

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

7-15-2021

Learning to Learn Variational Semantic Memory

Xiantong Zhen

AIM Lab, University of Amsterdam, Netherlands & Inception Institute of Artificial Intelligence, Abu Dhabi, UAE

Yingjun Du

AIM Lab, University of Amsterdam, Netherlands

Huan Xiong

Harbin Institute of Technology, Harbin, China & Mohamed bin Zayed University of Artificial Intelligence

Cees G.M. Snoek

AIM Lab, University of Amsterdam, Netherlands

Ling Shao

Inception Institute of Artificial Intelligence, Abu Dhabi, UAE & Mohamed bin Zayed University of Artificial Intelligence

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Computer Sciences Commons](#)

Preprint: arXiv

- Archived with thanks to arXiv
 - Preprint License: [CC by](#)
 - Uploaded 24 March 2022
-

Recommended Citation

X. Zhen, Y. Du, H. Xiong, Q. Qui, C.G.M. Snoek, and L. Shao, "Learning to learn variational semantic memory," 2021, arXiv:2010.10341

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Machine Learning Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Learning to Learn Variational Semantic Memory

Xiantong Zhen^{1,2*}, Yingjun Du^{1*}, Huan Xiong^{3,4}, Qiang Qiu⁵, Cees G. M. Snoek¹, Ling Shao^{2,4}

¹AIM Lab, University of Amsterdam, Netherlands

²Inception Institute of Artificial Intelligence, Abu Dhabi, UAE

³Harbin Institute of Technology, Harbin, China

⁴Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE

⁵Electrical and Computer Engineering, Purdue University, USA

Abstract

In this paper, we introduce *variational semantic memory* into meta-learning to acquire long-term knowledge for few-shot learning. The variational semantic memory accrues and stores semantic information for the probabilistic inference of class prototypes in a hierarchical Bayesian framework. The semantic memory is grown from scratch and gradually consolidated by absorbing information from tasks it experiences. By doing so, it is able to accumulate long-term, general knowledge that enables it to learn new concepts of objects. We formulate memory recall as the variational inference of a latent memory variable from addressed contents, which offers a principled way to adapt the knowledge to individual tasks. Our variational semantic memory, as a new long-term memory module, confers principled recall and update mechanisms that enable semantic information to be efficiently accrued and adapted for few-shot learning. Experiments demonstrate that the probabilistic modelling of prototypes achieves a more informative representation of object classes compared to deterministic vectors. The consistent new state-of-the-art performance on four benchmarks shows the benefit of variational semantic memory in boosting few-shot recognition.

1 Introduction

Memory plays an essential role in human intelligence, especially for aiding learning and reasoning in the present. In machine intelligence, neural memory [24, 78, 25] has been shown to enhance neural networks by augmentation with an external memory module. For instance, episodic memory storing past experiences helps reinforcement learning agents adapt more quickly and improve sample efficiency [7, 60, 26]. Memory is well-suited for few-shot learning by meta-learning in that it offers an effective mechanism to extract inductive bias [23] by accumulating prior knowledge from a set of previously observed tasks. One of the primary issues when designing a memory module is deciding what information should be memorized, which usually depends on the problems to solve. Though being highly promising, it is non-trivial to learn to store useful information in previous experience, which should be as non-redundant as possible. Existing few-shot learning works with external memory typically store the information from the support set of the current task [46, 75, 62, 45, 33], focusing on learning the access mechanism, which is assumed to be shared across tasks. The memory used in these works is short-term with limited capacity [20, 44] in that long-term information is not well retained, despite the importance for efficiently learning new tasks.

Semantic memory, also known as conceptual knowledge [48, 72, 64], refers to general facts and common world knowledge gathered throughout our lives [57]. It enables humans to quickly learn new concepts by recalling the knowledge acquired in the past [64]. Compared to episodic memory,

*Equal contribution.

semantic memory has been less studied [72, 64], despite its pivotal role in remembering the past and imagining the future [31]. By its very nature, semantic memory can provide conceptual context to facilitate novel event construction [30] and support a variety of cognitive activities, e.g., object recognition [5]. We draw inspiration from the cognitive function of semantic memory and introduce it into meta-learning to learn to collect long-term semantic knowledge for few-shot learning.

In this paper, we propose an external memory module to accrue and store long-term semantic information gained from past experiences, which we call variational semantic memory. The function of semantic memory closely matches that of prototypes [66, 1], which identify the semantics of objects in few-shot classification. The semantic knowledge accumulated in the memory helps build the new object concepts represented by prototypes typically obtained from only one or few samples [66]. We apply our variational semantic memory module to the probabilistic inference of class prototypes modelled as distributions. The probabilistic prototypes obtained are more informative and therefore better represent categories of objects compared to deterministic vectors [66, 1]. We formulate the memory recall as a variational inference of the latent memory, which is an intermediate stochastic variable. This offers a principled way to retrieve information from the external memory and incorporate it into the inference of class prototypes for each individual task. We cast the optimization as a hierarchical variational inference problem in the Bayesian framework; the parameters of the inference of prototypes are jointly optimized in conjunction with the memory recall and update. The semantic memory is gradually consolidated throughout the course of learning by updating the knowledge from new observations in each experienced task via an attention mechanism. The long-term semantic knowledge on seen object categories is acquired, maintained and enhanced during the learning process. This contrasts with existing works [62, 75] in which the memory stores data from the support set and therefore only considers the short term. In our memory, each entry stores semantics representing a distinct object category by summarizing feature representations of class samples. This reduces redundant information and saves storage overhead. More importantly it avoids collapsing memory reading and writing into single memory slots [24, 79], which ensures that rich context information is provided for better construction of new concepts.

To summarize our three contributions: *i)* We propose variational semantic memory, a long-term memory module, which learns to acquire semantic information and enables new concepts of object categories to be quickly learned for few-shot learning. *ii)* We formulate the memory recall as a variational inference problem by introducing the latent memory variable, which offers a principled way to retrieve relevant information that fits with specific tasks. *iii)* We introduce variational semantic memory into the probabilistic inference of prototypes modelled as distributions rather than deterministic vectors, which provides more informative representations of class prototypes.

2 Method

Few-shot classification is commonly learned by constructing T few-shot tasks from a large dataset and optimizing the model parameters on these tasks. A task, also called an episode, is defined as an N -way K -shot classification problem [75, 55]. An episode is drawn from a dataset by randomly sampling a subset of classes. Data points in an episode are partitioned into a support S and query Q set. We adopt the episodic optimization [75], which trains the model in an iterative way by taking one episode-update at a time. The update of the model parameters is defined by a variational learning objective, which is based on an evidence lower bound (ELBO) [6]. Different from traditional machine learning tasks, meta-learning for few-shot classification trains the model on the meta-training set, and evaluates on the meta-test set, whose classes are not seen during meta-training.

In this work, we develop our method based on the prototypical network (ProtoNet) [66]. Specifically, the prototype \mathbf{z}_n of an object class n is obtained by: $\mathbf{z}_n = \frac{1}{K} \sum_k \Phi(\mathbf{x}_{n,k})$ where $\Phi(\mathbf{x}_{n,k})$ is the feature embedding of the sample $\mathbf{x}_{n,k}$, which is usually obtained by a convolutional neural network. For each query sample \mathbf{x} , the distribution over classes is calculated based on the softmax over distances to the prototypes of all classes in the embedding space:

$$p(\mathbf{y}_n = 1 | \mathbf{x}) = \frac{\exp(-d(\Phi(\mathbf{x}), \mathbf{z}_n))}{\sum_{n'} \exp(-d(\Phi(\mathbf{x}), \mathbf{z}_{n'}))}, \quad (1)$$

where \mathbf{y} denotes a random one-hot vector, with \mathbf{y}_n indicating its n -th element, and $d(\cdot, \cdot)$ is some distance function. Due to its non-parametric nature, the ProtoNet enjoys high flexibility and efficiency, achieving great success in few-shot learning.

