), and $\boldsymbol{\alpha}$ is a vector of free weighting parameters, determining the importance of each d_i . This network structure can be seen in Fig. 6.

Koch et al. (2015) applied this technique to few-shot image recognition in two stages. In the first stage, they train the twin network on an *image verification* task, where the goal is to output whether two input images \mathbf{x}_1 and \mathbf{x}_2 have the same class. The network is thus stimulated to learn discriminative features. In the second stage, where the model is confronted with a new task, the network leverages its prior learning experience. That is, given a task $\mathcal{T}_j = (D_{\mathcal{T}_j}^r, D_{\mathcal{T}_j}^{est})$, and previously unseen input $\mathbf{x} \in D_{\mathcal{T}_j}^{est}$, the predicted class \hat{y} is equal to the label y_i of the example $(x_i, y_i) \in D_{\mathcal{T}_j}^r$ which yields the highest similarity score to \mathbf{x} . In contrast to other techniques mentioned further in this section, Siamese neural networks do not directly optimize for good performance across tasks (consisting of support and query sets). However, they do leverage learned knowledge from the verification task to learn new tasks quickly.

In summary, Siamese neural networks are a simple and elegant approach to perform few-shot learning. However, they are not readily applicable outside the supervised learning setting.

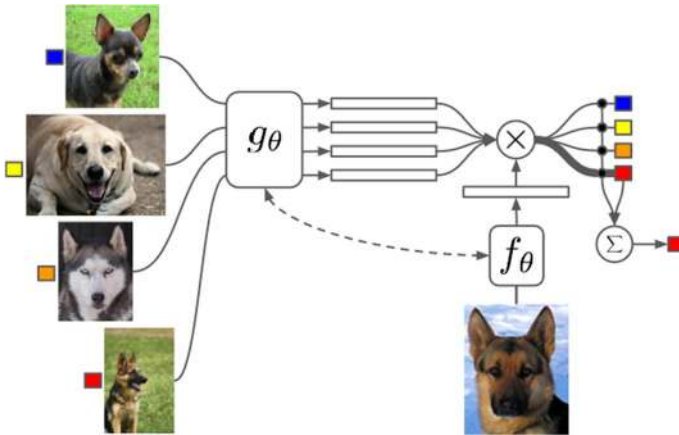


Fig. 7 The architecture of matching networks. Source: Vinyals et al. (2016)

3.3 Matching networks

Matching networks (Vinyals et al. 2016) build upon the idea that underlies Siamese neural networks (Koch et al. 2015). That is, they leverage pair-wise comparisons between the given support set $D_{\mathcal{T}_j}^r = \{(\mathbf{x}_i, y_i)\}_{i=1}^m$ (for a task \mathcal{T}_j), and new inputs $\mathbf{x} \in D_{\mathcal{T}_j}^{est}$ from the query set which we want to classify. However, instead of assigning the class y_i of the most similar example input \mathbf{x}_i , matching networks use a weighted combination of *all* example labels y_i in the support set, based on the similarity of inputs \mathbf{x}_i to new input \mathbf{x} . More specifically, predictions are computed as follows: $\hat{y} = \sum_{i=1}^m a(\mathbf{x}, \mathbf{x}_i)y_i$, where a is a non-parametric (non-trainable) attention mechanism, or similarity kernel. This classification process is shown in Fig. 7. In this figure, the input to f_θ has to be classified, using the support set $D_{\mathcal{T}_j}^r$ (input to g_θ).

The attention that is used consists of a softmax over the cosine similarity c between the input representations, i.e.,

$$a(\mathbf{x}, \mathbf{x}_i) = \frac{e^{c(f_\phi(\mathbf{x}), g_\phi(\mathbf{x}_i))}}{\sum_{j=1}^m e^{c(f_\phi(\mathbf{x}), g_\phi(\mathbf{x}_j))}}, \tag{8}$$

where f_ϕ and g_ϕ are neural networks, parameterized by ϕ and φ , that map raw inputs to a (lower-dimensional) latent vector, which corresponds to the output of the final hidden layer of a neural network. As such, neural networks act as embedding functions. The larger the cosine similarity between the embeddings of \mathbf{x} and \mathbf{x}_i , the larger $a(\mathbf{x}, \mathbf{x}_i)$, and thus the influence of label y_i on the predicted label \hat{y} for input \mathbf{x} .

Vinyals et al. (2016) propose two main choices for the embedding functions. The first is to use a single neural network, granting us $\theta = \phi = \varphi$ and thus $f_\phi = g_\phi$. This setup is the default form of matching networks, as shown in Fig. 7. The second choice is to make f_ϕ and g_ϕ dependent on the support set $D_{\mathcal{T}_j}^r$ using Long Short-Term Memory networks (LSTMs). In that case, f_ϕ is represented by an attention LSTM, and g_ϕ by a bidirectional one. This choice for embedding functions is called *Full Context Embeddings* (FCE), and yielded an accuracy improvement of roughly 2% on miniImageNet compared to the

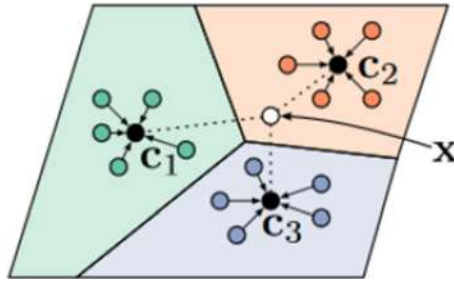


Fig. 8 Prototypical networks for the case of few-shot learning. The c_k are class prototypes for class k which are computed by averaging the representations of inputs (colored circles) in the support set. Note that the representation space is partitioned into three disjoint areas, where each area corresponds to one class. The class with the closest prototype to the new input x in the query set is then given as prediction. Source: Snell et al. (2017). (Color figure online)

regular matching networks, indicating that task-specific embeddings can aid the classification of new data points from the same distribution.

Matching networks learn a good feature space across tasks for making pair-wise comparisons between inputs. In contrast to Siamese neural networks (Koch et al. 2015), this feature space (given by weights θ) is learned across tasks, instead of on a distinct verification task.

In summary, matching networks are an elegant and simple approach to metric-based meta-learning. However, these networks are not readily applicable outside of supervised learning settings, and suffer from performance degradation when label distributions are biased (Vinyals et al. 2016).

3.4 Prototypical networks

Just like Matching networks (Vinyals et al. 2016), prototypical networks (Snell et al. 2017) base their class predictions on the entire support set $D_{T_j}^r$. However, instead of computing the similarity between new inputs and examples in the support set, prototypical networks only compare new inputs to *class prototypes* (centroids), which are single vector representations of classes in some embedding space. Since there are fewer (or equal) class prototypes than the number of examples in the support set, the amount of required pair-wise comparisons decreases, saving computational costs.

The underlying idea of class prototypes is that for a task T_j , there exists an embedding function that maps the support set onto a space where class instances cluster nicely around the corresponding class prototypes (Snell et al. 2017). Then, for a new input x , the class of the prototype nearest to that input will be predicted. As such, prototypical networks perform nearest centroid/prototype classification in a meta-learned embedding space. This is visualized in Fig. 8.

More formally, given a distance function $d : X \times X \rightarrow [0, +\infty)$ (e.g. Euclidean distance) and embedding function f_θ , parameterized by θ , prototypical networks compute class probabilities $p_\theta(Y|\mathbf{x}, D_{T_j}^r)$ as follows

$$p_\theta(y = k|\mathbf{x}, D_{T_j}^r) = \frac{\exp[-d(f_\theta(\mathbf{x}), \mathbf{c}_k)]}{\sum_{y_i} \exp[-d(f_\theta(\mathbf{x}), \mathbf{c}_{y_i})]}, \tag{9}$$

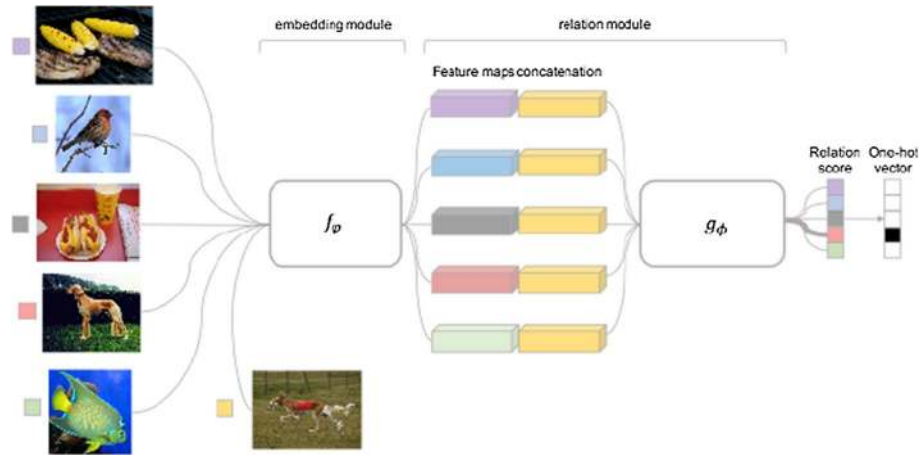


Fig. 9 Relation network architecture. First, the embedding network f_ϕ embeds all inputs from the support set D_T^r (the five example inputs on the left), and the query input (below the f_ϕ block). All support set embeddings $f_\phi(x_i)$ are then concatenated to the query embedding $f_\phi(x)$. These concatenated embeddings are passed into a relation network g_ϕ , which computes a relation score for every pair (x_i, x) . The class of the input x_i that yields the largest relation score $g_\phi([f_\phi(x), f_\phi(x_i)])$ is then predicted. Source: Sung et al. (2018)

where c_k is the prototype/centroid for class k and y_i are the classes in the support set D_T^r . Here, a class prototype for class k is defined as the average of all vectors x_i in the support set such that $y_i = k$. Thus, classes with prototypes that are nearer to the new input x obtain larger probability scores.

Snell et al. (2017) found that the squared Euclidean distance function as d gave rise to the best performance. With that distance function, prototypical networks can be seen as linear models. To see this, note that $-d(f_\theta(x), c_k) = -\|f_\theta(x) - c_k\|^2 = -f_\theta(x)^T f_\theta(x) + 2c_k^T f_\theta(x) - c_k^T c_k$. The first term does not depend on the class k , and does thus not affect the classification decision. The remainder can be written as $w_k^T f_\theta(x) + b_k$, where $w_k = 2c_k$ and $b_k = -c_k^T c_k$. Note that this is linear in the output of network f_θ , not linear in the input of the network x . Also, Snell et al. (2017) show that prototypical networks (coupled with Euclidean distance) are equivalent to matching networks in one-shot learning settings, as every example in the support set will be its own prototype.

In short, prototypical networks save computational costs by reducing the required number of pair-wise comparisons between new inputs and the support set, by adopting the concept of class prototypes. Additionally, prototypical networks were found to outperform matching networks (Vinyals et al. 2016) in 5-way, k -shot learning for $k = 1, 5$ on Omniglot (Lake et al. 2011) and miniImageNet (Vinyals et al. 2016), even though they do not use complex task-specific embedding functions. Despite these advantages, prototypical networks are not readily applicable outside of supervised learning settings.

3.5 Relation networks

In contrast to previously discussed metric-based techniques, Relation networks (Sung et al. 2018) employ a trainable similarity metric, instead of a pre-defined one (e.g. cosine similarity as used in matching networks (Vinyals et al. 2016)). More specifically, matching networks consist of two chained, neural network modules: the *embedding* network/module f_ϕ which is responsible for embedding inputs, and the *relation* network g_ϕ which computes similarity scores between new inputs \mathbf{x} and example inputs \mathbf{x}_i of which we know the labels. A classification decision is then made by picking the class of the example input which yields the largest *relation score* (or similarity). Note that Relation networks thus do not use the idea of class prototypes, and simply compare new inputs \mathbf{x} to all example inputs \mathbf{x}_i in the support set, as done by, e.g., matching networks (Vinyals et al. 2016).

More formally, we are given a support set $D_{\mathcal{T}_i}^r$ with some examples (\mathbf{x}_i, y_i) , and a new (previously unseen) input \mathbf{x} . Then, for every combination $(\mathbf{x}, \mathbf{x}_i)$, the Relation network produces a *concatenated* embedding $[f_\phi(\mathbf{x}), f_\phi(\mathbf{x}_i)]$, which is the vector obtained by concatenating the respective embeddings of \mathbf{x} and \mathbf{x}_i . This concatenated embedding is then fed into the *relation* module g_ϕ . Finally, g_ϕ computes the relation score between \mathbf{x} and \mathbf{x}_i as

$$r_i = g_\phi([f_\phi(\mathbf{x}), f_\phi(\mathbf{x}_i)]). \quad (10)$$

The predicted class is then $\hat{y} = y_{\arg \max_i r_i}$. This entire process is shown in Fig. 9. Remarkably enough, Relation networks use the Mean-Squared Error (MSE) of the relation scores, rather than the more standard cross-entropy loss. The MSE is then propagated backwards through the entire architecture (Fig. 9).

The key advantage of Relation networks is their expressive power, induced by the usage of a trainable similarity function. This expressivity makes this technique very powerful. As a result, it yields better performance than previously discussed techniques that use a fixed similarity metric.

3.6 Graph neural networks

Graph neural networks (Garcia and Bruna 2017) use a more general and flexible approach than previously discussed techniques for N -way, k -shot classification. As such, graph neural networks subsume Siamese (Koch et al. 2015) and prototypical networks (Snell et al. 2017). The graph neural network approach represents each task \mathcal{T}_i as a fully-connected graph $G = (V, E)$, where V is a set of nodes/vertices and E a set of edges connecting nodes. In this graph, nodes \mathbf{v}_i correspond to input embeddings $f_\theta(\mathbf{x}_i)$, concatenated with their one-hot encoded labels y_i , i.e., $\mathbf{v}_i = [f_\theta(\mathbf{x}_i), y_i]$. For inputs \mathbf{x} from the query set (for which we do not have the labels), a uniform prior over all N possible labels is used: $y = [\frac{1}{N}, \dots, \frac{1}{N}]$. Thus, each node contains an input and label section. Edges are weighted links that connect these nodes.

The graph neural network then propagates information in the graph using a number of local operators. The underlying idea is that label information can be transmitted from nodes of which we do have the labels, to nodes for which we have to predict labels. Which local operators are used, is out of scope for this paper, and the reader is referred to Garcia and Bruna (2017) for details.

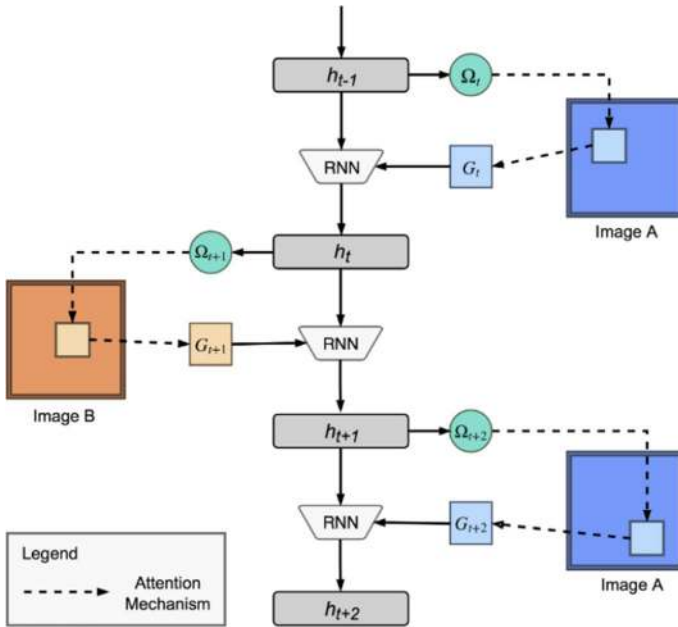


Fig. 10 Processing in an attentive recurrent comparator. At every time step, the model takes a glimpse of a part of an image and incorporates this information into the hidden state h_t . The final hidden state after taking various glimpses of a pair of images is then used to compute a class similarity score. Source: Shyam et al. (2017)

By exposing the graph neural network to various tasks \mathcal{T}_j , the propagation mechanism can be altered to improve the flow of label information in such a way that predictions become more accurate. As such, in addition to learning a good input representation function f_θ , graph neural networks also learn to propagate label information from labeled examples to unlabeled inputs.

Graph neural networks achieve good performance in few-shot settings (Garcia and Bruna 2017) and are also applicable in semi-supervised and active learning settings.

3.7 Attentive recurrent comparators

Attentive recurrent comparators (Shyam et al. 2017) differ from previously discussed techniques as they do not compare inputs as a whole, but by parts. This approach is inspired by how humans would make a decision concerning the similarity of objects. That is, we shift our attention from one object to the other, and move back and forth to take glimpses of different parts of both objects. In this way, information of two objects is fused from the beginning, whereas other techniques (e.g., matching networks (Vinyals et al. 2016) and graph neural networks (Garcia and Bruna 2017)) only combine information at the end (after embedding both images) (Shyam et al. 2017).

Given two inputs x_i and x , we feed them in interleaved fashion repeatedly into a recurrent neural network (controller): x_i, x, \dots, x_i, x . Thus, the image at time step t is given by $I_t = x_i$ if t is even else x . Then, at each time step t , the attention mechanism focuses on a square region of the current image: $G_t = \text{attend}(I_t, \Omega_t)$, where $\Omega_t = W_g h_{t-1}$ are attention

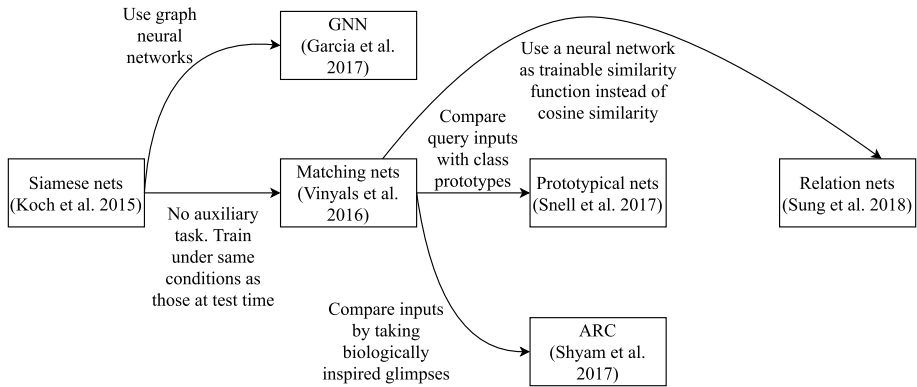


Fig. 11 The relationships between the covered metric-based meta-learning techniques

parameters, which are computed from the previous hidden state h_{t-1} . The next hidden state $h_{t+1} = \text{RNN}(G_t, h_{t-1})$ is given by the glimpse at time t , i.e., G_t , and the previous hidden state h_{t-1} . The entire sequence consists of g glimpses per image. After this sequence is fed into the recurrent neural network (indicated by $\text{RNN}(\circ)$), the final hidden state h_{2g} is used as a combined representation of x_i relative to x . This process is summarized in Fig. 10. Classification decisions can then be made by feeding the combined representations into a classifier. Optionally, the combined representations can be processed by bi-directional LSTMs before passing them to the classifier.

