
Learning to Learn Kernels with Variational Random Features

Xiantong Zhen^{*12} Haoliang Sun^{*3} Yingjun Du^{*2} Jun Xu⁴ Yilong Yin³ Ling Shao⁵¹ Cees Snoek²

Abstract

We introduce kernels with random Fourier features in the meta-learning framework for few-shot learning. We propose meta variational random features (MetaVRF) to learn adaptive kernels for the base-learner, which is developed in a latent variable model by treating the random feature basis as the latent variable. We formulate the optimization of MetaVRF as a variational inference problem by deriving an evidence lower bound under the meta-learning framework. To incorporate shared knowledge from related tasks, we propose a context inference of the posterior, which is established by an LSTM architecture. The LSTM-based inference network effectively integrates the context information of previous tasks with task-specific information, generating informative and adaptive features. The learned MetaVRF is able to produce kernels of high representational power with a relatively low spectral sampling rate and also enables fast adaptation to new tasks. Experimental results on a variety of few-shot regression and classification tasks demonstrate that MetaVRF can deliver much better, or at least competitive, performance compared to existing meta-learning alternatives.

1. Introduction

Learning to learn, or *meta-learning* (Schmidhuber, 1992; Thrun & Pratt, 2012), offers a promising tool for few-shot learning (Andrychowicz et al., 2016; Ravi & Larochelle, 2017; Finn et al., 2017) and has recently generated increasing popularity in machine learning. The crux of meta-learning for few-shot learning is to extract prior knowledge

from related tasks to enable fast adaptation to a new task with a limited amount of data. Generally speaking, existing meta-learning algorithms (Ravi & Larochelle, 2017; Bertinetto et al., 2019) design the meta-learner to extract meta-knowledge that improves the performance of the base-learner on individual tasks. Meta knowledge, like a good parameter initialization (Finn et al., 2017), or an efficient optimization update rule shared across tasks (Andrychowicz et al., 2016; Ravi & Larochelle, 2017) has been extensively explored in general learning framework, but how to define and use in few-shot learning remains an open question.

An effective base-learner should be powerful enough to solve individual tasks and able to absorb information provided by the meta-learner to improve its own performance. While potentially strong base-learners, kernels (Hofmann et al., 2008) have not yet been studied in the meta-learning scenario for few-shot learning. Learning adaptive kernels (Bach et al., 2004) in a data-driven way via random features (Rahimi & Recht, 2007) has demonstrated great success in regular learning tasks and remains of broad interest in machine learning (Sinha & Duchi, 2016; Hensman et al., 2017; Carratino et al., 2018; Bullins et al., 2018; Li et al., 2019). However, due to the limited availability of data, it is challenging for few-shot learning to establish informative and discriminant kernels. We thus explore the relatedness among distinctive but relevant tasks to generate rich random features to build strong kernels for base-learners, while still maintaining their ability to adapt quickly to individual tasks.

In this paper, we make three important contributions. First, we propose meta variational random features (MetaVRF), integrating, for the first time, kernel learning with random features and variational inference into the meta-learning framework for few-shot learning. We develop MetaVRF in a latent variable model by treating the random Fourier basis of translation-invariant kernels as the latent variable. Second, we formulate the optimization of MetaVRF as a variational inference problem by deriving a new evidence lower bound (ELBO) in the meta-learning setting, where the posterior over the random feature basis corresponds to the spectral distribution associated with the kernel. This formulation under probabilistic modeling provides a principled way of learning data-driven kernels with random Fourier features and more importantly, fits well in the meta-learning framework for few-shot learning allowing us to

^{*}Equal contribution ¹Inception Institute of Artificial Intelligence, UAE ²Informatics Institute, University of Amsterdam, The Netherlands ³School of Software, Shandong University, China ⁴College of Computer Science, Nankai University, China ⁵Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE. Correspondence to: X. Zhen <zhenxt@gmail.com>.