The ideal prototypical representation should be expressive and encompass enough intra-class variance, while being distinguishable between different classes. In the literature [66, 1], however, the prototypes are commonly modeled by a single or multiple deterministic vectors obtained by average pooling of only a few samples or clustering. Hence, they are not sufficiently representative of object categories. Moreover, uncertainty is inevitable due to the scarcity of data, which should also be encoded into the prototypical representations. In this paper, we derive a probabilistic latent variable model by modeling prototypes as distributions, which are learned by variational inference.

2.1 Variational Prototype Inference

We introduce the probabilistic modeling of class prototypes, in which we treat the prototype \mathbf{z} of each class as a distribution. In the few-shot learning scenario, to find \mathbf{z} is to infer the posterior $p(\mathbf{z}|\mathbf{x}, \mathbf{y})$, where (\mathbf{x}, \mathbf{y}) denotes the sample from the query set \mathcal{Q} . We derive a variational inference framework to solve \mathbf{z} by leveraging the support set \mathcal{S} .

Consider the conditional log-likelihood in a probabilistic latent variable model, where we incorporate the prototype \mathbf{z} as the latent variable

$$\log \left[\prod_{i=1}^{|\mathcal{Q}|} p(\mathbf{y}_i|\mathbf{x}_i) \right] = \log \left[\prod_{i=1}^{|\mathcal{Q}|} \int p(\mathbf{y}_i|\mathbf{x}_i, \mathbf{z}) p(\mathbf{z}|\mathbf{x}_i) d\mathbf{z} \right], \quad (2)$$

where $p(\mathbf{z}|\mathbf{x}_i)$ is the conditional prior in which we make the prototype dependent on \mathbf{x}_i . In general, it is intractable to directly solve the posterior, and usually we resort to a variational distribution to approximate the true posterior by minimizing the KL divergence:

$$D_{\text{KL}}[q(\mathbf{z}|\mathcal{S})||p(\mathbf{z}|\mathbf{x}, \mathbf{y})], \quad (3)$$

where $q(\mathbf{z}|\mathcal{S})$ is the variational posterior that makes the prototype \mathbf{z} dependent on the support set \mathcal{S} to leverage the meta-learning setting for few-shot classification. By applying the Baye’s rule, we obtain

$$\log \left[\prod_{i=1}^{|\mathcal{Q}|} p(\mathbf{y}_i|\mathbf{x}_i) \right] \geq \sum_{i=1}^{|\mathcal{Q}|} \left[\mathbb{E}_{q(\mathbf{z}|\mathcal{S})} [\log p(\mathbf{y}_i|\mathbf{x}_i, \mathbf{z})] - D_{\text{KL}}(q(\mathbf{z}|\mathcal{S})||p(\mathbf{z}|\mathbf{x}_i)) \right], \quad (4)$$

which is the ELBO of the conditional log-likelihood in (2). In practice, the variational posterior $q(\mathbf{z}|\mathcal{S})$ is implemented by a neural network that takes the average feature representations of samples in the support set \mathcal{S} and returns the mean and variance of the prototype \mathbf{z} . This can be directly adopted as the optimization objective for the variational inference of the prototype. While inheriting the flexibility of the prototype based few-shot learning [66, 1], our probabilistic inference enhances its class expressiveness by exploring higher-order information, i.e., variance, beyond a single or multiple deterministic mean vectors of samples in each class. More importantly, the probabilistic modeling provides a principled way of incorporating prior knowledge acquired from experienced tasks. In what follows, we introduce the external memory to augment the probabilistic latent model for enhanced variational inference of prototypes.

2.2 Variational Semantic Memory

We introduce the variational semantic memory to accumulate and store the semantic information from previous tasks for the inference of prototypes of new tasks. The knowledge on objects in the memory is consolidated episodically by seeing more object instances, which enables conceptual representations of new objects to be quickly built up for novel categories in tasks to come.

To be more specific, we deploy an external memory unit M which stores a key-value pair in each row of the memory array as [24]. The keys are the average feature representations of images from the same classes and the values are their corresponding class labels. The semantics of object categories in the memory provide context for quickly learning concepts of new object categories by seeing only a few examples in the current tasks. In contrast to most existing external memory modules [62, 52, 8], our variational semantic memory module stores semantic information by summarizing samples from individual categories, and therefore our memory module requires relatively light storage overhead, enabling more efficient retrieval of content from the memory.

Memory recall and inference It is pivotal to recall relevant information from the external memory and adapt it to learning new tasks when working with neural memory modules. When recalling a memory, it is not simply a read out; the content from the memory must be processed in order to fit the data in a specific task [52, 24, 78]. We regard memory recall as a decoding process of chosen content in the memory, which we accomplish via variational inference, instead of simply reading out the raw content from the external memory and directly incorporating it into specific tasks.

To this end, we introduce an intermediate stochastic variable, referred to as the latent memory \mathbf{m} . We cast the retrieval of memory into the inference of \mathbf{m} from the addressed memory M ; the memory addressing is based on the similarity between the content in the memory and the support set from the current task. The latent memory \mathbf{m} is inferred to connect the accrued semantic knowledge stored in the long-term memory to the current task, which is seamlessly coupled with the prototype inference under a hierarchical Bayesian framework.

From a Bayesian perspective, the prototype posterior can be inferred by marginalizing over the latent memory variable \mathbf{m} :

$$q(\mathbf{z}|S) = \int q(\mathbf{z}|\mathbf{m}, S)p(\mathbf{m}|S)d\mathbf{m}, \quad (5)$$

where $q(\mathbf{z}|\mathbf{m}, S)$ indicates that the prototype \mathbf{z} is now dependent on the support set S and the latent memory \mathbf{m} . To leverage the external memory M , we design a variational approximation $q(\mathbf{m}|M, S)$ to the posterior over the latent memory \mathbf{m} by inferring from M conditioned on S :

$$q(\mathbf{m}|M, S) = \sum_{a=1}^{|M|} p(\mathbf{m}|M_a, S)p(a|M, S). \quad (6)$$

Here, a is the addressed categorical variable, M_a denotes the corresponding memory content at address a , and $|M|$ represents the memory size, i.e., the number of memory entries.

We establish a hierarchical Bayesian framework for the variational inference of prototypes:

$$\tilde{q}(\mathbf{z}|M, S) = \sum_{a=1}^{|M|} p(a|M, S) \int q(\mathbf{z}|S, \mathbf{m})p(\mathbf{m}|M_a, S)d\mathbf{m}, \quad (7)$$

which is shown as a graphical model in Figure 1. We use the support set S and memory M to generate the categorical variable a to address the external memory, and then fetch the content M_a to infer the latent memory \mathbf{m} , which is incorporated as a conditional variable to assist S in the inference of the prototype \mathbf{z} . This offers a principled way to incorporate semantic knowledge and build up the prototypes of novel object categories. It mimics the cognitive mechanism of the human brain in learning new concepts by associating them with related concepts learned in the past [31]. Moreover, it naturally handles ambiguity and uncertainty when recalling memory better than the common strategy of using a deterministic transformation [24, 78].

When a is given, \mathbf{m} only depends on M_a and no longer relies on S . Therefore, we can attain $p(\mathbf{m}|M_a, S) = p(\mathbf{m}|M_a)$ by safely dropping S , which gives rise to:

$$q(\mathbf{m}|M, S) = \sum_{a=1}^{|M|} p(\mathbf{m}|M_a)p(a|M, S). \quad (8)$$

Since the memory size is finite, bounded by the number of seen classes, we further approximate $q(\mathbf{m}|M, S)$ empirically by

$$q(\mathbf{m}|M, S) = \sum_{a=1}^{|M|} \lambda_a p(\mathbf{m}|M_a), \quad \lambda_a = \frac{\exp(g(M_a, S))}{\sum_i \exp(g(M_i, S))}, \quad (9)$$

where M_a is the memory slot and stores the average feature representation of samples in each category that are seen in the learning stage, and $g(\cdot, \cdot)$ is a learnable similarity function, which we implement as a dot product for efficiency by taking the averages of samples in M_i and S , respectively.