The attention approach is biologically inspired, and biologically plausible. A downside of attentive recurrent comparators is the higher computational cost, while the performance is often not better than less biologically plausible techniques, such as graph neural networks (Garcia and Bruna 2017).

3.8 Metric-based techniques, in conclusion

In this section, we have seen various metric-based techniques. The metric-based techniques meta-learn an informative feature space that can be used to compute class predictions based on input similarity scores. Figure 11 shows the relationships between the various metric-based techniques that we have covered.

As we can see, Siamese networks (Koch et al. 2015) mark the beginning of metric-based, deep meta-learning techniques in few-shot learning settings. They are the first to use the idea of predicting classes by comparing inputs from the support and query sets. This idea was generalized in graph neural networks (GNNs) (Hamilton et al. 2017; Garcia and Bruna 2017) where the information flow between support and query inputs is parametric and thus more flexible. Matching networks (Vinyals et al. 2016) are directly inspired by Siamese networks as they use the same core idea (comparing inputs for making predictions), but directly train in the few-shot setting and use cosine similarity as a similarity function. Thus, the auxiliary, binary classification task used by Siamese networks is left out, and matching networks directly train on tasks. Prototypical networks (Snell et al. 2017) increase the robustness of input comparisons by comparing every query set input with a class prototype instead of individual support set examples. This reduces the number of required input comparisons for a single query input to N instead of $k \cdot N$. Relation

networks (Sung et al. 2018) replace the fixed, pre-defined similarity metrics used in matching and prototypical networks with a neural network, which allows for learning a domain-specific similarity function. Lastly, attentive recurrent comparators (ARCs) (Shyam et al. 2017) take a more biologically plausible approach by not comparing entire inputs but by taking multiple interleaved glimpses at various parts of the inputs that are being compared.

Key advantages of these metric-based techniques are that (i) the underlying idea of similarity-based predictions is conceptually simple, and (ii) they can be fast at test-time when tasks are small, as the networks do not need to make task-specific adjustments. However, when tasks at meta-test time become more distant from the tasks that were used at meta-train time, metric-learning techniques are unable to absorb new task information into the network weights. Consequently, performance may suffer.

Furthermore, when tasks become larger, pair-wise comparisons may become computationally expensive. Lastly, most metric-based techniques rely on the presence of labeled examples, which make them inapplicable outside of supervised learning settings.

4 Model-based meta-learning

A different approach to Deep Meta-Learning is the model-based approach. On a high level, model-based techniques rely upon an adaptive, internal state, in contrast to metric-based techniques, which generally use a fixed neural network at test-time.

More specifically, model-based techniques maintain a stateful, internal representation of a task. When presented with a task, a model-based neural network processes the support set in a sequential fashion. At every time step, an input enters and alters the internal state of the model. Thus, the internal state can capture relevant task-specific information, which can be used to make predictions for new inputs.

Because the predictions are based on internal dynamics that are hidden from the outside, model-based techniques are also called *black-boxes*. Information from previous inputs must be remembered, which is why model-based techniques have a memory component, either in- or externally.

Recall that the mechanics of metric-based techniques were limited to pair-wise input comparisons. This is not the case for model-based techniques, where the human designer has the freedom to choose the internal dynamics of the algorithm. As a result, model-based techniques are not restricted to meta-learning good feature spaces, as they can also learn internal dynamics, used to process and predict input data of tasks.

More formally, given a support set $D_{T_j}^r$ corresponding to a task T_j , model-based techniques compute a class probability distribution for a new input \mathbf{x} as

$$p_{\theta}(Y|\mathbf{x}, D_{T_j}^r) = f_{\theta}(\mathbf{x}, D_{T_j}^r), \quad (11)$$

where f represents the black-box neural network model, and θ its parameters.

4.1 Example

Using the same example as in Sect. 3, suppose we are given a task support set $D_{T_j}^r = \{([0, -4], 1), ([-2, -4], 2), ([-2, 4], 3), ([6, 0], 4)\}$, where a tuple denotes a pair (\mathbf{x}_i, y_i) . Furthermore, suppose our query set only contains one example $D_{T_j}^{est} = \{([4, 0.5], 4)\}$.

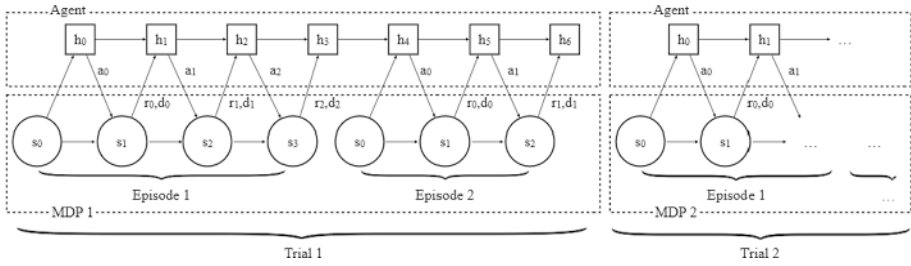


Fig. 12 Workflow of recurrent meta-learners in reinforcement learning contexts. As mentioned in Sect. 2.1.3, s_t , r_t , and d_t denote the state, reward, and termination flag at time step t . h_t refers to the hidden state at time t . Source: Duan et al. (2016)

This problem has been visualized in Fig. 5 (in Sect. 3). For the sake of the example, we do not use an input embedding function: our model will operate on the raw inputs of $D_{T_j}^r$ and $D_{T_j}^{test}$. As an internal state, our model uses an external *memory matrix* $M \in \mathbb{R}^{4 \times (2+1)}$, with four rows (one for each example in our support set), and three columns (the dimensionality of input vectors, plus one dimension for the correct label). Our model proceeds to process the support set in a sequential fashion, reading the examples from $D_{T_j}^r$ one by one, and by storing the i -th example in the i -th row of the memory module. After processing the support set, the memory matrix contains all examples, and as such, serves as internal task representation.

Given the new input [4, 0.5], our model could use many different techniques to make a prediction based on this representation. For simplicity, assume that it computes the dot product between x , and every memory $M(i)$ (the 2-D vector in the i -th row of M , ignoring the correct label), and predicts the class of the input which yields the largest dot product. This would produce scores $-2, -10, -6$, and 24 for the examples in $D_{T_j}^r$ respectively. Since the last example [6, 0] yields the largest dot product, we predict that class, i.e., 4.

Note that this example could be seen as a metric-based technique where the dot product is used as a similarity function. However, the reason that this technique is model-based is that it stores the entire task inside a memory module. This example was deliberately easy for illustrative purposes. More advanced and successful techniques have been proposed, which we will now cover.

4.2 Recurrent meta-learners

Recurrent meta-learners (Duan et al. 2016; Wang et al. 2016) are, as the name suggests, meta-learners based on recurrent neural networks. The recurrent network serves as dynamic task embedding storage. These recurrent meta-learners were specifically proposed for reinforcement learning problems, hence we will explain them in that setting.

The recurrence is implemented by e.g. an LSTM (Wang et al. 2016) or a GRU (Duan et al. 2016). The internal dynamics of the chosen Recurrent Neural Network (RNN) allows for fast adaptation to new tasks, while the algorithm used to train the recurrent network gradually accumulates knowledge about the task structure, where each task is modelled as an episode (or set of episodes).

The idea of recurrent meta-learners is quite simple. That is, given a task T_j , we simply feed the (potentially processed) environment variables $[s_{t+1}, a_t, r_t, d_t]$ (see Sect. 2.1.3) into

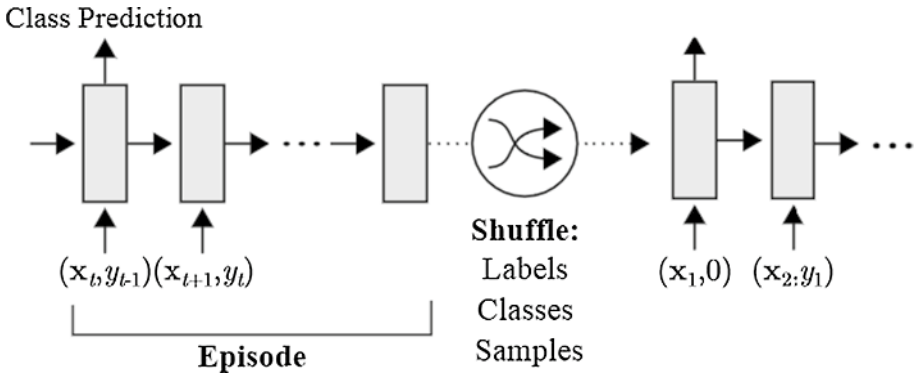


Fig. 13 Workflow of memory-augmented neural networks. Here, an episode corresponds to a given task T_j . After every episode, the order of labels, classes, and samples should be shuffled to minimize dependence on arbitrarily assigned orders. Source: Santoro et al. (2016)

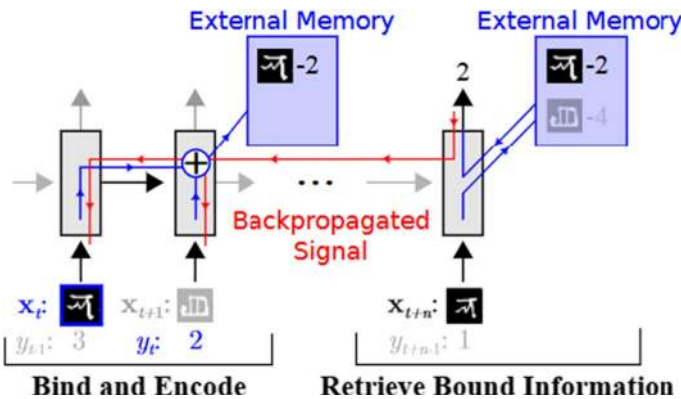


Fig. 14 Controller-memory interaction in memory-augmented neural networks. Source: Santoro et al. (2016)

an RNN at every time step t . Recall that s, a, r, d denote the state, action, reward, and termination flag respectively. At every time step t , the RNN outputs an action and a hidden state. Conditioned on its hidden state h_t , the network outputs an action a_t . The goal is to maximize the expected reward in each trial. See Fig. 12 for a visual depiction. From this figure, it also becomes clear why these techniques are model-based. That is, they embed information from previously seen inputs in the hidden state.

Recurrent meta-learners have been shown to perform almost as well as asymptotically optimal algorithms on simple reinforcement learning tasks (Wang et al. 2016; Duan et al. 2016). However, their performance degrades in more complex settings, where temporal dependencies can span a longer horizon. Making recurrent meta-learners better at such complex tasks is a direction for future research.

4.3 Memory-augmented neural networks (MANNs)

The key idea of memory-augmented neural networks (MANNs) (Santoro et al. 2016) is to enable neural networks to learn quickly with the help of an *external memory*. The main *controller* (the recurrent neural network interacting with the memory) then gradually accumulates knowledge across tasks, while the external memory allows for quick task-specific adaptation. For this, Santoro et al. (2016) used Neural Turing Machines (Graves et al. 2014). Here, the controller is parameterized by θ and acts as the long-term memory of the memory-augmented neural network, while the external memory module is the short-term memory.

The workflow of memory-augmented neural networks is displayed in Fig. 13. Note that the data from a task is processed as a sequence, i.e., data are fed into the network one by one. The support set is fed into the memory-augmented neural network first. Afterwards, the query set is processed. During the meta-train phase, training tasks can be fed into the network in arbitrary order. At time step t , the model receives input x_t with the label of the previous input, i.e., y_{t-1} . This was done to prevent the network from mapping class labels directly to the output (Santoro et al. 2016).

The interaction between the controller and memory is visualized in Fig. 14. The idea is that the external memory module, containing representations of previously seen inputs, can be used to make predictions for new inputs. In short, previously obtained knowledge is leveraged to aid the classification of new inputs. Note that neural networks also attempt to do this, however, their prior knowledge is slowly accumulated into the network weights, while an external memory module can directly store such information.

Given an input x_t at time t , the controller generates a key k_t , which can be stored in memory matrix M and can be used to retrieve previous representations from memory matrix M . When reading from memory, the aim is to produce a linear combination of stored keys in memory matrix M , giving greater weight to those which have a larger cosine similarity with the current key k_t . More specifically, a read vector w_t^r is created, in which each entry i denotes the cosine similarity between key k_t and the memory (from a previous input) stored in row i , i.e., $M_t(i)$. Then, the representation $r_t = \sum_i w_t^r(i)M(i)$ is retrieved, which is simply a linear combination of all keys (i.e., rows) in memory matrix M .

Predictions are made as follows. Given an input x_t , memory-augmented neural networks use the external memory to compute the corresponding representation r_t , which could be fed into a softmax layer, resulting in class probabilities. Across tasks, memory-augmented neural networks learn a good input embedding function f_θ and classifier weights, which can be exploited when presented with new tasks.

To write input representations to memory, Santoro et al. (2016) propose a new mechanism called Least Recently Used Access (LRUA). LRUA either writes to the least, or most recently used memory location. In the former case, it preserves recent memories, and in the latter, it updates recently obtained information. The writing mechanism works by keeping track of how often every memory location is accessed in a usage vector w_t^u , which is updated at every time step according to the following update rule: $w_t^u := \gamma w_{t-1}^u + w_t^r + w_t^w$, where superscripts u , w and r refer to usage, write and read vectors, respectively. In words, the previous usage vector is decayed (using parameter γ), while current reads (w_t^r) and writes (w_t^w) are added to the usage. Let n be the total number of reads to memory, and $\ell u(n)$ (ℓu for ‘least used’) be the n -th smallest value in the usage vector w_t^u . Then, the least-used weights are defined as follows:

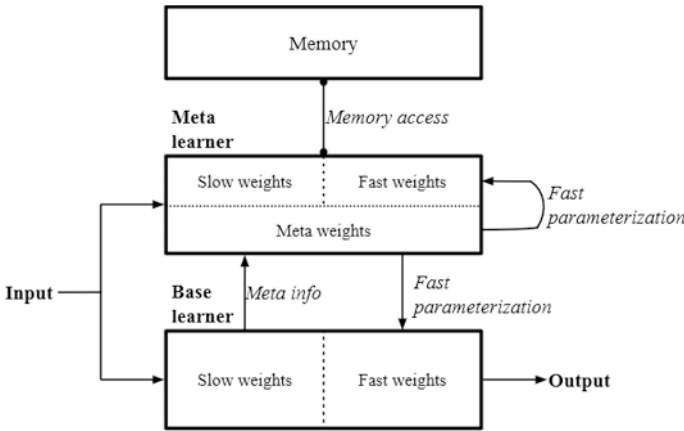


Fig. 15 The architecture of a Meta Network. Source: Munkhdalai and Yu (2017)

$$w_i^{\ell u}(i) = \begin{cases} 0 & \text{if } w_i^u(i) > \ell u(n) \\ 1 & \text{else} \end{cases}$$

Then, the write vector w_i^w is computed as $w_i^w = \sigma(\alpha)w_{i-1}^r + (1 - \sigma(\alpha))w_{i-1}^{\ell u}$, where α is a parameter that interpolates between the two weight vectors. As such, if $\sigma(\alpha) = 1$, we write to the most recently used memory, whereas when $\sigma(\alpha) = 0$, we write to the least recently used memory locations. Finally, writing is performed as follows: $M_i(i) := M_{i-1}(i) + w_i^w(i)k_i$, for all i .

In summary, memory-augmented neural networks (Santoro et al. 2016) combine external memory and a neural network to achieve meta-learning. The interaction between a controller, with long-term memory parameters θ , and memory M , may also be interesting for studying human meta-learning (Santoro et al. 2016). In contrast to many metric-based techniques, this model-based technique is applicable to both classification and regression problems. A downside of this approach is the architectural complexity.

4.4 Meta networks

Meta networks are divided into two distinct subsystems (consisting of neural networks), i.e., the base- and meta-learner (whereas in memory-augmented neural networks the base- and meta-components are intertwined). The base-learner is responsible for performing tasks, and for providing the meta-learner with meta-information, such as loss gradients. The meta-learner can then compute fast task-specific weights for itself and the base-learner, such that it can perform better on the given task $\mathcal{T}_j = (D_{\mathcal{T}_j}^r, D_{\mathcal{T}_j}^{est})$. This workflow is depicted in Fig. 15.

The meta-learner consists of neural networks u_ϕ, m_ϕ , and d_ψ . Network u_ϕ is used as an input representation function. Networks d_ψ and m_ϕ are used to compute task-specific weights ϕ^* and example-level fast weights θ^* . Lastly, b_θ is the base-learner which performs input predictions. Note that we used the term fast-weights throughout, which refers to task- or input-specific versions of slow (initial) weights.

In a similar fashion to memory-augmented neural networks (Santoro et al. 2016), meta networks (Munkhdalai and Yu 2017) also leverage the idea of an external memory module. However, meta networks use the memory for a different purpose. The memory stores for each observation \mathbf{x}_i in the support set two components, i.e., its representation \mathbf{r}_i and the fast weights θ_i^* . These are then used to compute an attention-based representation and fast weights for new inputs, respectively.

