



Meta-learning Approaches for Few-Shot Learning: A Survey of Recent Advances

HASSAN GHAROUN, Data Science Institute, University of Technology Sydney, Sydney, Australia

FERESHTEH MOMENIFAR, Western Sydney University, Sydney, Australia

FANG CHEN, Data Science Institute, University of Technology Sydney, Sydney, Australia

AMIR H. GANDOMI, Data Science Institute, University of Technology Sydney, Sydney, Australia and Obuda University, Budapest, Hungary

Despite its astounding success in learning deeper multi-dimensional data, the performance of deep learning declines on new unseen tasks mainly due to its focus on same-distribution prediction. Moreover, deep learning is notorious for poor generalization from few samples. Meta-learning is a promising approach that addresses these issues by adapting to new tasks with few-shot datasets. This survey first briefly introduces meta-learning and then investigates state-of-the-art meta-learning methods and recent advances in: (i) metric-based, (ii) memory-based, (iii), and learning-based methods. Finally, current challenges and insights for future researches are discussed.

CCS Concepts: • **Computing methodologies** → **Image representations; Neural networks; Learning paradigms;**

Additional Key Words and Phrases: Meta-learning, learning to learn, few-shot learning, representation learning

ACM Reference Format:

Hassan Gharoun, Fereshteh Momenifar, Fang Chen, and Amir H. Gandomi. 2024. Meta-learning Approaches for Few-Shot Learning: A Survey of Recent Advances. *ACM Comput. Surv.* 56, 12, Article 294 (July 2024), 41 pages. <https://doi.org/10.1145/3659943>

1 INTRODUCTION

Humans possess the extraordinary capability of learning a new concept even after minimal observation. To a greater extent, a child can distinguish a dog from a cat through a single picture [108]. This critical characteristic of human intelligence lies in the humans' ability to leverage obtained knowledge of prior experiences to unforeseen circumstances with small observation.

Authors' Contact Information: Hassan Gharoun, Data Science Institute, University of Technology Sydney, Sydney, New South Wales, Australia; e-mail: hassan.gharoun@student.uts.edu.au; Fereshteh Momenifar, Western Sydney University, Sydney, New South Wales, Australia; e-mail: 22046851@student.westernsydney.edu.au; Fang Chen, Data Science Institute, University of Technology Sydney, Sydney, New South Wales, Australia; e-mail: fang.chen@uts.edu.au; Amir H. Gandomi, Data Science Institute, University of Technology Sydney, Sydney, New South Wales, Australia and Obuda University, Budapest, Hungary; e-mail: gandomi@uts.edu.au.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2024 Copyright held by the owner/author(s).

ACM 0360-0300/2024/07-ART294

<https://doi.org/10.1145/3659943>

Unlike the human learning paradigm, traditional **machine learning (ML)** and **deep learning (DL)** models train a specific task from scratch through: (a) the training phase in which a model is initiated randomly and then updated, and (b) the test phase in which the model evaluates. While ML and DL have obtained remarkable success in a wide range of applications, they are notorious for requiring a huge number of samples to generalize. In many real-world problems, collecting more data is costly, time-consuming, and even might not be feasible due to physical system constraints [71]. Moreover, most ML and DL models presume that training and testing datasets have the same distribution [112]. Thus, their performance suffers under data distribution shifts [110]. While the model's generalization ability significantly depends on data distribution, training a model on a diverse dataset does not result in domain adaptation [25]. Domain adaptations aim to transfer obtained knowledge from prior experiences to learn new but relevant tasks [114]. Transfer learning, a well-established sub-field in machine learning, uses knowledge of related source domains with sufficient annotated data to improve the model's performance on related target domains with inadequate or a lack of labeled data [21]. Model parameters, feature representations, and instances are different forms of knowledge that can be transferred in transfer learning [83]. In addition, transfer learning has achieved tremendous success in a wide range of applications, such as computer vision [97] and natural language processing [92]. Arguably, transfer learning has obtained outstanding achievement for supervised pre-training of **convolutional neural networks (CNNs)** on ImageNet. However, the performance of the pre-trained network deteriorates during fine-tuning with few samples [88], or new samples become less similar [123]. One way to generalize unseen samples directly from few examples is by learning rich knowledge from similar events and re-using past knowledge to incrementally adapt to unseen examples. This idea, motivated by humans' intelligence, is known as meta-learning.

Meta-learning refers to a variety of techniques that aim to improve adaptation through transferring generic and accumulated knowledge (meta-data) from prior experiences with few data points to adapt to new tasks quickly without requiring training from scratch. Meta-learning is most commonly known as "*Learning to learn*". But, what does learning to learn mean? Mitchell [77] defined the term *learning* as the instance when a computer program's performance at a given task improves with experience. In conjunction with Mitchell's definition, Thrun and Pratt [105] defined "*Learning to learn*" as a capability of algorithms to improve their tasks' performance with experience and with a number of observed tasks. In this context, an algorithm's performance improvement depends on a family of tasks. In other words, if a learning algorithm's performance does not improve with the experience of other tasks, it is not said to learn to learn [105]. Accordingly, the learning level is elevated from data to tasks in meta-learning. Figure 1 better demonstrates the learning to learn paradigm.

Driven by the goal to scale AI to a real-world setting and approaches to the human intelligence paradigm, meta-learning has recently regained significant interest in few-shot problems using deep learning [113]. Currently, a number of surveys on meta-learning have been proposed [44, 62], however, the meta-learning methods progress rapidly. Accordingly, this paper examines meta-learning methods in detail and performs an extensive review of meta-learning to provide an up-to-date review. Additionally, a new organization from the algorithms' mechanics point of view is introduced. A summary of the benchmark comparison of review methods is also presented, followed by a comprehensive discussion of challenges to establish future directions based on current limitations and strengths.

The remainder of this paper is organized as follows. Section 2 reviews the preliminary concepts of meta-learning. Section 3 introduces the category of meta-learning methodologies and presents an in-depth review of state-of-the-art methods and their advances. Section 4 describes

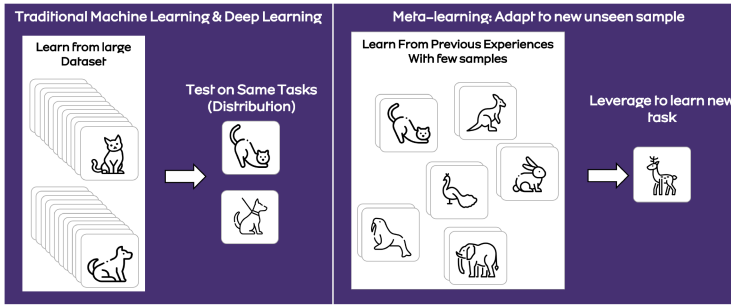


Fig. 1. Contrast between traditional machine learning and deep learning and the meta-learning approach. The left panel illustrates the conventional process where models learn from a large dataset and are tested on tasks within the same distribution. The right panel showcases the meta-learning paradigm, where models leverage previous experiences from a few samples to quickly adapt and learn new, unseen tasks, demonstrating the advantage of meta-learning in dynamic learning environments.

the challenges along with future research directions and finally Sections 5 and 6 briefly discuss applications of meta-learning and conclusions, respectively.

2 META-LEARNING PRELIMINARY TERMINOLOGY AND TRAINING STRUCTURE

2.1 Terminology and Training Structure

In the standard supervised machine learning setting, dataset D is divided into two separable and distinct datasets, namely train and test sets denoted by D_{train} and D_{test} , respectively. Conventional neural networks train during iterative cycles, termed “Epoch”. In every epoch, the entire D_{train} is feeding forward and backward. During each epoch, D_{train} may be partitioned into several packets termed “Batch”. Thus, each epoch may consist one or several “iterations”, which is the number of steps needed for one epoch that all batches (i.e., entire D_{train}) feed into the model.

In meta-learning, dataset D is divided with no overlapping into two meta sets, namely the meta train and meta test sets, denoted by D_{meta_train} and D_{meta_test} , respectively. Next, D_{meta_train} and D_{meta_test} are subdivided into train and test sets. To differentiate between these two forms of training and testing sets, the latter (inner) train and test sets are termed “Support set” and “Query set”, respectively. Let’s rewrite the meta train set and meta test set as $D_{meta_train} = \{(D_{meta_train}^{Support}, D_{meta_train}^{Query})\}$ and $D_{meta_test} = \{(D_{meta_test}^{Support}, D_{meta_test}^{Query})\}$. Support set is a N -way k -shot random samples and the query set is q random samples for each of the N categories of the support set. A meta-learner model trains in an “Episodic” mode. In each episode, support and query sets would created. A meta-learner model takes $D_{meta_train}^{Support}$ to learn, then applies on the $D_{meta_train}^{Query}$ to make predictions. The meta-learner updates based on the prediction error over episodes. Over a series of episodes, the meta-learner learns to learn from the small dataset. This phase is called meta-learning. Afterward, it takes $D_{meta_test}^{Support}$ to build a classifier. Here, the model concentrates on learning task-specific parameters quickly. Then, the performance of the classifier is evaluated through $D_{meta_test}^{Query}$, which is known as the adaptation phase. In other words, the model uses the prior knowledge learned in the meta-learning phase to quickly adapt to a new task in the adaptation phase. From a conventional neural networks view, episodic training corresponds to training on a series of batches [56]. A visualization of data preparation and episodic training is demonstrated in Figure 2.



Fig. 2. Overview of episodic training for few-shot learning, showing support and query sets for both meta-training and meta-testing phases.

2.2 Datasets

Most state-of-the-art meta-learning models have been evaluated and compared on several well-known benchmarks. While there several datasets exist, the three most widely used benchmark datasets are:

- Omniglot: an image dataset of different handwritten characters of different alphabets [57]. This dataset is comprised of 1,623 characters related to 50 different alphabets. In the few-shot learning setting, each character stands as classes/labels, and each character has 20 samples (or 20 shots) drawn by 20 different persons.
- MiniImageNet: the mini version of the popular large ImageNet dataset published in 2009, which revolutionized the artificial intelligence realm. ImageNet is a large colored image dataset annotated by humans for visual recognition competition. Now, ImageNet contains 14 million images and covers 20,000 objects/classes. For few-shot learning problems, [108] published MiniImageNet from ImageNet, a smaller dataset containing 100 classes with over 60,000 images.
- CUB-200-2011: the Caltech-USCD-Birds-200-2011 [109], in short CUB-200-2011, contains 11,788 images of 200 different bird species (200 classes). Each image has a textual description alongside labels describing various attributes, such as color and shape.

Oxford-102 (flowers dataset), CIFAR-10, and MNIST are other datasets that are less used in the literature. In this paper, the above three datasets are used to compare the performance of different methods.

3 META-LEARNING

This section provides a comprehensive overview of cutting-edge studies and their advances in meta-learning in three main categories (shown in Figure 3):

- Metric-based methods
- Memory-based methods
- Learning-based methods
 - Learning the initialization
 - Learning the parameters
 - Learning the optimizer

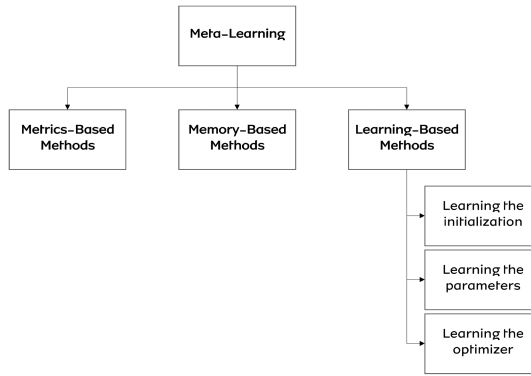


Fig. 3. An overview of the proposed taxonomy of meta-learning methods.

3.1 Metric-based Methods

The key concept in “learning the metric space” methods is comparing the query set samples to support set samples. Then, the labels of query samples are predicted based on proximity to the support set. Principally, the success of these methods depends on embedding learning. In this context, the aim of embedding learning is to convert each high dimensional sample to a lower dimensional sample in the form of vector representations in such a way that similar samples reside closely in a new lower dimension space. Given an N -way k -shot problem, let’s consider a support set $D_{support} = \{(x_1^s, y_1^s), \dots, (x_K^s, y_K^s)\}$, where x_i^s refers to input samples, $y_i^s \in N$ is predefined labels, and a query set $D_{query} = \{(x_1^q, y_1^q), \dots, (x_j^q, y_j^q)\}$, in which x_j^q is the test sample and $y_j^q \in N$ is the true label of test. In a generic form of learning the metric space methods, embedding function f and embedding function g apply on x_i^s and x_j^q , respectively, to generate embedding of the support set and query set. Next, the similarity between obtained embeddings, i.e., $f(x_i^s)$ and $g(x_j^q)$, is measured by similarity function $s(\cdot, \cdot)$. To predict the label of x_j^q , a decision is made based on the closest similarity between $f(x_i^s)$ and $g(x_j^q)$. In other words, the class $f(x_i^s)$ nearest to $g(x_j^q)$ is assigned to the sample test x_j^q . While the same embedding function can be used for both support and query sets, several researches have shown that different embedding functions lead to a better performance [10, 108]. Current state-of-the-art metric-based algorithms are:

3.1.1 Siamese Network. Siamese network is one of the simplest metric-based learners widely used in one-shot learning problems. The principal architecture of the Siamese network was proposed in 1993 by [6] for fingerprint similarity estimation. However, the term “Siamese” was first introduced by [12] for the signature verification problem. Principally, the Siamese network consists of two symmetrical neural networks joined with the similarity function called the “energy function.” The architecture of the Siamese network is shown by Figure 4.

As demonstrated in Figure 4, a input pair x_1 and x_2 from the training set is fed to two embedding functions, $f_1(\cdot)$ and $f_2(\cdot)$, sharing the same parameters. Then, embeddings $f_1(x_1)$ and $f_2(x_2)$ feed into energy function. Basically, the energy function could be any similarity measure, such as cosine similarity or Euclidean distance. The output of the energy function is the similarity score of the two embeddings. Thus, rather than a classification task, the Siamese network principally aims to learn the network for matching tasks between similar pairs [75]. The proposed network utilizes contrastive loss as a loss function for optimization. The traditional contrastive loss used in the

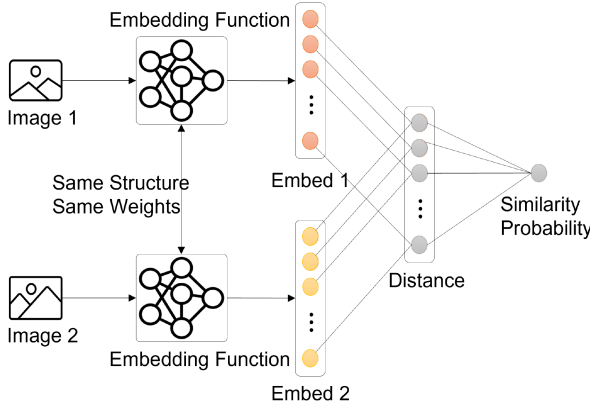


Fig. 4. Siamese network architecture for one-shot learning, optimizing pair-wise similarity for classification tasks with limited data.

generic Siamese network is expressed as Equation (1) [16, 45]:

$$\mathcal{L}(E, Y) = \sum_{i=1}^P (1 - Y) * L_G(E(Z_1, Z_2)_i) + Y * L_I(E(Z_1, Z_2)_i) \quad (1)$$

where $E(Z_1, Z_2)_i$ is the energy function that measures the similarity of the i -th pairs of embeddings Z_1 and Z_2 . Y is the true label that denotes whether the given pair is similar $Y = 0$ or dissimilar $Y = 1$. L_G and L_I are partial loss functions applied on similar and dissimilar pairs, respectively, and designed in a way that monotonically increases the loss function for a similar pair and decreases loss function for a dissimilar pair. Siamese networks are widely used in various applications where data are limited. The face verification task is one of the extensive applications of the Siamese network. In the state-of-the-art research, [16] studied the task of face verification by refining the loss function as Equation (2):

$$\mathcal{L}(E, Y) = (1 - Y) * E^2 + Y * \max(0, \text{margin} - E)^2 \quad (2)$$

where the margin is a hyper-parameter that specifies the lower bound distance between dissimilar pairs. Similar to Equation (1), E is the energy function, and Y will be 0 for the similar pair and 1 otherwise. This study was one of the earliest works that used CNNs as an embedding function. Later, [55] studied Siamese networks for a one-shot image recognition problem, which is known as the first study of a metric-based method extended Siamese networks for few-shot learning. Additionally, Siamese networks have been extensively studied in various applications. For further details on Siamese network applications, interested readers can refer to comprehensive reviews by [15, 45].