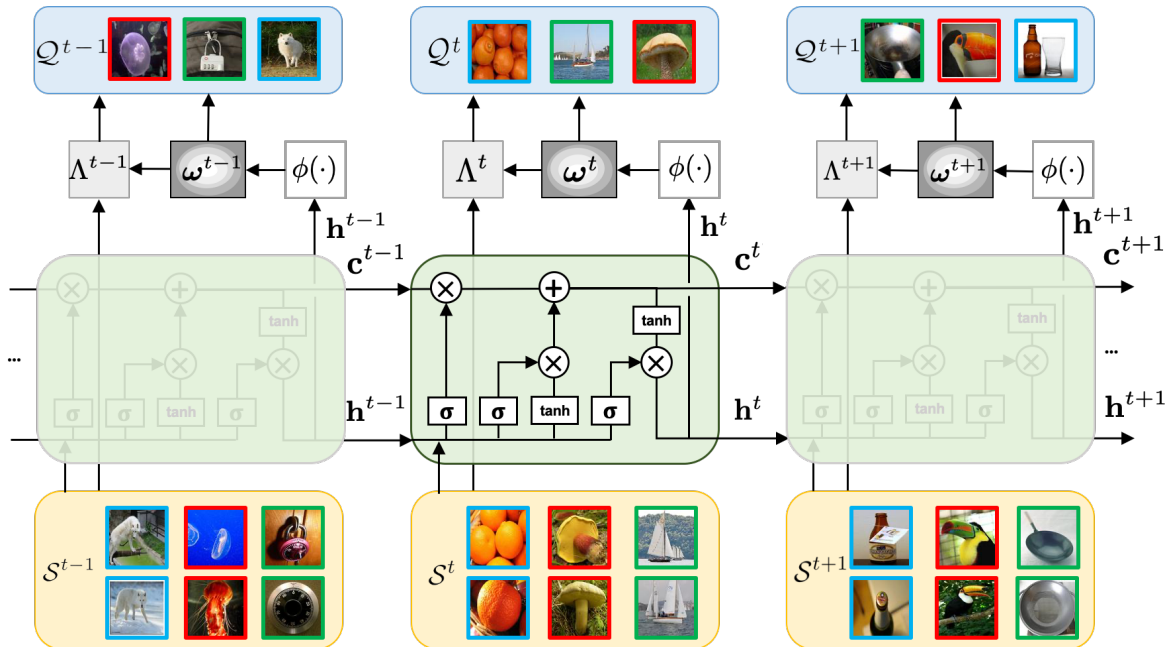


Figure 1. The learning framework of our meta variational random features (MetaVRF). The meta-learner employs an LSTM-based inference network () to infer the spectral distribution over Λ^t , the kernel from the support set S^t of the current task t , and the outputs \mathbf{h}^{t-1} and \mathbf{c}^{t-1} of the previous task. During the learning process, the cell state in the LSTM is deployed to accumulate the shared knowledge through experiencing a set of prior tasks. The *remember* and *forget* gates in the LSTM episodically refine the cell state by absorbing information from each experienced task. For each individual task, the task-specific information extracted from its support set is combined with distilled information from the previous tasks to infer the adaptive spectral distribution of the kernels.

flexibly customize the variational posterior to leverage the meta knowledge for inference. As the third contribution, we propose a context inference which puts the inference of random feature bases of the current task into the context of all previous, related tasks. The context inference provides a generalized way to integrate context information of related tasks with task-specific information for the inference of random feature bases. To establish the context inference, we introduce a recurrent LSTM architecture (Hochreiter & Schmidhuber, 1997), leveraging its innate capability of learning long-term dependencies, which can be adopted to explore shared meta-knowledge from a large set of previous tasks. The LSTM-based inference connects knowledge from previous tasks to the current task, gradually collecting and refreshing the knowledge across the course of learning. The learning process with an LSTM-based inference network is illustrated in Figure 1. Once learning ceases, the ultimate LSTM state gains meta-knowledge from related experienced tasks, which enables fast adaptation to new tasks.

We demonstrate the effectiveness of the proposed MetaVRF by extensive experiments on a variety of few-shot regression and classification tasks. Results show that our MetaVRF achieves better, or at least competitive, performance compared to previous methods. Moreover, we conduct further analysis on MetaVRF to demonstrate its ability to be integrated with deeper architectures and its efficiency with

relatively low sampling rates. We also apply MetaVRF to versatile and challenging settings with inconsistent training and test conditions, and it still delivers promising results, which further demonstrates its strong learning ability.

2. Method

We first describe the base-learner based on the kernel ridge regression in meta-learning for few-shot learning, and then introduce kernel learning with random features, based on which our meta variational random features are developed.

2.1. Meta-Learning with Kernels

We adopt the episodic training strategy commonly used for few-shot classification in meta-learning (Ravi & Larochelle, 2017), which involves *meta-training* and *meta-testing* stages. In the *meta-training* stage, a meta-learner is trained to enhance the performance of a base-learner on a *meta-training* set with a batch of few-shot learning tasks, where a task is usually referred as an episode (Ravi & Larochelle, 2017). In the *meta-test* stage, the base-learner is evaluated on a *meta-testing* set with different classes of data samples from the *meta-training* set.

