

EEML: Ensemble Embedded Meta-learning^{*}

Geng Li¹[0000-0002-9733-6891], Boyuan Ren²[0000-0002-1291-5377], and Hongzhi Wang³[0000-0002-7521-2871]

¹ Harbin Institute of Technology, Harbin, CHINA 21S003022@stu.hit.edu.cn

² Harbin Institute of Technology, Harbin, CHINA 1180300207@stu.hit.edu.cn

³ Harbin Institute of Technology, Harbin, CHINA wangzh@hit.edu.cn

Abstract. To accelerate learning process with few samples, meta-learning resorts to prior knowledge from previous tasks. However, the inconsistent task distribution and heterogeneity is hard to be handled through a global sharing model initialization. In this paper, based on gradient-based meta-learning, we propose an ensemble embedded meta-learning algorithm (EEML) that explicitly utilizes multi-model-ensemble to organize prior knowledge into diverse specific experts. We rely on a task embedding cluster mechanism to deliver diverse tasks to matching experts in training process and instruct how experts collaborate in test phase. As a result, the multi experts can focus on their own area of expertise and cooperate in upcoming task to solve the task heterogeneity. The experimental results show that the proposed method outperforms recent state-of-the-arts easily in few-shot learning problem, which validates the importance of differentiation and cooperation.

Keywords: Meta-learning · Ensemble-learning · Few-shot Learning.

1 Introduction

One of the Intelligent advantages of human being is acquiring new skill for unseen task from few samples. To achieve the fast learning capability, meta-learning also known as learning to learn, leverages what we learned from past tasks to improve learning in a new task. As a common practice for few-shot learning problem, meta-learning contains several series which differ in the way of employing prior knowledge [12]: metric-based [13, 29], model-based [19], gradient-based(a.k.a. optimization-based) meta-learning [8].

Though these methods earn relative improvement comparing with past methods, most of them hypothesize the prior knowledge learned from past should be shared indifferently for all tasks. As a consequence, they suffer from the tasks heterogeneity. On the other hand, a few research works try to fix the problem by customizing the prior knowledge to each task [1, 25, 35]. However, the defect of such methods lie in resorting to only a single expert who may not be sufficient for achieving the best result. In other words, instead of reasonably utilizing overall intelligence, these methods prefer to choose an expert as dictator.

^{*} Supported by Harbin Institute of Technology.

Hence, we are motivated to pursue a meta-learning framework both effectively diverge knowledge into multi experts and comprehensively utilizing over-all experts to aggregate an stable answer for new task to enhance the final performance. The inspiration comes from classical ensemble learning, bagging [4] or random forest [5]. In ensemble learning, though single learner may perform trivially or sometimes blindly on the whole task distribution, but the ensemble mechanism can help each expert reach a relatively good cooperative result through counteracting mutual bias.

Inspired by this, we propose a novel meta-learning algorithm named Ensemble Embedded Meta-learning(EEML). The key idea of EEML is to organize heterogeneous training tasks into matching experts and collecting all experts opinions to form a comprehensive answer for a new task. In this paper, we resorts to gradient-based meta-learning as the backbone which instantiates knowledge as parameter initialization. Specifically, we employ a task embedding cluster to cover the domain of each expert. The gradient of task would be considered as a strong embedding and we assign training task to the matching expert according to the similarity between the gradient and the expert domain. In predicting phase, all experts can utilize few samples to contribute its opinions, and the final result is composed as a weighted voting answer from all experts.

We would highlight the contributions of EEML as two points: 1) it employs multi-experts cooperation framework instead of single expert for uncertain tasks, so that it outperforms recent state-of-the-art meta-learning algorithms in few-shot learning. 2) as we know, EEML is the first generic meta-learning algorithm combined with ensemble learning, whose experiment results show great confidence in exploring further in composing meta-learning and ensemble learning for future work.

2 Related Work

2.1 Meta-learning

Meta-learning is understood as “learning to learn” [28, 32, 31, 33, 17], it starts with the form of learning to optimize model [2, 11] and now are mainly divided into three series, metric-based meta-learning [13, 29, 30, 34, 16], network-based meta-learning [19, 21, 23, 26, 20] and gradient-based meta-learning [36, 9, 10, 6, 15, 18]. The goal of metric-based algorithms is to measure the similarity between support samples and query samples with elaborate metric methodology. Network-based approaches emphasize finding structures effective in capturing the prior knowledge in training data. Gradient-based meta-learning aims to use different task gradients to optimize a base model with a batch form.