Algorithm 1 Meta networks, by Munkhdalai and Yu (2017)

```

1: Sample  $S = \{(\mathbf{x}_i, y_i) \sim D_{T_j}^{tr}\}_{i=1}^T$  from the support set
2: for  $(\mathbf{x}_i, y_i) \in S$  do
3:    $\mathcal{L}_i = \text{error}(u_\phi(\mathbf{x}_i), y_i)$ 
4: end for
5:  $\phi^* = d_\psi(\{\nabla_\phi \mathcal{L}_i\}_{i=1}^T)$ 
6: for  $(\mathbf{x}_i, y_i) \in D_{T_j}^{tr}$  do
7:    $\mathcal{L}_i = \text{error}(b_\theta(\mathbf{x}_i), y_i)$ 
8:    $\theta_i^* = m_\varphi(\nabla_\theta \mathcal{L}_i)$ 
9:   Store  $\theta_i^*$  in  $i$ -th position of example-level weight memory  $M$ 
10:   $\mathbf{r}_i = u_{\phi, \phi^*}(\mathbf{x}_i)$ 
11:  Store  $\mathbf{r}_i$  in  $i$ -th position of representation memory  $R$ 
12: end for
13:  $\mathcal{L}_{task} = 0$ 
14: for  $(\mathbf{x}, y) \in D_{T_j}^{test}$  do
15:   $\mathbf{r} = u_{\phi, \phi^*}(\mathbf{x})$ 
16:   $\mathbf{a} = \text{attention}(R, \mathbf{r})$  ▷  $a_k$  is the cosine similarity between  $\mathbf{r}$  and  $R(k)$ 
17:   $\theta^* = \text{softmax}(\mathbf{a})^T M$ 
18:   $\mathcal{L}_{task} = \mathcal{L}_{task} + \text{error}(b_{\theta, \theta^*}(\mathbf{x}), y)$ 
19: end for
20: Update  $\Theta = \{\theta, \phi, \psi, \varphi\}$  using  $\nabla_\Theta \mathcal{L}_{task}$ 

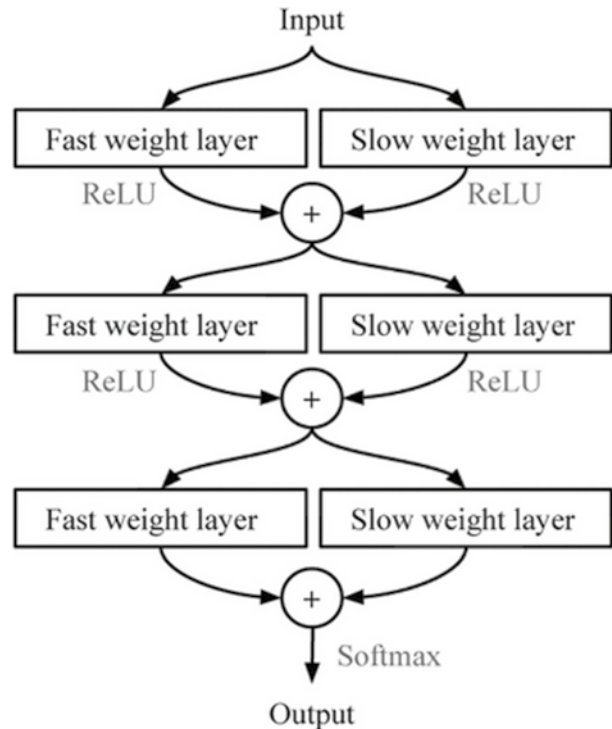
```

The pseudocode for meta networks is displayed in Algorithm 1. First, a sample of the support set is created (line 1), which is used to compute task-specific weights ϕ^* for the representation network u_ϕ (lines 2-5). Note that u_ϕ has two tasks, (i) it should compute a representation for inputs \mathbf{x}_i (line 10 and 15), and (ii) it needs to make predictions for inputs \mathbf{x}_i , in order to compute a loss (line 3). To achieve both goals, a conventional neural network can be used that makes class predictions. The states of the final hidden layer are then used as representations. Typically, the cross-entropy is calculated over the predictions of representation network u_ϕ . When there are multiple examples per class in the support set, an alternative is to use a contrastive loss function (Munkhdalai and Yu 2017).

Then, meta networks iterate over every example (\mathbf{x}_i, y_i) in the support set $D_{T_j}^{tr}$. The base-learner b_θ attempts to make class predictions for these examples, resulting in loss values \mathcal{L}_i (line 7). The gradients of these losses are used to compute fast weights θ^* for example i (line 8), which are then stored in the i -th row of memory matrix M (line 9). Additionally, input representations \mathbf{r}_i are computed and stored in memory matrix R (lines 10-11).

Now, meta networks are ready to address the query set $D_{T_j}^{test}$. They iterate over every example (\mathbf{x}, y) , and compute a representation \mathbf{r} of it (line 15). This representation is matched against the representations of the support set, which are stored in memory matrix R . This matching gives us a similarity vector \mathbf{a} , where every entry k denotes the similarity between input representation \mathbf{r} and the k -th row in memory matrix R , i.e., $R(k)$ (line 16). A softmax over this similarity vector is performed to normalize the entries. The resulting

Fig. 16 The layer augmentation setup used to combine slow and fast weights. Source: Munkhdalai and Yu (2017)



vector is used to compute a linear combination of weights that were generated for inputs in the support set (line 17). These weights θ^* are specific for input x in the query set, and can be used by the base-learner b to make predictions for that input (line 18). The observed error is added to the task loss. After the entire query set is processed, all involved parameters can be updated using backpropagation (line 20).

Note that some neural networks use both slow- and fast-weights at the same time. Munkhdalai and Yu (2017) use a so-called augmentation setup for this, as depicted in Fig. 16.

In short, meta networks rely on a reparameterization of the meta- and base-learner for every task. Despite the flexibility and applicability to both supervised and reinforcement learning settings, the approach is quite complex. It consists of many components, each with its own set of parameters, which can be a burden on memory usage and computation time. Additionally, finding the correct architecture for all the involved components can be time-consuming.

4.5 Simple neural attentive meta-learner (SNAIL)

Instead of an external memory matrix, SNAIL (Mishra et al. 2018) relies on a special model architecture to serve as memory. Mishra et al. (2018) argue that it is not possible to use Recurrent Neural Networks for this, as they have limited memory capacity, and cannot pinpoint specific prior experiences (Mishra et al. 2018). Hence, SNAIL uses a different architecture, consisting of 1D *temporal convolutions* (Oord et al. 2016) and a *soft attention*

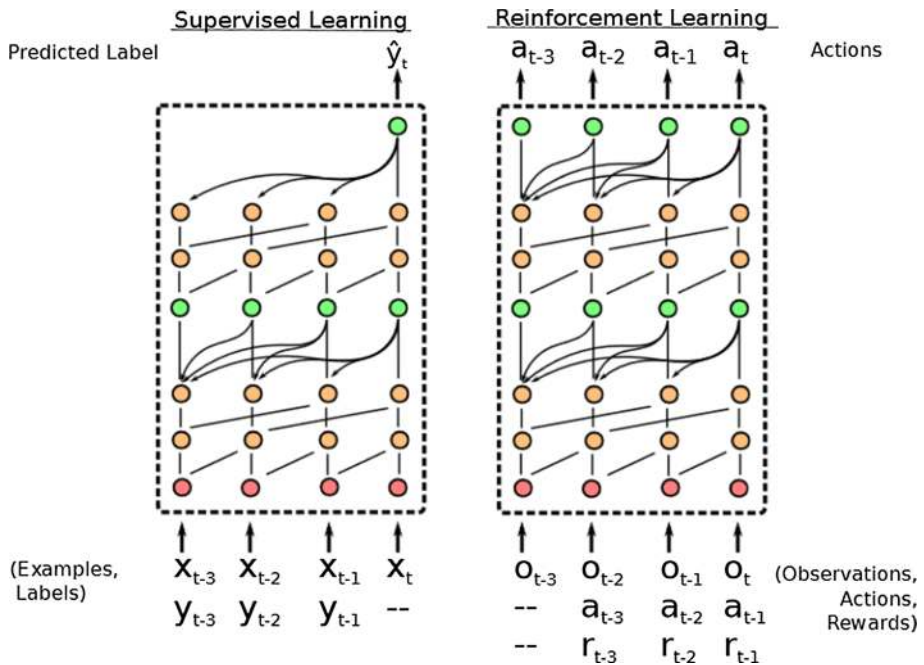


Fig. 17 Architecture and workflow of SNAIL for supervised and reinforcement learning settings. The input layer is red. Temporal Convolution blocks are orange; attention blocks are green. Source: Mishra et al. (2018). (Color figure online)

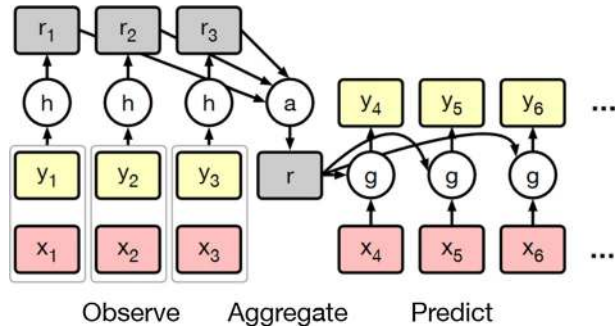
mechanism (Vaswani et al. 2017). The temporal convolutions allow for ‘high bandwidth’ memory access, and the attention mechanism allows one to pinpoint specific experiences. Figure 17 visualizes the architecture and workflow of SNAIL for supervised learning problems. From this figure, it becomes clear why this technique is model-based. That is, model outputs are based upon the internal state, computed from earlier inputs.

SNAIL consists of three building blocks. The first is the *DenseBlock*, which applies a single 1D convolution to the input, and concatenates (in the feature/horizontal direction) the result. The second is a *TCBlock*, which is simply a series of DenseBlocks with an exponentially increasing dilation rate of the temporal convolutions (Mishra et al. 2018). Note that the dilation is nothing but the temporal distance between two nodes in a network. For example, if we use a dilation of 2, a node at position p in layer L will receive the activation from node $p - 2$ from layer $L - 1$. The third block is the *AttentionBlock*, which learns to focus on the important parts of prior experience.

In a similar fashion to memory-augmented neural networks (Santoro et al. 2016) (Sect. 4.3), SNAIL also processes task data in sequence, as shown in Fig. 17. However, the input at time t is accompanied by the label at time t , instead of $t - 1$ (as was the case for memory-augmented neural networks). SNAIL learns internal dynamics from seeing various tasks so that it can make good predictions on the query set, conditioned upon the support set.

A key advantage of SNAIL is that it can be applied to both supervised and reinforcement learning tasks. In addition, it achieves good performance compared to previously discussed techniques. A downside of SNAIL is that finding the correct architecture of TCBlocks and DenseBlocks can be time-consuming.

Fig. 18 A schematic view of how conditional neural processes work. Here, h denotes a network outputting a representation for an observation, a denotes an aggregation function for these representations, and g denotes a neural network that makes predictions for unlabelled observations, based on the aggregated representation. Source: Garnelo et al. (2018)



4.6 Conditional neural processes (CNPs)

In contrast to previous techniques, a conditional neural process (CNP) (Garnelo et al. 2018) does not rely on an external memory module. Instead, it aggregates the support set into a single aggregated latent representation. The general architecture is shown in Fig. 18. As we can see, the conditional neural process operates in three phases on task \mathcal{T}_j . First, it observes the support set $D_{\mathcal{T}_j}^r$, including the ground-truth outputs y_i . Examples $(x_i, y_i) \in D_{\mathcal{T}_j}^r$ are embedded using a neural network h_θ into representations r_i . Second, these representations are aggregated using operator a to produce a single representation r of $D_{\mathcal{T}_j}^r$ (hence it is model-based). Third, a neural network g_ϕ processes this single representation r , new inputs x , and produces predictions \hat{y} .

Let the entire conditional neural process model be denoted by Q_Θ , where Θ is a set of all involved parameters $\{\theta, \phi\}$. The training process is different compared to other techniques. Let $x_{\mathcal{T}_j}$ and $y_{\mathcal{T}_j}$ denote all inputs and corresponding outputs in $D_{\mathcal{T}_j}^r$. Then, the first $\ell \sim U(0, \dots, k \cdot N - 1)$ examples in $D_{\mathcal{T}_j}^r$ are used as a conditioning set $D_{\mathcal{T}_j}^c$ (effectively splitting the support set into a true training set and a validation set). Given a value of ℓ , the goal is to maximize the log-likelihood (or minimize the negative log-likelihood) of the labels $y_{\mathcal{T}_j}$ in the entire support set $D_{\mathcal{T}_j}^r$

$$\mathcal{L}(\Theta) = -\mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} \left[\mathbb{E}_{\ell \sim U(0, \dots, k \cdot N - 1)} \left(Q_\Theta(y_{\mathcal{T}_j} | D_{\mathcal{T}_j}^c, x_{\mathcal{T}_j}) \right) \right]. \tag{12}$$

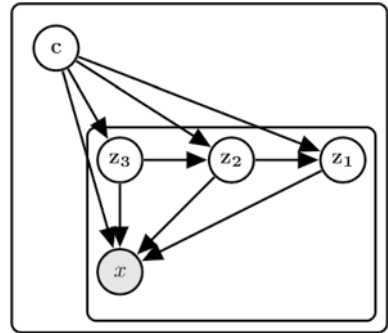
Conditional neural processes are trained by repeatedly sampling various tasks and values of ℓ , and propagating the observed loss backwards.

In summary, conditional neural processes use compact representations of previously seen inputs to aid the classification of new observations. Despite its simplicity and elegance, a disadvantage of this technique is that it is often outperformed in few-shot settings by other techniques such as matching networks (Vinyals et al. 2016) (see Sect. 3.3).

4.7 Neural statistician

A neural statistician (Edwards and Storkey 2017) differs from earlier approaches as it learns to compute *summary statistics*, or *meta-features*, of data sets in an unsupervised manner. These latent embeddings (making the approach model-based) can then later be

Fig. 19 Neural statistician architecture. Edges are neural networks. All incoming inputs to a node are concatenated



used for making predictions. Despite the broad applicability of the model, we discuss it in the context of Deep Meta-Learning.

A neural statistician performs both *learning* and *inference*. In the learning phase, the model attempts to produce generative models \hat{P}_i for every data set D_i . The key assumption that is made by Edwards and Storkey (2017) is that there exists a generative process P_i , which conditioned on a latent context vector c_i , can produce a data set D_i . At inference time, the goal is to infer a (posterior) probability distribution over the context $q(c|D)$.

The model uses a variational autoencoder, which consists of an encoder and decoder. The encoder is responsible for producing a distribution over latent vectors $z: q(z|x;\phi)$, where x is an input vector, and ϕ are the encoder parameters. The encoded input z , which is often of lower dimensionality than the original input x , can then be decoded by the decoder $p(x|z;\theta)$. Here, θ are the parameters of the decoder. To capture more complex patterns in data sets, the model uses multiple latent layers z_1, \dots, z_L , as shown in Fig. 19. Given this architecture, the posterior over c and z_1, \dots, z_L (shorthand $z_{1:L}$) is given by

$$q(c, z_{1:L}|D;\phi) = q(c|D;\phi) \prod_{x \in D} q(z_L|x, c;\phi) \prod_{i=1}^{L-1} q(z_i|z_{i+1}, x, c;\phi). \tag{13}$$

The neural statistician is trained to minimize a three-component loss function, consisting of the reconstruction loss (how well it models the data), context loss (how well the inferred context $q(c|D;\phi)$ corresponds to the prior $P(c)$, and latent loss (how well the inferred latent variables z_i are modelled).

This model can be applied to N -way, few-shot learning as follows. Construct N data sets for every of the N classes, such that one data set contains only examples of the same class. Then, the neural statistician is provided with a new input x , and has to predict its class. It computes a context posterior $N_x = q(c|x;\phi)$ depending on new input x . In a similar fashion, context posteriors are computed for all of the data sets $N_i = q(c|D_i;\phi)$. Lastly, it assigns the label i such that the difference between N_i and N_x is minimal.

In summary, the neural statistician (Edwards and Storkey 2017) allows for quick learning on new tasks through data set modeling. Additionally, it is applicable to both supervised and unsupervised settings. A downside is that the approach requires many data sets to achieve good performance (Edwards and Storkey 2017).

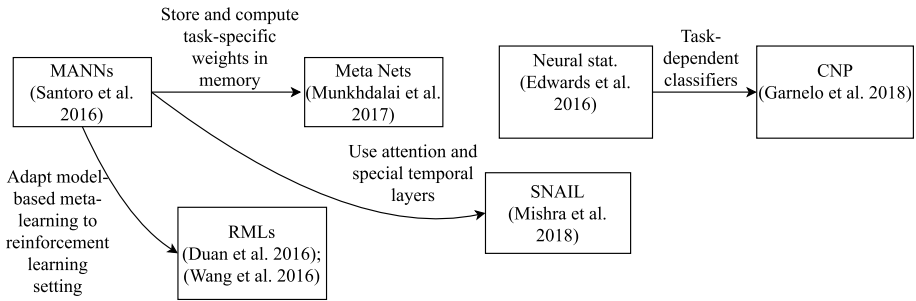


Fig. 20 The relationships between the covered model-based meta-learning techniques. The neural statistician and conditional neural processes form an island in the model-based approaches

4.8 Model-based techniques, in conclusion

In this section, we have discussed various model-based techniques. Despite apparent differences, they all build on the notion of task internalization. That is, tasks are processed and represented in the state of the model-based system. This state can then be used to make predictions. Figure 20 displays the relationships between the covered model-based techniques.