For the few-shot classification problem, we sample C -way k -shot classification tasks from the *meta-training* set, where k

is the number of labelled examples for each of the C classes. Given the t -th task with a support set $S^t = f(\mathbf{x}_i; \mathbf{y}_i) g_{i=1}^C$ and query set $Q^t = f(\mathbf{x}_i; \mathbf{y}_i) g_{i=1}^m (S^t; Q^t \subseteq X)$, we learn the parameters θ^t of the predictor f^t using a standard learning algorithm with kernel trick $\theta^t = \theta(\mathcal{X}; Y)$, where $S^t = f\mathcal{X}; Yg$. Here, θ is the base-learner and $\theta : X \rightarrow \mathbb{R}^X$ is a mapping function from X to a dot product space H . The similarity measure $k(\mathbf{x}; \mathbf{x}^\theta) = h(\mathbf{x}; \mathbf{x}^\theta)$ is usually called a kernel (Hofmann et al., 2008).

As in traditional supervised learning problems, the base-learner for the t -th single task can use a predefined kernel, e.g., radius base function, to map the input into a dot product space for efficient learning. Once the base-learner is obtained on the support set, its performance is evaluated on the query set by the following loss function:

$$\sum_{(\mathbf{x}; \mathbf{y}) \in Q^t} L(f^t(\mathbf{x}); \mathbf{y}); \quad (1)$$

where $L(\cdot)$ can be any differentiable function, e.g., cross-entropy loss. In the meta-learning setting for few-shot learning, we usually consider a batch of tasks. Thus, the meta-learner is trained by optimizing the following objective function *w.r.t.* the empirical loss on T tasks

$$\sum_{t=1}^T \sum_{(\mathbf{x}; \mathbf{y}) \in Q^t} L(f^t(\mathbf{x}); \mathbf{y}); \text{ s.t. } \theta^t = \theta(\mathcal{X}; Y); \quad (2)$$

where θ^t is the feature mapping function which can be obtained by learning a task-specific kernel k^t for each task t with data-driven random Fourier features.

In this work, we employ kernel ridge regression, which has an efficient closed-form solution, as the base-learner for few-shot learning. The kernel value in the Gram matrix $K \in \mathbb{R}^{Ck \times Ck}$ is computed as $k(\mathbf{x}; \mathbf{x}^\theta) = (\mathbf{x}; \mathbf{x}^\theta)^{\triangleright}$, where “ \triangleright ” is the transpose operation. The base-learner for a single task is obtained by solving the following objective *w.r.t.* the support set of this task,

$$\theta = \arg \min \text{Tr}[(Y - K)\theta]^{\triangleright} + \text{Tr}[K\theta]; \quad (3)$$

which admits a closed-form solution

$$\theta = Y(1 + K)^{-1} \quad (4)$$

The learned predictor is then applied to samples in the query set X :

$$\hat{Y} = f(X) = K\theta; \quad (5)$$

Here, $K = (X; X)^{\triangleright} \in \mathbb{R}^{Ck \times m}$, with each element as $k(\mathbf{x}; \mathbf{x})$ between the samples from the support and query sets. Note that we also treat θ in (3) as a trainable parameter by leveraging the meta-learning setting, and all these parameters are learned by the meta-learner.

Rather than using pre-defined kernels, we consider learning adaptive kernels with random Fourier features in a data-driven way. Moreover, we leverage the shared knowledge by exploring dependencies among related tasks to learn rich features for building up informative kernels.

2.2. Random Fourier Features

Random Fourier features (RFFs) were proposed to construct approximate translation-invariant kernels using explicit feature maps (Rahimi & Recht, 2007), based on Bochner’s theorem (Rudin, 1962).

Theorem 1 (Bochner’s theorem) (Rudin, 1962) *A continuous, real valued, symmetric and shift-invariant function $k(\mathbf{x}; \mathbf{x}^\theta) = k(\mathbf{x} - \mathbf{x}^\theta)$ on \mathbb{R}^d is a positive definite kernel if and only if it is the Fourier transform $p(\cdot)$ of a positive finite measure such that*

$$k(\mathbf{x}; \mathbf{x}^\theta) = \int_{\mathbb{R}^d} e^{i\mathbf{l}^\top(\mathbf{x} - \mathbf{x}^\theta)} dp(\mathbf{l}) = E_{\mathbf{l}} [f(\mathbf{l}; \mathbf{x}) f(\mathbf{l}; \mathbf{x}^\theta)];$$

where $f(\mathbf{l}; \mathbf{x}) = e^{i\mathbf{l}^\top \mathbf{x}};$ (6)

It is guaranteed that $f(\mathbf{l}; \mathbf{x}) f(\mathbf{l}; \mathbf{x}^\theta)$ is an unbiased estimation of $k(\mathbf{x}; \mathbf{x}^\theta)$ with sufficient RFF bases f, g drawn from $p(\cdot)$ (Rahimi & Recht, 2007).