Thus, the prototype inference can now be approximated by Monte Carlo sampling:

$$\tilde{q}(\mathbf{z}|M, S) \approx \frac{1}{J} \sum_{j=1}^J q(\mathbf{z}|\mathbf{m}^{(j)}, S), \quad \mathbf{m}^{(j)} \sim \sum_{a=1}^{|M|} \lambda_a p(\mathbf{m}|M_a), \quad (10)$$

where J is the number of Monte Carlo samples.

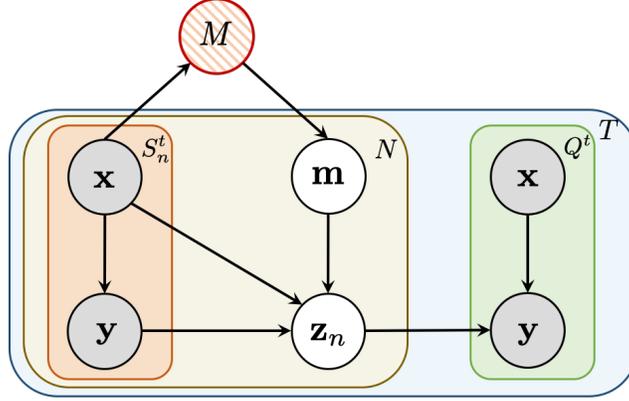


Figure 1: Graphical illustration of the proposed probabilistic prototype inference with variational semantic memory. M is the semantic memory module. S_n^t denotes the samples from the n -th class in the support set in each t task. Q^t is the query set. T is the number of tasks, and N is the number of classes in each task.

Memory update and consolidation The memory update is an important operation in the maintenance of memory, which should be able to effectively absorb new useful information to enrich memory content. We draw inspiration from the concept formation process in the human cognitive function [31]: the concept of an object category is formed and grown by seeing a set of similar objects of the same category. The memory is built from scratch and gradually consolidated by being episodically updated with knowledge observed from a series of related tasks. We adopt an attention mechanism to refresh content in the memory by taking into account the structural information of data.

To be more specific, the memory is empty at the beginning of the learning. When a new task arrives, we directly append the mean feature representation of data from a given category to the memory entries if this category is not seen. Otherwise, for seen categories, we update the memory content with new observed data from the current task using self-attention [73] similar to the graph attention mechanism [74]. This enables the structural information of data to be better explored for memory update. We first construct the graph with respect to the memory M_c to be updated. The nodes of the graph are a set of feature representations: $H_c = \{\mathbf{h}_c^0, \mathbf{h}_c^1, \mathbf{h}_c^2, \dots, \mathbf{h}_c^{N_c}\}$, where $\mathbf{h}_c^{N_c} \in \mathbb{R}^d$, $N_c = |S_c \cup Q_c|$, $\mathbf{h}_c^0 = M_c$, $\mathbf{h}_c^{i>0} = h_\phi(\mathbf{x}_c^i)$, $h_\phi(\cdot)$ is the convolutional neural network for feature extraction, and $\mathbf{x}_c^i \in \{S_c \cup Q_c\}$ contains all samples including both the support and query set from the c -th category in the current task.

We use the nodes H_c on the graph to generate a new representation of memory M_c , which better explores structural information of data. To do so, we need to compute a *similarity coefficient* between M_c and the nodes \mathbf{h}_c^i on the graph. We implement this by a single-layer feed-forward neural network parameterized by a weight vector $\mathbf{h} \in \mathbb{R}^{2d}$, that is, $e_{M_c}^i = \mathbf{w}^\top [M_c, \mathbf{h}_c^i]$ with $[\cdot, \cdot]$ being a concatenation operation. Here, $e_{M_c}^i$ indicates the importance of node i 's features to node M_c . In practice, we use the following normalized similarity coefficients [74]:

$$\beta_{M_c}^i = \text{softmax}_i(e_{M_c}^i) = \frac{\exp(\text{LeakyReLU}(\mathbf{w}^\top [M_c, \mathbf{h}_c^i]))}{\sum_{j=0}^{N_c} \exp(\text{LeakyReLU}(\mathbf{w}^\top [M_c, \mathbf{h}_c^j]))}. \quad (11)$$

We can now compute a linear combination of the feature representations of the nodes on the graph as the final output representation of \bar{M}_c :

$$\bar{M}_c = \sigma \left(\sum_{i=0}^{N_c} \beta_{M_c}^i \mathbf{h}_c^i \right), \quad (12)$$

where $\sigma(\cdot)$ is a nonlinear activation function, e.g., softmax. The graph attention operation can effectively find and assimilate the most useful information from the samples in the new task. We update the memory content with an attenuated weighted average,

$$M_c \leftarrow \alpha M_c + (1 - \alpha) \bar{M}_c, \quad (13)$$

where $\alpha \in (0, 1)$ is a hyperparameter. This operation allows useful information to be retained in the memory, while erasing less relevant or trivial information.

2.3 Objective

To train the model, we adopt stochastic gradient variational Bayes [36] and implement it using deep neural networks for end-to-end learning. By combining (4), (9) and (7), we obtain the following objective for the hierarchical variational inference:

$$\arg \min_{\{\phi, \theta, \psi, \varphi\}} - \sum_t^T \left[\sum_n^N \left[\sum_{(\mathbf{x}_i^t, \mathbf{y}_i^t) \in Q_n^t}^{|Q_n^t|} \left[\sum_\ell^L [\log p(\mathbf{y}_i^t | h_\phi(\mathbf{x}_i^t), \mathbf{z}_n^{(\ell)})] + D_{\text{KL}}\left(\frac{1}{J} \sum_{j=1}^J \tilde{q}_\varphi(\mathbf{z}_n | \mathbf{m}^{(j)}, \bar{\mathbf{h}}_{S_n^t}) || p_\theta(\mathbf{z}_n | h_\phi(\mathbf{x}_i^t))\right) + D_{\text{KL}}\left(\sum_a^{|M|} \lambda_a p_\psi(\mathbf{m} | M_a) || p_\psi(\mathbf{m} | \bar{\mathbf{h}}_{S_n^t})\right)\right] \right] \right], \quad (14)$$

where $\mathbf{z}^{(\ell)} \sim \frac{1}{J} \sum_{j=1}^J q_\varphi(\mathbf{z} | \mathbf{m}^{(j)}, S_n^t)$, $\mathbf{m}^{(j)} \sim \sum_{a=1}^{|M|} \lambda_a p_\psi(\mathbf{m} | M_a)$, L and J are numbers of Monte Carlo samples, $\bar{\mathbf{h}}_{S_n^t} = \frac{1}{|S_n^t|} \sum_{\mathbf{x} \in S_n^t} h_\phi(\mathbf{x})$, and n denotes the n -th class. To enable back propagation, we adopt the reparameterization trick [36] for sampling \mathbf{z} and \mathbf{m} . The third term in (14) essentially serves to constrain the inferred latent memory to ensure that it is relevant to the current task. Here, we make the parameters shared by the prior and the posterior for \mathbf{m} , and we also amortize the inference of prototypes across classes [22], which involves using the samples S_n^t from each class to infer their prototypes individually. In practice, the log-likelihood term is implemented as a cross entropy loss between predictions and ground-truth labels. The conditional probabilistic distributions are set to be diagonal Gaussian. We implement them using multi-layer perceptrons with the amortization technique and the reparameterization trick [36, 59], which take the conditionals as input and output the parameters of the Gaussian. In addition, we implement the model with the objective in (4), which we refer to as the variational prototypical network.