In this work, we focus on the discuss about gradient-based meta-learning. The most powerful representation should be MAML [8], which presents the “inner-loop and outer-loop” framework commonly used by plenty of followers. Reptile [22] uses accumulated first gradient replace the second derivative and earns speeding up but instability. MetaOptNet [14] and [29, 3] reset the inner optimization by dividing the model into feature extractor and head classifier and apply

on complex base learner like ResNet-12 without overfitting. LEO [25] introduces lower-dimensional latent space to consider task-conditioned initialization. L2F [1] uses extra MLP network to help attenuate the base learner parameters to gain a task-conditioned initialization.

2.2 Ensemble Learning

Ensemble learning exploits multiple base machine learning algorithms to obtain weak predictive results and fuse results with voting mechanism to achieve better performance than that obtained from any constituent algorithm alone [7]. The main classical ensemble learning algorithms are AdaBoost [27], Random Forest [5] and Bagging [4]. In this paper, to start with clarity we mainly rely on the bagging mechanism which is simple and efficient, besides, it's naturally adaptive to parallel training paradigm, which can effectively speedup our training process. However, with the precisely designed coefficients adopted during training and test phase, our method has inherent difference with original bagging.

3 Preliminaries

3.1 The Problem Definition

Supposed that a sequence of tasks $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N\}$ are sampled from a probability $p(\mathcal{T})$. Each task sampled from $p(\mathcal{T})$ is defined as $\{x_i, y_i\}_{i=1}^m$. Dividing parts samples into training set $\mathcal{D}^{tr} = \{x_i, y_i\}_{i=1}^{n^{tr}}$ and the rest as test set $\mathcal{D}^{te} = \{x_i, y_i\}_{i=n^{tr}+1}^m$, we then can also describe a task \mathcal{T} as $\{\mathcal{D}^{tr}, \mathcal{D}^{te}\}$. Given a model f_θ , the training aim is to obtain the optimal parameters θ to minimize the loss function:

$$Loss(f_\theta(\mathcal{D}^{tr}, \{x_i\}_{i=n^{tr}+1}^m, \{y_i\}_{i=n^{tr}+1}^m))$$

which can be simply noted as $\mathcal{L}_{\mathcal{T}}^{\mathcal{D}^{te}}(f_\theta)$.

Based on this, the meta-learning considers model f as a learning process, which earn θ' from θ and give final prediction based on θ' . To find a globally shared parameters θ_g , we follow:

$$\min_{\theta_g} \sum_{k=1}^N Loss(f_{\theta_g}(\mathcal{D}_k^{tr}, \{x_{k,i}\}_{i=n^{tr}+1}^m, \{y_{k,i}\}_{i=n^{tr}+1}^m)) \quad (1)$$

Different meta-learning algorithms mainly diverge in f , which is also utilized in test phase to map θ_g to θ'_g for new task t_i and gain predictions corresponding.

3.2 The Gradient-based Meta-learning

Here, we compactly give an overview of the representative algorithm of gradient-based meta-learning, i.e. model-agnostic meta-learning. With the aim of searching for a globally sharing parameters initialization for unseen tasks, MAML define

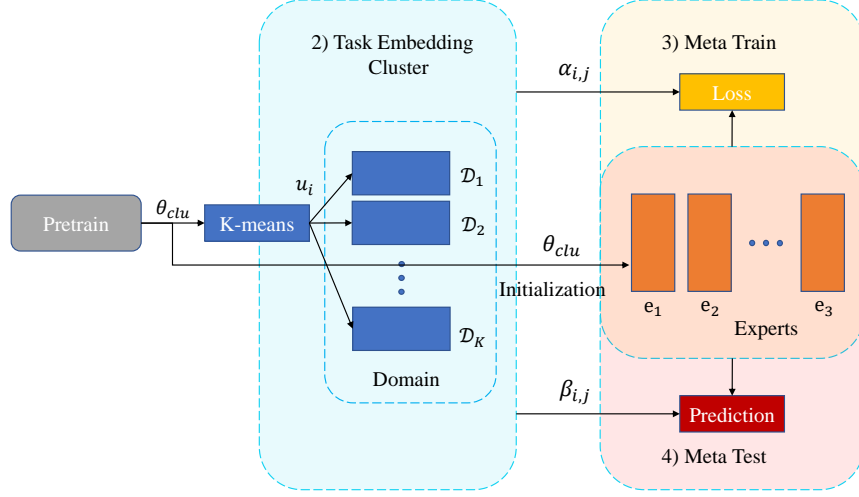


Fig. 1. The framework of EEML mainly concerns three stages: a) Task Embedding Cluster: We employ the θ_{clu} to gain task gradient embedding u_i , then leverage K-means to get Domains. b) Meta Train: via calculating coefficients $\alpha_{i,j}$, we train experts towards costumed loss function. c) Meta Test: Given coefficients $\beta_{i,j}$, we gain cooperative results through weighted voting mechanism.