3.1.2 Prototypical Networks. The prototypical network, proposed by [101], is basically learning an embedding function through clustering the embedding space. Principally, prototypical networks aim to generate a prototype of each class to use as an anchor for the classification task. Given the N -way k -task classification problem, the prototypical representation of each class is built based on the mean embedding of k data points in each N class. By using the same embedding function for both the support and query sets, the embedding of query points would be obtained. Mainly, Euclidean distance is used as a distance metric, measuring the distance of query embeddings and classes' prototypes. Each query point is classified based on the closest distance of the query embedding and classes' prototypes. Predictions probabilities are obtained by feeding query-prototypes distance




Fig. 5. Prototypical networks for few-shot learning, illustrating embedding space and prototype formation.

into the Soft-max function. The prototypical network trains by minimizing the average negative log-probability of the correct prediction, as expressed by Equation (3). Figure 5 shows the general flow of the prototypical networks.

$$-\frac{1}{T} \sum_i \log p(y_i^* | x_i^*, p_c) \quad (3)$$

Two well-known variants of prototypical networks are semi-supervised and Gaussian prototypical networks. [91] revised the prototypical network to deal with the semi-supervised setting, where the dataset contains both labeled and unlabeled data. In their model, similar to the generic prototypical network, class prototypes are first built with labeled data. Then, soft k-means algorithm is applied on the embeddings of both labeled and unlabeled data. However, inspired by seeded k-means, instead of initializing cluster centers randomly, the class prototypes are calculated from the labeled data, and are used as cluster centers. Then, unlabeled embeddings are assigned to each cluster based on their Euclidean distance to each cluster center, i.e., class prototypes. Finally, class prototypicals are refined by updating cluster centers by the usual k-means algorithm. Since some points that do not belong to any existing classes may exist in unlabeled data, these data would interfere with prototypes refinements. Thus, the authors proposed an additional class for these data called the distractor class and considered the origin as the center of the distractor class. Since distractor class members may belong to different subjects, a single distractor class will have a high variance. To solve this issue, the authors proposed a soft-masking mechanism, in which a threshold distance is set for all original class prototypicals. This threshold acts as an intra-cluster variance capturing distractors as exemplars outside the prototypical areas. This threshold is calculated by feeding statistics of prototypes, including min, max, skewness, and kurtosis, to a small **multi-layer perceptron (MLP)** network.

The Gaussian prototypical network is another novel architecture of the prototypical network proposed by [26]. In this model, a covariance matrix is constructed from the embedding function alongside the generated embedding vectors. The covariance matrix is used to build a confidence region around data points and adds weights to embedding vectors, then utilized to compute a total covariance matrix of each class. Afterward, the distance of each query from class prototypes is used to compute the linear Euclidean distance as follows (Equation (4)):

$$d_c^2(i) = (\vec{x}_i - \vec{p}_c)^T S_c (\vec{x}_i - \vec{p}_c) \quad (4)$$

where p_c is the cluster center or the prototype of class c ; S_c is the inverse covariance matrix; d_c is the distance of query point i from class prototype c ; and \vec{x}_i is query point i . Figure 6 illustrates the operation of Gaussian prototypical networks. Variable p_c is calculated using variance-weighted embedding vectors in Equation (5):

$$\vec{p}_c = \frac{\sum_i \vec{s}_i^c * \vec{x}_i^c}{\sum_i \vec{s}_i^c} \quad (5)$$




Fig. 6. Gaussian prototypical networks workflow; Solid circles are embedding vectors, dark ellipses represent each embedding covariance matrix, and light halos depict the class's total covariance matrix. Finally, stars illustrate classes' prototypicals.

where \vec{x}_i^c is the embedding vector of support point i and class c ; and \vec{s}_i^c is the covariance matrix of related support point i and class c . Class covariance matrix is computed as Equation (6):

$$\vec{S}_C = \sum_i \vec{s}_i^c \quad (6)$$

[84] proposed the **transferable prototypical network (TPN)** by remoulding prototypical networks to learn unsupervised domain adaptation. TPN uses labeled data as a source and unlabeled data as target domains and tries to learn a neural network that minimizes embedding shifts between source and target domains. In other words, TPN constructs an embedding function that generates domain-invariant representations. Technically, TPN first builds class prototypes from the source domain (i.e., labeled data), then the target domain (i.e., unlabeled data) assigned to the nearest class prototypes. The authors proposed three different classifiers regarding three datasets, namely source only data, target only data, and combined source and target data. The authors argued that if the source and target domain distributions are the same, the class prototypes on both domains should be identical. Thus, to achieve the representation invariant over domains, they compute class level discrepancy over domains by calculating **reproducing kernel Hilbert space (RKHS)**. Finally, TPN trains by minimizing combined classification loss and discrepancy loss.

With a similar belief and similar mechanics to the TPN [70, 84] introduced a new module called domain alignment component to prototypical networks to tackle the problem of domain shifts between test and train sets. Basically, the domain alignment component tries to minimize the **maximum mean discrepancy (MMD)** by calculating the similarity of two different domains, one from the train set and one from the test set, in kernel Hilbert space. Then, the networks are trained to minimize the combined negative log-probability loss (the generic prototypical networks loss) and MMD measure. Empirical analysis on benchmark datasets showed that their model outperforms the generic prototypical networks.

[47] argued that most metric-based meta-learning models forget intra-class distribution information. To tackle this problem, they proposed another tweak to prototypical networks called **improved prototypical network (IPN)**, which has three main components: (i) feature extraction, (ii) weight distribution, and (iii) distance scaling. The first component basically extracts feature embeddings and constructs class prototypes. In generic prototypical networks, samples of the same class have the same importance (weight) in class prototypes. The authors refined this idea by proposing an attention-analogous strategy. Hence, the second component, weight distribution, is responsible for computing weighted class prototypes. The weight distribution component is designed as a three-layer neural network that takes concatenated embeddings of the same class as the input and provides samples' weights as the output. The last component, distance scaling is a

three-layer neural network that maximizes inter-class differences and minimizes the intra-class differences. The results indicated that IPN outperforms the generic prototypical networks.

[67] expanded the prototypical network to multi-label few-shot problems. The authors argued that the main challenge with multi-label problems is that different labels can have the same support set leading to the same prototypes. To address this issue, they proposed a new prototypical network called **label-enhanced prototypical network (LPN)**. For a multi-label few-shot problems, the authors targeted the aspect category detection in sentiment analysis. Accordingly, their data points and class labels consisted of sentences and aspect categories. The LPN is comprised of four main modules: (i) feature extraction, (ii) label-enhanced prototypical network, (iii) contrastive learning, and (iv) adaptive multi-label inference. The authors used a pre-trained embedding function for language, such as BERT [18] as the feature extraction module. The second module is responsible for computing the relation between data points, here sentences, and labels. For this purpose, after extracting support sentence representations from the previous module, the same process is repeated for labels that are verbal phrases to get label representations. The authors used the low-rank bilinear model [125] to assign weight to each sentence's representations. In simple words, the authors calculated the sentence representations' weights by element-wise multiplication of sentences' representations and their label representations, normalized with the Soft-max function. Then, class prototypes are built based on the weighted sentences' representations. The label-enhanced prototypical network module is trained to minimize the cross-entropy loss. The third module is designed in a way to maximize intra-class similarity and minimize inter-class similarity by employing contrastive learning loss. Since in their problem, every sentence can have a different number of labels as a multi-label problem, the next challenge was determining the number of labels that should predict for each new query point. This challenge was addressed by proposing the fourth module, which is basically a multi-layer perception network that predicts the number of labels for each sentence. This network is trained to minimize the cross-entropy loss of predicted label counts. By integrating all modules and combining all losses, LPN outperforms the generic prototypical network. [17] proposed another variant to prototypical named the multi-prototypes prototypical network. In contrast to the generic prototypical networks that consider only one prototype per class, the proposed network investigates the idea of multiple prototypes per class. For this purpose, the authors partitioned the embedding space of each class into k clusters by employing the k -means clustering algorithm. Class prototypes are simply cluster centers, and the proposed network is trained to minimize the **COREL (clustering-oriented representation learning)** loss proposed by [52]. The results showed that the number of class clusters, i.e., classes prototypes affects model performance; in other words, the difficulty of determining the optimal number of prototypes per class limits the performance.

[54] argued that most metric-learner models are not able to distinguish classes out of support set classes and it is necessary for meta-learner models to be able to detect unseen classes during prediction. The authors proposed **dummy prototypical networks (D-ProtoNets)** to address this issue. Conceptually, D-ProtoNets follows the same mechanics of the generic prototypical networks, but in every episode, a dummy prototype is built by introducing a dummy generator function. In simple words, the dummy generator is a Maxpool layer followed by non-linear fully connected layers performed on the support classes' prototypes. The dummy prototypes are responsible for handling unseen classes. The authors targeted the keyword spotting detection problem and created and published splitGSC, a benchmark dataset for key spotting detection problems based on the Google speech commands dataset. Through various numerical analyses, the authors demonstrated that D-ProtoNets achieved state-of-the-art performance on splitGSC.

[42] argued that a single prototype for each class, like the generic prototypical network, is not able to capture the class distribution effectively, so they proposed a **local descriptor-based**




Fig. 7. Depiction of channel squeeze and spatial excitation. This schema illustrates the mechanism where $\sigma(v)$ signifies the varying positional relevance within the embeddings, enhancing feature representation through selective emphasis.

multi-prototype network (LMPNet). Technically, LMPNet utilizes a local descriptor to obtain feature embeddings. The local descriptor is basically the two-dimensional vector output from a convolutional layer [117]. LMPNet builds class prototypes based on the weighted average of the local descriptors. These prototypes were obtained through a novel mechanism by introducing channel squeeze and spatial excitation components, which flatten feature embeddings regarding their spatial position over channel dimensions (Figure 7). In summary, this attention module is designed to learn the output of multiple prototypes from local descriptors compressed by channel squeeze and spatial excitation. Using Resnet as a backbone, LMPNet outperforms the generic matching network.

[82] extended prototypical networks to a multi-modal prototypical network. The core novelty of this model is the use of textual descriptions attached to the image dataset to create more feature embeddings. For this purpose, the authors trained a **text-conditional generative adversarial network (tcGAN)**, which creates image features from the given textual description. Then, two prototypes per class were calculated – one from the original images’ embeddings and one from the generated images’ embeddings. Next, combined class prototypes are computed based on the weighted average of both original and generated prototypes. Finally, the multi-modal prototypical network was trained in a similar fashion to generic prototypical networks.

3.1.3 Matching Networks. Matching networks are another powerful metric-based meta-learner proposed by [108] that are able to generate labels for the unobserved class as well. In this work, a parametric neural network P was developed to produce a prediction for each query point from the support set $S = \{(x_i, y_i)\}_{i=1}^k$, (k number of sample per class), per Equation (7):

$$p(\hat{y}|\hat{x}, S) = \sum_{i=1}^k a(\hat{x}, x_i) y_i \quad (7)$$

where $p(\hat{y}|\hat{x}, S)$ is the predicted class for the query point \hat{x} given support set S ; x_i and y_i are the inputs and labels of samples from support set S ; and a is the attention mechanism between \hat{x} and x_i . The attention mechanism is simply the Soft-max function over the cosine distance between embeddings of \hat{x} and x_i , which are calculated via Equation (8):

$$a(\hat{x}, x_i) = \frac{e^{\cosine(f(\hat{x}), g(x_i))}}{\sum_{j=1}^k e^{\cosine(f(\hat{x}), g(x_j))}} \quad (8)$$

In Equation (7), the attention mechanism result is multiplied by the support set labels y_i . For this purpose, the support set labels are first converted into one-hot encoded values. The overall flow of the matching network is shown in Figure 8. As seen in Equation (8), the authors used two different embedding functions, f and g , to map data points into embedding for the query \hat{x} and support input x_i , respectively. Finally, the proposed network is trained to minimize the prediction error. [13] argued that support and query images could differentiate in local features and shapes. They tackled this problem by proposing a **cascaded feature matching network (CFMN)**, which conceptually

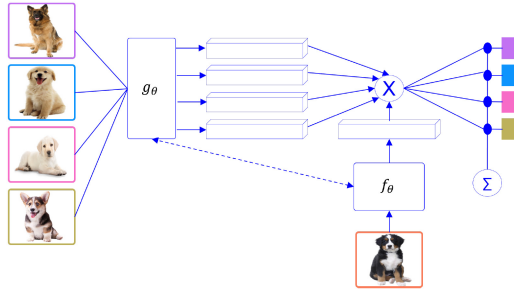


Fig. 8. Matching networks structure. Illustrates the transformation of input instances into a feature space for one-shot classification via embedding and comparison functions.

maps features of the support and query sets to a latent space then reshapes them into a two-dimensional array. Next, it calculates the element-wise similarity of support and query features. Additionally, the support set features maps to a higher space based on element-wise correlation calculated in the previous step. In other words, the final support set features are constructed by taking into account the relations of feature positions, the so-called spatial attention mechanism. This approach allows relevant features to be considered and less important features to be ignored. Finally, CFMN predicts a label for a new sample based on the distance-based similarity score of the support set and a new sample point. [74] proposed another matching network variant called the **attentive matching network (AMN)**, which exhibits two novelties. First, it incorporates a new loss function named **complementary cosine loss (CCL)** that is a combination of the **hardest-category discernment loss (HDL)** and cosine-distance metric, aiming to minimize the intra-class difference and maximize the inter-class differences. The second novelty is a new attention mechanism that focuses on the feature level. This module calculates the importance of features based on the standard deviation over classes. In other words, a higher standard deviation over classes means that features are more distinctive between classes. For this purpose, the authors first built class prototypes by class-wise mean of embeddings, then calculated the inter-class standard deviation from each class prototypes. In order for the feature attention module to be differentiable, they proposed a convolutional operation over the attention calculation.

[126] argued that a generic matching network randomly selects labels from the base dataset to create training samples, which leads to an over-easy label set. For example, the authors compared a label set cat, tree, water to another label set cat, dog, lion that is easy and low value in a classification task. To address this issue, they proposed schedule sampling for the matching network, which first samples a random label set, then calculates the difficulty value for the selected label set. The difficulty value is computed based on similarity metrics. The training of matching net started with a low difficulty label set and then continued into a more difficult label set. The authors showed that the performance of the classifier directly relates to the difficulty of the label set.