3 Related Work

Meta-Learning Meta-learning, or *learning to learn* [65, 69], for few-shot learning [37, 55, 62, 17, 75, 66] addresses the fundamental challenge of generalizing across tasks with limited labelled data. Meta-learning approaches for few-shot learning differ in the way they acquire inductive biases and adopt them for individual tasks [23]. They can be roughly categorized into four groups. Those in the first group are based on distance metrics and generally learn a shared/adaptive embedding space in which query images can be accurately matched to support images for classification [75, 66, 63, 47, 83, 1, 9]. Those based on optimization try to learn an optimization algorithm that is shared across tasks, and can be adapted to new tasks, enabling learning to be conducted efficiently and effectively [55, 2, 17, 18, 23, 71, 61, 82, 54]. The third group explicitly learns a base-learner that incorporates knowledge acquired by the meta-learner and effectively solves individual tasks [22, 4, 86]. In the fourth group, a memory mechanism has been incorporated. Usually, an external memory module is deployed to rapidly assimilate new data of unseen tasks, which is used for quick adaptation or to make decisions [62, 45, 46, 43]. The methods from different groups are not necessarily exclusive, and they can be combined to improve performance [71]. In addition, meta-learning has also been explored for reinforcement learning [42, 7, 77, 14, 16, 60] and other tasks [3, 81].

Prototypes The prototypical network is one of most successful meta-learning models for few-shot learning [75, 53, 27, 19, 34]. It learns to project samples into a metric space in which classification is conducted by computing the distance from query samples to class prototypes. Allen et al. [1] introduced an infinite mixture of prototypes that represents each category of objects by multiple clusters. The number of clusters is inferred from data by non-parametric Bayesian methods [53, 27]. Recently, Triantafillou et al. [71] combined the complementary strengths of prototypical networks and MAML [17] by leveraging their respective effective inductive bias and flexible adaptation mechanism for few-shot learning. Our work improves the prototypical network by probabilistic modeling of prototypes, inheriting the effectiveness and flexibility of the ProtoNet and further enriching the expressiveness of prototypes by the external memory mechanism.

Memory It has been shown that neural networks with memory, such as the long-short term memory [28] model, are capable of meta-learning [29]. Recent works augment neural networks with an external memory module to improve their learning capability [62, 49, 78, 25, 33, 45, 46, 52, 79, 80]. In few-shot learning, existing work with external memory mainly store the information contained in

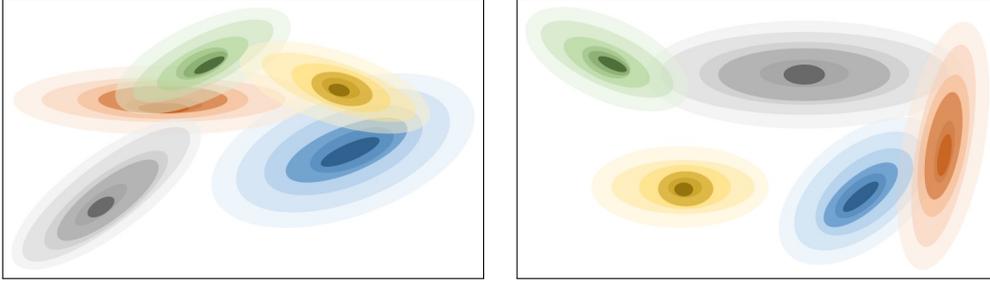


Figure 2: Prototype distributions of our variational prototype network without (left) and with variational semantic memory (right), where different colors indicate different categories. With the memory the prototypes become more distinctive and distant from each other, with less overlap.

Table 1: Benefit of variational prototype network over ProtoNet [66] in (%) on *miniImageNet*, *tieredImageNet* and CIFAR-FS.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
ProtoNet	47.40 \pm 0.60	65.41 \pm 0.52	53.31 \pm 0.89	72.69 \pm 0.74	55.50 \pm 0.70	72.01 \pm 0.60
Variational prototype network	52.11 \pm 1.70	66.13 \pm 0.83	55.13 \pm 1.88	73.71 \pm 0.84	61.35 \pm 1.60	75.72 \pm 0.90

Table 2: Benefit of variational semantic memory in our variational prototype network over alternative memory modules [24, 78] in (%) on *miniImageNet*, *tieredImageNet* and CIFAR-FS.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Variational prototype network	52.11 \pm 1.70	66.13 \pm 0.83	55.13 \pm 1.88	73.71 \pm 0.84	61.35 \pm 1.60	75.72 \pm 0.90
w/ Rote Memory	53.15 \pm 1.81	66.92 \pm 0.78	55.98 \pm 1.73	74.12 \pm 0.88	62.71 \pm 1.71	76.17 \pm 0.81
w/ Transformed Memory	53.85 \pm 1.71	67.23 \pm 0.89	56.15 \pm 1.70	74.33 \pm 0.80	62.97 \pm 1.88	76.97 \pm 0.77
w/ Variational semantic memory	54.73 \pm 1.60	68.01 \pm 0.90	56.88 \pm 1.71	74.65 \pm 0.81	63.42 \pm 1.90	77.93 \pm 0.80

the support set of the current task [46, 75, 45], focusing on learning the access mechanism shared across tasks. In these works, the external memory is wiped from episode to episode [20, 44]. Hence, it fails to maintain long-term information that has been shown to be crucial for efficiently learning new tasks [52, 20]. Memory has also been incorporated into generative models [8, 41, 79] and sequence modeling [39] by conditioning on the context information provided in the external memory. To store minimal amounts of data, Ramalho and Garnelo proposed a surprise-based memory module, which deploys a memory controller to select minimal samples to write into the memory [52]. In contrast to [8], our variational semantic memory adopts deterministic soft addressing, which enables us to leverage the full context of memory content by picking up multiple entries instead of a single one [8]. Our variational semantic memory is able to accrue long-term knowledge that provides rich context information for quickly learning novel tasks. Rather than directly using specific raw content or deploying a deterministic transformation [24, 78], we introduce the latent memory as an intermediate stochastic variable to be inferred from the addressed content in the memory. This enables the most relevant information to be retrieved from the memory and adapted to the data in specific tasks.

4 Experiments

Datasets and settings We evaluate our model on four standard few-shot classification tasks: *miniImageNet* [76], *tieredImageNet* [58], CIFAR-FS [4] and *Omniglot* [38]. For fair comparison with previous works, we experiment with both shallow convolutional neural networks with the same architecture as in [22] and a deep ResNet-12 [47, 40, 43, 56] architecture for feature extraction. More implementation details, including optimization settings and network architectures, are given in the supplementary material. We also provide a state-of-the-art comparison on the *Omniglot* dataset and more results on the performance with other deep architectures, e.g., WRN-28-10 [84].

Benefit of variational prototype network We compare against the ProtoNet [66] as our baseline model in which the prototypes are obtained by averaging the feature representations of each class. These results are obtained with shallow networks. As shown in Table 1, the proposed variational prototype network consistently outperforms the ProtoNet demonstrating the benefit brought by

Table 3: Advantage of memory update with attention mechanism.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
w/o Attention	53.97 ± 1.80	67.13 ± 0.76	56.05 ± 1.73	74.27 ± 0.85	62.93 ± 1.76	76.79 ± 0.80
w/ Attention	54.73 ± 1.60	68.01 ± 0.90	56.88 ± 1.71	74.65 ± 0.81	63.42 ± 1.90	77.93 ± 0.80

Table 4: Comparison with other memory models.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
MANN [62]	41.38 ± 1.70	61.73 ± 0.80	44.27 ± 1.69	67.15 ± 0.70	54.31 ± 1.91	67.98 ± 0.80
KM [79]	53.84 ± 1.70	67.35 ± 0.80	55.73 ± 1.65	73.36 ± 0.70	62.58 ± 1.80	77.11 ± 0.80
Variational semantic memory	54.73 ± 1.60	68.01 ± 0.90	56.88 ± 1.71	74.65 ± 0.81	63.42 ± 1.90	77.93 ± 0.80

probabilistic modeling. The probabilistic prototypes provide more informative representations of classes, which are able to encompass large intra-class variations and therefore improve performance.