learning process f like below. Formally, given the task distribution $p(\mathcal{T})$, we sample a task \mathcal{T}_i from it and define the base learner as g_θ then we can get:

$$\theta'_i = \theta - \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}^{\mathcal{D}_i^{tr}}(g_\theta) \quad (2)$$

$$f_\theta(\mathcal{D}_i^{tr}, \{x_i\}_{i=n^{tr}+1}^m) = g_{\theta'_i}(\{x_i\}_{i=n^{tr}+1}^m) \quad (3)$$

To update θ , MAML then follows:

$$\theta \leftarrow \theta - \nabla_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}^{\mathcal{D}_i^{te}}(f_{\theta_i}) \quad (4)$$

For simplicity of description, We define gradient of task \mathcal{T}_i , $\nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}^{\mathcal{D}_i^{te}}(f_{\theta_i})$, as u_i in following sections.

4 Methodology

In this section, we detail the proposed EEML algorithm like the Fig. 1. The EEML aims to raise a group of experts covering heterogeneous knowledge domains in meta-training, then collect all experts opinions in a weighted voting schema to achieve better results in test phase. The following parts of this section are organized in task embedding cluster, meta-training and meta-test.

4.1 Task Embedding Cluster

Instead of abstracting task embedding from only labeled data, we argue that gradient u_i calculated in (4) can consider dynamical training process which implies optimization direction information ignored by static task embedding mechanism like auto-encoder or rnn-based encoder. With gradient u_i considered as task \mathcal{T}_i embedding, we resorts to a classical cluster algorithm, K-means to divide sampled tasks into different groups. The distance function leveraged in K-means is cosine distance which focuses on the difference in direction instead of embedding norm. Given hyperparameter K , we then obtain K clusters which correspond to different optimization direction for specific model parameters θ_{clu} . We collect the centers of different clusters as $\{c_1, c_2, \dots, c_K\}$. To measure the distance between future coming task and existing clusters, we follow the cosine distance as the standard metric.

4.2 Meta Train

Unlike training single dictator model, raising a group of experts should carefully maintain the balance between diversity and generalization. To start with a generic initialization, we initialize the experts $\{g_{e_1}, g_{e_2}, \dots, g_{e_K}\}$ as K copies from $g_{\theta_{clu}}$, whose generalization then can be controlled through parameter initialization θ_{clu} . To get an effective initialization without introducing too many operations, we simply exploit MAML to earn a relatively proper θ_{clu} as the initialization parameters for clustering and experts. We can adjust the pretraining via limiting the total training epochs.

In training stage, to strengthen the diversity of experts in training paradigm, we utilize the loss function below as the target training loss.

$$\mathcal{Loss} = \sum_{i=1}^N \sum_{j=1}^K \alpha_{i,j} \mathcal{L}_{\mathcal{T}_i}^{\mathcal{D}_i^{te}}(g_{e_j}) \quad (5)$$

The coefficient $\alpha_{i,j}$ would induce the final tendency of expert e_i towards specific tasks. In our assumption, a qualified expert e_i should be familiar with a task from a specific domain D_i while has no strict demanding about the competence for other tasks outside. We define the D_i as $\{\mathcal{T}_j | \text{Dist}(\mathcal{T}_j, c_i) = \min_{p=1:K} \text{Dist}(\mathcal{T}_j, c_p)\}$, which indicates a task \mathcal{T}_j belonging to D_i should be nearest to the domain center c_i comparing with other domain centers. To expand the generalization, we do not assign coefficient $\alpha_{i,j}$ from the indicator of whether task \mathcal{T}_i is in D_j , instead, we choose the cosine similarity metric after softmax layer as final $\alpha_{i,j}$ to ensure the probability sum as one.