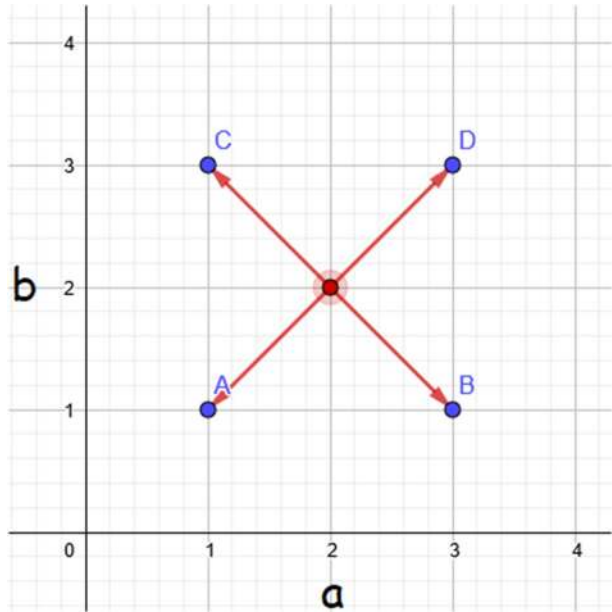
Memory-augmented neural networks (Santoro et al. 2016) mark the beginning of deep model-based meta-learning techniques. They use the idea of feeding the entire support set in a sequential fashion into the model and then making predictions for the query set inputs using the internal state of the model. Such a model-based approach, where inputs sequentially enter the model was also taken by recurrent meta-learners (RMLs) (Duan et al. 2016; Wang et al. 2016) in the reinforcement learning setting. Meta networks (Munkhdalai and Yu 2017) also use a large black-box solution but generate task-specific weights for every task that is encountered. SNAIL (Mishra et al. 2018) tries to improve the memory capacity and ability to pinpoint memories, which is limited in recurrent neural networks, by using attention mechanisms coupled with special temporal layers. Lastly, the neural statistician and conditional neural processes are two techniques that try to learn meta-features of data sets in an end-to-end fashion. The neural statistician uses the distance between meta-features to make class predictions, while the conditional neural process conditions classifiers on these features.

Advantages of model-based approaches include the flexibility of the internal dynamics of the systems, and their broader applicability compared to most metric-based techniques. However, model-based techniques are often outperformed by metric-based techniques in supervised settings (e.g. graph neural networks (Garcia and Bruna 2017); Sect. 3.6), may not perform well when presented with larger data sets (Hospedales et al. 2020), and generalize less well to more distant tasks than optimization-based techniques (Finn and Levine 2018). We discuss this optimization-based approach next.

5 Optimization-based meta-learning

Optimization-based techniques adopt a different perspective on meta-learning than the previous two approaches. They explicitly optimize for fast learning. Most optimization-based techniques do so by approaching meta-learning as a bi-level optimization problem. At the

Fig. 21 Example of an optimization-based technique, inspired by Finn et al. (2017)



inner-level, a base-learner makes task-specific updates using some optimization strategy (such as gradient descent). At the outer-level, the performance across tasks is optimized.

More formally, given a task $\mathcal{T}_j = (D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test})$ with new input $\mathbf{x} \in D_{\mathcal{T}_j}^{test}$ and base-learner parameters θ , optimization-based meta-learners return

$$P(Y|\mathbf{x}, D_{\mathcal{T}_j}^{tr}) = f_{g_{\varphi}(\theta, D_{\mathcal{T}_j}^{tr}, \mathcal{L}_{\mathcal{T}_j})}(\mathbf{x}), \tag{14}$$

where f is the base-learner, g_{φ} is a (learned) optimizer that makes task-specific updates to the base-learner parameters θ using the support data $D_{\mathcal{T}_j}^{tr}$, and loss function $\mathcal{L}_{\mathcal{T}_j}$.

5.1 Example

Suppose we are faced with a linear regression problem, where every task is associated with a different function $f(x)$. For this example, suppose our model only has two parameters: a and b , which together form the function $\hat{f}(x) = ax + b$. Suppose further that our meta-training set consists of four different tasks, i.e., A, B, C, and D. Then, according to the optimization-based view, we wish to find a single set of parameters $\{a, b\}$ from which we can quickly learn the optimal parameters for each of the four tasks, as displayed in Fig. 21. In fact, this is the intuition behind the popular optimization-based technique MAML (Finn et al. 2017). By exposing our model to various meta-training tasks, we can update parameters a and b to facilitate quick adaptation.

We will now discuss the core optimization-based techniques in more detail.

5.2 LSTM optimizer

Standard gradient update rules have the form

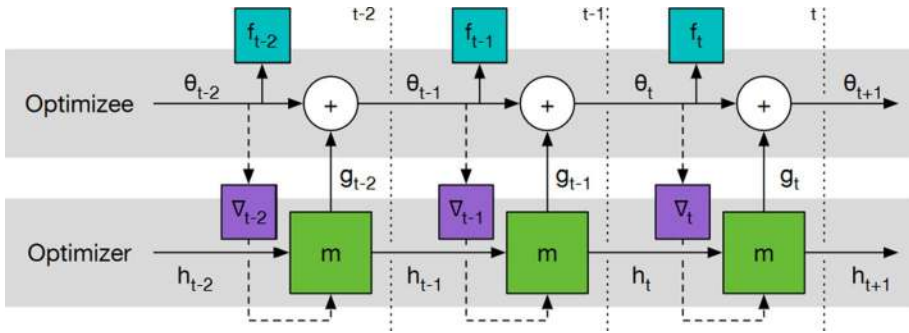


Fig. 22 Workflow of the LSTM optimizer. Gradients can only propagate backwards through solid edges. f_t denotes the observed loss at time step t . Source: Andrychowicz et al. (2016)

$$\theta_{t+1} := \theta_t - \alpha \nabla_{\theta_t} \mathcal{L}_{\mathcal{T}_j}(\theta_t), \tag{15}$$

where α is the learning rate, and $\mathcal{L}_{\mathcal{T}_j}(\theta_t)$ is the loss function with respect to task \mathcal{T}_j and network parameters at time t , i.e., θ_t . The key idea underlying LSTM optimizers (Andrychowicz et al. 2016) is to replace the update term $(-\alpha \nabla \mathcal{L}_{\mathcal{T}_j}(\theta_t))$ by an update proposed by an LSTM g with parameters φ . Then, the new update becomes

$$\theta_{t+1} := \theta_t + g_{\varphi}(\nabla_{\theta_t} \mathcal{L}_{\mathcal{T}_j}(\theta_t)). \tag{16}$$

This new update allows the optimization strategy to be tailored to a specific family of tasks. Note that this is meta-learning, i.e., the LSTM learns to learn. As such, this technique learns an update policy.

The loss function used to train an LSTM optimizer is:

$$\mathcal{L}(\varphi) = \mathbb{E}_{\mathcal{L}_{\mathcal{T}_j}} \left[\sum_{t=1}^T w_t \mathcal{L}_{\mathcal{T}_j}(\theta_t) \right], \tag{17}$$

where T is the number of parameter updates that are made, and w_t are weights indicating the importance of performance after t steps. Note that generally, we are only interested in the final performance after T steps. However, the authors found that the optimization procedure was better guided by equally weighting the performance after each gradient descent step. As is often done, second-order derivatives (arising from the dependency between the updated weights and the LSTM optimizer) were ignored due to the computational expenses associated with the computation thereof. This loss function is fully differentiable and thus allows for training an LSTM optimizer (see Fig. 22). To prevent a parameter explosion, the same network is used for every *coordinate*/weight in the base-learner’s network, causing the update rule to be the same for every parameter. Of course, the updates depend on their prior values and gradients.

The key advantage of LSTM optimizers is that they can enable faster learning compared to hand-crafted optimizers, also on different data sets than those used to train the optimizer. However, Andrychowicz et al. (2016) did not apply this technique to few-shot learning. In fact, they did not apply it across tasks at all. Thus, it is unclear whether this technique

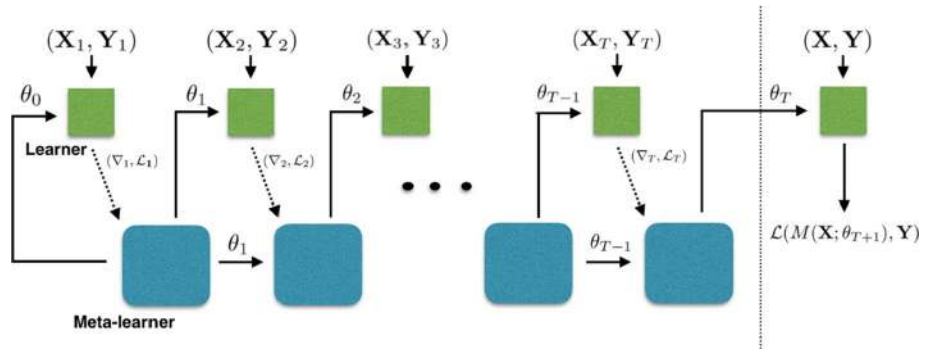


Fig. 23 LSTM meta-learner computation graph. Gradients can only propagate backwards through solid edges. The base-learner is denoted as M . (X_t, Y_t) are training sets, whereas (X, Y) is the test set. Source: Ravi and Larochelle (2017)

can perform well in few-shot settings, where few data per class are available for training. Furthermore, the question remains whether it can scale to larger base-learner architectures.

5.3 LSTM meta-learner

Instead of having an LSTM predict gradient updates, Ravi and Larochelle (2017) embed the weights of the base-learner parameters into the cell state (long-term memory component) of the LSTM, giving rise to LSTM meta-learners. As such, the base-learner parameters θ are literally inside the LSTM memory component (cell state). In this way, cell state updates correspond to base-learner parameter updates. This idea was inspired by the resemblance between the gradient and cell state update rules. Gradient updates often have the form as shown in Eq. 15. The LSTM cell state update rule, in contrast, looks as follows

$$c_t := f_t \odot c_{t-1} + \alpha_t \odot \bar{c}_t, \tag{18}$$

where f_t is the forget gate (which determines which information should be forgotten) at time t , \odot represents the element-wise product, c_t is the cell state at time t , and \bar{c}_t the candidate cell state for time step t , and α_t the learning rate at time step t . Note that if $f_t = \mathbf{1}$ (vector of ones), $\alpha_t = \alpha$, $c_{t-1} = \theta_{t-1}$, and $\bar{c}_t = -\nabla_{\theta_{t-1}} \mathcal{L}_{T_t}(\theta_{t-1})$, this update is equivalent to the one used by gradient-descent. This similarity inspired Ravi and Larochelle (2017) to use an LSTM as a meta-learner that learns to make updates for a base-learner, as shown in Fig. 23.

More specifically, the cell state of the LSTM is initialized with $c_0 = \theta_0$, which will be adjusted by the LSTM to a good common initialization point across different tasks. Then, to update the weights of the base-learner for the next time step $t + 1$, the LSTM computes c_{t+1} , and sets the weights of the base-learner equal to that. There is thus a one-to-one correspondence between c_t and θ_t . The meta-learner’s learning rate α_t (see Eq. 18), is set equal to $\sigma(\mathbf{w}_\alpha \cdot [\nabla_{\theta_{t-1}} \mathcal{L}_{T_t}(\theta_{t-1}), \mathcal{L}_{T_t}(\theta_t), \theta_{t-1}, \alpha_{t-1}] + \mathbf{b}_\alpha)$, where σ is the sigmoid function. Note that the output is a vector, with values between 0 and 1, which denote the the learning rates for the corresponding parameters. Furthermore, \mathbf{w}_α and \mathbf{b}_α are trainable parameters that part of the LSTM meta-learner. In words, the learning rate at any time depends on the loss gradients, the loss value, the previous parameters, and the previous learning rate. The

forget gate, f_t , determines what part of the cell state should be forgotten, and is computed in a similar fashion, but with different weights.

To prevent an explosion of meta-learner parameters, weight-sharing is used, in a similar fashion to LSTM optimizers proposed by Andrychowicz et al. (2016) (Sect. 5.2). This implies that the same update rule is applied to every weight at a given time step. The exact update, however, depends on the history of that specific parameter in terms of the previous learning rate, loss, etc. For simplicity, second-order derivatives were ignored, by assuming the base-learner's loss does not depend on the cell state of the LSTM optimizer. Batch normalization was applied to stabilize and speed up the learning process.

In short, LSTM optimizers can learn to optimize a base-learner by maintaining a one-to-one correspondence over time between the base-learner's weights and the LSTM cell state. This allows the LSTM to exploit commonalities in the tasks, allowing for quicker optimization. However, there are simpler approaches (e.g. MAML Finn et al. 2017) that outperform this technique.

5.4 Reinforcement learning optimizer

Li and Malik (2018) proposed a framework that casts optimization as a reinforcement learning problem. Optimization can then be performed by existing reinforcement learning techniques. At a high-level, an optimization algorithm g takes as input an initial set of weights θ_0 and a task \mathcal{T}_j with the corresponding loss function $\mathcal{L}_{\mathcal{T}_j}$, and produces a sequence of new weights $\theta_1, \dots, \theta_T$, where θ_T is the final solution found. On this sequence of proposed new weights, we can define a loss function \mathcal{L} that captures unwanted properties (e.g. slow convergence, oscillations, etc.). The goal of learning an optimizer can then be formulated more precisely as follows. We wish to learn an optimal optimizer

$$g^* = \underset{g}{\operatorname{argmin}} \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T}), \theta_0 \sim p(\theta_0)} [\mathcal{L}(g(\mathcal{L}_{\mathcal{T}_j}, \theta_0))] \quad (19)$$

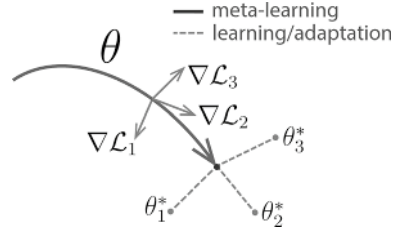
The key insight is that the optimization can be formulated as a Partially Observable Markov Decision Process (POMDP). Then, the state corresponds to the current set of weights θ_t , the action to the proposed update at time step t , i.e., $\Delta\theta_t$, and the policy to the function that computes the update. With this formulation, the optimizer g can be learned by existing reinforcement learning techniques. In their paper, they used a recurrent neural network as an optimizer. At each time step, they feed it observation features, which depend on the previous set of weights, loss gradients, and objective functions, and use guided policy search to train it.

In summary, Li and Malik (2018) made the first step towards general optimization through reinforcement learning optimizers, which were shown able to generalize across network architectures and data sets. However, the base-learner architecture that was used was quite small. The question remains whether this approach can scale to larger architectures.

5.5 MAML

Model-agnostic meta-learning (MAML) (Finn et al. 2017) uses a simple gradient-based inner optimization procedure (e.g. stochastic gradient descent), instead of more complex LSTM procedures or procedures based on reinforcement learning. The key idea of MAML is to explicitly optimize for fast adaptation to new tasks by learning a good set

Fig. 24 MAML learns an initialization point from which it can perform well on various tasks. Source: Finn et al. (2017)



of initialization parameters θ . This is shown in Fig. 24: from the learned initialization θ , we can quickly move to the best set of parameters for the task T_j , i.e., θ_j^* for $j = 1, 2, 3$. The learned initialization can be seen as the *inductive bias* of the model, or simply the set of assumptions (encapsulated in θ) that the model makes with respect to the overall task structure.

More formally, let θ denote the initial model parameters of a model. The goal is to quickly learn new concepts, which is equivalent to achieving a minimal loss in few gradient update steps. The amount of gradient steps s has to be specified upfront, such that MAML can explicitly optimize for achieving good performance within that number of steps. Suppose we pick only one gradient update step, i.e., $s = 1$. Then, given a task $T_j = (D_{T_j}^r, D_{T_j}^{test})$, gradient descent would produce updated parameters (fast weights)

$$\theta'_j = \theta - \alpha \nabla_{\theta} \mathcal{L}_{D_{T_j}^r}(\theta), \tag{20}$$

specific to task j . The *meta-loss* of quick adaptation (using $s = 1$ gradient steps) across tasks can then be formulated as

$$ML := \sum_{T_j \sim p(T)} \mathcal{L}_{D_{T_j}^{test}}(\theta'_j) = \sum_{T_j \sim p(T)} \mathcal{L}_{D_{T_j}^{test}}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{D_{T_j}^r}(\theta)), \tag{21}$$

where $p(T)$ is a probability distribution over tasks. This expression contains an inner gradient ($\nabla_{\theta} \mathcal{L}_{T_j}(\theta_j)$). As such, by optimizing this meta-loss using gradient-based techniques, we have to compute second-order gradients. One can easily see this in the computation below

$$\begin{aligned} \nabla_{\theta} ML &= \nabla_{\theta} \sum_{T_j \sim p(T)} \mathcal{L}_{D_{T_j}^{test}}(\theta'_j) \\ &= \sum_{T_j \sim p(T)} \nabla_{\theta} \mathcal{L}_{D_{T_j}^{test}}(\theta'_j) \\ &= \sum_{T_j \sim p(T)} \mathcal{L}'_{D_{T_j}^{test}}(\theta'_j) \nabla_{\theta}(\theta'_j) \\ &= \sum_{T_j \sim p(T)} \mathcal{L}'_{D_{T_j}^{test}}(\theta'_j) \nabla_{\theta}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{D_{T_j}^r}(\theta)) \\ &= \underbrace{\sum_{T_j \sim p(T)} \mathcal{L}'_{D_{T_j}^{test}}(\theta'_j)}_{\text{FOMAML}} (\nabla_{\theta} \theta - \alpha \nabla_{\theta}^2 \mathcal{L}_{D_{T_j}^r}(\theta)), \end{aligned} \tag{22}$$

where we used $\mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j)$ to denote the derivative of the loss function with respect to the query set, evaluated at the post-update parameters θ'_j . The term $\alpha \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta)$ contains the second-order gradients. The computation thereof is expensive in terms of time and memory costs, especially when the optimization trajectory is large (when using a larger number of gradient updates s per task). Finn et al. (2017) experimented with leaving out second-order gradients, by assuming $\nabla_{\theta} \theta'_j = I$, giving us First-Order MAML (FOMAML, see Eq. 22). They found that FOMAML performed reasonably similarly to MAML. This means that updating the initialization using only first-order gradients $\sum_{\mathcal{T}_j \sim p(\mathcal{T})} \mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j)$ is roughly equal to using the full gradient expression of the meta-loss in Eq. 22. One can extend the meta-loss to incorporate multiple gradient steps by substituting θ'_j by a multi-step variant.