Benefit of variational semantic memory We compare with two alternative methods of memory recall: rote memory and transformed memory [24, 78] (The implementation details of are provided in the supplementary material). As shown in Table 2, our variational semantic memory surpasses alternatives on all three benchmarks. The advantage over rote memory indicates the benefit of introducing the intermediate latent memory variable; the advantage over transformed memory demonstrates the benefit of formulating the memory recall as the variational inference of the latent memory, which is treated as a stochastic variable. To understand the empirical benefit, we visualize the distributions of prototypes obtained with/without variational semantic memory in Figure 2 on *miniImageNet*. The variational semantic memory enables the prototypes of different classes to be more distinctive and distant from each other, with less overlap, which enables larger intra-class variations to be encompassed, resulting in improved performance.

Benefit of attentional memory update We investigate the benefit of the attention mechanism for memory update. Specifically, we replace the attention-based update with a mean-based one; that is, we use $\bar{M} = \frac{1}{N_c} \sum_i h(\mathbf{x}_c^i)$. The experimental results are reported in Table 3. We can see that the memory update with attention mechanism performs consistently better than that using the mean-based update. This is because that with the attention mechanism, we are able to better absorb more informative knowledge from the data of new tasks by exploring the structural information.

Effect of memory size We conduct this experiment on *miniImageNet*. From Figure 3, we can see that the performance increases along with the increase in memory size. This is reasonable since larger memory provides more context information for building better prototypes. Moreover, we observe that the memory module plays a more significant role in the 1-shot setting. In this case, the prototype inferred from only one example might be insufficiently representative of the object class. Leveraging context information provided by the memory, however, compensates for the limited number of examples.

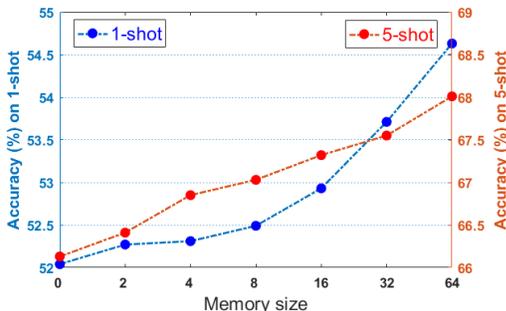


Figure 3: Effect of memory size.

Comparison with other memory models To demonstrate the effectiveness of our variational memory mechanism, we compare with two other representative memory models, i.e., the memory augmented neural network (MANN) [62] and the Kanerva machine (KM) [79]. MANN adopts an architecture with augmented memory capacities similar to neural Turing machines [24] while the KM deploys Kanerva’s sparse distributed memory mechanism and introduces learnable addresses and reparameterized latent variables. The KM was originally proposed for generative models, but we adopt its reading and writing mechanism to our semantic memory in the meta-learning setting for few-shot classification. The results are shown in Table 4. Our variational semantic memory consistently outperforms MANN and the KM on all three datasets.

Table 5: Comparison (%) on *miniImageNet*, *tieredImageNet* and CIFAR-FS using a shallow feature extractor.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Matching Net [75]	43.56 ± 0.84	55.31 ± 0.73	-	-	-	-
MAML [17]	48.70 ± 1.84	63.11 ± 0.92	51.67 ± 1.81	70.30 ± 1.75	58.90 ± 1.91	71.52 ± 1.10
Relation Net [68]	50.44 ± 0.82	65.32 ± 0.70	54.48 ± 0.93	65.32 ± 0.70	55.00 ± 1.01	69.30 ± 0.80
SNAIL (32C) by [4]	45.10 ± 0.85	55.20 ± 0.80	-	-	-	-
GNN [21]	50.31 ± 0.83	66.42 ± 0.90	-	-	61.90 ± 1.03	75.30 ± 0.91
PLATIPUS [18]	50.10 ± 1.90	-	-	-	-	-
VERSA [22]	53.31 ± 1.80	67.30 ± 0.91	-	-	62.51 ± 1.70	75.11 ± 0.91
R2-D2 (64C) [4]	49.50 ± 0.20	65.40 ± 0.20	-	-	62.30 ± 0.20	77.40 ± 0.20
R2-D2 [12]	51.70 ± 1.80	63.31 ± 0.91	-	-	60.20 ± 1.80	70.91 ± 0.91
CAVIA [87]	51.80 ± 0.70	65.61 ± 0.60	-	-	-	-
iMAML [51]	49.30 ± 1.90	-	-	-	-	-
VSM (This paper)	54.73 ± 1.60	68.01 ± 0.90	56.88 ± 1.71	74.65 ± 0.81	63.42 ± 1.90	77.93 ± 0.80

Table 6: Comparison (%) on *miniImageNet* and *tieredImageNet* using a deep feature extractor.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way	
	1-shot	5-shot	1-shot	5-shot
SNAIL [43]	55.71 ± 0.99	68.88 ± 0.92	-	-
AdaResNet [46]	56.88 ± 0.62	71.94 ± 0.57	-	-
TADAM [47]	58.50 ± 0.30	76.70 ± 0.30	-	-
Shot-Free [56]	59.04 ± n/a	77.64 ± n/a	63.52 ± n/a	82.59 ± n/a
TEWAM [50]	60.07 ± n/a	75.90 ± n/a	-	-
MTL [67]	61.20 ± 1.80	75.50 ± 0.80	-	-
Variational FSL [85]	61.23 ± 0.26	77.69 ± 0.17	-	-
MetaOptNet [40]	62.64 ± 0.61	78.63 ± 0.46	65.99 ± 0.72	81.56 ± 0.53
Diversity w/ Cooperation [15]	59.48 ± 0.65	75.62 ± 0.48	-	-
Meta-Baseline [11]	63.17 ± 0.23	79.26 ± 0.17	-	-
Tian et al. [70]	64.82 ± 0.60	82.14 ± 0.43	71.52 ± 0.69	86.03 ± 0.49
VSM (This paper)	65.72 ± 0.57	82.73 ± 0.51	72.01 ± 0.71	86.77 ± 0.44

State-of-the-art comparison As shown in Tables 5 and 6, our variational semantic memory (VSM) sets a new state-of-the-art on all few-shot learning benchmarks. On *miniImageNet*, our model using either a shallow or deep network achieves high recognition accuracy, surpassing the second best method, i.e., VERSA [22], by a margin of 1.43% on the 5-way 1-shot using a shallow network. On *tieredImageNet*, our model again outperforms previous methods using shallow networks, e.g., MAML [17] and Relation Net [68], and deep networks, e.g., [70]. On CIFAR-FS, our model delivers 63.42% on the 5-way 1-shot setting, surpassing the second best R2D2 [4] by 1.12%. The consistent state-of-the-art results on all benchmarks using either shallow or deep feature extraction networks validate the effectiveness of our model for few-shot learning.

5 Conclusion

In this paper, we introduce a new long-term memory module, named *variational semantic memory*, into meta-learning for few-shot learning. We apply it as an external memory for the probabilistic modelling of prototypes in a hierarchical Bayesian framework. The memory episodically learns to accrue and store semantic information by experiencing a set of related tasks, which provides semantic context that enables new object concepts to be quickly learned in individual tasks. The memory recall is formulated as the variational inference of a latent memory variable from the addressed content in the external memory. The memory is established from scratch and gradually consolidated by updating with knowledge absorbed from data in each task using an attention mechanism. Extensive experiments on four benchmarks demonstrate the effectiveness of variational semantic memory in learning to accumulate long-term knowledge. Our model achieves new state-of-the-art performance on four benchmark datasets, consistently surpassing previous methods. More importantly, the findings in this work demonstrate the benefit of semantic knowledge accrued through long-term memory in effectively learning novel concepts of object categories, and therefore highlight the pivotal role of semantic memory in few-shot recognition.