For a predefined kernel, e.g., radius basis function (RBF), we use Monte Carlo sampling to draw bases from the spectral distribution, which gives rise to the explicit feature map:

$$\mathbf{z}(\mathbf{x}) = \frac{1}{D} [\cos(\mathbf{l}_1^\top \mathbf{x} + b_1); \dots; \cos(\mathbf{l}_D^\top \mathbf{x} + b_D)]; \quad (7)$$

where $f, g; \mathbf{l}_1; \dots; \mathbf{l}_D g$ are the random bases sampled from $p(\cdot)$, and $[b_1; \dots; b_D]$ are D biases sampled from a uniform distribution with a range of $[0; 2\pi]$. Finally, the kernel values $k(\mathbf{x}; \mathbf{x}^\theta) = \mathbf{z}(\mathbf{x})\mathbf{z}(\mathbf{x}^\theta)^{\triangleright}$ in K are computed as the dot product of their random feature maps with the same bases.

3. Meta Variational Random Features

We introduce our MetaVRF using a latent variable model in which we treat random Fourier bases as latent variables inferred from data. Learning kernels with random Fourier features is tantamount to finding the posterior distribution over random bases in a data-driven way. It is naturally cast into a variational inference problem, where the optimization objective is derived from an evidence lower bound (ELBO) under the meta-learning framework.

3.1. Evidence Lower Bound

From a probabilistic perspective, under the meta-learning setting for few-shot learning, the random feature basis

is obtained by maximizing the conditional predictive log-likelihood of samples from the query set \mathcal{Q} .

$$\begin{aligned} & \max_p \prod_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}} \log p(\mathbf{y}|\mathbf{x}; S) \\ & = \max_p \prod_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}} \int \log p(\mathbf{y}|\mathbf{x}; S; \mathbf{!}) p(\mathbf{!}|\mathbf{x}; S) d\mathbf{!} : \end{aligned} \quad (8)$$

We adopt a conditional prior distribution $p(\mathbf{!}|\mathbf{x}; S)$ over the base $\mathbf{!}$ as in the conditional variational auto-encoder (CVAE) (Sohn et al., 2015) rather than an uninformative prior (Kingma & Welling, 2013; Rezende et al., 2014). By depending on the input \mathbf{x} , we infer the bases that can specifically represent the data, while leveraging the context of the current task by conditioning on the support set S .

In order to infer the posterior $p(\mathbf{!}|\mathbf{y}; \mathbf{x}; S)$ over $\mathbf{!}$, which is generally intractable, we resort to using a variational distribution $q(\mathbf{!}|S)$ to approximate it, where the base is conditioned on the support set S by leveraging meta-learning. We obtain the variational distribution by minimizing the Kullback-Leibler (KL) divergence

$$D_{\text{KL}}[q(\mathbf{!}|S) \text{jj} p(\mathbf{!}|\mathbf{y}; \mathbf{x}; S)] : \quad (10)$$

By applying the Bayes' rule to the posterior $p(\mathbf{!}|\mathbf{y}; \mathbf{x}; S)$, we derive the ELBO as

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x}; S) &= E_{q(\mathbf{!}|S)} \log p(\mathbf{y}|\mathbf{x}; S; \mathbf{!}) \\ &\quad - D_{\text{KL}}[q(\mathbf{!}|S) \text{jj} p(\mathbf{!}|\mathbf{x}; S)] : \end{aligned} \quad (11)$$

The first term of the ELBO is the predictive log-likelihood conditioned on the observation \mathbf{x} , S and the inferred RFF bases $\mathbf{!}$. Maximizing it enables us to make an accurate prediction for the query set by utilizing the inferred bases from the support set. The second term in the ELBO minimizes the discrepancy between the meta variational distribution $q(\mathbf{!}|S)$ and the meta prior $p(\mathbf{!}|\mathbf{x}; S)$, which encourages samples from the support and query sets to share the same random Fourier bases. The full derivation of the ELBO is provided in the supplementary material.

We now obtain the objective by maximizing the ELBO with respect to a batch of T tasks:

$$\begin{aligned} L &= \frac{1}{T} \sum_{t=1}^T \prod_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t} E_{q(\mathbf{!}^t|S^t)} \log p(\mathbf{y}|\mathbf{x}; S^t; \mathbf{!}^t) \\ &\quad - D_{\text{KL}}[q(\mathbf{!}^t|S^t) \text{jj} p(\mathbf{!}^t|\mathbf{x}; S^t)] ; \end{aligned} \quad (12)$$

where S^t is the support set of the t -th task associated with its specific bases $\mathbf{!}^t$ $g_{d=1}^D$ and $(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t$ is the sample from the query set of the t -th task. Directly optimizing the above objective does not take into account the task dependency. Thus, we introduce context inference by conditioning the posterior on both the support set of the current task and the shared knowledge extracted from previous tasks.