MAML is trained as follows. The initialization weights θ are updated by continuously sampling a batch of m tasks $B = \{\mathcal{T}_j \sim p(\mathcal{T})\}_{i=1}^m$. Then, for every task $\mathcal{T}_j \in B$, an *inner update* is performed to obtain θ'_j , granting an observed loss $\mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta'_j)$. These losses across a batch of tasks are used in the *outer update*

$$\theta := \theta - \beta \nabla_{\theta} \sum_{\mathcal{T}_j \in B} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta'_j). \tag{23}$$

The complete training procedure of MAML is displayed in Algorithm 2. At test-time, when presented with a new task \mathcal{T}_j , the model is initialized with θ , and performs a number of gradient updates on the task data. Note that the algorithm for FOMAML is equivalent to Algorithm 2, except for the fact that the update on line 8 is done differently. That is, FOMAML updates the initialization with the rule $\theta = \theta - \beta \sum_{\mathcal{T}_j \sim p(\mathcal{T})} \mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j)$.

Algorithm 2 One-step MAML for supervised learning, by Finn et al. (2017)

- 1: Randomly initialize θ
 - 2: **while** not done **do**
 - 3: Sample batch of J tasks $B = \mathcal{T}_1, \dots, \mathcal{T}_J \sim p(\mathcal{T})$
 - 4: **for** $\mathcal{T}_j = (D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test}) \in B$ **do**
 - 5: Compute $\nabla_{\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta)$
 - 6: Compute $\theta'_j = \theta - \alpha \nabla_{\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta)$
 - 7: **end for**
 - 8: Update $\theta = \theta - \beta \nabla_{\theta} \sum_{\mathcal{T}_j \in B} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta'_j)$
 - 9: **end while**
-

Antoniou et al. (2019), in response to MAML, proposed many technical improvements that can improve training stability, performance, and generalization ability. Improvements include (i) updating the initialization θ after every inner update step (instead of after all steps are done) to increase gradient propagation, (ii) using second-order gradients only after 50 epochs to increase the training speed, (iii) learning layer-wise learning rates to improve flexibility, (iv) annealing the meta-learning rate β over time, and (v) some Batch Normalization tweaks (keep running statistics instead of batch-specific ones, and using per-step biases).

MAML has obtained great attention within the field of Deep Meta-Learning, perhaps due to its (i) simplicity (only requires two hyperparameters), (ii) general applicability, and (iii) strong performance. A downside of MAML, as mentioned above, is that it can be quite expensive in terms of running time and memory to optimize a base-learner for every task and compute higher-order derivatives from the optimization trajectories.

5.6 iMAML

Instead of ignoring higher-order derivatives (as done by FOMAML), which potentially decreases the performance compared to regular MAML, iMAML (Rajeswaran et al. 2019) approximates these derivatives in a way that is less memory-consuming.

Let \mathcal{A} denote an inner optimization algorithm (e.g., stochastic gradient descent), which takes a support set $D_{\mathcal{T}_j}^{tr}$ corresponding to task \mathcal{T}_j and initial model weights θ , and produces new weights $\theta'_j = \mathcal{A}(\theta, D_{\mathcal{T}_j}^{tr})$. MAML has to compute the derivative

$$\nabla_{\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta'_j) = \mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j) \nabla_{\theta}(\theta'_j), \tag{24}$$

where $D_{\mathcal{T}_j}^{test}$ is the query set corresponding to task \mathcal{T}_j . This equation is a simple result of applying the chain rule. Importantly, note that $\nabla_{\theta}(\theta'_j)$ differentiates through $\mathcal{A}(\theta, D_{\mathcal{T}_j}^{tr})$, while $\mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j)$ does not, as it represents the gradient of the loss function evaluated at θ'_j .

Rajeswaran et al. (2019) make use of the following lemma.

If $(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j))$ is invertible (i.e., $(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j))^{-1}$ exists), then

$$\nabla_{\theta}(\theta'_j) = \left(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j) \right)^{-1}. \tag{25}$$

Here, λ is a regularization parameter. The reason for this is discussed below.

Combining Eqs. 24 and 25, we have that

$$\nabla_{\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta'_j) = \mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j) \left(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j) \right)^{-1}. \tag{26}$$

The idea is to obtain an approximate gradient vector \mathbf{g}_j that is close to this expression, i.e., we want the difference to be small

$$\mathbf{g}_j - \mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j) \left(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j) \right)^{-1} = \epsilon, \tag{27}$$

for some small tolerance vector ϵ . If we multiply both sides by the inverse of the inverse, i.e., $\left(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j) \right)$, we get

$$\mathbf{g}_j^T \left(\mathbf{I} + \frac{1}{\lambda} \nabla_{\theta}^2 \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta'_j) \right) \mathbf{g}_j - \mathbf{g}_j^T \mathcal{L}'_{D_{\mathcal{T}_j}^{test}}(\theta'_j) = \epsilon', \tag{28}$$

where ϵ' absorbed the multiplication factor. We wish to minimize this expression for \mathbf{g}_j , and that can be performed using optimization techniques such as the conjugate gradient algorithm (Rajeswaran et al. 2019). This algorithm does not need to store Hessian matrices,

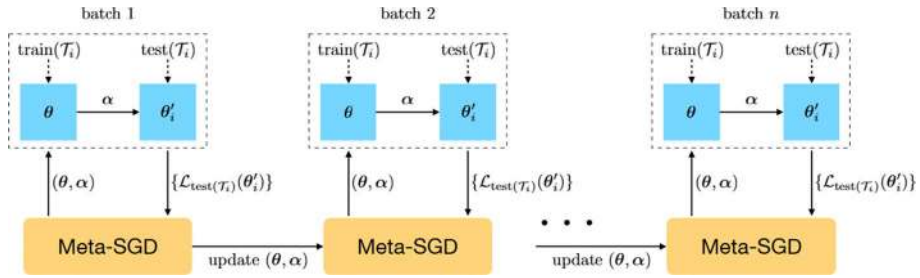


Fig. 25 Meta-SGD learning process. Source: Li et al. (2017)

which decreases the memory cost significantly. In turn, this allows iMAML to work with more inner gradient update steps. Note, however, that one needs to perform explicit regularization in that case to avoid overfitting. The conventional MAML did not require this, as it uses only a few number of gradient steps (equivalent to an early stopping mechanism).

At each inner loop step, iMAML computes the meta-gradient g_j . After processing a batch of tasks, these gradients are averaged and used to update the initialization θ . Since it does not differentiate through the optimization process, we are free to use any other (non-differentiable) inner-optimizer.

In summary, iMAML reduces memory costs significantly as it need not differentiate through the optimization trajectory, also allowing for greater flexibility in the choice of the inner optimizer. Additionally, it can account for larger optimization paths. The computational costs stay roughly the same compared to MAML (Finn et al. 2017). Future work could investigate more inner optimization procedures (Rajeswaran et al. 2019).

5.7 Meta-SGD

Meta-SGD (Li et al. 2017), or meta-stochastic gradient descent, is similar to MAML (Finn et al. 2017) (Sect. 5.5). However, on top of learning an initialization, Meta-SGD also learns learning rates for every model parameter in θ , building on the insight that the optimizer can be seen as trainable entity.

The standard SGD update rule is given in Eq. 15. The meta-SGD optimizer uses a more general update, namely

$$\theta'_j \leftarrow \theta - \alpha \odot \nabla_{\theta} \mathcal{L}_{D_{\mathcal{T}_j}^r}(\theta), \tag{29}$$

where \odot is the element-wise product. Note that this means that alpha (learning rate) is now a vector—hence the bold font— instead of scalar, which allows for greater flexibility in the sense that each parameter has its own learning rate. The goal is to learn the initialization θ , and learning rate vector α , such that the generalization ability is as large as possible. More mathematically precise, the learning objective is

$$\min_{\alpha, \theta} \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} [\mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta'_j)] = \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} [\mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\theta - \alpha \odot \nabla_{\theta} \mathcal{L}_{D_{\mathcal{T}_j}^r}(\theta))], \tag{30}$$

where we used a simple substitution for θ'_j . $\mathcal{L}_{D_{\mathcal{T}_j}^r}$ and $\mathcal{L}_{D_{\mathcal{T}_j}^{test}}$ are the losses computed on the support and query set respectively. Note that this formulation stimulates generalization

ability (as it includes the query set loss $\mathcal{L}_{D_{T_j}^{test}}$, which can be observed during the meta-training phase). The learning process is visualized in Fig. 25. Note that the meta-SGD optimizer is trained to maximize generalization ability after only one update step. Since this learning objective has a fully differentiable loss function, the meta-SGD optimizer itself can be trained using standard SGD.

In summary, Meta-SGD is more expressive than MAML as it does not only learn an initialization but also learning rates per parameter. This, however, does come at the cost of an increased number of hyperparameters.

5.8 Reptile

Reptile (Nichol et al. 2018) is another optimization-based technique that, like MAML (Finn et al. 2017), solely attempts to find a good set of initialization parameters θ . The way in which Reptile attempts to find this initialization is quite different from MAML. It repeatedly samples a task, trains on the task, and moves the model weights towards the trained weights (Nichol et al. 2018). Algorithm 3 displays the pseudocode describing this simple process.

Algorithm 3 Reptile, by Nichol et al. (2018)

- 1: Initialize θ
 - 2: **for** $i = 1, 2, \dots$ **do**
 - 3: Sample task $\mathcal{T}_j = (D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test})$ and corresponding loss function $\mathcal{L}_{\mathcal{T}_j}$
 - 4: $\theta'_j = \text{SGD}(\mathcal{L}_{D_{\mathcal{T}_j}^{tr}}, \theta, k)$ ▷ Perform k gradient update steps to get θ'_j
 - 5: $\theta := \theta + \epsilon(\theta'_j - \theta)$ ▷ Move initialization point θ towards θ'_j
 - 6: **end for**
-

Nichol et al. (2018) note that it is possible to treat $(\theta - \theta'_j)/\alpha$ as gradients, where α is the learning rate of the inner stochastic gradient descent optimizer (line 4 in the pseudocode), and to feed that into a meta-optimizer (e.g. Adam). Moreover, instead of sampling one task at a time, one could sample a batch of n tasks, and move the initialization θ towards the average update direction $\bar{\theta} = \frac{1}{n} \sum_{j=1}^n (\theta'_j - \theta)$, granting the update rule $\theta := \theta + \epsilon \bar{\theta}$.

The intuition behind Reptile is that updating the initialization weights towards updated parameters will grant a good inductive bias for tasks from the same family. By performing Taylor expansions of the gradients of Reptile and MAML (both first-order and second-order), Nichol et al. (2018) show that the expected gradients differ in their direction. They argue, however, that in practice, the gradients of Reptile will also bring the model towards a point minimizing the expected loss over tasks.

A mathematical argument as to why Reptile works goes as follows. Let θ denote the initial parameters, and θ_j^* the optimal set of weights for the task \mathcal{T}_j . Lastly, let d be the Euclidean distance function. Then, the goal is to minimize the distance between the initialization point θ and the optimal point θ_j^* , i.e.,

$$\min_{\theta} \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} \left[\frac{1}{2} d(\theta, \theta_j^*)^2 \right]. \tag{31}$$

The gradient of this expected distance with respect to the initialization θ is given by

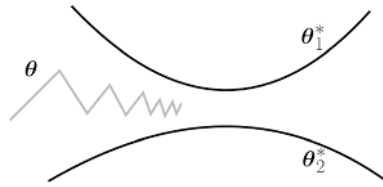


Fig. 26 Schematic visualization of Reptile’s learning trajectory. Here, θ_1^* and θ_2^* are the optimal weights for tasks \mathcal{T}_1 and \mathcal{T}_2 respectively. The initialization parameters θ oscillate between these. Adapted from Nichol et al. (2018)

$$\begin{aligned} \nabla_{\theta} \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} \left[\frac{1}{2} d(\theta, \theta_j^*)^2 \right] &= \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} \left[\frac{1}{2} \nabla_{\theta} d(\theta, \theta_j^*)^2 \right] \\ &= \mathbb{E}_{\mathcal{T}_j \sim p(\mathcal{T})} [\theta - \theta_j^*], \end{aligned} \tag{32}$$

where we used the fact that the gradient of the squared Euclidean distance between two points \mathbf{x}_1 and \mathbf{x}_2 is the vector $2(\mathbf{x}_1 - \mathbf{x}_2)$. Nichol et al. (2018) go on to argue that performing gradient descent on this objective would result in the following update rule

$$\begin{aligned} \theta &= \theta - \epsilon \nabla_{\theta} \frac{1}{2} d(\theta, \theta_j^*)^2 \\ &= \theta - \epsilon (\theta_j^* - \theta). \end{aligned} \tag{33}$$

Since we do not know $\theta_{\mathcal{T}_j}^*$, one can approximate this by term by k steps of gradient descent $SGD(\mathcal{L}_{\mathcal{T}_j}, \theta, k)$. In short, Reptile can be seen as gradient descent on the distance minimization objective given in Eq. 31. A visualization is shown in Fig. 26. The initialization θ is moving towards the optimal weights for tasks 1 and 2 in an interleaved fashion (hence the oscillations).

In conclusion, Reptile is an extremely simple meta-learning technique, which does not need to differentiate through the optimization trajectory like, e.g., MAML (Finn et al. 2017), saving time and memory costs. However, the theoretical foundation is a bit weaker due to the fact that it does not directly optimize for fast learning as done by MAML, and performance may be a bit worse than that of MAML in some settings.

5.9 LEO

Latent Embedding Optimization (LEO) was proposed by Rusu et al. (2018) to combat an issue of gradient-based meta-learners, such as MAML (Finn et al. 2017) (see Sect. 5.5), in few-shot settings (N -way, k -shot). These techniques operate in a high-dimensional parameter space using gradient information from only a few examples, which could lead to poor generalization.

LEO alleviates this issue by learning a lower-dimensional latent embedding space, which indirectly allows us to learn a good set of initial parameters θ . Additionally, the embedding space is conditioned upon tasks, allowing for more expressivity. In theory, LEO could find initial parameters for the entire base-learner network, but the authors only experimented with setting the parameters for the final layers.

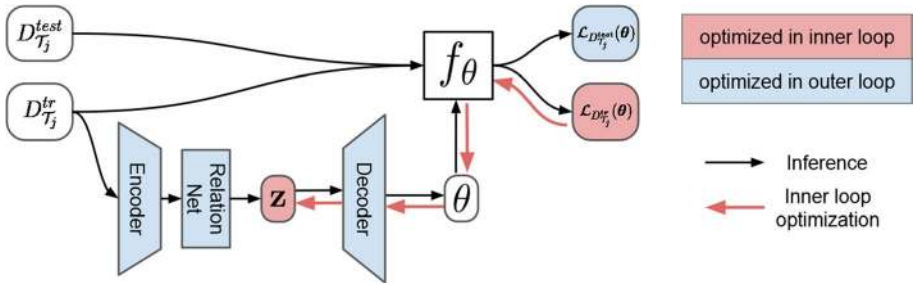


Fig. 27 Workflow of LEO, adapted from Rusu et al. (2018)

The complete workflow of LEO is shown in Fig. 27. As we can see, given a task \mathcal{T}_j , the corresponding support set $D_{\mathcal{T}_j}^{tr}$ is fed into an encoder, which produces hidden codes for each example in that set. These hidden codes are paired and concatenated in every possible manner, granting us $(Nk)^2$ pairs, where N is the number of classes in the training set, and k the number of examples per class. These paired codes are then fed into a relation network (Sung et al. 2018) (see Sect. 3.5). The resulting embeddings are grouped by class, and parameterize a probability distribution over latent codes z_n (for class n) in a low dimensional space \mathcal{Z} . More formally, let x_n^ℓ denote the ℓ -th example of class n in $D_{\mathcal{T}_j}^{tr}$. Then, the mean μ_n^e and variance σ_n^e of a Gaussian distribution over latent codes for class n are computed as

$$\mu_n^e, \sigma_n^e = \frac{1}{Nk^2} \sum_{\ell_p=1}^k \sum_{m=1}^N \sum_{\ell_q=1}^k g_{\phi_r} \left(g_{\phi_e}(x_n^{\ell_p}), g_{\phi_e}(x_m^{\ell_q}) \right), \tag{34}$$

where ϕ_r, ϕ_e are parameters for the relation network and encoder respectively. Intuitively, the three summations ensure that every example with class n in $D_{\mathcal{T}_j}^{tr}$ is paired with every example from all classes n . Given μ_n^e , and σ_n^e , one can sample a latent code $z_n \sim N(\mu_n^e, \text{diag}(\sigma_n^{e2}))$ for class n , which serves as a latent embedding of the task training data.

The decoder can then generate a task-specific initialization θ_n for class n as follows. First, one computes a mean and variance for a Gaussian distribution using the latent code

$$\mu_n^d, \sigma_n^d = g_{\phi_d}(z_n). \tag{35}$$

These are then used to sample initialization weights $\theta_n \sim N(\mu_n^d, \text{diag}(\sigma_n^{d2}))$. The loss from the generated weights can then be propagated backwards to adjust the embedding space. In practice, generating such high-dimensional set of parameters from a low-dimensional embedding can be quite problematic. Therefore, LEO uses pre-trained models, and only generates weights for the final layer, which limits the expressivity of the model.

A key advantage of LEO is that it optimizes in a lower-dimensional latent embedding space, which aids generalization performance. However, the approach is more complex than e.g. MAML (Finn et al. 2017), and its applicability is limited to few-shot learning settings.