Broader Impact

This work introduces the concept of semantic memory from cognitive science into the machine learning field. We use it to augment a probabilistic model for few-shot learning. The developed variational framework offers a principled way to achieve memory recall, which could also be applied to other learning scenarios, e.g., continual learning. The empirical findings indicate the potential role of neural semantic memory as a long-term memory module in enhancing machine learning models. Finally, this work will not cause any foreseeable ethical issue or societal consequence.

Acknowledgements This work is financially supported by the Inception Institute of Artificial Intelligence, the University of Amsterdam and the allowance Top consortia for Knowledge and Innovation (TKIs) from the Netherlands Ministry of Economic Affairs and Climate Policy. The authors would like to thank all anonymous reviewers for their constructive feedback and valuable suggestions. Thanks to the partial support from national natural science foundation of China (Grants: 61976060, 61871016).

References

- [1] K. R. Allen, E. Shelhamer, H. Shin, and J. B. Tenenbaum. Infinite mixture prototypes for few-shot learning. In *ICML*, 2019.
- [2] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. de Freitas. Learning to learn by gradient descent by gradient descent. In *NeurIPS*, 2016.
- [3] Y. Balaji, S. Sankaranarayanan, and R. Chellappa. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*, 2018.
- [4] L. Bertinetto, J. F. Henriques, P. H. Torr, and A. Vedaldi. Meta-learning with differentiable closed-form solvers. In *ICLR*, 2019.
- [5] J. R. Binder and R. H. Desai. The neurobiology of semantic memory. *Trends in cognitive sciences*, 15(11):527–536, 2011.
- [6] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [7] C. Blundell, B. Uria, A. Pritzel, Y. Li, A. Ruderman, J. Z. Leibo, J. Rae, D. Wierstra, and D. Hassabis. Model-free episodic control. *arXiv preprint arXiv:1606.04460*, 2016.
- [8] J. Bornschein, A. Mnih, D. Zoran, and D. J. Rezende. Variational memory addressing in generative models. In *NeurIPS*, 2017.
- [9] T. Cao, M. Law, and S. Fidler. A theoretical analysis of the number of shots in few-shot learning. In *ICLR*, 2020.
- [10] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019.
- [11] Y. Chen, X. Wang, Z. Liu, H. Xu, and T. Darrell. A new meta-baseline for few-shot learning. *arXiv preprint arXiv:2003.04390*, 2020.
- [12] A. Devos, S. Chatel, and M. Grossglauser. Reproducing meta-learning with differentiable closed-form solvers. In *ICLR Workshop*, 2019.
- [13] G. S. Dhillon, P. Chaudhari, A. Ravichandran, and S. Soatto. A baseline for few-shot image classification. In *ICLR*, 2020.
- [14] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel. R^2 : Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [15] N. Dvornik, C. Schmid, and J. Mairal. Diversity with cooperation: Ensemble methods for few-shot classification. In *ICCV*, 2019.
- [16] R. Fakoor, P. Chaudhari, S. Soatto, and A. J. Smola. Meta-q-learning. *arXiv preprint arXiv:1910.00125*, 2020.
- [17] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [18] C. Finn, K. Xu, and S. Levine. Probabilistic model-agnostic meta-learning. In *NeurIPS*, 2018.
- [19] S. Fort. Gaussian prototypical networks for few-shot learning on omniglot. *arXiv preprint arXiv:1708.02735*, 2017.

- [20] M. Fortunato, M. Tan, R. Faulkner, S. Hansen, A. P. Badia, G. Buttimore, C. Deck, J. Z. Leibo, and C. Blundell. Generalization of reinforcement learners with working and episodic memory. In *NeurIPS*, 2019.
- [21] V. Garcia and J. Bruna. Few-shot learning with graph neural networks. In *ICLR*, 2018.
- [22] J. Gordon, J. Bronskill, M. Bauer, S. Nowozin, and R. E. Turner. Meta-learning probabilistic inference for prediction. In *ICLR*, 2019.
- [23] E. Grant, C. Finn, S. Levine, T. Darrell, and T. Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *ICLR*, 2018.
- [24] A. Graves, G. Wayne, and I. Danihelka. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [25] A. Graves, G. Wayne, M. Reynolds, T. Harley, I. Danihelka, A. Grabska-Barwińska, S. G. Colmenarejo, E. Grefenstette, T. Ramalho, J. Agapiou, et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471–476, 2016.
- [26] S. Hansen, A. Pritzel, P. Sprechmann, A. Barreto, and C. Blundell. Fast deep reinforcement learning using online adjustments from the past. In *NeurIPS*, 2018.
- [27] N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker. *Bayesian nonparametrics*, volume 28. Cambridge University Press, 2010.
- [28] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [29] S. Hochreiter, A. S. Younger, and P. R. Conwell. Learning to learn using gradient descent. In *ICANN*. Springer, 2001.
- [30] M. Irish, D. R. Addis, J. R. Hodges, and O. Piguet. Considering the role of semantic memory in episodic future thinking: evidence from semantic dementia. *Brain*, 135(7):2178–2191, 2012.
- [31] M. Irish and O. Piguet. The pivotal role of semantic memory in remembering the past and imagining the future. *Frontiers in behavioral neuroscience*, 7:27, 2013.
- [32] M. A. Jamal and G.-J. Qi. Task agnostic meta-learning for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11719–11727, 2019.
- [33] L. Kaiser, O. Nachum, A. Roy, and S. Bengio. Learning to remember rare events. In *ICLR*, 2016.
- [34] J. Kim, T.-H. Oh, S. Lee, F. Pan, and I. S. Kweon. Variational prototyping-encoder: One-shot learning with prototypical images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9462–9470, 2019.
- [35] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [36] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [37] G. Koch. Siamese neural networks for one-shot image recognition. In *ICML Workshop*, 2015.
- [38] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [39] H. Le, T. Tran, T. Nguyen, and S. Venkatesh. Variational memory encoder-decoder. In *NeurIPS*, 2018.
- [40] K. Lee, S. Maji, A. Ravichandran, and S. Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, 2019.
- [41] C. Li, J. Zhu, and B. Zhang. Learning to generate with memory. In *ICML*, 2016.
- [42] Z. Li, F. Zhou, F. Chen, and H. Li. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- [43] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. A simple neural attentive meta-learner. In *ICLR*, 2018.
- [44] A. Miyake and P. Shah. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press, 1999.
- [45] T. Munkhdalai and H. Yu. Meta networks. In *ICML*, 2017.
- [46] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler. Rapid adaptation with conditionally shifted neurons. In *ICML*, 2018.
- [47] B. Oreshkin, P. R. López, and A. Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018.
- [48] K. Patterson, P. J. Nestor, and T. T. Rogers. Where do you know what you know? the representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12):976–987, 2007.
- [49] A. Pritzel, B. Urias, S. Srinivasan, A. P. Badia, O. Vinyals, D. Hassabis, D. Wierstra, and C. Blundell. Neural episodic control. In *ICML*, 2017.