Figure 2. Illustration of MetaVRF in a directed graphical model, where (\mathbf{x}, \mathbf{y}) is a sample in the query set \mathcal{Q}^t . The base $\mathbf{!}^t$ of the t -th task is dependent on the support set S^t of the current task and the context \mathcal{C} of related tasks.

3.2. Context Inference

We propose a context inference which puts the inference of random feature bases for the current task in the context of related tasks. We replace the variational distribution in (10) with a conditional distribution $q(\mathbf{!}^t|S^t; \mathcal{C})$ that makes the bases $\mathbf{!}^t$ $g_{d=1}^D$ of the current t -th task conditioned also on the context \mathcal{C} of related tasks.

The context inference gives rise to a new ELBO, as follows:

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x}; S^t) &= E_{q(\mathbf{!}^t|S^t; \mathcal{C})} \log p(\mathbf{y}|\mathbf{x}; S^t; \mathbf{!}^t) \\ &\quad - D_{\text{KL}}[q(\mathbf{!}^t|S^t; \mathcal{C}) \text{jj} p(\mathbf{!}^t|\mathbf{x}; S^t)] : \end{aligned} \quad (13)$$

which can be represented in a directed graphical model as shown in Figure 2. In a practical sense, the KL term in (13) encourages the model to extract useful information from previous tasks for inferring the spectral distribution associated with each individual sample \mathbf{x} of the query set in the current task.

The context inference integrates the knowledge shared across tasks with the task-specific knowledge to build up adaptive kernels for individual tasks. The inferred random features are highly informative due to the absorbed information from prior knowledge of experienced tasks. The base-learner built on the inferred kernel with the informative random features can effectively solve the current task.

However, since there is usually a huge number of related tasks, it is non-trivial to model all these tasks simultaneously. We consider using recurrent neural networks to gradually accumulate information episodically along with the learning process by organizing tasks in a sequence. We propose an LSTM-based inference network by leveraging its innate capability of remembering long-term information (Gers & Schmidhuber, 2000). The LSTM offers a well-suited structure to implement the context inference. The cell state \mathbf{c} stores and accrues the meta knowledge shared among related tasks, which can also be updated when experiencing a new task in each episode during the course of learning; the output \mathbf{h} is used to adapt to each specific task.

To be more specific, we model the variational posterior $q(\mathbf{h}^t; S^t; \mathcal{C})$ through $q(\mathbf{h}^t)$ which is parameterized as a multi-layer perceptron (MLP) (\mathbf{h}^t). Note that \mathbf{h}^t is the output from an LSTM that takes S^t and \mathcal{C} as inputs. We implement the inference network with both vanilla and bidirectional LSTMs (Schuster & Paliwal, 1997; Graves & Schmidhuber, 2005). For a vanilla LSTM, we have

$$[\mathbf{h}^t; \mathbf{c}^t] = g_{\text{LSTM}}(S^t; \mathbf{h}^{t-1}; \mathbf{c}^{t-1}); \quad (14)$$

where $g_{\text{LSTM}}(\cdot)$ is a vanilla LSTM network that takes the current support set, the output \mathbf{h}^{t-1} and the cell state \mathbf{c}^{t-1} as the input. S^t is the average over the feature representation vectors of samples in the support set (Zaheer et al., 2017). The feature representation is obtained by a shared convolutional network (\cdot). To incorporate more context information, we also implement the inference with a bidirectional LSTM, and we have $\mathbf{h}^t = [\mathbf{h}^t; \mathbf{h}^t]$, where \mathbf{h}^t and \mathbf{h}^t are the outputs from forward and backward LSTMs, respectively, and $[\cdot; \cdot]$ indicates a concatenation operation.

Therefore, the optimization objective with the context inference is:

$$L = \frac{1}{T} \sum_{t=1}^T \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t} \mathbb{E}_{q(\mathbf{h}^t)} \log p(\mathbf{y}|\mathbf{x}; S^t; \mathbf{h}^t) \quad (15)$$

$$D_{\text{KL}}[q(\mathbf{h}^t) \parallel p(\mathbf{h}^t; S^t)] :$$

where the variational approximate posterior $q(\mathbf{h}^t)$ is taken as a multivariate Gaussian with a diagonal co-variance. Given the support set as input, the mean μ and standard deviation σ are output from the inference network (\cdot). The conditional prior $p(\mathbf{h}^t; S^t)$ is implemented with a prior network which takes an aggregated representation by using the cross attention (Kim et al., 2019) between \mathbf{x} and S^t . The details of the prior network are provided in the supplementary material. To enable back-propagation with the sampling operation during training, we adopt the reparametrization trick (Rezende et al., 2014; Kingma & Welling, 2013) as $\mathbf{h}^t = \mu + \sigma \cdot \mathbf{z}$, where $\mathbf{z} \sim \mathcal{N}(0; 1)$:

During the course of learning, the LSTMs accumulate knowledge in the cell state by updating their cells using information extracted from each task. For the current task t , the knowledge stored in the cell is combined with the task-specific information from the support set to infer the spectral distribution for this task. To accrue the information across all the tasks in the meta-training set, the output and the cell state of LSTMs are passed down across batches. As a result, the final cell state contains the distilled prior knowledge from all those experienced tasks in the meta-training set.