In this paper, instead of considering static task embedding without any information from model, we decide to leverage the gradient u_i as the embedding of task \mathcal{T}_i . With both model optimization information and labelled data contained, the u_i help EEML achieve better performance.

$$\begin{aligned} \alpha &= \text{Softmax}(\text{Sim}(\mathcal{T}_j, \mathbf{c})) \\ &= \text{Softmax}\left(\frac{u_i \cdot \mathbf{c}}{\|u_i\| \|\mathbf{c}\|}\right) \end{aligned} \quad (6)$$

4.3 Meta Test

Facing unseen task from heterogeneity distribution, single dictator model is insufficient and easy to be blind. In our paradigm, all experts are able to join the test phase through two steps. First, the experts fine-tune the parameters with training set \mathcal{D}_i^{tr} from \mathcal{T}_{test} .

$$e'_i = e_i - \nabla_{\theta} \mathcal{L}_{\mathcal{T}_{test}}^{\mathcal{D}_i^{tr}}(e_i) \quad (7)$$

With fine-tuned experts $\mathbf{e}' = e'_1, e'_2, \dots, e'_n$. We leverage a weighted voting mechanism inspired by Bagging like:

$$p_j = \sum_{i=1}^K \beta_{i,j} g_{e'_i}(\mathcal{T}_j) \quad (8)$$

The coefficient $\beta_{i,j}$ is designed as:

$$\beta_{i,j} = \frac{Sim(\mathcal{T}_j, c_i)}{err_{i,j}} \quad (9)$$

The $err_{i,j}$ is the error value of training set \mathcal{D}_j^{tr} with model g_{e_i} , which implies the adaptation level of expert e_i . Like $\alpha_{i,j}$ we also take the softmax function as filter to maintain probability property.

$$err_j = Softmax(\mathcal{L}_{\mathcal{T}_j}^{\mathcal{D}_j^{tr}}(g_{\mathbf{e}})) \quad (10)$$

$\mathcal{L}_{\mathcal{T}_j}^{\mathcal{D}_j^{tr}}(g_{\mathbf{e}})$ is a vector composed of $\{\mathcal{L}_{\mathcal{T}_j}^{\mathcal{D}_j^{tr}}(g_{e_1}), \mathcal{L}_{\mathcal{T}_j}^{\mathcal{D}_j^{tr}}(g_{e_2}), \dots, \mathcal{L}_{\mathcal{T}_j}^{\mathcal{D}_j^{tr}}(g_{e_K})\}$. In our paradigm, the final prediction of all experts would be considered based on the expert similarity with the task and the final performance. And the dictator mechanism which consists of only one final model can be considered as a specific instance when $K = 1$.

5 Experiments

In this section, we evaluate the effectiveness of our proposed EEML on both toy regression and complex image few-shot classification cases. This section is organized as Datasets and Implementation Details, Toy Regression, and Image Few-shot Classification.

5.1 Datasets and Implementation Details

Datasets for regression and classification For toy regression, to compare with [35, 8]. We follow the task design from four function families. They are (1) Sinusoids (2) Line (3) Cubic (4) Quadratic. See Table 1 for details of function definition. We also take MSE as the final evaluation metric. As for image few-shot classification, we use the widely accepted dataset, i.e. MiniImageNet, as benchmark to validate the effectiveness. MiniImageNet is a larger benchmark containing 100 classes randomly chosen from ImageNet ILSVRC-2012 challenge [24] with 600 images of size 84×84 pixels per class. It is split into 64 base classes, 16 validation classes and 20 novel classes.

Table 1. Toy regression function family definition details.

function family	formulation ($y(x)$)	parameters range			
		p_1	p_2	p_3	p_4
Sinusoids	$p_1 \sin(p_2 x + p_3)$	$[0.1, 5.0]$	$[0.8, 1.2]$	$[0, 2\pi]$	—
Line	$p_1 x + p_2$	$[-3.0, 3.0]$	$[-3.0, 3.0]$	—	—
Quadratic	$p_1 x^2 + p_2 x + p_3$	$[-0.2, 0.2]$	$[-2.0, 2.0]$	$[-3.0, 3.0]$	—
Cubic	$p_1 x^3 + p_2 x^2 + p_3 x + p_4$	$[-0.1, 0.1]$	$[-0.2, 0.2]$	$[-2.0, 2.0]$	$[-3.0, 3.0]$

Models for regression and classification Following [8, 35], we take 2 hidden layer mlp and 4-conv network as backbones for toy regression and image classification respectively. The related hyperparameters are listed in Table 2.

Table 2. Hyperparameters of the experiments.