- [50] S. Qiao, C. Liu, W. Shen, and A. L. Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, 2018.
- [51] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine. Meta-learning with implicit gradients. In *NeurIPS*, 2019.
- [52] T. Ramalho and M. Garnelo. Adaptive posterior learning: few-shot learning with a surprise-based memory module. In *ICLR*, 2019.
- [53] C. E. Rasmussen. The infinite gaussian mixture model. In *NeurIPS*, 2000.
- [54] S. Ravi and A. Beaton. Amortized bayesian meta-learning. In *ICLR*, 2019.
- [55] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- [56] A. Ravichandran, R. Bhotika, and S. Soatto. Few-shot learning with embedded class models and shot-free meta training. In *ICCV*, 2019.
- [57] D. Reisberg. "*Semantic Memory*" in *The Oxford handbook of cognitive psychology*. Oxford University Press, 2013.
- [58] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018.
- [59] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [60] S. Ritter, J. Wang, Z. Kurth-Nelson, S. Jayakumar, C. Blundell, R. Pascanu, and M. Botvinick. Been there, done that: Meta-learning with episodic recall. In *ICML*, 2018.
- [61] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell. Meta-learning with latent embedding optimization. In *ICLR*, 2019.
- [62] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, 2016.
- [63] V. G. Satorras and J. B. Estrach. Few-shot learning with graph neural networks. In *ICLR*, 2018.
- [64] D. Saumier and H. Chertkow. Semantic memory. *Current neurology and neuroscience reports*, 2(6):516–522, 2002.
- [65] J. Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.
- [66] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [67] Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele. Meta-transfer learning for few-shot learning. In *CVPR*, 2019.
- [68] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018.
- [69] S. Thrun and L. Pratt. *Learning to learn*. Springer Science & Business Media, 2012.
- [70] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola. Rethinking few-shot image classification: a good embedding is all you need? *arXiv preprint arXiv:2003.11539*, 2020.
- [71] E. Triantafillou, T. Zhu, V. Dumoulin, P. Lamblin, U. Evci, K. Xu, R. Goroshin, C. Gelada, K. Swersky, P.-A. Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *ICLR*, 2020.
- [72] E. Tulving et al. Episodic and semantic memory. *Organization of memory*, 1:381–403, 1972.
- [73] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [74] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *ICLR*, 2018.
- [75] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NeurIPS*, 2016.
- [76] O. Vinyals, M. Fortunato, and N. Jaitly. Pointer networks. In *NeurIPS*. 2015.
- [77] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.
- [78] J. Weston, S. Chopra, and A. Bordes. Memory networks. In *ICLR*, 2014.
- [79] Y. Wu, G. Wayne, A. Graves, and T. Lillicrap. The kanerva machine: A generative distributed memory. In *ICLR*, 2018.
- [80] Y. Wu, G. Wayne, K. Gregor, and T. Lillicrap. Learning attractor dynamics for generative memory. In *NeurIPS*, 2018.

- [81] T. Yang, X. Zhang, Z. Li, W. Zhang, and J. Sun. Metaanchor: Learning to detect objects with customized anchors. In *NeurIPS*, 2018.
- [82] J. Yoon, T. Kim, O. Dia, S. Kim, Y. Bengio, and S. Ahn. Bayesian model-agnostic meta-learning. In *NIPS*, 2018.
- [83] S. W. Yoon, J. Seo, and J. Moon. Tapnet: Neural network augmented with task-adaptive projection for few-shot learning. In *ICML*, 2019.
- [84] S. Zagoruyko and N. Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [85] J. Zhang, C. Zhao, B. Ni, M. Xu, and X. Yang. Variational few-shot learning. In *ICCV*, 2019.
- [86] X. Zhen, H. Sun, Y. Du, J. Xu, Y. Yin, L. Shao, and C. Snoek. Learning to learn kernels with variational random features. *ICML*, 2020.
- [87] L. Zintgraf, K. Shiarli, V. Kurin, K. Hofmann, and S. Whiteson. Fast context adaptation via meta-learning. In *ICML*, 2019.

6 More Results

We provide more experimental results on the *Omniglot* dataset, under the 20-way settings on *miniImageNet* and *tieredImageNet* datasets, comparisons with other deep architectures, e.g., WRN-10-28 [84], and further ablation studies on effect of the hyperparameter (α) in the memory update and the effect of the Gumbel-softmax for approximating the addressing vectors in the memory recall.

On the *Omniglot* dataset, as shown in Table 7 our model consistently achieves high performance, exceeding other competitive methods (98.3% - up 0.07%) under the 20-way 1-shot setting. The results are consistent with the findings on the other datasets.

The performance under the 20-way 1-shot setting on the *miniImageNet* and *tieredImageNet* datasets are reported in Table 8. We compare against state-of-the-art methods, i.e., [32] and Baseline++ [10] under this setting. The proposed variational semantic memory consistently outperforms these two methods on *miniImageNet*. We do not observe previous results under the 20-way 1-shot setting on *tieredImageNet* and we provide our results for future comparison.

We have also experimented with one more deep architecture, i.e., WRN-28-10 [84], to compare with previous methods using the same architecture. The results are reported in Table 9. Again, our model outperforms those methods, which is consistent with the results using other architectures.

In addition, we test the impact of α in (13) the memory update. The value of α control how much information in the memory will be kept during the update with new information. The experimental results on the *miniImageNet* dataset under both 1-shot and 5-shot setting are shown in Table 10. We can see that the performance achieves the best when the values of α are 0.7 and 0.8. This means that in each update we need to keep the majority of old information in the memory.

Finally, in the memory recall, we use softmax to generate the addressing vector, and compare it against the Gumbel-softmax approximation. The results on three datasets are shown in Table 11. We can see that Gumbel-softmax achieves comparable performance with the regular softmax. It is worth mentioning that the memory size needs to be pre-fixed when using the Gumbel-softmax approximation, while the regular softmax can deal with dynamic memories with growing size.

7 Datasets

miniImageNet. The *miniImageNet* is originally proposed in [76], has been widely used for evaluating few-shot learning algorithms. It consists of 60,000 color images from 100 classes with 600 examples per class. The images have dimensions of 84×84 pixels. We follow the train/val/ test split introduced in [55], which uses 64 classes for meta-training, 16 classes for meta-validation, and the remaining 20 classes for meta-testing.

tieredImageNet. The *tieredImageNet* dataset [58] is a larger subset of ImageNet with 608 classes (779, 165 images) grouped into 34 higher-level categories in the ImageNet human-curated hierarchy. These categories are further divided into 20 training categories (351 classes), 6 validation categories (97 classes), and 8 testing categories (160 classes). This construction near the root of the ImageNet hierarchy results in a more challenging, yet realistic regime with test classes that are distinctive enough from training classes.

CIFAR-FS. The CIFAR-FS is proposed in [4], which is randomly sampled from CIFAR-100 by using the same standard with which *miniImageNet* has been generated. The original resolution of 32×32 pixels makes the task harder.

Omniglot. The Omniglot [38] is a few-shot learning dataset consisting of 1623 handwritten characters (each with 20 instances) derived from 50 alphabets. We follow a pre-processing and training procedure defined in [75]. We first resize images to 28×28 and then character classes are augmented with rotations of 90 degrees. The training, validation and test sets consist of a random split of 1100, 100, and 423 characters, respectively.

8 Implementation details

We provide more implementation details. For the feature extraction networks, we do not use any fully connected layer after the convolutional layers. All of our models were trained via SGD with

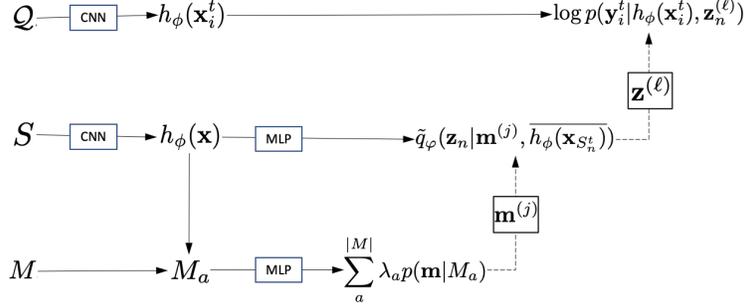


Figure 4: The sketch of the implementation.

the Adam [35] optimizer. For the 5-way 5-shot model, we train using the setting of 8 tasks per batch for 100,000 iterations and use a constant learning rate of 0.0001. For the 5-way 1-shot model, we train with the setting of 8 tasks per batch for 150,000 iterations and use a constant learning rate of 0.00025. No regularization was used other than batch normalization. In the Monte Carlo sampling, we set the number $J = 150$ for \mathbf{m} and to $L = 100$ for \mathbf{z} , which are chosen by using the validation set. The architectures of inference networks $\tilde{q}_\varphi(\cdot)$, $p_\psi(\cdot)$, the prior network $p_\theta(\cdot)$ and feature extraction networks $h_\phi(\cdot)$ are provided in Tables 12 and 13. The sketch of the implementation is shown in Figure 4. We implemented all models in the Tensorflow framework and tested on an NVIDIA Tesla V100.