5.10 Online MAML (FTML)

Online MAML (Finn et al. 2019) is an extension of MAML (Finn et al. 2017) to make it applicable to *online learning* settings (Anderson 2008). In the online setting, we are presented with a sequence of tasks \mathcal{T}_t with corresponding loss functions $\{\mathcal{L}_{\mathcal{T}_t}\}_{t=1}^T$, for some potentially infinite time horizon T . The goal is to pick a sequence of parameters $\{\theta_t\}_{t=1}^T$ that performs well on the presented loss functions. This objective is captured by the $Regret_T$ over the entire sequence, which is defined by Finn et al. (2019) as follows

$$Regret_T = \sum_{t=1}^T \mathcal{L}_{\mathcal{T}_t}(\theta'_t) - \min_{\theta} \sum_{t=1}^T \mathcal{L}_{\mathcal{T}_t}(\theta'_t), \quad (36)$$

where θ are the initial model parameters (just as MAML), and θ'_t are parameters resulting from a one-step gradient update (starting from θ) on task t . Here, the left term reflects the updated parameters chosen by the agent (θ_t), whereas the right term presents the minimum obtainable loss (in hindsight) from a single fixed set of parameters θ . Note that this setup assumes that the agent can make updates to its chosen parameters (transform its initial choice at time t from θ_t to θ'_t).

Finn et al. (2019) propose FTML (Follow The Meta Leader), inspired by FTL (Follow The Leader) (Hannan 1957; Kalai and Vempala 2005), to minimize the regret. The basic idea is to set the parameters for the next time step ($t + 1$) equal to the best parameters in hindsight, i.e.,

$$\theta_{t+1} := \operatorname{argmin}_{\theta} \sum_{k=1}^t \mathcal{L}_{\mathcal{T}_k}(\theta'_k). \quad (37)$$

The gradient to perform meta-updates is then given by

$$g_t(\theta) := \nabla_{\theta} \mathbb{E}_{\mathcal{T}_k \sim p_t(\mathcal{T})} \mathcal{L}_{\mathcal{T}_k}(\theta'_k), \quad (38)$$

where $p_t(\mathcal{T})$ is a uniform distribution over tasks $1, \dots, t$ (at time t).

Algorithm 4 contains the full pseudocode for FTML. In this algorithm, *MetaUpdate* performs a few (N_{meta}) meta-steps. In each meta-step, a task is sampled from B , together with train and test mini-batches to compute the gradient g_t in Eq. 37. The initialization θ is then updated ($\theta := \theta - \beta g_t(\theta)$), where β is the meta-learning rate. Note that the memory usage keeps increasing over time, as at every time step t , we append tasks to the buffer B , and keep task data sets in memory.

Algorithm 4 FTML by Finn et al. (2019)

Require: Performance threshold γ

- 1: Initialize empty task buffer B
- 2: **for** $t = 1, \dots$ **do**
- 3: Initialize data set $D_t = \emptyset$
- 4: Append \mathcal{T}_t to B
- 5: **while** $|D_t| < N$ **do**
- 6: Append batch of data $\{(x_i, y_i)\}_{i=1}^n$ to D_t
- 7: $\theta_t = \text{MetaUpdate}(\theta_t, B, t)$
- 8: Compute θ'_t
- 9: **if** $\mathcal{L}_{D_{\mathcal{T}_t}^{\text{test}}}(\theta'_t) < \gamma$ **then**
- 10: Save $|D_t|$ as the efficiency for task \mathcal{T}_t
- 11: **end if**
- 12: **end while**
- 13: Save final performance $\mathcal{L}_{D_{\mathcal{T}_t}^{\text{test}}}(\theta'_t)$
- 14: $\theta_{t+1} = \theta_t$
- 15: **end for**

In summary, Online MAML is a robust technique for online-learning (Finn et al. 2019). A downside of this approach is the computational costs that keep growing over time, as all encountered data are stored. Reducing these costs is a direction for future work. Also, one could experiment with how well the approach works when more than one inner gradient update steps per task are used, as mentioned by Finn et al. (2019).

5.11 LLAMA

Grant et al. (2018) mold MAML into a probabilistic framework, such that a probability distribution over task-specific parameters θ'_j is learned, instead of a single one. In this way, multiple potential solutions can be obtained for a task. The resulting technique is called LLAMA (Laplace Approximation for Meta-Adaptation). Importantly, LLAMA is only developed for supervised learning settings.

A key observation is that a neural network $f_{\theta'_j}$, parameterized by updated parameters θ'_j (obtained from few gradient updates using $D_{\mathcal{T}_j}^r$), outputs class probabilities $P(y_i | x_i, \theta'_j)$. To minimize the error on the query set $D_{\mathcal{T}_j}^{\text{test}}$, the model must output large probability scores for the true classes. This objective is captured in the maximum log-likelihood loss function

$$\mathcal{L}_{D_{\mathcal{T}_j}^{\text{test}}}(\theta'_j) = - \sum_{x_i, y_i \in D_{\mathcal{T}_j}^{\text{test}}} \log P(y_i | x_i, \theta'_j). \tag{39}$$

Simply put, if we see a task j as a probability distribution over examples $p_{\mathcal{T}_j}$, we wish to maximize the probability that the model predicts the correct class y_i , given an input x_i . This can be done by plain gradient descent, as shown in Algorithm 5, where β is the meta-learning rate. Line 4 refers to ML-LAPLACE, which is a subroutine that computes task-specific updated parameters θ'_j , and estimates the negative log-likelihood (loss function) which is used to update the initialization θ , as shown in Algorithm 6. Grant et al. (2018) approximated the quadratic curvature matrix \hat{H} using K-FAC (Martens and Grosse 2015).

The trick is that the initialization θ defines a distribution $p(\theta'_j|\theta)$ over task-specific parameters θ'_j . This distribution was taken to be a diagonal Gaussian (Grant et al. 2018). Then, to sample solutions for a new task \mathcal{T}_j , one can simply generate possible solutions θ'_j from the learned Gaussian distribution.

Algorithm 5 LLAMA by Grant et al. (2018)

- 1: Initialize θ randomly
 - 2: **while** not converged **do**
 - 3: Sample a batch of J tasks: $B = \mathcal{T}_1, \dots, \mathcal{T}_J \sim p(\mathcal{T})$
 - 4: Estimate $\mathbb{E}_{(\mathbf{x}_i, y_i) \sim p_{\mathcal{T}_j}} [-\log P(y_i|\mathbf{x}_i, \theta)] \forall \mathcal{T}_j \in B$ using ML-LAPLACE
 - 5: $\theta = \theta - \beta \nabla_{\theta} \sum_j \mathbb{E}_{(\mathbf{x}_i, y_i) \sim p_{\mathcal{T}_j}} [-\log P(y_i|\mathbf{x}_i, \theta)]$
 - 6: **end while**
-

Algorithm 6 ML-LAPLACE (Grant et al., 2018)

- 1: $\theta'_j = \theta$
 - 2: **for** $k=1, \dots, K$ **do**
 - 3: $\theta'_j = \theta'_j + \alpha \nabla_{\theta'_j} \log P(y_i \in D_{\mathcal{T}_j}^{tr} | \theta'_j, \mathbf{x}_i \in D_{\mathcal{T}_j}^{tr})$
 - 4: **end for**
 - 5: Compute curvature matrix $\hat{H} = \nabla_{\theta'_j}^2 [-\log P(y_i \in D_{\mathcal{T}_j}^{test} | \theta'_j, \mathbf{x}_i \in D_{\mathcal{T}_j}^{test})] + \nabla_{\theta'_j}^2 [-\log P(\theta'_j | \theta)]$
 - 6: **return** $-\log P(y_i \in D_{\mathcal{T}_j}^{test} | \theta'_j, \mathbf{x}_i \in D_{\mathcal{T}_j}^{test}) + \eta \log[\det(\hat{H})]$
-

In short, LLAMA extends MAML in a probabilistic fashion, such that one can obtain multiple solutions for a single task, instead of one. This does, however, increase the computational costs. On top of that, the used Laplace approximation (in ML-LAPLACE) can be quite inaccurate (Grant et al. 2018).

5.12 PLATIPUS

PLATIPUS (Finn et al. 2018) builds upon the probabilistic interpretation of LLAMA (Grant et al. 2018) but learns a probability distribution over initializations θ , instead of task-specific parameters θ'_j . Thus, PLATIPUS allows one to sample an initialization $\theta \sim p(\theta)$, which can be updated with gradient descent to obtain task-specific weights (fast weights) θ'_j .

Algorithm 7 PLATIPUS training algorithm by Finn et al. (2018)

```

1: Initialize  $\Theta = \{\mu_\theta, \sigma_\theta^2, v_q, \gamma_p, \gamma_q\}$ 
2: while Not done do
3:   Sample batch of tasks  $B = \{\mathcal{T}_j \sim p(\mathcal{T})\}_{i=1}^m$ 
4:   for  $\mathcal{T}_j \in B$  do
5:      $D_{\mathcal{T}_j}^{tr}, D_{\mathcal{T}_j}^{test} = \mathcal{T}_j$ 
6:     Compute  $\nabla_{\mu_\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\mu_\theta)$ 
7:     Sample  $\theta \sim q = N(\mu_\theta - \gamma_q \nabla_{\mu_\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\mu_\theta), v_q)$ 
8:     Compute  $\nabla_\theta \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta)$ 
9:     Compute fast weights  $\theta'_i = \theta - \alpha \nabla_\theta \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\theta)$ 
10:   end for
11:    $p(\theta|D_{\mathcal{T}_j}^{tr}) = N(\mu_\theta - \gamma_p \nabla_{\mu_\theta} \mathcal{L}_{D_{\mathcal{T}_j}^{tr}}(\mu_\theta), \sigma_\theta^2)$ 
12:   Compute  $\nabla_\Theta \left[ \sum_{\mathcal{T}_j} \mathcal{L}_{D_{\mathcal{T}_j}^{test}}(\phi_i) + D_{KL}(q(\theta|D_{\mathcal{T}_j}^{test}), p(\theta|D_{\mathcal{T}_j}^{tr})) \right]$ 
13:   Update  $\Theta$  using the Adam optimizer
14: end while

```

The approach is best explained by its pseudocode, as shown in Algorithm 7. In contrast to the original MAML, PLATIPUS introduces five more parameter vectors (line 1). All of these parameters are used to facilitate the creation of Gaussian distributions over prior initializations (or simply priors) θ . That is, μ_θ represents the vector mean of the distributions. σ_q^2 and v_q represent the covariances of train and test distributions respectively. γ_x for $x = q, p$ are learning rate vectors for performing gradient steps on distributions q (lines 6 and 7) and P (line 11).

The key difference with the regular MAML is that instead of having a single initialization point θ , we now learn distributions over priors: q and P , which are based on query and support data sets of task \mathcal{T}_j respectively. Since these data sets come from the same task, we want the distributions $q(\theta|D_{\mathcal{T}_j}^{test})$, and $p(\theta|D_{\mathcal{T}_j}^{tr})$ to be close to each other. This is enforced by the Kullback–Leibler divergence (D_{KL}) loss term on line 12, which measures the distance between the two distributions. Importantly, note that q (line 7) and P (line 11) use vector means which are computed with one gradient update step using the query and support data sets respectively. The idea is that the mean of the Gaussian distributions should be close to the updated mean μ_θ , because we want to enable fast learning. As one can see, the training process is very similar to that of MAML (Finn et al. 2017) (Sect. 5.5), with some small adjustments to allow us to work with the probability distributions over θ .

At test-time, one can simply sample a new initialization θ from the prior distribution $p(\theta|D_{\mathcal{T}_j}^{tr})$ (note that q cannot be used at test-time as we do not have access to $D_{\mathcal{T}_j}^{test}$), and apply a gradient update on the provided support set $D_{\mathcal{T}_j}^{tr}$. Note that this allows us to sample multiple potential initializations θ for the given task.

The key advantage of PLATIPUS is that it is aware of its uncertainty, which greatly increases the applicability of Deep Meta-Learning in critical domains such as medical diagnosis (Finn et al. 2018). Based on this uncertainty, it can ask for labels of some inputs it is unsure about (active learning). A downside to this approach, however, is the increased computational costs, and the fact that it is not applicable to reinforcement learning.

5.13 Bayesian MAML (BMAML)

Bayesian MAML (Yoon et al. 2018) is another probabilistic variant of MAML that can generate multiple solutions. However, instead of learning a distribution over potential solutions, BMAML simply keeps M possible solutions and jointly optimizes them. Recall that probabilistic MAMLs (e.g., PLATIPUS) attempt to maximize the data likelihood of task \mathcal{T}_j , i.e., $p(\mathbf{y}_j^{test}|\theta'_j)$, where θ'_j are task-specific fast weights obtained by one or more gradient updates. Yoon et al. (2018) model this likelihood using Stein Variational Gradient Descent (SVGD) (Liu and Wang 2016).

To obtain M solutions, or equivalently, parameter settings θ^m , SVGD keeps a set of M particles $\Theta = \{\theta^m\}_{m=1}^M$. At iteration t , every $\theta_t \in \Theta$ is updated as follows

$$\theta_{t+1} = \theta_t + \epsilon(\phi(\theta_t)) \tag{40}$$

$$\text{where } \phi(\theta_t) = \frac{1}{M} \sum_{m=1}^M \left[k(\theta_t^m, \theta_t) \nabla_{\theta_t^m} \log p(\theta_t^m) + \nabla_{\theta_t^m} k(\theta_t^m, \theta_t) \right]. \tag{41}$$

Here, $k(\mathbf{x}, \mathbf{x}')$ is a similarity kernel between \mathbf{x} and \mathbf{x}' . The authors used a radial basis function (RBF) kernel, but in theory, any other kernel could be used. Note that the update of one particle depends on the other gradients of particles. The first term in the summation ($k(\theta_t^m, \theta_t) \nabla_{\theta_t^m} \log p(\theta_t^m)$) moves the particle in the direction of the gradients of other particles, based on particle similarity. The second term ($\nabla_{\theta_t^m} k(\theta_t^m, \theta_t)$) ensures that particles do not collapse (repulsive force) (Yoon et al. 2018).

These particles can then be used to approximate the probability distribution of the test labels

$$p(\mathbf{y}_j^{test}|\theta'_j) \approx \frac{1}{M} \sum_{m=1}^M p(\mathbf{y}_j^{test}|\theta_{\mathcal{T}_j}^m), \tag{42}$$

where $\theta_{\mathcal{T}_j}^m$ is the m -th particle obtained by training on the support set $D_{\mathcal{T}_j}^{tr}$ of task \mathcal{T}_j .

Yoon et al. (2018) proposed a new meta-loss to train BMAML, called the *Chaser Loss*. This loss relies on the insight that we want the approximated parameter distribution (obtained from the support set $p_{\mathcal{T}_j}^n(\theta_{\mathcal{T}_j}|D_{\mathcal{T}_j}^{tr}, \Theta_0)$) and true distribution $p_{\mathcal{T}_j}^\infty(\theta_{\mathcal{T}_j}|D_{\mathcal{T}_j}^{tr} \cup D_{\mathcal{T}_j}^{test})$ to be close to each other (since the task is the same). Here, n denotes the number of SVGD steps, and Θ_0 is the set of initial particles, in a similar fashion to the initial parameters θ seen by MAML. Since the true distribution is unknown, Yoon et al. (2018) approximate it by running SVGD for s additional steps, granting us the *leader* $\theta_{\mathcal{T}_j}^{n+s}$, where the s additional steps are performed on the combined support and query set. The intuition is that as the number of updates increases, the obtained distributions become more like the true ones. $\theta_{\mathcal{T}_j}^n$ in this context is called the *chaser* as it wants to get closer to the leader. The proposed meta-loss is then given by

$$\mathcal{L}_{BMAML}(\Theta_0) = \sum_{\mathcal{T}_j \in \mathcal{B}} \sum_{m=1}^M \|\theta_{\mathcal{T}_j}^{n,m} - \theta_{\mathcal{T}_j}^{n+s,m}\|_2^2. \tag{43}$$

The full pseudocode of BMAML is shown in Algorithm 8. Here, $\Theta_{\mathcal{T}_j}^n(\Theta_0)$ denotes the set of particles after n updates on task \mathcal{T}_j , and SG means “stop gradients” (we do not want the leader to depend on the initialization, as the leader must lead).

Algorithm 8 BMAML by Yoon et al. (2018)

```

1: Initialize  $\Theta_0$ 
2: for  $t=1, \dots$  until convergence do
3:   Sample a batch of tasks  $B$  from  $p(\mathcal{T})$ 
4:   for task  $\mathcal{T}_j \in B$  do
5:     Compute chaser  $\Theta_{\mathcal{T}_j}^n(\Theta_0) = SVGD_n(\Theta_0; D_{\mathcal{T}_j}^{tr}, \alpha)$ 
6:     Compute leader  $\Theta_{\mathcal{T}_j}^{n+s}(\Theta_0) = SVGD_s(\Theta_{\mathcal{T}_j}^n(\Theta_0); D_{\mathcal{T}_j}^{tr} \cup D_{\mathcal{T}_j}^{test}, \alpha)$ 
7:   end for
8:    $\Theta_0 = \Theta_0 - \beta \nabla_{\Theta_0} \sum_{\mathcal{T}_j \in B} d(\Theta_{\mathcal{T}_j}^n(\Theta_0), SG(\Theta_{\mathcal{T}_j}^{n+s}(\Theta_0)))$ 
9: end for
    
```

In summary, BMAML is a robust optimization-based meta-learning technique that can propose M potential solutions to a task. Additionally, it is applicable to reinforcement learning by using Stein Variational Policy Gradient instead of SVGD. A downside of this approach is that one has to keep M parameter sets in memory, which does not scale well. Reducing the memory costs is a direction for future work (Yoon et al. 2018). Furthermore, SVGD is sensitive to the selected kernel function, which was pre-defined in BMAML. However, Yoon et al. (2018) point out that it may be beneficial to learn the kernel function instead. This is another possibility for future research.

5.14 Simple differentiable solvers

Bertinetto et al. (2019) take a quite different approach. That is, they pick simple base-learners that have an analytical closed-form solution. The intuition is that the existence of a closed-form solution allows for good learning efficiency. They propose two techniques using this principle, namely **R2-D2** (Ridge Regression Differentiable Discriminator), and **LR-D2** (Logistic Regression Differentiable Discriminator). We cover both in turn.