3.1.4 Relation Networks. Relation network is another simple yet effective metric-based meta-learner presented by [103]. This model has two main components: (i) embedding function denoted by f_ϕ and (ii) relation function g_ϕ . According to the flow of the relation network illustrated in Figure 9, the embedding function f_ϕ maps the support and query sets into the embedding environment. For the k -shot problem where $k > 1$, embeddings of the same class in the support set are aggregated through the element-wise summation to construct each class embedding map. The obtained consolidated classes embeddings are then combined with embedding of query samples and fed into the relation function g_ϕ . The output of the relation function g_ϕ is the N (number of

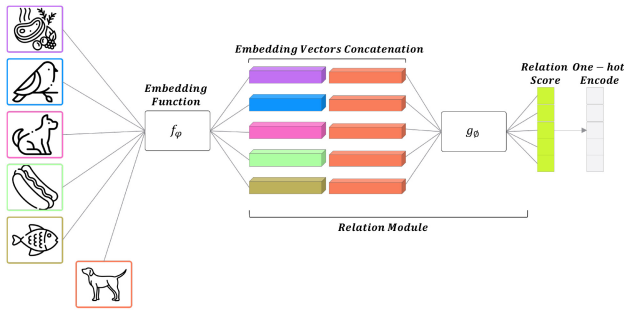


Fig. 9. Schematic of a relation network for few-shot learning.

classes) similarity score called the relation score that ranges between 0 and 1. While the proposed model is a classification task, the authors trained their model using the **mean square error (MSE)** loss function since they aim to predict a relation score scale between 0 and 1.

Reference [39], sharing a similar belief to [47], introduced the memory-augmented relation network (MRN), which aimed to investigate the interactions between instances. MRN is composed of four components: (i) feature embedding function, (ii) information generation component, (iii) relation module, and (iv) metric-based classifier. The feature embedding function simply converts data points into latent space. The information generation component adds information to samples basically from its k -nearest samples' embeddings. This component is equipped with memory to control embedding propagation. The relation module, similar to a generic relation network, is a convolutional neural network that takes input samples and computes their distances. The last component, the metric-based classifier, predicts a label for an unlabeled query point based on the distance from class centroids, which are defined as the mean of the same class feature representations. This classifier trains to minimize cross-entropy loss. [69] proposed a prototype-relation network by basically combining prototypical networks and relation net. This model has two main components: (1) the prototype component that generates class prototypes, and (2) the relation component, similar to the generic relation net, that computes relations between query points and class prototypes. Another novelty of this work is the introduction of a new loss function that considers inter and intra-class distances. [19] developed another variant of relation networks called **multi-scale relation networks (MSRN)** that has two main components: a feature extractor and metric learner. In CNNs, low-level layers generate detailed information while the higher-level layers produce more semantic information. The authors argued that combining the feature outputs of different layers may improve representations. The second component, the metric learner, is basically a generic relation network; however, the authors subtracted these two components instead of concatenating two multi-scale features. They claimed that subtraction generates a better representation of feature differences. Their model outperforms the generic relation net. [89] proposed a **multi-local feature relation network (MLFRNet)** to improve relation networks. They used random cropping and random dropping to get local features, then mapped to embedding space. Finally, class prototypes are calculated by the mean of each class's local feature embeddings. The authors also introduced a local feature attention module to calculate the relation between local feature prototypes and query samples. Basically, the local feature attention module is responsible for computing the contribution of local features in query sample prediction. Additionally, the authors proposed a new strategy called support-support similarity to dynamically calculate the margin loss of local features. Support-support similarity was designed to compute pair-wise cosine similarity of classes that is used to generate a margin. The authors suggest that adding the margin

is necessary for different classes that might have local features with high similarity in order to separate different classes effectively. [1] argued that the original relation network is not robust to spatial positions of feature maps and may perform poorly in scale variations of compared objects. To address this issue, they extended a new variation of the relation network called the **multi-scale kronecker-product relation network (MsKPRN)**, which consists of three main modules. The first module is the encoding module that is responsible for converting different scales of support and query sets into feature space. The output of this module is feature maps with different scales that is organized in a way to feed the same scale support and query feature maps into the next module. The next module is the **kronecker-product** (in short **KP**) module, which calculates the spatial correlation from pair representations by operating the inner product on the same scale support and query feature maps. The obtained spatial correlation is concatenated with the representations to feed into the next module. The third module is the relation network module, which is responsible for learning the relation between the same scale support and query feature maps and producing similarity scores in the range between 0 and 1. Their results demonstrated that the kronecker-product module improves few-shot classification tasks by incorporating spatial correlation maps into feature maps.

3.1.5 Other Advanced Methods. The next advanced metric-based meta-learning algorithm reviewed in this section is the **graph neural network (GNN)** proposed by [28]. The authors argued that their proposed GNN for the few-shot classification problem generalizes Siamese networks, prototypical networks, and matching networks based on their architectures. Conceptually, in their model, each task T_j is approached as a fully-connected graph $G = (V, E)$, where V and E are denoted as nodes and edges, respectively. Nodes are associated with support set embeddings concatenated with one hot encoded label. Edges correspond to tasks' similarities, which are trainable. For a new query point, i.e., an unlabeled data point, a uniform distribution over all labels is used as label concatenation to query embedding. Finally, GNN was trained to minimize the cross-entropy loss. Results show that GNN outperforms Siamese networks, prototypical networks, and matching networks on several benchmark problems. [59] proposed the **global class representation (GCR)** model, which achieved the best performance on benchmarks and outperformed many well-known meta-learning algorithms. The idea of GCR is to engage novel classes with base classes at the first stages of training, thus the reason for naming the model global class. Concretely, global class representation is first created from the whole class set, called C_{Total} , which is composed of both base and novel classes. Next, in each episode, N random classes and n_s random samples per each class are selected to form the support set. Then, episodic class representation is built based on the support set. Here, class representation is analogous to class prototypes in prototypical networks. As one novelty of this study, the authors proposed a new component called the registration module. Basically, the registration module is responsible for calculating the similarity between each episodic class representation and global class representations, or the so-called registration loss. GCR minimizes both registration loss and classification cross-entropy loss to achieve optimal global classes. GCR predicts a label of a new sample by comparing it to global class representations. [38] disputed that most metric-based meta-learners use standard similarity measures, such as cosine similarity and Euclidean distance. They support that those metrics could not capture non-linear relations between samples. To address this issue, they introduce a **multi-scale feature network (MSFN)**, which has three main components. The first component is the feature extractor that extracts multi-scale features from the support set. The next component is the label feature that is responsible for learning classes representations (analogous to class prototypes). The label feature basically is a network that takes a concatenation of support features from the previous component as the input and generates class representations as the output. The last component is the metric-component,

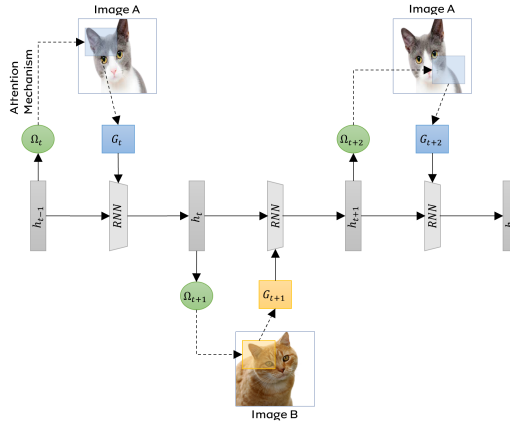


Fig. 10. Attentive recurrent comparators (ARCs) operation. Depicts the ARCs mechanism focusing sequentially on image regions via an RNN controller and attention mechanism, mirroring human feature observation strategies.

which principally is the relation network that computes the matching degree between a new query feature map and class representations. Through numerical analysis, the authors demonstrated that multi-scale features and non-metric similarity measures improve the overall model performance.

[99] argued that humans differentiate two objectives by recurrent multiple observations, and during each observation, a specific feature is inferred. This human-wise approach is quite different from Siamese networks. [99] proposed **attentive recurrent comparators (ARCs)**, mimicking humans. ARCs consist of two components: (i) **recurrent neural network (RNN)** controller, and (ii) attention mechanism. The attention mechanism is designed in a way that concentrates on a specific region of the image. Figure 10 demonstrates the scheme of ARCs, in which the attention mechanism gets a glimpse G_t from the image I_t at each step (Equation (9)):

$$G_t \leftarrow (I_t, \Omega_t) \quad (9)$$

where Ω_t is the attention glimpse parameter, denoting the location and size of the glimpse. Basically, Ω_t is computed based on the previous hidden state (h_{t-1}) of the RNN controller (Equation (10)):

$$\Omega_t = w_t h_{t-1} \quad (10)$$

where w_t is the projection matrix necessary to map h_t to the number of attention parameters. Next, the hidden state is obtained via Equation (11):

$$h_t \leftarrow RNN(G_t, h_{t-1}) \quad (11)$$

where $RNN(\cdot)$ is the RNN controller update function. In short, ARCs recurrently observe different partitions of two images over a finite number of representations, so the model can focus better on the important context of each image. While the presented model outperforms Siamese networks, it can be quite computationally expensive due to sequential operations. Similarly, [119] believed that humans differentiate objects based on partitioning images and comparing them with the prototypical part of each image category. Accordingly, they proposed a new metric-based method called the **region comparison network (RCN)**. The aim of RCN, as demonstrated in Figure 11, is to compare the whole query image with different partitions of support samples. RCN consists of three main components: (i) feature extractor, which is responsible for producing image representations;

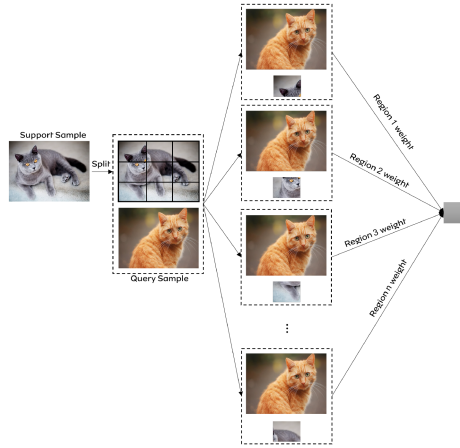


Fig. 11. Comparison support sample's parts with query sample in Region comparison network (RCN).

(ii) region matching, which is designed in a way to calculate the similarity scores between support and query samples; and (iii) explain network, which is used to obtain the final classification. To be more specific, the output of the first step is decomposed into several region vectors. Then, cosine similarity between support set region vectors and query set representation is computed. Subsequently, global max-pooling is applied to pick the most important information in the region similarity vector. The explain network takes region similarity maps from the previous step and region weight to make a classification. In each task, the important region may change. For example, in classifying dogs and birds, the important regions are the dog's head and bird's wings. The explain network uses task-specific region weights produced by the region meta-learner, which is designed to take a concatenation of support set representations and query set representations and produce task-specific region weights. The most salient aspects of RCN are its interpretability and quantitative explainability of classification tasks in the few-shot setting.

[40] aimed to explore the effects of pre-defined similarity metrics, such as Euclidean and cosine distances, on existing metric-based methods and proposed a new method called **metric-agnostic conditional embeddings (MACO)**. The architecture of MACO is demonstrated in Figure 12. Basically, MACO is composed of four stages: (i) feature stage, (ii) relation stage, (iii) conditioning stage, and (iv) classifier stage. The first stage is responsible for obtaining a feature vector from input images, which is simply a CNN-based network with a fixed structure for both the support and query sets. The next stage is designed to produce a single vector for each class. This class-specific single vector is analogous to the class prototype. The authors used the relational network (proposed by [96]) to generate the single vector class representations. The relational network is comprised of fully connected layers (instead of convolutional layers) followed by batch normalization and ELU activation function using pair-wise image comparisons within each class. The single vector class representation would be the average output across all comparisons (Equation (12)):

$$R_{\rho}(s_k) = \frac{1}{\binom{n}{2}} \sum_{(x_i, x_j) \in s_k} g_{\rho}(x_i, x_j), \quad i \neq j \quad (12)$$

where R_{ρ} is the single vector of support set class k ; n is the number of images in class k ; g_{ρ} is the relational network parameterized with ρ ; and x_i, x_j are images representations within support class k . The conditioning stage uses a network similar to a relational network that takes the concatenation of class vectors and query images. The output of this stage is the modified class

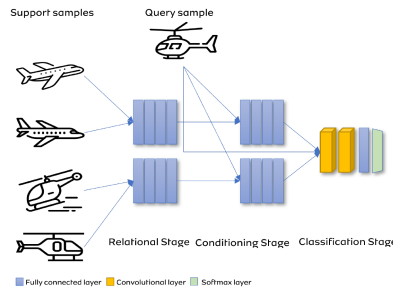


Fig. 12. Metric-agnostic conditional embeddings (MACO) architecture. This flowchart outlines the four-stage process of MACO, integrating CNN-based feature extraction with relational networks for class vector generation and conditioning, culminating in classification.

representation in the context of the query image. The last stage is a neural network that takes the modified class vectors and query images as input and produces Softmax classification. This network is composed of convolutional layers followed by fully connected layers. In short, MACO is a novel model in which a learnable classifier supersedes the pre-defined metrics and modified class vectors (or prototypes) conditioned to the query image context. Similarly, [120] developed a learnable network instead of pre-defined similarity metrics termed **relative position and map network (RPMN)**. Basically, RPMN consists of three components: (i) feature extractor, (ii) relation position network, and (iii) relative map network. The relative position network aims to produce different weights to the position of feature embeddings, and the relative map network is analogous to the original relation network but has been modified in such a way to compare two single maps from support and query feature representations. The relative map network is trained to learn task-specific embedding models. Moreover, the relative map network is equipped with a fully connected head, which is responsible for computing the final similarity score between the support and query sets. An important aspect of RPMN is the introduction of the relation position network as a new attention mechanism, allowing for more comparisons of effective feature maps, and the relative map network as a learnable similarity metric network.

3.2 Memory-based Methods

Memory-based meta-learning techniques have a more dynamic mechanism since they are adaptive to the presented tasks. Principally, memory-based meta-learners own the memory component (internally or externally) that enables them to retrieve information from previous inputs. The **memory-augmented neural networks (MANN)** is a well-known and effective memory-based meta-learner proposed by [95], which aims to use external memory to improve task adaptation utilizing the **Neural Turing Machine (NTM)** [34]. On a high level view, NTM is a neural network equipped with external memory to retrieve information. More specifically, NTM contains two main components: (i) controller and (ii) memory bank. Figure 13 shows the interaction of the controller and memory bank. Similar to a standard neural network, NTM receives inputs and predict outputs. In contrast to a standard neural network, NTM communicates with the memory bank through selective read and write operations, which the authors named as *head* analogous to the Turing machine. In order for the NTM to be trainable, including the way memory bank is utilized, the authors developed blurry read and write operations that should be differentiable to use backpropagation. For this purpose, read and write operations interact with distribution over memory locations instead of a single location. Specifically, the read operation simply reads values from the memory bank, where the value is basically the neural network parameters (i.e., weights). The memory bank is simply a two-dimension matrix. The write operation is more complex than




Fig. 13. Overview of Neural Turing Machine (NTM). Illustrates the controller's interaction with the memory bank via read and write heads.

read and includes two modules called erase and add, which erases information that is no longer required from a memory cell and adds a new operation to the memory, respectively.

To produce the parameters of the NTM neural network, the authors proposed two different addressing mechanisms. First, the content-based approach uses the similarity of the current value and controller's value. While the content-based approach is quite straightforward, the location-based approach is employed to iterate over different locations rather than content, which is essential for generalization.