Fast Adaptation. Once meta-training ceases, the output and the cell state are directly used for a new incoming task in the meta-test set to achieve fast adaptation with a simple

feed-forward computation operation. To be more specific, for a task with the support set S in the meta-test set, we draw D samples $f^{(d)} g_{f=1}^D$ as the bases: $f^{(d)} = q(\mathbf{h}^t)$, where \mathbf{h}^t is output from either a vanilla LSTM or a bidirectional LSTM, depending on which is used during the meta-training stage. The bases are adopted to compute the kernels on the support set and construct the classifier of the base-learner for the task, using (4). The classifier is then used to make predictions of samples in the query set for performance evaluation.

4. Related Work

Meta-learning, or learning to learn, endues machine learning models the ability to improve their performance by leveraging knowledge extracted from a number of prior tasks. It has received increasing research interest with breakthroughs in many directions (Finn et al., 2017; Rusu et al., 2019; Gordon et al., 2019; Aravind Rajeswaran, 2019). Gradient-based methods (e.g., MAML (Finn et al., 2017)) learn an appropriate initialization of model parameters and adapt it to new tasks with only a few gradient steps (Finn & Levine, 2018; Zintgraf et al., 2019; Rusu et al., 2019). Learning a shared optimization algorithm has also been explored in order to quickly learn of new tasks (Ravi & Larochelle, 2017; Andrychowicz et al., 2016; Chen et al., 2017).

Metric learning has been widely studied with great success for few-shot learning (Vinyals et al., 2016; Snell et al., 2017; Satorras & Estrach, 2018; Oreshkin et al., 2018; Allen et al., 2019). The basic assumption is that a common metric space is shared across related tasks. By extending the matching network (Vinyals et al., 2016) to few-shot scenarios, Snell et al. (Snell et al., 2017) constructed a prototype for each class by averaging feature representations of samples from the class in the metric space. The classification is conducted via matching the query samples to prototypes by computing their distances. To enhance the prototype representation, Allen et al. (Allen et al., 2019) proposed an infinite mixture of prototypes (IMP) to adaptively represent data distribution of each class by multiple clusters instead of using a single vector. In addition, Oreshkin et al. (Oreshkin et al., 2018) proposed task dependent adaptive metric for improved few-shot learning and established prototypes of classes conditioning on a task representation encoded by a task embedding network. Their results indicate the benefit of learning task-specific metric in few-shot learning.

While these meta-learning algorithms have made great progress in few-shot learning tasks, exploring prior knowledge from previous tasks remains an open challenge (Titsias et al., 2019). In this work, we introduce kernels based on random features as the base-learners, which enables us to acquire shared knowledge across tasks by modeling their dependency via the random feature basis of kernels.

Kernel learning with random Fourier features is a versatile and powerful tool in machine learning (Bishop, 2006; Hofmann et al., 2008; Shervashidze et al., 2011). Pioneering works (Bach et al., 2004; Gönen & Alpaydm, 2011; Duvenaud et al., 2013) learn to combine predefined kernels in a multi-kernel learning manner. Kernel approximation by random Fourier features (RFFs) (Rahimi & Recht, 2007) is an effective technique for efficient kernel learning (Gärtner et al., 2002), which has recently become increasingly popular (Sinha & Duchi, 2016; Carratino et al., 2018). Recent works (Wilson & Adams, 2013) learn kernels in the frequency domain by modeling the spectral distribution as a mixture of Gaussians and computing its optimal linear combination. Instead of modeling the spectral distribution with explicit density functions, other works focus on optimizing the random base sampling strategy (Yang et al., 2015; Sinha & Duchi, 2016). Nonetheless, it has been shown that accurate approximation of kernels does not necessarily result in high classification performance (Avron et al., 2016; Chang et al., 2017). This suggests that learning adaptive kernels with random features by data-driven sampling strategies (Sinha & Duchi, 2016) can improve the performance, even with a low sampling rate compared to using universal random features (Avron et al., 2016; Chang et al., 2017).