Hyperparameter	Toy Regression	MiniImageNet
task batch size	32	32
Inner loop learning rate	0.001	0.03
Outer loop learning rate	0.001	0.001
Inner loop adaptation steps	5	5
epoch for pretraining	15000	20000
epoch for meta-train	5000	10000
K selected for k-means	4	3

Table 3. Performance of MSE \pm 95% confidence intervals on toy regression tasks, averaged over 4,000 tasks. Both 5-shot and 10-shot results are reported.

Model	5-shot	10-shot
MAML	2.205 ± 0.121	0.761 ± 0.068
Meta-SGD	2.053 ± 0.117	0.836 ± 0.065
MT-net	2.435 ± 0.130	0.967 ± 0.056
HSML	0.856 ± 0.073	0.161 ± 0.021
EEML	0.765 ± 0.140	0.151 ± 0.032

5.2 Toy Regression

We compare our methods with recent state-of-the-art works in meta-learning, such as MAML, Meta-SGD, HSML. We can see from Table 3 that our method achieves better results which indicates the effectiveness of our ensemble mechanism.

Table 4. Accuracy of image few-shot classification problem on MiniImageNet. Comparing with gradient-based meta-learning on 1-shot and 5-shot cases with 4-conv model as backbone. The results are with $\pm 95\%$ confidence intervals.

Model	MiniImageNet	
	1-shot	5-shot
MAML	48.07 ± 1.75	63.15 ± 0.91
Meta-SGD	50.47 ± 1.87	64.03 ± 0.94
ANIL	46.70 ± 0.40	61.50 ± 0.50
BOIL	49.61 ± 0.16	66.45 ± 0.37
MT-net	51.70 ± 1.84	—
sparse-MAML	50.35 ± 0.39	67.03 ± 0.74
HSML	50.38 ± 1.85	—
EEML	52.42 ± 1.75	68.40 ± 0.95

5.3 Image Few-shot Classification

We keep comparing with recent SOTA methods in much more complex images classification problem. As the Table 4 show, our method improves at least %4 for original MAML, which easily outperform recent SOTA like Meta-SGD, MT-net, HSML. Both 1-shot and 5-shot cases earn stable improvement which validates the effectiveness of our method.

6 Conclusion

In this paper, we propose a novel ensemble embedded meta-learning(EEML) algorithm, and comparing with recent meta-learning on both toy regression and image few-shot classification tasks. The experiment results show that our method possesses great potential not only for simple task like function regression, but also complex task like image classification. Through replacing the single dictatorial expert with an cluster based ensemble voting mechanism, our method earn easily improvement comparing with recent state-of-the-arts.

Although our method already earns explicit effects on two benchmarks, the potential of combining ensemble learning with meta-learning still calls for more exploration. And in reality employment case, the hyperparameters should be searched carefully to gain better performance.

References

1. Baik, S., Hong, S., Lee, K.M.: Learning to Forget for Meta-Learning. arXiv:1906.05895 [cs, stat] (Jun 2020), <http://arxiv.org/abs/1906.05895>, arXiv: 1906.05895
2. Bengio, S., Bengio, Y., Cloutier, J., Gecsei, J.: On the optimization of a synaptic learning rule. In: Preprints Conf. Optimality in Artificial and Biological Neural Networks. vol. 2 (1992)