[95] proposed a new addressing mechanism called **least recently used access (LRUA)**. In MANN, the read operation is performed through the content-based addressing mechanism, and the write operation uses LRUA to write the most recently used memory location. MANN accumulates information of a coupled representation-class label in external memory through the explained mechanism, which enables it to retrieve the information later for classification. [76] proposed a new memory-based meta-learner model called **simple neural attentive meta-learner (SNAIL)** that combines temporal convolution and soft attention mechanism that acts as a memory compared to the external memory of MANN. The temporal convolution serves as high band-width memory access, and the soft attention allows access to a specific experience. The combination of these two components leads to better leveraging of information from past experiences. [29] proposed **conditional neural processes (CNPs)** that, in a similar fashion to SNAIL, do not employ external memory. On a high level, CNPs consists of a meta-learner and task-learner. Figure 14 shows the general architecture of CNPs. In this figure, h is an embedding function that generates representations r_i from the input and ground truth (labels) from the support set S . Operator a is the aggregation function that produces a single representation from r_i , and g is the neural network (task learner) that predicts labels for a new sample from the query set based on the output of a . As seen in Figure 14, the meta-learner produces a memory value by aggregating representations of the support set through operator a . Then, the label of new inputs is predicted by processing the aggregated representations. In other words, CNPs use concise representations of seen classes to make predictions.

[106] proposed a memory augmented matching network by combining MANN and matching network. The authors argued that if support data distribution is skewed, constructing a prototype by mean of embeddings leads to biased class prototypes. To support this argument, they developed weighted class prototypes by incorporating the distances of class-wise samples. The weight of sample i in each class is calculated as the inverse of the total distance of sample i to all other samples in that class. The results showed that weighted class prototypes lead to remarkable improvement.




Fig. 14. Conditional neural processes (CNP) architecture synthesis. Showcases the CNP’s embedding and aggregation stages for memory-efficient prediction without external memory reliance.

3.3 Learning-based Methods

This work categorizes the three groups of *learning the initialization*, *learning the parameters*, and *learning the optimizer* as a subset of a larger group *learning-based methods*. In previous surveys, three groups were introduced as model-based methods or optimization-based methods, but these methods are mechanically different. This work refers to these three groups as learning-based methods because a network (meta-learner) is trained to learn shared global initializer (learning the initialization), to learn the learning rate (learning the optimizer), or to learn the network’s parameters for the base learner (i.e., the model that predicts for unseen tasks). The following subsections discuss each group in detail.

3.3.1 Learning the Initialization. In a standard neural network, weights are initialized randomly, then after calculating loss, optimal weights are obtained by minimizing loss through the gradient descent process. The learning an initialization algorithms seeks to generate a better shared global initializer learned from the support set so that a new unseen task can learn quickly with fewer data points and less optimization steps. The **model-agnostic meta-learning (MAML)** is a pioneering learning the initialization algorithm proposed by [22]. Intuitively, MAML aims to learn a model’s parameters in such a way that a gradient-based learning rule on a new unseen task T progresses quickly. Figure 15 shows how MAML can lead to improvement on the direction of the gradient. In this figure, three different related tasks are represented with optimal parameters θ'_1 , θ'_2 and θ'_3 , where θ is the randomly initialized parameter. By shifting θ to a new position closer to θ'_1 , θ'_2 and θ'_3 , these optimal parameters can be achieved quickly with fewer steps. The general architecture of MAML consists of a model denoted by function f_θ with parameter θ . The model parameter θ is updated by Equation (13):

$$\theta'_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{T_i}(f_\theta) \quad (13)$$

where $\mathcal{L}_{T_i}(f_\theta)$ is a gradient of loss for task T_i ; and α is a hyper-parameter. In this step (inner loop), the optimal parameter θ'_i for each task T_i is obtained. In the next step (outer loop), the randomly initialized parameter θ is updated on a new set of tasks by Equation (14). The flow of inner and outer loops update rule is shown in Figure 16. Thus, parameter θ , initialized randomly, shifts to an optimal position where fewer gradient steps are required for the next batch of tasks.

$$\theta = \theta - \beta \nabla \sum_{T_i \sim p(T)} \mathcal{L}_{T_i}(f_\theta) \quad (14)$$

where β is a hyper-parameter; $p(T)$ is distribution over tasks T_i and $\mathcal{L}_{T_i}(f_\theta)$ is the gradient for each new task T_i .




Fig. 15. Efficient learning initialization via MAML: fast adaptation to new tasks with fewer gradient steps.




Fig. 16. Streamlined workflow for MAML: initializing and adapting model parameters across multiple tasks.

Continuing their previous work, [24] proposed probabilistic MAML, inspired by [33] who interpreted MAML as an approximate inference for the posterior distribution of parameters of the support set. The proposed model, termed **probabilistic latent model for incorporating priors and uncertainty (PLATIPUS)**, was able to capture uncertainty via variational inference. Numerical experiments indicated that PLATIPUS can effectively be used for a diverse family of tasks with multi-modal task distributions. In another tweak to MAML, [122] proposed **Bayesian Model-Agnostic Meta-Learning (BMAML)** to learn uncertainty, which principally combines the **Stein variational gradient descent (SVGD)** with the MAML algorithm. Their experiments demonstrated that this model is robust to overfitting during the meta-update phase. To alleviate the difficulties of learning in high-dimensional parameter space such as those faced by MAML, [93] proposed a novel algorithm called **latent embedding optimization (LEO)**. Similar to MAML, LEO consists of an inner loop where the task-specific values are obtained and an outer loop where global shared initializations are updated. Theoretically, LEO improves generalization performance by lowering the dimension of embedding space by using a combination of an encoder and relation network. The workflow of LEO is shown by Figure 17. The encoder produces hidden codes from the support set. Then, hidden codes are concatenated pairwise and fed into a relation network, resulting in a probability distribution over latent codes in a lower dimension. Finally, the decoder generates task-specific initial parameters that are differentiable to backpropagate for adaptation.

While meta-learning algorithms have a remarkable performance on benchmark few-shot problems, they are not robust to the adversarial samples [32]. [131] proposed a modified MAML model called **context adaptation for meta-learning (CAML)** by introducing context parameters. Context parameters essentially are task-specific parameters updated in an inner loop. In each iteration of the adaptation process, context parameters are initialized with zero value. This modification leads to better memory management and faster adaptation. Additionally, through the numerical examples, the authors demonstrated that CAML outperforms MAML in regression tasks. [121] proposed a network robust to adversarial samples by training a MAML model on both clean and adversarial samples. In the developed **ADversarial Meta-Learner (ADML)** model, inner and outer gradient updating is performed regarding correlations between clean and adversarial




Fig. 17. LEO algorithm for task-specific parameter initialization: enhancing generalization with encoder-relation network coupling.

examples. This research opens up a new direction in meta-learning. However, the potential problem of meta-learning models, such as MAML, with learning the initialization is that they only learn initializers, which can be biased toward some tasks and leads to poor generalization, specifically on unseen tasks. Addressing this issue, [46] proposed a **task-agnostic meta-learner (TAML)** model, in which two approaches are used to train a model unbiased over tasks. The first approach, “entropy maximization reduction,” makes a random guess over predicted labels, so the initial model will have a large entropy. Then, entropy is combined with the meta-learner loss to minimize the entropy, or in other words make the model more confident about the predicted labels. The limitation of the entropy-based approach is that it can only apply to discrete labels. Thus, the authors proposed the second approach, “inequality-minimization,” inspired by economic inequalities theory. They used a large group of statistics, including the Gini-coefficient, Theil-index, and variance of logarithm, to compute the tasks bias. They presumed the loss of the initial model of each task T_i as income inequality, which was minimized by the meta-learner.

As can be seen in Equation (13) and Equation (14), MAML has two hyperparameters, α and β , that are required to be tuned, which can be expensive. [8] alleviate the tuning burden of MAML hyperparameters by utilizing the **hypergradient descent (HD)** algorithm (proposed by [7]). Theoretically, HD adapts hyperparameters by backpropagating on the hyperparameters (here, α and β) in conjunction with the original gradient descent iterations. [81] believed that representation change plays an important role in the performance of MAML. Theoretically, the body of the embedding function in the MAML inner loop is almost fixed and the embedding function’s head is responsible for updating the features. This mechanic led to only small changes in representations during MAML inner loop task updates, which authors termed as representation reuse. The authors proposed a new variant of the agnostic meta-learner model called **body-only update in inner loop (BOIL)**. Contrary to the previous belief of MAML, in their model, the body of the embedding function updates features, leading to remarkable changes in representations, phrased representation change. Empirical results on benchmark datasets showed that the learning body of the embedding function in the inner loop improves MAML performance. [5] investigated the possible disagreement of the initialization location between tasks as well as layers of neural networks. They argued that sharing an initialization between conflicting tasks, i.e., tasks with unaligned gradient directions, degenerates the optimization landscape. To solve this issue, they proposed a new technique called **learn to forget (L2F)**, which is designed to selectively forget prior knowledge for each task. For this purpose, they proposed a learnable parameter for each layer that is computed through gradients of the support set batch of tasks. Through numerical analysis, the authors showed that the proposed attenuation parameters lead to a stable performance improvement of MAML. [33] empirically showed that MAML can be cast within a hierarchical Bayesian




Fig. 18. Implicit MAML (iMAML) approach. While MAML differentiates through the optimization path, iMAML estimates local curvature for the exact meta-gradient

model to estimate task-specific parameters and, thus, proposed another variation of MAML called **Laplace approximation for meta-adaptation (LLAMA)**. More specifically, LLAMA is an extension of MAML that incorporates uncertainty about task-specific parameters, so instead of point estimation, Laplace approximation is applied to obtain the probability distribution over parameter tasks T_i . While LLAMA resulted in a more accurate estimate than MAML's point estimation, Laplace approximation might be quite inaccurate. [87] proposed another extension to MAML called **implicit MAML (iMAML)**. Figure 18 illustrates the difference between MAML and iMAML. While MAML differentiates the meta-gradient through the optimization path, the proposed iMAML estimates local curvature for the exact meta-gradient and is also more memory-efficient since there is no need to differentiate throughout the optimization path.

3.3.2 Learning the Parameters. In this group of algorithms, the parameters of task-specific networks (i.e., base learner) are generated by another learner (i.e., meta-learner) instead of being initialized randomly, leading to a faster adaptation for new tasks. One early work in this category was conducted by [20], who proposed neural statistician, a novel extension of the variational auto-encoder that summarizes dataset and learn dataset statistics. In simple words, a neural statistician generates summary statistics, including mean and variance, to provide a generative model that can be trained unsupervised and is used to estimate the posterior distribution of task-specific parameters. However, the major limitation of the neural statistician is that its performance depends on a large dataset. [79] proposed a meta-learner model called MetaNet to quickly parameterize a predictive model (i.e., a neural network) by processing meta information. MetaNet consists of two main modules: the base learner operates at the level of tasks and provides the meta-learner with information including loss gradients; and the meta-learner principally generates weights quickly in a task-agnostic setting by processing received information. These generated weights are employed to leverage both the meta- and base learners. In other words, the meta-learner parameterizes itself and the base learner. Here, the meta-learner benefits an external memory to store the support set inputs, their representations, and generated example level weights. Prediction over the query set is made based on comparing the query sample's representation against stored representation in memory. In other words, weights of the base learner would compute based on the similar support set information in memory. Finally, the base learner makes predictions based on the computed weights and seen query points. [10] targeted the one-shot classification problem and proposed non-iterative feed-forward neural nets (called learnet) to predict the parameters of a task-specific predictor (named pupil network). While their model architect may be reminiscent of Siamese networks, the meta-learner (i.e., learnet) is dynamic and changes regarding the output of mappings. [61] proposed a novel model called LGM-Net consisting of two main components: (i) TargetNet, which operates in task-specific environment; and (ii) MetaNet, which is responsible for generating parameters for TargetNet. This MetaNet is different from the Meta-Net model presented by [79]

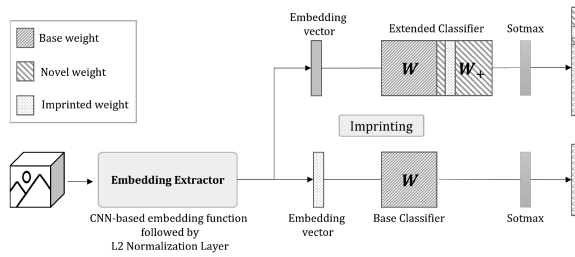


Fig. 19. An overview of the weight imprinting [86] process. Feature vectors of latent space are utilized as imprinted weights in a fully connected layer of the extended classifier which is used for the prediction of new query point. Source: [86]

and consists of two modules: (i) task context encoder and (ii) weight generator. The task context encoder is responsible for learning representations, while the weight generator is designed to compute the conditional distributions of weights. The authors used matching networks as TargetNet which use generated weights to compute matching probability score as well as classification loss. Another novelty of this study is the introduction of intertask normalization, a strategy that allows similar tasks in a batch of tasks to interact with each other.

[30] proposed a dynamic algorithm by introducing an attention-based classification weight generator and a cosine-similarity based ConvNet recognition model. The former is responsible for producing weights vectors for unseen classes by seeing a few examples from those classes. The authors argued that it is not feasible for a ConvNet model with a typical linear layer classifier to handle classification weights of both base and novel classes. They address this issue by proposing the second component that provides prediction by computing the cosine similarity between the generated weights and samples' representations. The combination of these two components allowed the authors to achieve fast adaptation for novel classes without forgetting the base classes. Continuing the previous work and with the same belief of developing the meta-learner without forgetting base classes, [31] proposed a **denoising auto-encoder network (DAE)** in the form of **graph neural network (GNN)**, which is responsible for regularizing the weight generating meta-model. Basically, DAE takes contaminated classification weights with Gaussian noise and generates classification weights based on the conditional distribution of seen data from both novel and base classes. Results showed that this model outperforms their previous model in [30].

[86] argued that embedding vectors are comparable with the weights of the last linear layer of the ConvNet classifier. Accordingly, instead of using two different components of meta-learner and base learner, they proposed a new process called weight imprinting that basically sets final layer weights based on normalized embeddings of the novel class's training samples. Figure 19 shows the overall flow of their model, in which an embedding function first trains on base classes then tunes on novel classes where imprint weights are used in the final layers. With a similar belief, [90] addressed the issue of forgetting the base classes by proposing incremental learning. Formally, the classifier trains on base classes and learns a set of embeddings. Then, the meta-learner trains on novel classes. This study also introduced an additional regularizer that is optimized by combining both novel and base classes and is computed by an attention attractor network, which solves the catastrophic forgetting problem.

Until now, the reviewed studies in this section focused on generating parameters for a classifier network. [115] proposed a new model called **task-aware feature embeddings for low-shot learning (TAFE-Net)** that concentrates on generating parameters of feature layers to construct feature embeddings tuned for a given task. TAFE-Net consists of two components: (i) meta-learner

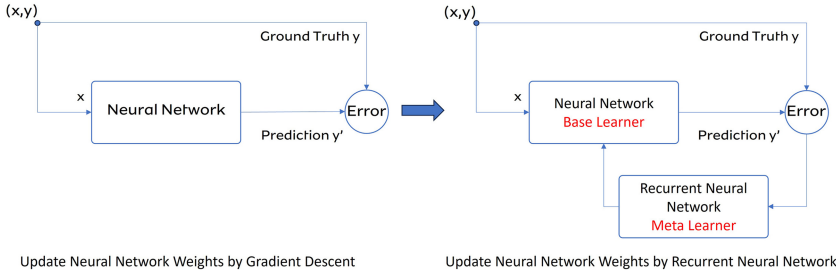


Fig. 20. Optimization via RNN meta-learner: learning gradient descent for faster neural network training.

that is designed to generate parameters for feature layers in the prediction network; and (ii) prediction network, which generates a binary label that indicates that the input image x is compliant to task description T . The authors used well-known pre-trained deep neural networks, such as ResNet, to get the features of input images. These generated features feed to the dynamic feature layers. The output of this layer is called task-aware feature embedding since feature embedding of each image may vary depending on the task description. Similarly, [102] introduced the **meta-transfer learner (MTL)** model, focusing on developing task-specific feature extractors by leveraging both transfer learning and meta-learning. Conceptually, they trained a deep neural network on a large dataset, such as MiniImageNet dataset, and froze the feature extractor by scaling and shifting operations on pre-trained feature embeddings. This work also introduced a novel hard task meta-batch process, which increases the meta-learner focus on hard tasks through sampling cases that the classifier failed.