Let $g_\phi : X \rightarrow \mathbb{R}^e$ be a pre-trained input embedding model (e.g. a CNN), which outputs embeddings with a dimensionality of e . Furthermore, assume that we use a linear predictor function $f(g_\phi(x_i)) = g_\phi(x_i)W$, where W is a $e \times o$ weight matrix and o is the output dimensionality (of the label). When using (regularized) Ridge Regression (done by R2-D2), one uses the optimal W , i.e.,

$$\begin{aligned}
 W^* &= \arg \min_W ||XW - Y||_2^2 + \gamma ||W||^2 \\
 &= (X^T X + \gamma I)^{-1} X^T Y,
 \end{aligned}
 \tag{44}$$

where $X \in \mathbb{R}^{n \times e}$ is the input matrix, containing n rows (one for each embedded input $g_\phi(x_i)$), $Y \in \mathbb{R}^{n \times o}$ is the output matrix with correct outputs corresponding to the inputs, and γ is a regularization term to prevent overfitting. Note that the analytical solution contains the term $(X^T X) \in \mathbb{R}^{e \times e}$, which is quadratic in the size of the embeddings. Since e can become quite large when using deep neural networks, Bertinetto et al. (2019) use Woodbury’s identity

$$W^* = X^T (XX^T + \gamma I)^{-1} Y, \tag{45}$$

where $XX^T \in \mathbb{R}^{n \times n}$ is linear in the embedding size, and quadratic in the number of examples, which is more manageable in few-shot settings, where n is very small. To make predictions with this Ridge Regression model, one can compute

$$\hat{Y} = \alpha X_{test} W^* + \beta, \tag{46}$$

where α and β are hyperparameters of the base-learner that can be learned by the meta-learner, and $X_{test} \in \mathbb{R}^{m \times e}$ corresponds to the m test inputs of a given task. Thus, the meta-learner needs to learn α, β, γ , and ϕ (embedding weights of the CNN).

The technique can also be applied to iterative solvers when the optimization steps are differentiable (Bertinetto et al. 2019). LR-D2 uses the Logistic Regression objective and Newton’s method as solver. Outputs $\mathbf{y} \in \{-1, +1\}^n$ are now binary. Let \mathbf{w} denote a parameter row of our linear model (parameterized by W). Then, the i -th iteration of Newton’s method updates \mathbf{w}_i as follows

$$\mathbf{w}_i = (X^T \text{diag}(s_i)X + \gamma I)^{-1} X^T \text{diag}(s_i)\mathbf{z}_i, \tag{47}$$

where $\mu_i = \sigma(\mathbf{w}_{i-1}^T X)$, $s_i = \mu_i(1 - \mu_i)$, $\mathbf{z}_i = \mathbf{w}_{i-1}^T X + (\mathbf{y} - \mu_i)/s_i$, and σ is the sigmoid function. Since the term $X^T \text{diag}(s_i)X$ is a matrix of size $e \times e$, and thus again quadratic in the embedding size, Woodburry’s identity is also applied here to obtain

$$\mathbf{w}_i = X^T (XX^T + \lambda \text{diag}(s_i)^{-1})^{-1} \mathbf{z}_i, \tag{48}$$

making it quadratic in the input size, which is not a big problem since n is small in the few-shot setting. The main difference compared to R2-D2 is that the base-solver has to be run for multiple iterations to obtain W .

In the few-shot setting, the base-level optimizers compute the weight matrix W for a given task \mathcal{T}_i . The obtained loss on the query set of a task $\mathcal{L}_{D_{test}}$ is then used to update the parameters ϕ of the input embedding function (e.g. CNN) and the hyperparameters of the base-learner.

Lee et al. (2019) have done similar work to Bertinetto et al. (2019), but with linear Support Vector Machines (SVMs) as base-learner. Their approach is dubbed **MetaOptNet**, and achieved state-of-the-art performance on few-shot image classification.

In short, simple differentiable solvers are simple, reasonably fast in terms of computation time, but limited to few-shot learning settings. Investigating the use of other simple base-learners is a direction for future work.

5.15 Optimization-based techniques, in conclusion

Optimization-based aim to learn new tasks quickly through (learned) optimization procedures. Note that this closely resembles base-level learning, which also occurs through optimization (e.g., gradient descent). However, in contrast to base-level techniques, optimization-based meta-learners can learn the optimizer and/or are exposed to multiple tasks, which allows them to learn to learn new tasks quickly. Figure 28 shows the relationships between the covered optimization-based techniques.

As we can see, the LSTM optimizer (Andrychowicz et al. 2016), which replaces hand-crafted optimization procedures such as gradient descent with a trainable LSTM, can be seen as the starting point for these optimization-based meta-learning techniques. Li and

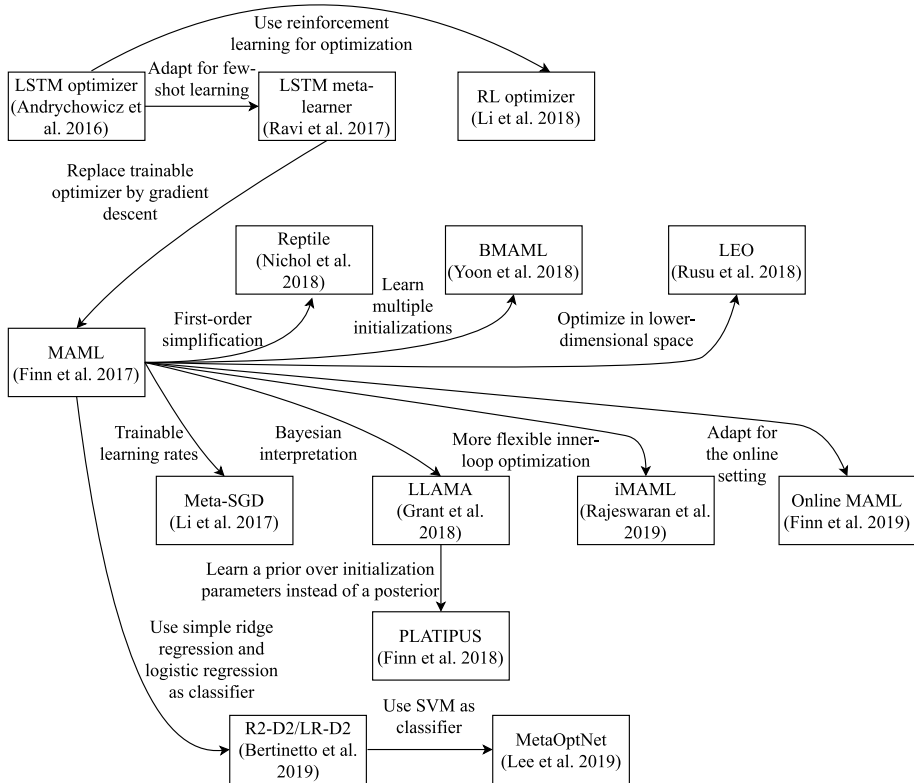


Fig. 28 The relationships between the covered optimization-based meta-learning techniques. As one can see, MAML has a central position in this network of techniques as it has inspired many other works

Malik (2018) also aim to learn the optimization procedure with reinforcement learning instead of gradient-based methods. The LSTM meta-learner (Ravi and Larochelle 2017) extends the LSTM optimizer to the few-shot setting by not only learning the optimization procedure but also a good set of initial weights. This way, it can be used across tasks. MAML (Finn et al. 2017) is a simplification of the LSTM meta-learner as it replaces the trainable LSTM optimizer by hand-crafted gradient descent. MAML has received considerable attention within the field of deep meta-learning, and has, as one can see, inspired many other works.

Meta-SGD is an enhancement of MAML that not only learns the initial parameters, but also the learning rates (Li et al. 2017). LLAMA (Grant et al. 2018), PLATIPUS (Finn et al. 2018), and online MAML (Finn et al. 2019) extend MAML to the active and online learning settings. LLAMA and PLATIPUS are probabilistic interpretations of MAML, which allow them to sample multiple solutions for a given task and quantify their uncertainty. BMAML (Yoon et al. 2018) takes a more discrete approach as it jointly optimizes a discrete set of M initializations. iMAML (Rajeswaran et al. 2019) aims to overcome the computational expenses associated with the computation of second-order derivatives, which is needed by MAML. Through implicit differentiation, they also allow for the use of non-differentiable inner loop optimization procedures. Reptile (Nichol et al. 2018) is an elegant first-order meta-learning algorithm for finding a set of initial parameters and removes the

need of computing higher-order derivatives. LEO (Rusu et al. 2018) tries to improve the robustness of MAML by optimizing in lower-dimensional parameter space through the use of an encoder-decoder architecture. Lastly, R2-D2, LR-D2 (Bertinetto et al. 2019), and Lee et al. (2019) use simple classical machine learning methods (ridge regression, logistic regression, SVM, respectively) as a classifier on top of a learned feature extractor.

A key advantage of optimization-based approaches is that they can achieve better performance on wider task distributions than, e.g., model-based approaches (Finn and Levine 2018). However, optimization-based techniques optimize a base-learner for every task that they are presented with and/or learn the optimization procedure, which is computationally expensive (Hospedales et al. 2020).

Optimization-based meta-learning is a very active area of research. We expect future work to be done in order to reduce the computational demands of these methods and improve the solution quality and level of generalization. We think that benchmarking and reproducibility research will play an important role in these improvements.

6 Concluding remarks

In this section, we give a helicopter view of all that we discussed, and the field of Deep Meta-Learning in general. We will also discuss challenges and future research.

6.1 Overview

In recent years, there has been a shift in focus in the broad meta-learning community. Traditional algorithm selection and hyperparameter optimization for classical machine learning techniques (e.g. Support Vector Machines, Logistic Regression, Random Forests, etc.) have been augmented by Deep Meta-Learning, or equivalently, the pursuit of self-improving neural networks that can leverage prior learning experience to learn new tasks more quickly. Instead of training a new model from scratch for different tasks, we can use the same (meta-learning) model across tasks. As such, meta-learning can widen the applicability of powerful deep learning techniques to domains where less data is available and computational resources are limited.

Deep Meta-Learning techniques are characterized by their meta-objective, which allows them to maximize performance across various tasks, instead of a single one, as is the case in base-level learning objectives. This meta-objective is reflected in the training procedure of meta-learning methods, as they learn on a set of different meta-training tasks. The few-shot setting lends itself nicely towards this end, as tasks consist of few data points. This makes it computationally feasible to train on many different tasks, and it allows us to evaluate whether a neural network can learn new concepts from few examples. Task construction for training and evaluation does require some special attention. That is, it has been shown beneficial to match training and test conditions (Vinyals et al. 2016), and perhaps train in a more difficult setting than the one that will be used for evaluation (Snell et al. 2017).

On a high level, there are three categories of Deep Meta-Learning techniques, namely (i) metric-, (ii) model-, and (iii) optimization-based ones, which rely on (i) computing input similarity, (ii) task embeddings with states, and (iii) task-specific updates, respectively. Each approach has strengths and weaknesses. Metric-learning techniques are simple and effective (Garcia and Bruna 2017), but are not readily applicable outside of the supervised learning setting (Hospedales et al. 2020). Model-based techniques, on the other hand,

can have very flexible internal dynamics, but lack generalization ability to more distant tasks than the ones used at meta-train time (Finn and Levine 2018). Optimization-based approaches have shown greater generalizability, but are in general computationally expensive, as they optimize a base-learner for every task (Finn and Levine 2018; Hospedales et al. 2020).

Table 2 provides a concise, tabular overview of these approaches. Many techniques have been proposed for each one of the categories, and the underlying ideas may vary greatly, even within the same category. Table 3, therefore, provides an overview of all methods and key ideas that we have discussed in this work, together with their applicability to supervised learning (SL) and reinforcement learning (RL) settings, key ideas, and benchmarks that were used for testing them. Table 5 displays an overview of the 1- and 5-shot classification performances (reported by the original authors) of the techniques on the frequently used miniImageNet benchmark. Moreover, it displays the used backbone (feature extraction module) as well as the final classification mechanism. From this table, it becomes clear that the 5-shot performance is typically better than the 1-shot performance, indicating that data scarcity is a large bottleneck for achieving good performance. Moreover, there is a strong relationship between the expressivity of the backbone and the performance. That is, deeper backbones tend to give rise to better classification performance. The best performance is achieved by MetaOptNet, yielding a 1-shot accuracy of 64.09% and a 5-shot accuracy of 80.00%. Note however that MetaOptNet used a deeper backbone than most of the other techniques.

6.2 Open challenges and future work

Despite the great potential of Deep Meta-Learning techniques, there are still open challenges, which we discuss here.

Figure 1 in Sect. 1 displays the accuracy scores of the covered meta-learning techniques on 1-shot miniImageNet classification. Techniques that were not tested in this setting by the original authors are omitted. As we can see, the performance of the techniques is related to the expressivity of the used backbone (ordered in increasing order on the x-axis). For example, the best-performing techniques, LEO and MetaOptNet, use the largest network architectures. Moreover, the fact that different techniques use different backbones poses a problem as it is difficult to fairly compare their classification performance. An obvious question arises to which degree the difference in performance is due to methodological improvements, opposed to the fact that a better backbone architecture was chosen. For this reason, we think that it would be useful to perform a large-scale benchmark test where techniques are compared when they use the same backbones. This would also allow us to get a more clear idea of how the expressivity of the feature extraction module affects the performance.

Another challenge of Deep Meta-Learning techniques is that they can be susceptible to the *memorization problem* (meta-overfitting), where the neural network has memorized tasks seen at meta-training time and fails to generalize to new tasks. More research is required to better understand this problem. Clever task design and meta-regularization may prove useful to avoid such problems (Yin et al. 2020).

Another problem is that most of the meta-learning techniques discussed in this work are evaluated on narrow benchmark sets. This means that the data that the meta-learner used for training are not too distant from the data used for evaluating its performance. As such, one may wonder how well these techniques are able to adapt to more distant tasks.

Table 5 Comparison of the accuracy scores of the covered meta-learning techniques on 1- and 5-shot mini-ImageNet classification

Name	Backbone	Classifier	1-shot	5-shot
<i>Metric-based</i>				
Siamese nets	–	–	–	–
Matching nets	64-64-64-64	Cosine sim.	43.56 ± 0.84	55.31 ± 0.73
Prototypical nets	64-64-64-64	Euclidean dist.	49.42 ± 0.78	68.20 ± 0.66
Relation nets	64-96-128-256	Sim. network	50.44 ± 0.82	65.32 ± 0.70
ARC	–	64-1 dense	49.14 ± –	–
GNN	64-96-128-256	Softmax	50.33 ± 0.36	66.41 ± 0.63
<i>Model-based</i>				
Recurrent ml.	–	–	–	–
MANNs	–	–	–	–
Meta nets	64-64-64-64-64	64-Softmax	49.21 ± 0.96	–
SNAIL	Adj. ResNet-12	Softmax	55.71 ± 0.99	68.88 ± 0.92
CNP	–	–	–	–
Neural stat.	–	–	–	–
<i>Opt.-based</i>				
LSTM optimizer	–	–	–	–
LSTM ml.	32-32-32-32	Softmax	43.44 ± 0.77	60.60 ± 0.71
RL optimizer	–	–	–	–
MAML	32-32-32-32	Softmax	48.70 ± 1.84	63.11 ± 0.92
iMAML	64-64-64-64	Softmax	49.30 ± 1.88	–
Meta-SGD	64-64-64-64	Softmax	50.47 ± 1.87	64.03 ± 0.94
Reptile	32-32-32-32	Softmax	48.21 ± 0.69	66.00 ± 0.62
LEO	WRN-28-10	Softmax	61.76 ± 0.08	77.59 ± 0.12
Online MAML	–	–	–	–
LLAMA	64-64-64-64	Softmax	49.40 ± 1.83	–
PLATIPUS	–	–	–	–
BMAML	64-64-64-64-64	Softmax	53.80 ± 1.46	–
Diff. solvers				
R2-D2	96-192-384-512	Ridge regr.	51.8 ± 0.2	68.4 ± 0.2
LR-D2	96-192-384-512	Log. regr.	51.90 ± 0.20	68.70 ± 0.20
MetaOptNet	ResNet-12	SVM	64.09 ± 0.62	80.00 ± 0.45

Scores are taken from the original papers. The \pm indicates the 95% confidence interval. The backbone is the used feature extraction module. The classifier column shows the final layer(s) that were used to transform the features into class predictions. Used abbreviations: “sim.”: similarity, “Adj.”: adjusted, and “dist.”: distance, “log.”: logistic, “regr.”: regression, “ml.”: meta-learner, “opt.”: optimization

Chen et al. (2019) showed that the ability to adapt to new tasks decreases as they become more distant from the tasks seen at training time. Moreover, a simple non-meta-learning baseline (based on pre-training and fine-tuning) can outperform state-of-the-art meta-learning techniques when meta-test tasks come from a different data set than the one used for meta-training.

In reaction to these findings, Triantafillou et al. (2020) have recently proposed the Meta-Dataset benchmark, which consists of various previously used meta-learning benchmarks

such as Omniglot (Lake et al. 2011) and ImageNet (Deng et al. 2009). This way, meta-learning techniques can be evaluated in more challenging settings where tasks are diverse. Following Hospedales et al. (2020), we think that this new benchmark can prove to be a good means for the investigation and development of meta-learning algorithms for such challenging scenarios.

As mentioned earlier in this section, Deep Meta-Learning has the appealing prospect of widening the applicability of deep learning techniques to more real-world domains. For this, increasing the generalization ability of these techniques is very important. Additionally, the computational costs associated with the deployment of meta-learning techniques should be small. While these techniques can learn new tasks quickly, meta-training can be quite computationally expensive. Thus, decreasing the required computation time and memory costs of Deep Meta-Learning techniques remains an open challenge.

Some real-world problems demand systems that can perform well in online, or active learning settings. The investigation of Deep Meta-Learning in these settings (Finn et al. 2018; Yoon et al. 2018; Finn et al. 2019; Munkhdalai and Yu 2017; Vuorio et al. 2018) remains an important direction for future work.

Yet another direction for future research is the creation of *compositional* Deep Meta-Learning systems, which instead of learning flat and associative functions $x \rightarrow y$, organize knowledge in a *compositional* manner. This would allow them to decompose an input x into several (already learned) components $c_1(x), \dots, c_n(x)$, which in turn could help the performance in low-data regimes (Tokmakov et al. 2019).