In the experiments of comparison with alternative methods of memory recall, For rote memory, we put concatenation of the addressed memory contents with mean feature representation of the support set as input to the inference network: $q(\mathbf{z}|\bar{M}, \bar{h}_\phi(\mathbf{x}_{S_n^t}))$, where $\bar{M} = \sum_a^{|M|} \lambda_a M_a$, and $\lambda_a = \frac{g(M_a, S)}{\sum_i g(M_i, S)}$; for transformed memory, we follow the strategy in [24, 78] and pass the addressed memory \bar{M} through a parameterized transformation before feeding into the inference network: $q(\mathbf{z}|T(\bar{M}), \bar{h}_\phi(\mathbf{x}_{S_n^t}))$, where $T(\cdot)$ is the transformation implemented as a multi-layer perceptron (MLP).

Table 7: Comparison (%) on *Omniglot* using a shallow feature extractor.

	<i>Omniglot</i> , 5-way		<i>Omniglot</i> , 20-way	
	1-shot	5-shot	1-shot	5-shot
Siamese Net [37]	96.7	98.4	88.0	96.5
Matching net [75]	98.1	98.9	93.8	98.5
MAML [17]	98.7 \pm 0.4	99.9 \pm 0.1	95.8 \pm 0.3	98.9 \pm 0.2
SNAIL [43]	99.1 \pm 0.2	99.8 \pm 0.1	97.6 \pm 0.3	99.4 \pm 0.2
GNN [21]	99.2	99.7	97.4	99.0
VERSA [22]	99.7 \pm 0.2	99.8 \pm 0.1	97.7 \pm 0.3	98.8 \pm 0.2
R2-D2 [4]	98.6	99.7	94.7	98.9
IMP [1]	98.4 \pm 0.3	99.5 \pm 0.1	95.0 \pm 0.1	98.6 \pm 0.1
ProtoNet [66]	98.5 \pm 0.2	99.5 \pm 0.1	95.3 \pm 0.2	98.7 \pm 0.1
<i>This paper</i>	99.8 \pm 0.1	99.9 \pm 0.1	98.3 \pm 0.3	99.4 \pm 0.2

Table 8: Performance comparison under 20-way settings on the *miniImageNet* and *tieredImageNet* datasets.

	<i>miniImageNet</i>		<i>tieredImageNet</i>	
	1-shot	5-shot	1-shot	5-shot
TAML [32]	19.73 \pm 0.65	29.81 \pm 0.35	n/a	n/a
Baseline++ [10]	n/a	38.03 \pm 0.24	n/a	n/a
<i>This paper</i>	22.07 \pm 0.53	39.98 \pm 0.27	24.76 \pm 0.51	41.84 \pm 0.31

Table 9: Comparison (%) on *miniImageNet* and *tieredImageNet* using WRN-28-10 feature extractor.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way	
	1-shot	5-shot	1-shot	5-shot
Activation to Parameter [50]	59.60 \pm 0.41	73.74 \pm 0.19	-	-
Fine-Tuning [13]	57.73 \pm 0.62	78.17 \pm 0.49	66.58 \pm 0.70	85.55 \pm 0.48
LEO [61]	61.76 \pm 0.08	77.59 \pm 0.12	66.33 \pm 0.05	81.44 \pm 0.09
<i>This paper</i>	63.45 \pm 0.39	78.99 \pm 0.:	68.54 \pm 0.61	86.25 \pm 0.39

Table 10: Performance comparison by using various α on the *miniImageNet* dataset.

	1-shot	5-shot
$\alpha = 0.$	51.28 \pm 1.70	65.77 \pm 0.70
$\alpha = 0.1$	51.93 \pm 1.80	65.71 \pm 0.90
$\alpha = 0.2$	52.15 \pm 1.70	65.99 \pm 0.80
$\alpha = 0.3$	52.11 \pm 1.70	65.96 \pm 0.70
$\alpha = 0.4$	52.10 \pm 1.90	66.11 \pm 0.90
$\alpha = 0.5$	53.15 \pm 1.60	66.93 \pm 0.70
$\alpha = 0.6$	53.61 \pm 1.70	67.15 \pm 0.80
$\alpha = 0.7$	54.73 \pm 1.60	67.37 \pm 0.80
$\alpha = 0.8$	53.94 \pm 1.70	68.01 \pm 0.90
$\alpha = 0.9$	53.77 \pm 1.80	67.53 \pm 0.90
$\alpha = 1.0$	53.53 \pm 1.70	67.05 \pm 0.80

Table 11: Performance comparison with Gumbel-softmax for memory addressing.

	<i>miniImageNet</i> , 5-way		<i>tieredImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Gumbel-softmax	53.27 \pm 1.70	68.37 \pm 0.80	55.25 \pm 1.80	74.57 \pm 0.80	62.11 \pm 1.60	77.69 \pm 0.80
Softmax	54.73 \pm 1.60	68.01 \pm 0.90	56.88 \pm 1.71	74.65 \pm 0.81	63.42 \pm 1.90	77.93 \pm 0.80

Table 12: The architectures of inference networks and prior network.

The inference network $q_\varphi(\cdot)$ for Omniglot, <i>miniImageNet</i> , CIFAR-FS.	
Output size	Layers
$J \times 512$	concatenate $\mathbf{m}^{(j)}$ and $\overline{h_\phi(\mathbf{x}_{S_n^i})}$
256	fully connected, ELU
256	fully connected, ELU
256	linear fully connected to $\mu_z, \log \sigma_z^2$
The prior network $p_\theta(\cdot)$ for Omniglot, <i>miniImageNet</i> , CIFAR-FS	
Output size	Layers
256	Input query feature
256	fully connected, ELU
256	fully connected, ELU
256	fully connected to $\mu_z, \log \sigma_z^2$
The inference network $p_\psi(\cdot)$ for Omniglot, <i>miniImageNet</i> , CIFAR-FS.	
Output size	Layers
256	Input memory feature
256	fully connected, ELU
256	fully connected, ELU
256	linear fully connected to $\mu_m, \log \sigma_m^2$

Table 13: The architectures of CNN for different datasets.

The CNN architecture $h_\phi(\cdot)$ for Omniglot.	
Output size	Layers
$28 \times 28 \times 1$	Input images
$14 \times 14 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.9, <i>pool</i> (2×2, stride=2, SAME)
$7 \times 7 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.9, <i>pool</i> (2×2, stride=2, SAME)
$4 \times 4 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.9, <i>pool</i> (2×2, stride=2, SAME)
$2 \times 2 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.9, <i>pool</i> (2×2, stride=2, SAME)
256	flatten
The shallow architecture $h_\phi(\cdot)$ for CIFAR-FS	
Output size	Layers
$32 \times 32 \times 3$	Input images
$16 \times 16 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$8 \times 8 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$4 \times 4 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$2 \times 2 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
256	flatten
The CNN architecture $h_\phi(\cdot)$ for <i>miniImageNet</i> and <i>tieredImageNet</i>	
Output size	Layers
$84 \times 84 \times 3$	Input images
$42 \times 42 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$21 \times 21 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$10 \times 10 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$5 \times 5 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
$2 \times 2 \times 64$	<i>conv2d</i> (3×3, stride=1, SAME, RELU), dropout 0.5, <i>pool</i> (2×2, stride=2, SAME)
256	flatten