3.3.3 Learning the Optimizer. A standard neural network trains by calculating loss and using gradient descent to minimize the loss. The key idea of learning the optimizer is to learn gradient descent by replacing it with the RNN. Figure 20 shows the high-level mechanics of learning the optimizer algorithms. The base learner's parameters are updated through an RNN network (meta-learner). Since RNN optimizes itself through gradient descent, these models are also known as learn gradient descent by gradient descent. [88] proposed an LSTM-based meta-learner to learn the optimization process used to train task-specific learner.

Equation (15) denotes the standard gradient descent update rule:

$$\theta_t = \theta_{t-1} - \alpha_t \nabla_{\theta_{t-1}} \mathcal{L}_{Tj}(\theta_{t-1}) \quad (15)$$

where θ_{t-1} is the parameter of the learner at $t - 1$ step; α_t is the learning rate; and $\nabla_{\theta_{t-1}} \mathcal{L}_{Tj}$ is the gradient of loss for task T parameterized with θ at step $t - 1$. The authors argued that Equation (15) resembles the update cell in LSTM and replaced term $-\alpha_t \nabla_{\theta_{t-1}} \mathcal{L}_{Tj}(\theta_{t-1})$ by LSTM network denoted by g and parameterized with φ . Thus, Equation (15) becomes Equation (16):

$$\theta_{t+1} = \theta_t + g_{\varphi}(\nabla_{\theta_t} \mathcal{L}_{Tj}(\theta_t)) \quad (16)$$

With this new update rule (Equation (16)), the meta-learner (i.e., LSTM) is able to learn the learning rate θ , which leads to task adaptation quickly. **Meta stochastic gradient descent (Meta-SGD)** is another learning the optimizer algorithm proposed by [64], which is able to learn the learning rate on top of learning initialization too. Recall the standard gradient descent update rule (Equation (15)), Meta-SGD proposed a new update rule as Equation (17):

$$\theta_t' = \theta_{t-1} - \alpha \odot \nabla_{\theta_{t-1}} \mathcal{L}_{Tj}(\theta_{t-1}) \quad (17)$$

where α is a vector instead of a scalar with the same shape as θ . The operator \odot denotes the element-wise product. In an inner loop, task-specific parameters are updated, and in an outer loop,



Fig. 21. Reptile scheme shows the movement of the random parameter θ towards tasks optimal parameters.

meta-related optimization is performed on both θ and α by Equation (18) and (19).

$$\theta_t = \theta_{t-1} - \beta \nabla_{\theta, t-1} \sum_{T_i \sim p(T)} \mathfrak{L}_{T_j}(f(\theta_{i, t-1}')) \quad (18)$$

$$\alpha_t = \alpha_{t-1} - \beta \nabla_{\alpha, t-1} \sum_{T_i \sim p(T)} \mathfrak{L}_{T_j}(f(\theta_{i, t-1}')) \quad (19)$$

Another learning the optimizer algorithm called Reptile proposed by [80] tries to generate a task-specific learner close to optimal values. In this regard, the Reptile concept is similar to MAML but quite simpler. Figure 21 shows the mechanism of Reptile. Assume two tasks, T_1 and T_2 , sampled from task distribution $p(T_i)$ randomly, where each task is initialized randomly with parameter θ . For each task, T_1 and T_2 , optimal parameters θ_1' and θ_2' will be obtained through gradient descent after n steps. Then, a random parameter θ will shift in the direction closer to both obtained parameters, θ_1' and θ_2' , by minimizing the Euclidean distance of parameter θ and task-specific optimal parameters θ_1' and θ_2' , as shown in Equation (20):

$$\text{minimize } E \left[\frac{1}{2} D(\theta, \theta_i') \right] \quad (20)$$

Finally, gradient descent is performed on Equation (20) to obtain the following update rule (Equation (21)):

$$\theta = \theta - \epsilon (\theta_i' - \theta) \quad (21)$$

where ϵ is a step size parameter. [9] introduced a new idea about the existing learning the optimizer methods, that is using a non-deep machine learning algorithm to adapt new tasks from few samples. Let's consider a pre-trained network (typically, a CNN-based network) $\phi(x) : \mathbb{R}^m \rightarrow \mathbb{R}^e$ (where \mathbb{R}^m and \mathbb{R}^e denote the input and embedding space, respectively); and episode-specific predictor $f(\phi(x), w_\epsilon) : \mathbb{R}^e \times \mathbb{R}^p \rightarrow \mathbb{R}^o$ (where \mathbb{R}^p and \mathbb{R}^o refer to the predictor parameters and output spaces, respectively; and $w_\epsilon \in \mathbb{R}^p$ is a set of task-specific predictor parameters). Given the support set of episode, ϵ , $z_\epsilon = \{(x_i, y_i) \sim \epsilon\}$ and query set of episode ϵ , $z'_\epsilon = \{(x'_i, y'_i) \sim \epsilon\}$, f can be considered as a final layer of embedding function denoted by $f(\phi(x)) = \phi(x)W$, ($w_\epsilon \sim W$). The authors argued that f can be trained to obtain parameter w_ϵ through a learning algorithm Λ , which can be considered as a typical machine learning algorithm applied in an episodically training fashion and incorporating meta-parameters w and ρ . These meta-parameters allow knowledge transfer between episodes. Accordingly, the authors proposed two algorithms, namely **Ridge Regression Differentiable Discriminator (R2-D2)** and **Logistic Regression Differentiable Discriminator (LR-D2)**. In R2-D2, the optimal W is obtained by Equation (22):

$$\begin{aligned} \Lambda(z) &= \text{argmin} \|XW - Y\|^2 + \lambda \|W\|^2 \\ &= (X^T X + \lambda I)^{-1} X^T Y \end{aligned} \quad (22)$$

where X , Y , and λ are embeddings, outputs, and regularization terms, respectively. Considering the high dimensionality of embeddings, $X^T X \in \mathbb{R}^{e \times e}$ can be computationally expensive since the

matrix $X^T X$ grows quadratically. The authors addressed this issue by using the Woodbury formula, as given in Equation (23):

$$W = \Lambda(z) = X^T (XX^T + \lambda I)^{-1} Y \quad (23)$$

With small number of n in the n -shot problem, the computation cost of Equation (23) is generally small. Although Equation (22) originally works for regression tasks, the authors modified it to Equation (24) for classification tasks:

$$\hat{Y} = \alpha X' W + \beta \quad (24)$$

The meta-learner learns α , β , and λ (hyper-parameters of base learner) along with CNN parameters W . The authors asserted that other similar solvers can be employed for differentiable operations. Thus, they proposed the LR-D2 method, which is basically logistic regression that uses Newton's method as an iterative solver. LR-D2, unlike R2-D2, is designed to apply to classification tasks directly. Considering inputs $X \in \mathbb{R}^{n \times e}$ and binary label $y \in \{-1, +1\}^n$, parameters $w_i \in \mathbb{R}^e$ are updated by Equation (25):

$$w_i = (X^T \text{diag}(s_i) X + \lambda I)^{-1} X^T \text{diag}(s_i) z_i \quad (25)$$

where w_i denotes the parameters at iteration i ; and I stands for the identity matrix. Moreover, $s_i = \mu_i(1 - \mu_i)$, $z_i = w_{i-1}^T X + (y - \mu_i)/s_i$ and $\mu_i = \sigma(w_{i-1}^T X)$, where σ is the sigmoid function. Similarly, the authors applied Woodbury's identity on $X^T \text{diag}(s_i) X$ and modified Equation (25) to Equation (26):

$$w_i = X^T (XX^T + \lambda \text{diag}(s_i)^{-1})^{-1} z_i \quad (26)$$

With the same idea of using a machine learning classifier as a base learner, [58] proposed a new method called MetaOptNet. In other words, a **support vector machine (SVM)** is applied as a base learner incorporated with a differentiable **quadratic programming (QP)** solver to learn parameters of the embedding function f_ϕ . This method was shown to outperform both R2-D2 and LR-D2.

4 OVERVIEW, CHALLENGES AND FUTURE DIRECTIONS

Along with the considerable success of deep learning, its impotence in domain generalization and avaricious craving for a large dataset for training have increased the demand for few-shot meta-learning algorithms. Despite the fact that in recent years, meta-learning has attracted remarkable interest from researchers, meta-learning is still in its infancy. Although the Siamese network was originally proposed in 1993, the paper "Siamese neural networks for one-shot image recognition" in 2015 [55] triggered a new era of meta-learning and metric-based algorithms. While Siamese networks are limited to binary classification tasks, other metric-based methods resolve this limitation by training on tasks in an episodic training manner. Prototypical networks bring the idea of using class prototypes to compare every new query point instead of individual comparisons in Siamese networks. This idea was the starting point of many metric-based methods. Relation networks employ a RNN instead of a fixed similarity function, leading to a dynamic nature of learning domain-related similarity functions. A general view on recent metric-based meta-learners demonstrates that the concentration on inter-class and intra-class relationships and a proposed weighted prototype has been increased. Compared to other non-metric based models, metric-based methods are simpler but more innovative in proposing novel similarity measures. Given the pair-wise comparisons – whether a comparison with class prototypes or direct comparison with support data points – metric-based methods can be computationally expensive, especially with large tasks. The next group, memory-based meta-learning models, have received less attention than other groups, but show strong performance. The key advantage of memory-based models is their robustness to forget base classes as the fruit of external/internal memory. Finally, the last large group, the

learning-based methods, comprises the major area of research. The LSTM-based meta-learner was the first to materialize the idea of learning gradient descent by gradient descent. It will not be an overstatement if MAML is considered the most inspiring method and commencement mark of learning the initialization methods. While the learning-based methods outperform other methods on a broader range of tasks [23], they can be computationally expensive since they optimize a base learner for every task [23]. Table 1 compares the performance of the discussed methods on benchmark datasets extracted from the papers. The experimental results from the Omniglot and MiniImageNet datasets offer a rich tapestry for understanding the nuances of few-shot learning across different meta-learning methods. Figures 22, 23, 24, 25, 26, 27, summarize the Table 1 results for the Omniglot, MiniImageNet and CUB200-2011 5-way problems.

Starting with the Omniglot dataset, the clear trend is the superior performance of models in the 5-shot setting compared to the 1-shot setting. This is intuitive as the 5-shot setting provides more examples for the model to learn from, which is particularly beneficial in few-shot learning scenarios where data scarcity is a primary challenge. When examining specific models, Meta-SGD (99.93%), Prototype-relation networks (99.91%), MAML (99.90%), and Global class representation (99.90%) stand out with the highest accuracy in the 5-way 5-shot problem on the Omniglot dataset.

Metric-based methods like Prototypical Networks and their extensions show strong performance on the Omniglot dataset, which is known for its simpler tasks with less intra-class variance. These methods effectively leverage class prototypes and attention to local features, which seems to be particularly effective for the character recognition tasks in Omniglot. However, when faced with the more complex and visually diverse MiniImageNet dataset, the performance of these methods may not be as strong, indicating a potential limitation in handling tasks with greater intra-class variability and more complex feature distributions.

The holistic approach of GNN to task structure seems to provide an edge on datasets like Omniglot. However, the performance gap between GNNs and other methods on MiniImageNet suggests that while GNNs can capture complex interrelationships between data points, they may require more sophisticated adaptations to handle high intra-class variation effectively.

Memory-based methods such as MANN and SNAIL show a consistent performance across different datasets, which may be attributed to their ability to leverage past experiences and focus on relevant features. Their robustness to forgetting and ability to draw on a rich memory of past tasks could explain their reliable performance across different levels of task complexity.

Learning the Initialization methods like MAML are designed to find a suitable starting point for new tasks, which should, in theory, enable them to perform well on both the Omniglot and MiniImageNet datasets. However, the results indicate that while MAML performs well on simpler tasks, it may struggle with the more complex MiniImageNet tasks, potentially due to overfitting to the simpler structures of the Omniglot dataset or an initialization that doesn't generalize well across diverse tasks.

Learning the Parameters methods are expected to adapt quickly to new tasks due to task-specific parameter generation. Yet, the results show that while they perform competently, they do not always outperform simpler metric-based methods, suggesting that the generation of parameters may not always capture the nuances required for optimal performance on complex tasks.

The experimental results from CUB200-2011 offer metric-based methods like the MPNet and MACO perform moderately well, indicating their capability to capture distinct features even with a single example. However, their performance is not as high as more sophisticated methods, likely due to the complexity of the CUB200 dataset which has higher intra-class variance compared to Omniglot. Notably, the RCN significantly outperforms other methods, suggesting that focusing on specific regions and comparing them can be particularly effective for fine-grained tasks where subtle differences between classes are important. The MSFN performs notably better than in

Table 1. Accuracy Results on Benchmarks Datasets Reported in Original Papers, i: Metric-Based; ii: Memory-Based; iii: Learning the Initialization; iv: Learning the Parameter; v: Learning the Optimizer; The \pm Indicates the 95% Confidence Interval