The question has been raised whether contemporary Deep Meta-Learning techniques actually learn how to perform rapid learning, or simply learn a set of robust high-level features, which can be (re)used for many (new) tasks. Raghu et al. (2020) investigated this question for the most popular Deep Meta-Learning technique MAML and found that it largely relies on feature reuse. It would be interesting to see whether we can develop techniques that rely more upon fast learning, and what the effect would be on performance.

Lastly, it may be useful to add more meta-abstraction levels, giving rise to, e.g., meta-meta-learning, meta-meta-...-learning (Hospedales et al. 2020; Schmidhuber 1987).

Acknowledgements Thanks to Herke van Hoof for an insightful discussion on LLAMA. Thanks to Pavel Brazdil for his encouragement and feedback on a preliminary version of this work.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson T (2008) The theory and practice of online learning. AU Press, Athabasca University
- Andrychowicz M, Denil M, Colmenarejo SG, Hoffman MW, Pfau D, Schaul T, Shillingford B, de Freitas N (2016) Learning to learn by gradient descent by gradient descent. In: Advances in neural information processing systems 29, Curran Associates Inc., NIPS' 16, pp 3988–3996
- Antoniou A, Edwards H, Storkey A (2019) How to train your MAML. In: International conference on learning representations, ICLR' 19

- Barrett DG, Hill F, Santoro A, Morcos AS, Lillicrap T (2018) Measuring abstract reasoning in neural networks. In: Proceedings of the 35th international conference on machine learning, JMLR.org, ICML'18, pp 4477–4486
- Bengio S, Bengio Y, Cloutier J, Gecsei J (1997) On the optimization of a synaptic learning rule. In: Optimality in artificial and biological neural networks, Lawrence Erlbaum Associates, Inc
- Bengio Y, Bengio S, Cloutier J (1991) Learning a synaptic learning rule. In: International joint conference on neural networks, IEEE, IJCNN'91, vol 2
- Bertinetto L, Henriques JF, Torr PHS, Vedaldi A (2019) Meta-learning with differentiable closed-form solvers. In: International conference on learning representations, ICLR'19
- Brazdil P, Carrier CG, Soares C, Vilalta R (2008) Metalearning: applications to data mining. Springer, Berlin
- Chen WY, Liu YC, Kira Z, Wang YC, Huang JB (2019) A closer look at few-shot classification. In: International conference on learning representations, ICLR'19
- Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L (2009) ImageNet: a large-scale hierarchical image database. In: Proceedings of the IEEE conference on computer vision and pattern recognition, IEEE, pp 248–255
- Duan Y, Schulman J, Chen X, Bartlett PL, Sutskever I, Abbeel P (2016) RL²: fast reinforcement learning via slow reinforcement learning. arXiv preprint [arXiv:161102779](https://arxiv.org/abs/1611.02779)
- Edwards H, Storkey A (2017) towards a neural statistician. In: International conference on learning representations, ICLR'17
- Finn C, Levine S (2018) Meta-Learning and Universality: deep representations and gradient descent can approximate any learning algorithm. In: international conference on learning representations, ICLR'18
- Finn C, Abbeel P, Levine S (2017) Model-agnostic meta-learning for fast adaptation of deep networks. In: Proceedings of the 34th international conference on machine learning, JMLR.org, ICML'17, pp 1126–1135
- Finn C, Xu K, Levine S (2018) Probabilistic model-agnostic meta-learning. In: Advances in neural information processing systems 31, Curran Associates Inc., NIPS'18, pp 9516–9527
- Finn C, Rajeswaran A, Kakade S, Levine S (2019) Online meta-learning. In: Chaudhuri K, Salakhutdinov R (eds) Proceedings of the 36th international conference on machine learning, JMLR.org, ICML'19, pp 1920–1930
- Garcia V, Bruna J (2017) Few-shot learning with graph neural networks. In: International conference on learning representations, ICLR'17
- Garnelo M, Rosenbaum D, Maddison C, Ramalho T, Saxton D, Shanahan M, Teh YW, Rezende D, Eslami SMA (2018) Conditional neural processes. In: Dy J, Krause A (eds) Proceedings of the 35th international conference on machine learning, JMLR.org, ICML'18, vol 80, pp 1704–1713
- Goceri E (2019a) Capsnet topology to classify tumours from brain images and comparative evaluation. IET Image Proc 14(5):882–889
- Goceri E (2019b) Challenges and recent solutions for image segmentation in the era of deep learning. In: 2019 ninth international conference on image processing theory, tools and applications (IPTA), IEEE, pp 1–6
- Goceri E (2020) Convolutional neural network based desktop applications to classify dermatological diseases. In: 2020 IEEE 4th international conference on image processing applications and systems (IPAS), IEEE, pp 138–143
- Goceri E, Karakas AA (2020) Comparative evaluations of CNN based networks for skin lesion classification. In: 14th International conference on computer graphics. visualization, computer vision and image processing (CGVCVIP), Zagreb, Croatia, pp 1–6
- Grant E, Finn C, Levine S, Darrell T, Griffiths T (2018) Recasting gradient-based meta-learning as hierarchical bayes. In: International conference on learning representations, ICLR'18
- Graves A, Wayne G, Danihelka I (2014) Neural turing machines. arXiv preprint [arXiv:14105401](https://arxiv.org/abs/1410.5401)
- Gupta A, Mendonca R, Liu Y, Abbeel P, Levine S (2018) Meta-reinforcement learning of structured exploration strategies. In: Advances in neural information processing systems 31, Curran Associates Inc., NIPS'18, pp 5302–5311
- Hamilton W, Ying Z, Leskovec J (2017) Inductive representation learning on large graphs. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) Advances in neural information processing systems, Curran Associates, Inc., vol 30, <https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7e9-Paper.pdf>
- Hannan J (1957) Approximation to bayes risk in repeated play. *Contribut Theory Games* 3:97–139
- Hastie T, Tibshirani R, Friedman J (2009) The elements of statistical learning: data mining, inference, and prediction, 2nd edn. Springer, New York

- He K, Zhang X, Ren S, Sun J (2015) Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE international conference on computer vision, pp 1026–1034
- Hinton GE, Plaut DC (1987) Using fast weights to deblur old memories. In: Proceedings of the 9th annual conference of the cognitive science society, pp 177–186
- Hochreiter S, Younger AS, Conwell PR (2001) Learning to learn using gradient descent. In: International conference on artificial neural networks, Springer, pp 87–94
- Hospedales T, Antoniou A, Micaelli P, Storkey A (2020) Meta-learning in neural networks: a survey. arXiv preprint [arXiv:200405439](https://arxiv.org/abs/2004.05439)
- Iqbal MS, Luo B, Khan T, Mehmood R, Sadiq M (2018) Heterogeneous transfer learning techniques for machine learning. *Iran J Comput Sci* 1(1):31–46
- Iqbal MS, El-Ashram S, Hussain S, Khan T, Huang S, Mehmood R, Luo B (2019a) Efficient cell classification of mitochondrial images by using deep learning. *J Opt* 48(1):113–122
- Iqbal MS, Luo B, Mehmood R, Alrige MA, Alharbey R (2019b) Mitochondrial organelle movement classification (fission and fusion) via convolutional neural network approach. *IEEE Access* 7:86570–86577
- Iqbal MS, Ahmad I, Bin L, Khan S, Rodrigues JJ (2020) Deep learning recognition of diseased and normal cell representation. *Trans Emerg Telecommun Technol* p e4017
- Jankowski N, Duch W, Grąbczewski K (2011) Meta-learning in computational intelligence, vol 358. Springer, Berlin
- Kalai A, Vempala S (2005) Efficient algorithms for online decision problems. *J Comput Syst Sci* 71(3):291–307
- Koch G, Zemel R, Salakhutdinov R (2015) Siamese neural networks for one-shot image recognition. In: Proceedings of the 32nd international conference on machine learning, JMLR.org, ICML'15, vol 37
- Krizhevsky A (2009) Learning multiple layers of features from tiny images. University of Toronto, Tech. rep
- Krizhevsky A, Sutskever I, Hinton GE (2012) ImageNet classification with deep convolutional neural networks. In: Advances in neural information processing systems, pp 1097–1105
- Lake B, Salakhutdinov R, Gross J, Tenenbaum J (2011) One shot learning of simple visual concepts. In: Proceedings of the annual meeting of the cognitive science society vol 33, pp 2568–2573
- Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ (2017) Building machines that learn and think like people. *Behav Brain Sci* 40
- LeCun Y, Cortes C, Burges C (2010) MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist>, accessed: 7-10-2020
- Lee K, Maji S, Ravichandran A, Soatto S (2019) Meta-learning with differentiable convex optimization. In: Proceedings of the IEEE conference on computer vision and pattern recognition, IEEE, pp 10657–10665
- Li K, Malik J (2018) Learning to Optimize Neural Nets. arXiv preprint [arXiv:170300441](https://arxiv.org/abs/1703.00441)
- Li Z, Zhou F, Chen F, Li H (2017) Meta-SGD: learning to learn quickly for few-shot learning. arXiv preprint [arXiv:170709835](https://arxiv.org/abs/1707.09835)
- Liu Q, Wang D (2016) Stein variational gradient descent: a general purpose bayesian inference algorithm. In: Advances in neural information processing systems 29, Curran Associates Inc., NIPS'16, pp 2378–2386
- Martens J, Grosse R (2015) Optimizing neural networks with kronecker-factored approximate curvature. In: Proceedings of the 32th international conference on machine learning, JMLR.org, ICML'15, pp 2408–2417
- Miconi T, Stanley K, Clune J (2018) Differentiable plasticity: training plastic neural networks with back-propagation. In: Dy J, Krause A (eds) Proceedings of the 35th international conference on machine learning, JMLR.org, ICML'18, pp 3559–3568
- Miconi T, Rawal A, Clune J, Stanley KO (2019) Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity. In: International conference on learning representations, ICLR'19
- Mishra N, Rohaninejad M, Chen X, Abbeel P (2018) A simple neural attentive meta-learner. In: International conference on learning representations, ICLR'18
- Mitchell TM (1980) The need for biases in learning generalizations. Tech. Rep. CBM-TR-117, Rutgers University
- Mnih V, Kavukcuoglu K, Silver D, Graves A, Antonoglou I, Wierstra D, Riedmiller M (2013) Playing atari with deep reinforcement learning. arXiv preprint [arXiv:13125602](https://arxiv.org/abs/1312.5602)
- Munkhdalai T, Yu H (2017) Meta networks. In: Proceedings of the 34th international conference on machine learning, JMLR.org, ICML'17, pp 2554–2563

- Nagabandi A, Clavera I, Liu S, Fearing RS, Abbeel P, Levine S, Finn C (2019) Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. In: International conference on learning representations, ICLR'19
- Naik DK, Mammone RJ (1992) Meta-neural networks that learn by learning. In: International joint conference on neural networks, IEEE, IJCNN'92, vol 1, pp 437–442
- Nichol A, Achiam J, Schulman J (2018) On first-order meta-learning algorithms. arXiv preprint [arXiv:180302999](https://arxiv.org/abs/180302999)
- Oord Avd, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A, Kavukcuoglu K (2016) WaveNet: a generative model for raw audio. arXiv preprint [arXiv:160903499](https://arxiv.org/abs/160903499)
- Oreshkin B, López PR, Lacoste A (2018) Tadam: Task dependent adaptive metric for improved few-shot learning. In: Advances in neural information processing systems 31, Curran Associates Inc., NIPS'18, pp 721–731
- Pan SJ, Yang Q (2009) A survey on transfer learning. *IEEE Trans Knowl Data Eng* 22(10):1345–1359
- Peng Y, Flach PA, Soares C, Brazdil P (2002) Improved Dataset Characterisation for Meta-learning. In: International conference on discovery science, Springer, Lecture Notes in Computer Science vol 2534, pp 141–152
- Raghu A, Raghu M, Bengio S, Vinyals O (2020) Rapid learning or feature reuse? Towards understanding the effectiveness of MAML. In: International conference on learning representations, ICLR'20
- Rajeswaran A, Finn C, Kakade SM, Levine S (2019) Meta-learning with implicit gradients. In: Advances in neural information processing systems 32, Curran Associates Inc., NIPS'19, pp 113–124
- Ravi S, Larochelle H (2017) Optimization as a model for few-shot learning. In: International conference on learning representations, ICLR'17
- Ren M, Triantafillou E, Ravi S, Snell J, Swersky K, Tenenbaum JB, Larochelle H, Zemel RS (2018) Meta-learning for semi-supervised few-shot classification. In: International conference on learning representations, ICLR'18
- Rusu AA, Rao D, Sygnowski J, Vinyals O, Pascanu R, Osindero S, Hadsell R (2018) Meta-learning with latent embedding optimization. In: International conference on learning representations, ICLR'18
- Santoro A, Bartunov S, Botvinick M, Wierstra D, Lillicrap T (2016) Meta-learning with memory-augmented neural networks. In: Proceedings of the 33rd international conference on international conference on machine learning, JMLR.org, ICML'16, pp 1842–1850
- Schmidhuber J (1987) Evolutionary principles in self-referential learning
- Schmidhuber J (1993) A neural network that embeds its own meta-levels. In: IEEE international conference on neural networks, IEEE, pp 407–412
- Schmidhuber J, Zhao J, Wiering M (1997) Shifting inductive bias with success-story algorithm, adaptive levin search, and incremental self-improvement. *Mach Learn* 28(1):105–130
- Shyam P, Gupta S, Dukkipati A (2017) Attentive recurrent comparators. In: Proceedings of the 34th international conference on machine learning, JMLR.org, ICML'17, pp 3173–3181
- Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M, Dieleman S, Grewe D, Nham J, Kalchbrenner N, Sutskever I, Lillicrap T, Leach M, Kavukcuoglu K, Graepel T, Hassabis D (2016) Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587):484
- Snell J, Swersky K, Zemel R (2017) Prototypical networks for few-shot learning. In: Advances in neural information processing systems 30, Curran Associates Inc., NIPS'17, pp 4077–4087
- Sun C, Shrivastava A, Singh S, Gupta A (2017) Revisiting unreasonable effectiveness of data in deep learning era. In: Proceedings of the IEEE international conference on computer vision, pp 843–852
- Sung F, Yang Y, Zhang L, Xiang T, Torr PH, Hospedales TM (2018) Learning to compare: relation network for few-shot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition, IEEE, pp 1199–1208
- Sutton RS, Barto AG (2018) Reinforcement learning: an introduction, 2nd edn. MIT Press, Cambridge
- Taylor ME, Stone P (2009) Transfer learning for reinforcement learning domains: a survey. *J Mach Learn Res* 10(7):1633–1685
- Thrun S (1998) Lifelong learning algorithms. In: Learning to learn, Springer, pp 181–209
- Tokmakov P, Wang YX, Hebert M (2019) Learning compositional representations for few-shot recognition. In: Proceedings of the IEEE international conference on computer vision, pp 6372–6381
- Triantafillou E, Zhu T, Dumoulin V, Lamblin P, Evci U, Xu K, Goroshin R, Gelada C, Swersky K, Manzagol PA, Larochelle H (2020) Meta-dataset: a dataset of datasets for learning to learn from few examples. In: International conference on learning representations, ICLR'20
- Vanschoren J (2018) Meta-learning: a survey. arXiv preprint [arXiv:181003548](https://arxiv.org/abs/181003548)
- Vanschoren J, van Rijn JN, Bischl B, Torgo L (2014) OpenML: networked science in machine learning. *SIGKDD Expl.* 15(2):49–60

- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: *Advances in neural information processing systems* 30, Curran Associates Inc., NIPS'17, pp 5998–6008
- Vinyals O (2017) Talk: model vs optimization meta learning. <http://metalearning-symposium.ml/files/vinyals.pdf>, *neural information processing systems (NIPS'17)*. Accessed 6 June 2020
- Vinyals O, Blundell C, Lillicrap T, Kavukcuoglu K, Wierstra D (2016) Matching networks for one shot learning. In: *Advances in neural information processing systems* 29, Curran Associates Inc., NIPS'16, pp 3637–3645
- Vuorio R, Cho DY, Kim D, Kim J (2018) Meta continual learning. arXiv preprint [arXiv:180606928](https://arxiv.org/abs/1806.06928)
- Wah C, Branson S, Welinder P, Perona P, Belongie S (2011) The Caltech-UCSD Birds-200-2011 Dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology
- Wang JX, Kurth-Nelson Z, Tirumala D, Soyer H, Leibo JZ, Munos R, Blundell C, Kumaran D, Botvinick M (2016) Learning to reinforcement learn. arXiv preprint [arXiv:161105763](https://arxiv.org/abs/1611.05763)
- Wu Y, Schuster M, Chen Z, Le QV, Norouzi M, Macherey W, Krikun M, Cao Y, Gao Q, Macherey K, Klingner J, Shah A, Johnson M, Liu X, Łukasz Kaiser, Gouws S, Kato Y, Kudo T, Kazawa H, Stevens K, Kurian G, Patil N, Wang W, Young C, Smith J, Riesa J, Rudnick A, Vinyals O, Corrado G, Hughes M, Dean J (2016) Google's neural machine translation system: bridging the gap between human and machine translation. arXiv preprint [arXiv:160908144](https://arxiv.org/abs/1609.08144)
- Yin M, Tucker G, Zhou M, Levine S, Finn C (2020) Meta-learning without memorization. In: *International conference on learning representations, ICLR'20*
- Yoon J, Kim T, Dia O, Kim S, Bengio Y, Ahn S (2018) Bayesian model-agnostic meta-learning. In: *Advances in neural information processing systems* 31, Curran Associates Inc., NIPS'18, pp 7332–7342
- Younger AS, Hochreiter S, Conwell PR (2001) Meta-learning with backpropagation. In: *International joint conference on neural networks, IEEE, IJCNN'01*, vol 3
- Yu T, Quillen D, He Z, Julian R, Hausman K, Finn C, Levine S (2019) Meta-world: a benchmark and evaluation for multi-task and meta reinforcement learning. arXiv preprint [arXiv:191010897](https://arxiv.org/abs/1910.10897)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.