Our MetaVRF is the first work to introduce kernel learning with random features to the meta-learning framework for few-shot learning. The optimization of MetaVRF is naturally cast as a variational inference and the context inference offers a principled way to incorporate prior knowledge and achieve informative and adaptive kernels.

5. Experiments

We evaluate our MetaVRF on several few-shot learning problems for both regression and classification. We demonstrate the benefit of exploring task dependency by implementing a baseline MetaVRF (12) without using the LSTM, which infers the random base solely from the support set. We also conduct further analysis to validate the effectiveness of our MetaVRF by showing its performance with deep embedding architectures, different numbers of bases, and under versatile and challenging settings with inconsistent training and test conditions.

5.1. Few-Shot Regression

We conduct regression tasks with different numbers of shots k , and compare our MetaVRF with MAML (Finn et al., 2017), a representative meta-learning algorithm. We follow the MAML work (Finn et al., 2017) to fit a target sine function $y = A \sin(wx + b)$, with only a few annotated samples. $A \in [0.1; 5]$, $w \in [0.8; 1.2]$, and $b \in [0; \pi]$ denote the amplitude, frequency, and phase, respectively, which follow a uniform distribution within the corresponding interval.




Figure 3. Performance (MSE) comparison for few-shot regression. Our MetaVRF fits the target function well, even with only three shots, and consistently outperforms regular RFFs and the counterpart MAML. (--- MetaVRF with bi-LSTM; --- MetaVRF with LSTM; --- MetaVRF w/o LSTM; --- MAML; — Ground Truth; ▲ Support Samples.)

The goal is to estimate the target sine function given only n randomly sampled data points. In our experiments, we consider the input in the range of $x \in [-5; 5]$, and conduct three tests under the conditions of $k = 3; 5; 10$. For a fair comparison, we compute the feature embedding using a small multi-layer perceptron (MLP) with two hidden layers of size 40, following the same settings used in MAML.

The results in Figure 3 show that our MetaVRF fits the function well with only three shots but performs better with an increasing number of shots, almost entirely fitting the target function with ten shots. Moreover, the results demonstrate the advantage of exploring task dependency by LSTM-based inference. MetaVRF with bi-LSTM performs better than regular LSTM since more context tasks are incorporated by bi-LSTM. In addition, we observe that MetaVRF performs better than MAML for all three settings with varying numbers of shots. We provide more results on few-shot regression tasks in the supplementary material.

5.2. Few-Shot Classification

The classification experiments are conducted on three commonly-used benchmark datasets, i.e., Omniglot (Lake et al., 2015), miniImageNet (Vinyals et al., 2016) and CIFAR-FS (Krizhevsky et al., 2009); for more details, please refer to the supplementary material. We extract image features using a shallow convolutional neural network with the same architecture as in (Gordon et al., 2019). We do not use any fully connected layers for these CNNs. The dimension of all feature vectors is 256. We also evaluate the random Fourier features (RFFs) and the radial basis function (RBF) kernel in which we take the bandwidth as the mean of pair-wise distances between samples in the support set of each task. The inference network () is a three-layer MLP with 256 units in the hidden layers and rectifier non-linearity where input sizes are 256 and 512 for the vanilla and bidirectional LSTMs, respectively.

Table 1. Performance (%) on *miniImageNet* and CIFAR-FS.

Method	<i>miniImageNet</i> , 5-way		CIFAR-FS, 5-way	
	1-shot	5-shot	1-shot	5-shot
MATCHING NET (Vinyals et al., 2016)	44.2	57	—	—
MAML (Finn et al., 2017)	48.7 1.8	63.1 0.9	58.9 1.9	71.5 1.0
MAML (64C)	46.7 1.7	61.1 0.1	58.9 1.8	71.5 1.1
META-LSTM (Ravi & Larochelle, 2017)	43.4 0.8	60.6 0.7	—	—
PROTO NET (Snell et al., 2017)	47.4 0.6	65.4 0.5	55.5 0.7	72.0 0.6
RELATION NET (Sung et al., 2018)	50.4 0.8	65.3 0.7	55.0 1.0	69.3 0.8
SNAIL (32C) by (Bertinetto et al., 2019)	45.1	55.2	—	—
GNN (Garcia & Bruna, 2018)	50.3	66.4	61.9	75.3
PLATIPUS (Finn et al., 2018)	50.1 1.9	—	—	—
VERSA (Gordon et al., 2019)	53.3 1.8	67.3 0.9	62.5 1.7	75.1 0.9
R2-D2 (64C) (Bertinetto et al., 2019)	49.5 0.2	65.4 0.2	62.3 0.2	77.4 0.2
R2-D2 (Devos et al., 2019)	51.7 1.8	63.3 0.9	60.2 1.8	70.9 0.9
CAVIA (Zintgraf et al., 2019)	51.8 0.7	65.6 0.6	—	—
IMAML (Aravind Rajeswaran, 2019)	49.3 1.9	—	—	—
RBF KERNEL	42.1 1.2	54.9 1.1	46.0 1.2	59.8 1.0
RFFs (2048d)	52.8 0.9	65.4 0.9	61.1 0.8	74.7 0.9
METAVRF (w/o LSTM, 780d)	51.3 0.8	66.1 0.7	61.1 0.7	74.3 0.9
METAVRF (vanilla LSTM, 780d)	53.1 0.9	66.8 0.7	62.1 0.8	76.0 0.8
METAVRF (bi-LSTM, 780d)	54.2 0.8	67.8 0.7	63.1 0.7	76.5 0.9