Model	Category	Omniglot				MiniImageNet				CUB-200-2011			
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Siamese [55]	i	97.3%	98.4%	88.2%	97%	-	-	-	-	-	-	-	-
Prototypical network [101]	i	98.8%	99.7%	96%	98.9%	49.42 \pm 0.78%	68.20 \pm 0.66%	-	-	-	-	-	-
Semi Supervised PN [91]	i	97.52 \pm 0.07%	-	-	-	50.41 \pm 0.31	64.39 \pm 0.24	-	-	-	-	-	-
Gaussian PN [26]	i	99.07 \pm 0.03%	99.73 \pm 0.02%	96.94 \pm 0.14%	99.29 \pm 0.04%	-	-	-	-	-	-	-	-
Transferable prototypical network (TPN) [84]	i	-	-	-	-	-	-	-	-	-	-	-	-
Domain alignment PN [70]	i	-	-	-	-	50.56 \pm 0.85%	69.62 \pm 0.76%	-	-	-	-	-	-
Improved prototypical networks (IPN) [47]	i	-	-	-	-	56.18 \pm 0.85%	74.60 \pm 0.83%	-	-	-	-	-	86.81%
Label enhanced prototypical network (LPN) [67]	i	-	-	-	-	-	-	-	-	-	-	-	-
Multi-prototypes PN [17]	i	-	-	-	-	-	-	-	-	-	-	-	-
Dummy prototypical networks [54]	i	-	-	-	-	-	-	-	-	-	-	-	-
Multi-prototype network (LMPNet) [42]	i	-	-	-	-	62.74 \pm 0.11%	80.23 \pm 0.52%	65.59 \pm 0.13%	68.19 \pm 0.23%	75.01 \pm 0.81%	85.30 \pm 0.54%	-	-
Multi-modal prototypical network [82]	i	-	-	-	-	-	-	-	-	-	-	-	-
Matching Networks [108]	i	98.1%	98.9%	93.8%	98.7%	46.6%	60%	-	-	-	-	-	-
Cascaded feature matching network (CFMN) [13]	i	99.7 \pm 0.2%	99.8 \pm 0.1%	98.0 \pm 0.2%	99.2 \pm 0.1%	52.98 \pm 0.84%	68.33 \pm 0.70%	-	-	-	-	-	-
Attentive matching network (AMN) [74]	i	99.44 \pm 0.09%	99.86 \pm 0.06%	98.06 \pm 0.11%	99.50 \pm 0.07%	55.30 \pm 0.89%	71.80 \pm 0.78%	-	-	-	-	-	-
Schedule sampling for matching network [126]	i	-	-	-	-	48.3%	65.3%	-	-	-	-	-	-
Relation networks [103]	i	99.6 \pm 0.2%	99.8 \pm 0.1%	97.6 \pm 0.2%	99.1 \pm 0.1%	50.44 \pm 0.82%	65.32 \pm 0.70%	-	-	-	-	-	-
Memory-augmented relation networks (MRN) [39]	i	-	-	-	-	57.83 \pm 0.69%	71.13 \pm 0.50%	-	-	-	-	-	-
Prototype-relation networks [69]	i	99.27 \pm 0.23%	99.91 \pm 0.06%	95.97 \pm 0.56%	99.32 \pm 0.13%	49.54 \pm 0.09%	68.34 \pm 0.06%	-	-	-	-	-	-
Multi-scale relation networks (MSRN) [19]	i	-	-	-	-	50.21 \pm 1.08%	65.89 \pm 0.32%	-	-	-	-	-	-
Multi-local feature relation network (MLFRNet) [89]	i	-	-	-	-	66.79 \pm 0.47%	83.16 \pm 0.39%	-	-	-	-	-	-
Graph neural networks (GNN) [27]	i	99.2%	99.7%	97.4%	99%	50.33 \pm 0.36%	66.41 \pm 0.63%	-	-	-	-	-	-
Global class representation (GCR) [59]	i	99.72 \pm 0.06%	99.90 \pm 0.10%	99.63 \pm 0.09%	99.32 \pm 0.04%	53.21 \pm 0.40%	72.34 \pm 0.32%	-	-	-	-	-	-
Multi-scale feature network (MSFN) [38]	i	99.7%	99.8%	98.1%	99.2%	54.48 \pm 1.23%	69.06 \pm 0.69%	62.4%	79.14%	-	-	-	-
Attentive recurrent comparators [99]	i	-	-	-	-	49.14%	-	-	-	-	-	-	-
Relative position and map networks (RPMN) [120]	i	-	-	-	-	53.35%	69.35%	-	-	-	-	-	-
Region comparison network (RCN) [119]	i	-	-	-	-	57.40 \pm 0.86%	75.19 \pm 0.64%	74.65 \pm 0.86%	88.81 \pm 0.57%	-	-	-	-
Multi-scale kronecker-product relation networks [1]	i	-	-	-	-	57.02 \pm 0.88%	72.06 \pm 0.68%	69.49 \pm 0.95%	82.94 \pm 0.65%	-	-	-	-
Metric-agnostic conditional embeddings (MACO) [40]	ii	-	-	-	-	41.09%	58.32%	60.76%	74.96%	-	-	-	-
MANN [95]	ii	82.8%	94.9%	-	-	-	-	-	-	-	-	-	-
SNAIL [76]	ii	99.07 \pm 0.16%	99.78 \pm 0.09%	97.64 \pm 0.30%	99.36 \pm 0.18%	55.71 \pm 0.99%	68.88 \pm 0.92%	-	-	-	-	-	-
CNFs [29]	ii	95.3%	98.5%	89.9%	96.8%	-	-	-	-	-	-	-	-
Memory augmented matching network [106]	ii	98.9%	99.7%	96.3%	98.9%	49.8%	66.5%	-	-	-	-	-	-
MAML [22]	iii	98.7 \pm 0.4%	99.9 \pm 0.1%	95.8 \pm 0.3%	98.9 \pm 0.2%	48.70 \pm 1.84%	63.11 \pm 0.92%	-	-	-	-	-	-
PLATIPUS [24]	iii	-	-	-	-	-	-	-	-	-	-	-	-

(Continued)

Table 1. Continued

Model	Category	Omniglot				MiniImageNet				CUB-200-2011	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
BMAML [122]	iii	-	-	-	-	53.80 ± 1.46%	-	-	-	-	-
LEO [93]	iii	-	-	-	-	61.76 ± 0.08%	77.59 ± 0.12%	-	-	-	-
CAML [131]	iii	-	-	-	-	51.82 ± 0.65%	65.85 ± 0.55%	-	-	-	-
ADML [121]	iii	-	-	-	-	44.00 ± 1.83%	-	-	-	-	-
TAML [46]	iii	99.47 ± 0.25 %	99.83 ± 0.09%	95.62 ± 0.5%	98.64 ± 0.13%	51.73 ± 1.88%	66.05 ± 0.85%	-	-	-	-
Alpha MAML [8]	iii	-	-	-	-	-	-	-	-	-	-
BOIL [81]	iii	-	-	-	-	49.61 ± 0.16%	66.45 ± 0.37%	-	-	-	-
MAML+L2F [5]	iii	-	-	-	-	62.12 ± 0.13%	78.13 ± 0.15%	-	-	-	-
iMAML [87]	iii	99.50 ± 0.26%	99.74 ± 0.11%	96.18 ± 0.36%	99.14 ± 0.1%	49.30 ± 1.88 %	-	-	-	-	-
LLAMA [33]	iii	-	-	-	-	49.40 ± 1.83%	-	-	-	-	-
MetaNet [79]	iv	98.45%	-	95.92%	-	49.21 ± 0.96%	-	-	-	-	-
Pupil network [10]	iv	-	-	-	-	-	-	-	-	-	-
LGM-Net [61]	iv	99%	99.4%	96.5%	98.5%	69.13 ± 0.35%	71.18 ± 0.68%	-	-	-	-
Dynamic learning without forgetting [30]	iv	-	-	-	-	74.92 ± 0.36%	58.55 ± 0.50%	-	-	-	-
DAE [31]	iv	-	-	-	-	61.07 ± 0.15%	76.75 ± 0.11%	-	-	-	-
Weight imprinting [86]	iv	-	-	-	-	-	-	-	-	-	-
Incremental learning with attention attractor networks [90]	iv	-	-	-	-	54.95 ± 0.30%	63.04 ± 0.30%	-	-	-	-
TAFE-Net [115]	iv	-	-	-	-	-	-	-	-	-	-
Meta-transfer learner (MTL) [102]	iv	98.1%	99.5%	93.2%	98.1%	61.2 ± 1.8%	75.5 ± 0.8%	-	-	-	-
Neural statistician [20]	v	-	-	-	-	-	-	-	-	-	-
LSTM-based meta-learner [88]	v	-	-	-	-	43.44 ± 0.77%	60.60 ± 0.71%	-	-	-	-
Meta-SGD [64]	v	99.53 ± 0.26%	99.93 ± 0.09%	95.93 ± 0.38%	98.97 ± 0.19%	50.47 ± 1.87%	64.03 ± 0.94%	-	-	-	-
Reptile [80]	v	97.68 ± 0.04%	99.48 ± 0.06%	89.43 ± 0.14%	97.12 ± 0.32%	49.97 ± 0.32%	65.99 ± 0.58%	-	-	-	-
MetaOptNet [58]	v	-	-	-	-	64.09 ± 0.62%	80.00 ± 0.45%	-	-	-	-
R2-D2 [9]	v	98.91 ± 0.05%	99.74 ± 0.02%	96.24 ± 0.05%	99.20 ± 0.02%	51.8 ± 0.2%	68.4 ± 0.2%	-	-	-	-
LR-D2 [9]	v	-	-	-	-	51.9 ± 0.2%	68.7 ± 0.2%	-	-	-	-

“-” indicates not reported. The MANN [95] performance is extracted from [29].




Fig. 22. 5way-1shot Omniglot accuracy results.




Fig. 23. 5way-5shot Omniglot accuracy results.

the 1-shot scenario, which may be due to its ability to utilize the additional information from multiple examples to refine its understanding of each class’s features. Once again, the RCN shows the highest accuracy, affirming its effectiveness in fine-grained image recognition tasks. The additional examples in the 5-shot setting likely allow the RCN to make more accurate region comparisons, leading to better performance.

Having examined the performance of different methods on specific datasets, it’s important to note a broader trend observed across these methodologies. The performance of all methods generally improves from 1-shot to 5-shot learning, indicating that having more examples aids in the adaptation process for new tasks. This improvement highlights another crucial aspect. The gap in performance between methods on the 5-way and 20-way classification tasks suggests that as the task complexity increases (more classes to distinguish between), the ability of these methods to generalize effectively from few examples varies significantly. In line with these findings, let’s consider the MiniImageNet dataset, which exemplifies this complexity. The top-performing methods on the more complex MiniImageNet dataset, such as MLFRNet, indicate the importance of




Fig. 27. 5way-5shot MinilImageNet accuracy results.

Despite the promising achievements of meta-learning, some challenges remain. These existing issues along with suggested future research directions are presented here.

- **Uncertainty of meta-learners’ performance.** As the name of few-shot learning implies, the main challenge is learning from a handful of samples. This issue has been addressed by meta-learning via sharing knowledge or generating knowledge for each unseen task from previously experienced tasks. In this context, experience is interpreted as fitting to data. However, during episodic training, unrelated tasks or even less related tasks may result in divergent optimization directions as well as disparate performance, which eventually may lead to uncertainty in unseen class prediction [46]. The disparity between Omniglot and more complex sets like MiniImageNet, raises concerns about the predictability and reliability of meta-learners. Addressing this uncertainty demands a nuanced approach to training, perhaps by cultivating a diverse task environment that mirrors the variance seen in real-world applications. By clustering similar tasks or increasing heterogeneity within episodic training, meta-learners can be developed to exhibit enhanced adaptability and reduced performance volatility. In light of this challenge, a promising direction for future research involves the integration of **Bayesian Active Learning (BAL)** and **Variational Autoencoders (VAE)** in the realm of meta-learning. Bayesian Active Learning, with its focus on strategically selecting information-rich data points, could be invaluable in refining the episodic training process for meta-learners. It is hypothesized that this selection mechanism can substantially mitigate the uncertainty in unseen class prediction. The implementation of Variational Autoencoders is suggested for the analysis and encoding of task complexities. It is posited that by capturing task distributions in a latent space, VAEs could enable a deeper understanding of task variability. This understanding is critical for the effective clustering of similar tasks and the generation of nuanced task heterogeneity representations. Consequently, an enhancement in the adaptability and a reduction in the performance volatility of meta-learners are expected.
- **Effective use of episodic training.** The improvement in model performance from 1-shot to 5-shot learning reveals the potential of episodic training in meta-learning. However, episodic training is prone to catastrophic forgetting, which results in model underfitting in base

classes. While several researches have addressed this issue [30, 31, 47], enhancing model performance on both base and unseen classes remains an important direction for future work. Innovations that amalgamate the resilience of memory-based methods with the acuity of metric-based methods, as evidenced by their application in complex datasets like CUB200-2011, could offer a path forward. Such hybrid approaches may refine the episodic training paradigm, balancing the retention of foundational knowledge with the flexibility required for novel tasks.

- **Improving stability.** Although GCR, a metric-based meta-learner, achieved the best performance on the Omniglot dataset, metric-based models cannot compete with other non-metric-based in MiniImageNet a colored dataset that is more complex than Omniglot. In other words, metric-based meta-learning algorithms are sensitive to the dataset. Approaches like the RCN have shown promise in fine-grained image recognition tasks, suggesting a future research direction for developing more stable meta-learning algorithms that are less sensitive to dataset complexity. Moreover, it has been shown that other existing meta-learning algorithms, such as MAML, are not robust to adversarial samples [32]. While ADML investigates the adversarial samples to improve meta-learner performance, more exploration will be valuable.
- **Training multi-domain, multi-modal and cross-domain meta-learning.** In theory, D_{Base} and D_{Novel} can be from separate domains. However, in practice, as the difference between D_{Base} and D_{Novel} increases, most meta-learning performance will decline. This can be seen by the poor performance of metric-based methods on the MiniImageNet dataset. Additionally, the current methods are trained on a single domain D_{Base} . As a suggestion for future research exploration, meta-learning can be developed on multi-domain and cross-domain performance. While training on a multi-domain D_{Base} is expected to increase generalization ability, it can be quite challenging due to divergent optimization direction or the possibility of forgetting D_{Base} tasks. Therefore, using memory-based methods or grouping related tasks based on statistical information is worth exploring. Although [107] recently proposed a new benchmark dataset by combining Omniglot and ImageNet, which is useful for multi-domain and cross-domain evaluation, other benchmark datasets from a combination of real-world applications, such as remote sensing scene images and agricultural pest images, and other real-world domains should be investigated. Last, only a few researches have studied multi-modal domains [24, 82, 125]. Therefore, the development of multi-modal benchmark datasets and methods on multi-modal datasets will be valuable to study. Further research on the multi-modal datasets and real-world multi-domain and cross-domains datasets is pertinent to determine if the current methods are successful in learning from a few samples or simply re-using learned features.
- **Representation learning effects.** Different meta-learning methods use embedding functions with different architectures. While most methods use four-layer CNNs as a backbone, few studies have applied larger networks, such as resnet, and multi-scale networks as backbone [19, 38]. Accordingly, it may be insightful to analyze the effect of embedding functions, like multi-scale networks, or changing the architecture of the backbone or pre-trained neural networks on state-of-the-art methods such as prototypical networks performance as well as computational costs. Another intriguing extension of this exploration would be the integration of **Monte Carlo (MC) Dropout** within these advanced architectures. This technique, primarily recognized for its robustness in uncertainty estimation and regularization, could provide significant advancements in learning the optimizer methods. In the context of Meta-SGD, the integration of MC Dropout could be explored to enhance the robustness of the learned optimization process. By applying dropout during both the training and inference phases of the meta-learner (RNN or LSTM), it could provide a more nuanced adap-

tation to task-specific optimization challenges. This might result in a more effective learning rate adaptation and improved convergence properties, particularly in scenarios with high variability or uncertainty in the data.

- **Improvement in computational cost.** As mentioned earlier in this section, learning-based methods can be quite computationally expensive in spite of learning from a few samples. Future research must address the optimization of computational resources, seeking efficiency gains that do not compromise the model's learning capability. This endeavor is critical not only for academic advancement but also for the democratization of meta-learning applications across varied and resource-constrained environments. Incorporating **Neural Architecture Search (NAS)** into the development of meta-learning models presents a promising avenue for future research, aiming to optimize model architectures for both performance and computational efficiency. NAS automates the process of designing neural network architectures, which can lead to the discovery of more efficient models tailored for specific meta-learning tasks. NAS can be used to develop lightweight models that maintain high performance with fewer parameters and less computational overhead by exploring a vast space of architectural possibilities.
- **Enhancing transparency and trust in meta-learning.** The imperatives of **Explainable Artificial Intelligence (XAI)** within the domain of meta-learning are recognized as being multifaceted. In meta-learning algorithms, known for their capability to make accurate predictions from limited data, complexity is inherent and often obscures the inner workings of these models, thereby necessitating the application of XAI principles. In realms where rapid model adaptability is required, such as healthcare or finance, the trustworthiness and transparency of meta-learning models are considered non-negotiable. In meta-learning, the pivotal role of XAI is underscored due to the reliance on these methods for high-stakes decisions. For instance, it is noted that in medical diagnostics, meta-learning has the capability to identify rare conditions from scarce examples, necessitating that practitioners understand the rationale behind diagnostic recommendations for trust and actionable insights. Similarly, in automated financial systems where meta-learning might be employed to detect fraudulent patterns from minimal data, the necessity for algorithms to provide explanations justifying their alerts is emphasized.