3. Bertinetto, L., Henriques, J.F., Torr, P.H., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. arXiv preprint arXiv:1805.08136 (2018)
4. Breiman, L.: Bagging predictors. *Machine learning* **24**(2), 123–140 (1996)
5. Breiman, L.: Random forests. *Machine Learning* **45**(1), 5–32 (2001). <https://doi.org/10.1023/A:1010933404324>, <http://dx.doi.org/10.1023/A%3A1010933404324>
6. Denevi, G., Ciliberto, C., Grazi, R., Pontil, M.: Learning-to-learn stochastic gradient descent with biased regularization. In: *International Conference on Machine Learning*. pp. 1566–1575. PMLR (2019)
7. Dong, X., Yu, Z., Cao, W., Shi, Y., Ma, Q.: A survey on ensemble learning. *Frontiers Comput. Sci.* **14**(2), 241–258 (2020), <http://dblp.uni-trier.de/db/journals/fcsc/fcsc14.html#DongYCSM20>
8. Finn, C., Abbeel, P., Levine, S.: Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. arXiv:1703.03400 [cs] (Jul 2017), <http://arxiv.org/abs/1703.03400>, arXiv: 1703.03400
9. Finn, C., Xu, K., Levine, S.: Probabilistic model-agnostic meta-learning. arXiv preprint arXiv:1806.02817 (2018)
10. Grant, E., Finn, C., Levine, S., Darrell, T., Griffiths, T.: Recasting gradient-based meta-learning as hierarchical bayes. arXiv preprint arXiv:1801.08930 (2018)
11. Hochreiter, S., Younger, A.S., Conwell, P.R.: Learning to learn using gradient descent. In: *International Conference on Artificial Neural Networks*. pp. 87–94. Springer (2001)
12. Hospedales, T., Antoniou, A., Micaelli, P., Storkey, A.: Meta-learning in neural networks: A survey (2020). <https://doi.org/10.48550/ARXIV.2004.05439>, <https://arxiv.org/abs/2004.05439>
13. Koch, G., Zemel, R., Salakhutdinov, R.: Siamese neural networks for one-shot image recognition. In: *ICML deep learning workshop*. vol. 2. Lille (2015)
14. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10657–10665 (2019)
15. Lee, Y., Choi, S.: Gradient-based meta-learning with learned layerwise metric and subspace. In: *International Conference on Machine Learning*. pp. 2927–2936. PMLR (2018)
16. Li, H., Eigen, D., Dodge, S., Zeiler, M., Wang, X.: Finding task-relevant features for few-shot learning by category traversal. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1–10 (2019)
17. Li, S., Xue, L., Feng, L., Wang, Y., Wang, D.: Object detection network pruning with multi-task information fusion. *World Wide Web* **25**(4), 1667–1683 (Jul 2022). <https://doi.org/10.1007/s11280-021-00991-3>, <https://doi.org/10.1007/s11280-021-00991-3>
18. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few-shot learning. arXiv preprint arXiv:1707.09835 (2017)
19. Munkhdalai, T., Yu, H.: Meta networks. In: *International Conference on Machine Learning*. pp. 2554–2563. PMLR (2017)
20. Munkhdalai, T., Yuan, X., Mehri, S., Trischler, A.: Rapid adaptation with conditionally shifted neurons. In: *International Conference on Machine Learning*. pp. 3664–3673. PMLR (2018)
21. Munkhdalai, T., Yuan, X., Mehri, S., Wang, T., Trischler, A.: Learning rapid-temporal adaptations. *CoRR* **abs/1712.09926** (2017), <http://arxiv.org/abs/1712.09926>

22. Nichol, A., Achiam, J., Schulman, J.: On First-Order Meta-Learning Algorithms. arXiv:1803.02999 [cs] (Oct 2018), <http://arxiv.org/abs/1803.02999>, arXiv: 1803.02999
23. Oreshkin, B.N., Rodriguez, P., Lacoste, A.: Tadam: Task dependent adaptive metric for improved few-shot learning. arXiv preprint arXiv:1805.10123 (2018)
24. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
25. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-Learning with Latent Embedding Optimization. arXiv:1807.05960 [cs, stat] (Mar 2019), <http://arxiv.org/abs/1807.05960>, arXiv: 1807.05960
26. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T.: Meta-learning with memory-augmented neural networks. In: *International conference on machine learning*. pp. 1842–1850. PMLR (2016)
27. Schapire, R.E.: Explaining adaboost. In: *Empirical inference*, pp. 37–52. Springer (2013)
28. Schmidhuber, J.: Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation* **4**(1), 131–139 (1992). <https://doi.org/10.1162/neco.1992.4.1.131>
29. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. arXiv preprint arXiv:1703.05175 (2017)
30. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1199–1208 (2018)
31. Thrun, S.: Lifelong learning algorithms. In: *Learning to learn*, pp. 181–209. Springer (1998)
32. Thrun, S.B., Pratt, L.: *Learning to learn: introduction and overview* (1998)
33. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. *Artificial intelligence review* **18**(2), 77–95 (2002)
34. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. arXiv preprint arXiv:1606.04080 (2016)
35. Yao, H., Wei, Y., Huang, J., Li, Z.: Hierarchically structured meta-learning (2019). <https://doi.org/10.48550/ARXIV.1905.05301>, <https://arxiv.org/abs/1905.05301>
36. Yoon, J., Kim, T., Dia, O., Kim, S., Bengio, Y., Ahn, S.: Bayesian model-agnostic meta-learning. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. pp. 7343–7353 (2018)