The incorporation of existing XAI methodologies into meta-learning frameworks is suggested as a means to bridge the gap between high-performance predictive models and the need for interpretability, which is inherently human-centric. Current XAI methods, ranging from feature attribution to model visualization techniques, are identified as capable of pinpointing influential factors in a model's decision-making process and illuminating decision boundaries or structures within data. These methods are proposed to be tailored to meet the specific requirements of meta-learning models, which often operate with simpler architectures and reduced complexity due to data scarcity.

In the realm of meta-learning, specifically within algorithms like LSTM-based Meta-Learners and Meta-SGD, the adaptation of feature attribution methods such as **LIME (Local Interpretable Model-Agnostic Explanations)** and **SHAP (SHapley Additive exPlanations)** [3] is suggested. These methods are crucial for elucidating how different aspects of input data influence the updates made by the LSTM in the optimization process. For instance, in medical diagnostics tasks, understanding which features of a dataset drive the LSTM's updates can enhance trust in the model's decisions. Additionally, employing model visualization techniques can illustrate the effects of changes in the optimization algorithm, such as learning rate adjustments in Meta-SGD, on the base learner's learning trajectory.

For algorithms like Reptile, the application of global explanation methods is proposed to comprehend how the strategy of moving towards optimal parameters for different tasks impacts overall model performance and decision-making. Local explanation methods can offer valuable insights into specific instances of successful task adaptation by the Reptile algorithm, which is particularly important in high-stakes domains like healthcare and finance.

In the context of algorithms like R2-D2, LR-D2, and MetaOptNet, integrating local explanation techniques can shed light on how these models arrive at specific predictions. This is especially relevant in scenarios requiring high interpretability, such as medical image analysis. Techniques like **Concept Activation Vectors (CAVs)** are recommended for use in understanding and visualizing the high-level concepts employed by models like MetaOptNet in making predictions. Moreover, exploring knowledge distillation, where simpler 'student' models learn to mimic the complex optimization processes of 'teacher' models, can offer a pathway to comprehend the decision-making process in a more interpretable manner.

However, it is acknowledged that when adapting XAI methods to meta-learning, the unique challenges posed by limited data environments must be considered. Models trained on sparse datasets are noted to potentially exhibit behaviors distinct from those trained on abundant data, necessitating novel approaches in explanation generation. Knowledge distillation techniques are suggested, wherein simpler, interpretable 'student' models are employed to mimic complex 'teacher' meta-learning models, thus offering a pathway to comprehend the decision-making process.

5 APPLICATIONS

In light of the ability of meta-learning to adapt and generalize to new domains with limited data, its application is rendered highly pertinent in real-world scenarios characterized by data scarcity. This distinctive feature of meta-learning within the framework of few-shot learning positions it as a crucial tool in fields where acquiring extensive datasets is challenging, costly, or impractical. Originally, the foundational applications of meta-learning were predominantly centered around the field of computer vision. However, over time, its utility has transcended to encompass a broader spectrum of domains. In a broader view, the applications of meta-learning can now be categorized into three primary areas:

- **Computer vision:** Meta-learning has found significant application in the field of computer vision, particularly due to its effectiveness in addressing few-shot learning challenges. Character recognition (frequently utilizing the Omniglot benchmark dataset) and image classification (MiniImageNet dataset,) have emerged as the two most prevalent areas of focus, as evidenced by the reviewed literature. Meta-learning has been pivotal in advancing remote sensing scene classification which represents another significant domain of image classification where the distinct nuances of aerial and satellite imagery are categorized [2, 60, 63, 104, 127]. The scope of meta-learning in image classification extends beyond objects and scenes, as typically represented in datasets like MiniImageNet and scene classification. It has also been effectively applied in specialized domains, such as medical image classification [48, 53, 72]. Object detection [116, 118, 128] and object recognition [66, 129] have also benefited from meta-learning advancements. Meta-learning emerges as a particularly useful approach for image segmentation when compared to traditional deep learning methods, primarily due to its ability to effectively learn from a limited amount of pixel-wise labeled data [73]. Moving beyond static images, meta-learning can facilitate visual tracking [68], object tracking [50, 111], video classification [41, 130].

- **Language and speech:** In language modeling, the adaptability of meta-learners has been showcased through a variety of tasks. Its applications span a diverse array of tasks including, but not limited to text classification [36, 37, 49], sentiment analysis [65, 85], intent detection [11], machine translation [35, 98]. In speech recognition meta-learning is emerging as a key tool for addressing few-shot adaptation challenges. These include training models for speech recognition of low-resource languages [4, 78, 100], and adapting to various accents [43]. The integration of meta-learning in these areas represents a significant step forward, enabling more personalized and adaptable language and speech processing systems.
- **Reinforcement learning and robotics:** Meta-learning has also been pivotal in **reinforcement learning (RL)**, with various meta-representations being explored. These include adjustments in initialization [22], and hyperparameters [124], facilitating gradient-based learning of neural policies with fewer environmental interactions. Another unique aspect of meta-learning is exploration in RL by focusing on how agents acquire data for learning [14, 64]. Other trends include the successful application of meta-learning on physical robots. This is evident in the role of few-shot learning in tasks like learning robot arm movements from minimal imitation learning [94] and acquiring manipulation actions like quickly adapting to diverse real-world situations like motor failures with minimal guidance [51].

6 CONCLUSION

This study embarked on an exploration of meta-learning, contrasting it with traditional machine and deep learning paradigms. Meta-learning transcends the conventional reliance on extensive data, pivoting towards efficient, task-oriented learning. This approach aligns closely with human cognitive processes, enabling rapid adaptation and learning from limited examples. Such capabilities address the shortcomings of traditional models, particularly in their limited ability to generalize effectively across diverse domains.

In delving into the foundational concepts of meta-learning, the study revealed its efficiency in harnessing prior learning experiences via meticulously reviewing and comparing cutting-edge and transformative techniques in the field. Unlike deep learning models, which learn tasks from scratch often requiring substantial datasets, meta-learning adapts quickly to new scenarios, resonating with intuitive human learning patterns.

The investigation also identified key challenges in meta-learning, notably the uncertainty in performance due to limited samples and varying optimization directions. Innovative approaches, such as Bayesian Active Learning and Variational Autoencoders, are proposed for these challenges, alongside hybrid methods that blend memory-based and metric-based techniques. Despite these advancements, issues like catastrophic forgetting and dataset complexity remain, calling for more stable algorithms and further exploration in multi-domain meta-learning. Integrating Explainable Artificial Intelligence into meta-learning is suggested to enhance model transparency, especially in critical areas like healthcare and finance.

Finally, the study highlighted the applications of meta-learning in fields like computer vision, natural language processing, and robotics, particularly in few-shot learning scenarios. This exploration not only sheds light on the current state of meta-learning but also opens avenues for future research, fostering novel approaches in the dynamic realm of machine learning and artificial intelligence.

REFERENCES

- [1] Mounir Abdelaziz and Zuping Zhang. 2022. Multi-scale kronecker-product relation networks for few-shot learning. *Multimedia Tools and Applications* 81, 5 (2022), 6703–6722.

- [2] Dalal A. Alajaji and Haikel Alhichri. 2020. Few shot scene classification in remote sensing using meta-agnostic machine. In *2020 6th Conference on Data Science and Machine Learning Applications (CDMA)*. IEEE, 77–80.
- [3] Sajid Ali, Tamer Abuhmed, Shaker El-Sappagh, Khan Muhammad, Jose M. Alonso-Moral, Roberto Confalonieri, Riccardo Guidotti, Javier Del Ser, Natalia Diaz-Rodríguez, and Francisco Herrera. 2023. Explainable artificial intelligence (XAI): What we know and what is left to attain trustworthy artificial intelligence. *Information Fusion* 99 (2023), 101805.
- [4] Md Rabiul Awal, Roy Ka-Wei Lee, Eshaan Tanwar, Tanmay Garg, and Tanmoy Chakraborty. 2023. Model-agnostic meta-learning for multilingual hate speech detection. *IEEE Transactions on Computational Social Systems* (2023).
- [5] Sungyong Baik, Seokil Hong, and Kyoung Mu Lee. 2020. Learning to forget for meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2379–2387.
- [6] Pierre Baldi and Yves Chauvin. 1993. Neural networks for fingerprint recognition. *Neural Computation* 5, 3 (1993), 402–418.
- [7] Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. 2017. Online learning rate adaptation with hypergradient descent. *arXiv preprint arXiv:1703.04782* (2017).
- [8] Harkirat Singh Behl, Atilim Güneş Baydin, and Philip H. S. Torr. 2019. Alpha MAML: Adaptive model-agnostic meta-learning. *arXiv preprint arXiv:1905.07435* (2019).
- [9] Luca Bertinetto, Joao F. Henriques, Philip Torr, and Andrea Vedaldi. 2018. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*.
- [10] Luca Bertinetto, João F. Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. 2016. Learning feed-forward one-shot learners. *Advances in Neural Information Processing Systems* 29 (2016).
- [11] Hemanthage S. Bhatthiya and Uthayasanker Thayasivam. 2020. Meta learning for few-shot joint intent detection and slot-filling. In *Proceedings of the 2020 5th International Conference on Machine Learning Technologies*. 86–92.
- [12] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. Signature verification using a “Siamese” time delay neural network. *Advances in Neural Information Processing Systems* 6 (1993).
- [13] Mengting Chen, Xinggang Wang, Heng Luo, Yifeng Geng, and Wenyu Liu. 2021. Learning to focus: Cascaded feature matching network for few-shot image recognition. *Science China Information Sciences* 64, 9 (2021), 1–13.
- [14] Xi Chen, Ali Ghadirzadeh, Márten Björkman, and Patric Jensfelt. 2019. Meta-learning for multi-objective reinforcement learning. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 977–983.
- [15] Davide Chicco. 2021. Siamese neural networks: An overview. *Artificial Neural Networks* (2021), 73–94.
- [16] Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, Vol. 1. IEEE, 539–546.
- [17] Jessica Deuschel, Daniel Firmbach, Carol I. Geppert, Markus Eckstein, Arndt Hartmann, Volker Bruns, Petr Kuritcyn, Jakob Dexe, David Hartmann, Dominik Perrin, Thomas Wittenberg, and Michaela Benz. 2021. Multi-prototype few-shot learning in histopathology. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 620–628.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [19] Yueming Ding, Xia Tian, Lirong Yin, Xiaobing Chen, Shan Liu, Bo Yang, and Wenfeng Zheng. 2019. Multi-scale relation network for few-shot learning based on meta-learning. In *International Conference on Computer Vision Systems*. Springer, 343–352.
- [20] Harrison Edwards and Amos Storkey. 2016. Towards a neural statistician. *arXiv preprint arXiv:1606.02185* (2016).
- [21] Abolfazl Farahani, Sahar Voghoei, Khaled Rasheed, and Hamid R. Arabnia. 2021. A brief review of domain adaptation. *Advances in Data Science and Information Engineering* (2021), 877–894.
- [22] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*. PMLR, 1126–1135.
- [23] Chelsea Finn and Sergey Levine. 2018. Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm. In *International Conference on Learning Representations*.
- [24] Chelsea Finn, Kelvin Xu, and Sergey Levine. 2018. Probabilistic model-agnostic meta-learning. *Advances in Neural Information Processing Systems* 31 (2018).
- [25] Chelsea B. Finn. 2018. *Learning to Learn with Gradients*. University of California, Berkeley.
- [26] Stanislav Fort. 2017. Gaussian prototypical networks for few-shot learning on Omniglot. *arXiv preprint arXiv:1708.02735* (2017).
- [27] Victor Garcia and Joan Bruna. 2017. Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043* (2017).
- [28] Victor Garcia and Joan Bruna. 2018. Few-shot learning with graph neural networks. In *6th International Conference on Learning Representations, ICLR 2018*.

- [29] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and S. M. Ali Eslami. 2018. Conditional neural processes. In *International Conference on Machine Learning*. PMLR, 1704–1713.
- [30] Spyros Gidaris and Nikos Komodakis. 2018. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4367–4375.
- [31] Spyros Gidaris and Nikos Komodakis. 2019. Generating classification weights with GNN denoising autoencoders for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21–30.
- [32] Micah Goldblum, Liam Fowl, and Tom Goldstein. 2020. Adversarially robust few-shot learning: A meta-learning approach. *Advances in Neural Information Processing Systems* 33 (2020), 17886–17895.
- [33] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. 2018. Recasting gradient-based meta-learning as hierarchical Bayes. *arXiv preprint arXiv:1801.08930* (2018).
- [34] Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural Turing machines. *arXiv preprint arXiv:1410.5401* (2014).
- [35] Jiatao Gu, Yong Wang, Yun Chen, Kyunghyun Cho, and Victor OK Li. 2018. Meta-learning for low-resource neural machine translation. *arXiv preprint arXiv:1808.08437* (2018).
- [36] Chengcheng Han, Zeqiu Fan, Dongxiang Zhang, Minghui Qiu, Ming Gao, and Aoying Zhou. 2021. Meta-learning adversarial domain adaptation network for few-shot text classification. *arXiv preprint arXiv:2107.12262* (2021).
- [37] Chengcheng Han, Yuhe Wang, Yingnan Fu, Xiang Li, Minghui Qiu, Ming Gao, and Aoying Zhou. 2023. Meta-learning Siamese network for few-shot text classification. In *International Conference on Database Systems for Advanced Applications*. Springer, 737–752.
- [38] Mengya Han, Ronggui Wang, Juan Yang, Lixia Xue, and Min Hu. 2020. Multi-scale feature network for few-shot learning. *Multimedia Tools and Applications* 79, 17 (2020), 11617–11637.
- [39] Jun He, Richang Hong, Xueliang Liu, Mingliang Xu, Zheng-Jun Zha, and Meng Wang. 2020. Memory-augmented relation network for few-shot learning. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1236–1244.
- [40] Nathan Hilliard, Lawrence Phillips, Scott Howland, Artëm Yankov, Courtney D. Corley, and Nathan O. Hodas. 2018. Few-shot learning with metric-agnostic conditional embeddings. *arXiv preprint arXiv:1802.04376* (2018).
- [41] Yufan Hu, Junyu Gao, and Changsheng Xu. 2020. Learning dual-pooling graph neural networks for few-shot video classification. *IEEE Transactions on Multimedia* 23 (2020), 4285–4296.
- [42] Hongwei Huang, Zhangkai Wu, Wenbin Li, Jing Huo, and Yang Gao. 2021. Local descriptor-based multi-prototype network for few-shot learning. *Pattern Recognition* 116 (2021), 107935.
- [43] Kuan-Po Huang, Yuan-Kuei Wu, and Hung-yi Lee. 2021. Multi-accent speech separation with one shot learning. *arXiv preprint arXiv:2106.11713* (2021).
- [44] Mike Huisman, Jan N. van Rijn, and Aske Plaat. 2021. A survey of deep meta-learning. *Artificial Intelligence Review* 54, 6 (2021), 4483–4541.
- [45] Olga Ilina, Vadim Ziyadinov, Nikolay Klenov, and Maxim Tereshonok. 2022. A survey on symmetrical neural network architectures and applications. *Symmetry* 14, 7 (2022), 1391.
- [46] Muhammad Abdullah Jamal and Guo-Jun Qi. 2019. Task agnostic meta-learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11719–11727.
- [47] Zhong Ji, Xingliang Chai, Yunlong Yu, Yanwei Pang, and Zhongfei Zhang. 2020. Improved prototypical networks for few-shot learning. *Pattern Recognition Letters* 140 (2020), 81–87.
- [48] Hongyang Jiang, Mengdi Gao, Heng Li, Richu Jin, Hanpei Miao, and Jiang Liu. 2022. Multi-learner based deep meta-learning for few-shot medical image classification. *IEEE Journal of Biomedical and Health Informatics* 27, 1 (2022), 17–28.
- [49] Xiang Jiang, Mohammad Havaei, Gabriel Chartrand, Hassan Chouaib, Thomas Vincent, Andrew Jesson, Nicolas Chapados, and Stan Matwin. 2018. Attentive task-agnostic meta-learning for few-shot text classification. In *2nd Workshop on Meta-Learning at Advances in Neural Information Processing Systems* (2018).
- [50] Ilchae Jung, Kihyun You, Hyeonwoo Noh, Minsu Cho, and Bohyung Han. 2020. Real-time object tracking via meta-learning: Efficient model adaptation and one-shot channel pruning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 11205–11212.
- [51] Rituraj Kaushik, Timothée Anne, and Jean-Baptiste Mouret. 2020. Fast online adaptation in robotics through meta-learning embeddings of simulated priors. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 5269–5276.
- [52] Kian Kenyon-Dean, Andre Cianflone, Lucas Page-Caccia, Guillaume Rabusseau, Jackie Chi Kit Cheung, and Doina Precup. 2018. Clustering-oriented representation learning with attractive-repulsive loss. *arXiv preprint arXiv:1812.07627* (2018).

- [53] Rabindra Khadka, Debesh Jha, Steven Hicks, Vajira Thambawita, Michael A. Riegler, Sharib Ali, and Pål Halvorsen. 2022. Meta-learning with implicit gradients in a few-shot setting for medical image segmentation. *Computers in Biology and Medicine* 143 (2022), 105227.
- [54] Byeonggeun Kim, Seunghan Yang, Inseop Chung, and Simyung Chang. 2022. Dummy prototypical networks for few-shot open-set keyword spotting. *arXiv preprint arXiv:2206.13691* (2022).
- [55] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, Vol. 2. Lille, 0.
- [56] Steinar Laenen and Luca Bertinetto. 2021. On episodes, prototypical networks, and few-shot learning. *Advances in Neural Information Processing Systems* 34 (2021), 24581–24592.
- [57] Brenden Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua Tenenbaum. 2011. One shot learning of simple visual concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 33.
- [58] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. 2019. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10657–10665.
- [59] Aoxue Li, Tiange Luo, Tao Xiang, Weiran Huang, and Liwei Wang. 2019. Few-shot learning with global class representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9715–9724.
- [60] Haifeng Li, Zhenqi Cui, Zhiqing Zhu, Li Chen, Jiawei Zhu, Haozhe Huang, and Chao Tao. 2020. RS-MetaNet: Deep meta metric learning for few-shot remote sensing scene classification. *arXiv preprint arXiv:2009.13364* (2020).
- [61] Huaiyu Li, Weiming Dong, Xing Mei, Chongyang Ma, Feiyue Huang, and Bao-Gang Hu. 2019. LGM-Net: Learning to generate matching networks for few-shot learning. In *International Conference on Machine Learning*. PMLR, 3825–3834.
- [62] Xiaoxu Li, Zhuo Sun, Jing-Hao Xue, and Zhanyu Ma. 2021. A concise review of recent few-shot meta-learning methods. *Neurocomputing* 456 (2021), 463–468.
- [63] Yong Li, Zhenfeng Shao, Xiao Huang, Bowen Cai, and Song Peng. 2021. Meta-FSEO: A meta-learning fast adaptation with self-supervised embedding optimization for few-shot remote sensing scene classification. *Remote Sensing* 13, 14 (2021), 2776.
- [64] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-SGD: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835* (2017).
- [65] Bin Liang, Xiang Li, Lin Gui, Yonghao Fu, Yulan He, Min Yang, and Ruifeng Xu. 2023. Few-shot aspect category sentiment analysis via meta-learning. *ACM Transactions on Information Systems* 41, 1 (2023), 1–31.
- [66] Bo Liu, Hao Kang, Haoxiang Li, Gang Hua, and Nuno Vasconcelos. 2020. Few-shot open-set recognition using meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8798–8807.
- [67] Han Liu, Feng Zhang, Xiaotong Zhang, Siyang Zhao, Junjie Sun, Hong Yu, and Xianchao Zhang. 2022. Label-enhanced prototypical network with contrastive learning for multi-label few-shot aspect category detection. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1079–1087.
- [68] Shuai Liu, Shichen Huang, Shuai Wang, Khan Muhammad, Paolo Bellavista, and Javier Del Ser. 2023. Visual tracking in complex scenes: A location fusion mechanism based on the combination of multiple visual cognition flows. *Information Fusion* 96 (2023), 281–296.
- [69] Xiaoqian Liu, Fengyu Zhou, Jin Liu, and Lianjie Jiang. 2020. Meta-learning based prototype-relation network for few-shot classification. *Neurocomputing* 383 (2020), 224–234.
- [70] Jiang Lu, Zhong Cao, Kailun Wu, Gang Zhang, and Changshui Zhang. 2018. Boosting few-shot image recognition via domain alignment prototypical networks. In *2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 260–264.
- [71] Jiang Lu, Pinghua Gong, Jieping Ye, and Changshui Zhang. 2020. Learning from very few samples: A survey. *arXiv preprint arXiv:2009.02653* (2020).
- [72] Liangfu Lu, Xudong Cui, Zhiyuan Tan, and Yulei Wu. 2023. MedOptNet: Meta-learning framework for few-shot medical image classification. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* (2023).
- [73] Shuai Luo, Yujie Li, Pengxiang Gao, Yichuan Wang, and Seiichi Serikawa. 2022. Meta-Seg: A survey of meta-learning for image segmentation. *Pattern Recognition* 126 (2022), 108586.
- [74] Sijie Mai, Haifeng Hu, and Jia Xu. 2019. Attentive matching network for few-shot learning. *Computer Vision and Image Understanding* 187 (2019), 102781.
- [75] Iaroslav Melekhov, Juho Kannala, and Esa Rahtu. 2016. Siamese network features for image matching. In *2016 23rd International Conference on Pattern Recognition (ICPR)*. IEEE, 378–383.
- [76] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2017. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141* (2017).
- [77] Tom M. Mitchell. 1997. *Machine Learning*. Vol. 1. McGraw-Hill New York.

- [78] Marzieh Mozafari, Reza Farahbakhsh, and Noel Crespi. 2022. Cross-lingual few-shot hate speech and offensive language detection using meta learning. *IEEE Access* 10 (2022), 14880–14896.
- [79] Tsendsuren Munkhdalai and Hong Yu. 2017. Meta networks. In *International Conference on Machine Learning*. PMLR, 2554–2563.
- [80] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018).
- [81] Jaehoon Oh, Hyungjun Yoo, ChangHwan Kim, and Se-Young Yun. 2020. BOIL: Towards representation change for few-shot learning. *arXiv preprint arXiv:2008.08882* (2020).
- [82] Frederik Pahde, Mihai Puscas, Tassilo Klein, and Moin Nabi. 2021. Multimodal prototypical networks for few-shot learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 2644–2653.
- [83] Sinno Jialin Pan and Qiang Yang. 2009. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* 22, 10 (2009), 1345–1359.
- [84] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. 2019. Transferrable prototypical networks for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2239–2247.
- [85] Sumit Pandey and Srishti Sharma. 2023. Meta-learned word embeddings for few-shot sentiment classification. In *International Conference on Smart Trends in Computing and Communications*. Springer, 577–589.
- [86] Hang Qi, Matthew Brown, and David G. Lowe. 2018. Low-shot learning with imprinted weights. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5822–5830.
- [87] Aravind Rajeswaran, Chelsea Finn, Sham M. Kakade, and Sergey Levine. 2019. Meta-learning with implicit gradients. *Advances in Neural Information Processing Systems* 32 (2019).
- [88] Sachin Ravi and Hugo Larochelle. 2016. Optimization as a model for few-shot learning. *International Conference on Learning Representations* (2016).
- [89] Li Ren, Guiduo Duan, Tianxi Huang, and Zhao Kang. 2022. Multi-local feature relation network for few-shot learning. *Neural Computing and Applications* 34, 10 (2022), 7393–7403.
- [90] Mengye Ren, Renjie Liao, Ethan Fetaya, and Richard Zemel. 2019. Incremental few-shot learning with attention attractor networks. *Advances in Neural Information Processing Systems* 32 (2019).
- [91] Mengye Ren, Eleni Triantafyllou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. 2018. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676* (2018).
- [92] Sebastian Ruder, Matthew E. Peters, Swabha Swayamdipta, and Thomas Wolf. 2019. Transfer learning in natural language processing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*. 15–18.
- [93] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. 2018. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960* (2018).
- [94] Alok Ranjan Sahoo and Pavan Chakraborty. 2022. A study on position control of a continuum arm using MAML (model-agnostic meta-learning) for adapting different loading conditions. *IEEE Access* 10 (2022), 14980–14992.
- [95] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *International Conference on Machine Learning*. PMLR, 1842–1850.
- [96] Adam Santoro, David Raposo, David G. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. *Advances in Neural Information Processing Systems* 30 (2017).
- [97] Manali Shaha and Meenakshi Pawar. 2018. Transfer learning for image classification. In *2018 Second International Conference on Electronics, Communication and Aerospace Technology (ICECA)*. IEEE, 656–660.
- [98] Amr Sharaf, Hany Hassan, and Hal Daumé III. 2020. Meta-learning for few-shot NMT adaptation. *arXiv preprint arXiv:2004.02745* (2020).
- [99] Pranav Shyam, Shubham Gupta, and Ambedkar Dukkipati. 2017. Attentive recurrent comparators. In *International Conference on Machine Learning*. PMLR, 3173–3181.
- [100] Satwinder Singh, Ruili Wang, and Feng Hou. 2022. Improved meta learning for low resource speech recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4798–4802.
- [101] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems* 30 (2017).
- [102] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. 2019. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 403–412.
- [103] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1199–1208.

- [104] Xu Tang, Weiquan Lin, Chao Liu, Xiao Han, Wenjing Wang, Jingjing Ma, and Licheng Jiao. 2021. Multi-scale meta-learning-based networks for high-resolution remote sensing scene classification. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 4928–4931.
- [105] Sebastian Thrun and Loren Pratt. 2012. *Learning to Learn*. Springer Science & Business Media.
- [106] Kien Tran, Hiroshi Sato, and Masao Kubo. 2019. Memory augmented matching networks for few-shot learnings. *International Journal of Machine Learning and Computing* 9, 6 (2019).
- [107] Eleni Triantafyllou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre Antoine Manzagol, and Hugo Larochelle. 2019. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096* (2019).
- [108] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, and Koray Kavukcuoglu. 2016. Matching networks for one shot learning. *Advances in Neural Information Processing Systems* 29 (2016).
- [109] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The Caltech-UCSD Birds-200-2011 dataset. (2011).
- [110] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. 2020. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726* (2020).
- [111] Guangting Wang, Chong Luo, Xiaoyan Sun, Zhiwei Xiong, and Wenjun Zeng. 2020. Tracking by instance detection: A meta-learning approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6288–6297.
- [112] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip Yu. 2022. Generalizing to unseen domains: A survey on domain generalization. *IEEE Transactions on Knowledge and Data Engineering* (2022).
- [113] Jane X. Wang. 2021. Meta-learning in natural and artificial intelligence. *Current Opinion in Behavioral Sciences* 38 (2021), 90–95.
- [114] Mei Wang and Weihong Deng. 2018. Deep visual domain adaptation: A survey. *Neurocomputing* 312 (2018), 135–153.
- [115] Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E. Gonzalez. 2019. TAFE-Net: Task-aware feature embeddings for low shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1831–1840.
- [116] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. 2019. Meta-learning to detect rare objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9925–9934.
- [117] Xiu-Shen Wei, Jian-Hao Luo, Jianxin Wu, and Zhi-Hua Zhou. 2017. Selective convolutional descriptor aggregation for fine-grained image retrieval. *IEEE Transactions on Image Processing* 26, 6 (2017), 2868–2881.
- [118] Xiongwei Wu, Doyen Sahoo, and Steven Hoi. 2020. Meta-RCNN: Meta learning for few-shot object detection. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1679–1687.
- [119] Zhiyu Xue, Lixin Duan, Wen Li, Lin Chen, and Jiebo Luo. 2020. Region comparison network for interpretable few-shot image classification. *arXiv preprint arXiv:2009.03558* (2020).
- [120] Zhiyu Xue, Zhenshan Xie, Zheng Xing, and Lixin Duan. 2020. Relative position and map networks in few-shot learning for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 932–933.
- [121] Chengxiang Yin, Jian Tang, Zhiyuan Xu, and Yanzhi Wang. 2018. Adversarial meta-learning. *arXiv preprint arXiv:1806.03316* (2018).
- [122] Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. 2018. Bayesian model-agnostic meta-learning. *Advances in Neural Information Processing Systems* 31 (2018).
- [123] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How transferable are features in deep neural networks? *Advances in Neural Information Processing Systems* 27 (2014).
- [124] Kenny Young, Baoxiang Wang, and Matthew E. Taylor. 2018. Metatrace actor-critic: Online step-size tuning by meta-gradient descent for reinforcement learning control. *arXiv preprint arXiv:1805.04514* (2018).
- [125] Zhou Yu, Jun Yu, Chenchao Xiang, Jianping Fan, and Dacheng Tao. 2018. Beyond bilinear: Generalized multimodal factorized high-order pooling for visual question answering. *IEEE Transactions on Neural Networks and Learning Systems* 29, 12 (2018), 5947–5959.
- [126] Lingling Zhang, Jun Liu, Minnan Luo, Xiaojun Chang, Qinghua Zheng, and Alexander G. Hauptmann. 2019. Scheduled sampling for one-shot learning via matching network. *Pattern Recognition* 96 (2019), 106962.
- [127] Pei Zhang, Yunpeng Bai, Dong Wang, Bendu Bai, and Ying Li. 2020. Few-shot classification of aerial scene images via meta-learning. *Remote Sensing* 13, 1 (2020), 108.
- [128] Pei Zhang, Yunpeng Bai, Dong Wang, Bendu Bai, and Ying Li. 2021. A meta-learning framework for few-shot classification of remote sensing scene. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4590–4594.

- [129] Yaohui Zhu, Chenlong Liu, and Shuqiang Jiang. 2020. Multi-attention meta learning for few-shot fine-grained image recognition. In *IJCAI*. 1090–1096.
- [130] Zhenxi Zhu, Limin Wang, Sheng Guo, and Gangshan Wu. 2021. A closer look at few-shot video classification: A new baseline and benchmark. *arXiv preprint arXiv:2110.12358* (2021).
- [131] Luisa M. Zintgraf, Kyriacos Shiarlis, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. 2018. Caml: Fast context adaptation via meta-learning. *arXiv preprint arXiv:1810.03642* (2018).

Received 27 January 2023; revised 12 December 2023; accepted 2 April 2024