The key hyperparameter for the number of bases D in (7) is set to $D = 780$ for MetaVRF in all experiments, while we use RFFs with $D = 2048$ as this produces the best performance. The sampling rate in our MetaVRF is much lower than in previous works using RFFs, in which D is usually set to be 5 to 10 times the dimension of the input features (Yu et al., 2016; Rahimi & Recht, 2007). We adopt a similar meta-testing protocol as (Gordon et al., 2019; Finn et al., 2017), but we test on 3,000 episodes rather than 600 and present the results with 95% confidence intervals. All reported results are produced by models trained from scratch. We compare with previous methods that use the same training procedures and similar shallow conventional CNN architectures as ours. The comparison results on three benchmark datasets are reported in Tables 1 and 2.

On all benchmark datasets, MetaVRF delivers the state-of-the-art performance. Even with a relatively low sampling rate, MetaVRF produces consistently better performance compared with RBF kernels, RFFs. MetaVRF with bi-LSTM outperforms the one with vanilla LSTM since it can leverage more information. It is worth mentioning that MetaVRF with bi-LSTM achieves good performance (54.2%) under the 5-way 1-shot setting on the *miniImageNet* dataset, surpassing the second best model by 1%. The MetaVRFs with bi-LSTM and vanilla LSTM consistently outperform the one without the LSTM, which demonstrates the effectiveness of using LSTM to explore task dependency. The much better performance over RBF kernels on all three datasets indicates the great benefit of adaptive kernels based on random Fourier features. The relatively inferior performance produced by RBF kernels is due to that

the mean of pair-wise distances of support samples would not be able to provide a proper estimate of the kernel bandwidth since we only have a few samples in each task, for instance, five samples under the 5-way 1-shot setting in the support set.

Note that on Omniglot, the performance of existing methods saturates and MetaVRF with bi-LSTM achieves the best performance for most settings, including 5-way 1-shot, 5-way 5-shot, and 20-way 1-shot. It is also competitive under the 20-way 5-shot setting falling within the error bars of the state-of-the-arts. In Table 1, we also implement a MAML (64C) with 64 channels in each convolutional layer. However, while it obtains modest performance, we believe the increased model size leads to overfitting. Since in the original SNAIL, a very deep ResNet-12 network is used for embedding, we cite the result of SNAIL reported in Bertinetto et al. (2019) using similar shallow networks as ours. For fair comparison, we also cite the original results of R2-D2 (Bertinetto et al., 2019) using 64 channels.

5.3. Further Analysis

Deep embedding. Our MetaVRF is independent of the convolutional architectures for feature extraction and can work with deeper embeddings either pre-trained or trained from scratch. In general, the performance improves with more powerful feature extraction architectures. We evaluate our method using pre-trained embeddings in order to compare with existing methods using deep embedding architectures. To benchmark with those methods, we adopt the pre-trained embeddings from a 28-layer wide residual network (WRN-28-10) (Zagoruyko & Komodakis, 2016), in a similar fash-

Table 2. Performance (%) on Omniglot.

Method	Omniglot, 5-way		Omniglot, 20-way	
	1-shot	5-shot	1-shot	5-shot
SIAMESE NET (Koch, 2015)	96.7	98.4	88	96.5
MATCHING NET (Vinyals et al., 2016)	98.1	98.9	93.8	98.5
MAML (Finn et al., 2017)	98.7 0.4	99.9 0.1	95.8 0.3	98.9 0.2
PROTO NET (Snell et al., 2017)	98.5 0.2	99.5 0.1	95.3 0.2	98.7 0.1
SNAIL (Mishra et al., 2018)	99.1 0.2	99.8 0.1	97.6 0.3	99.4 0.2
GNN (Garcia & Bruna, 2018)	99.2	99.7	97.4	99.0
VERSA (Gordon et al., 2019)	99.7 0.2	99.8 0.1	97.7 0.3	98.8 0.2
R2-D2 (Bertinetto et al., 2019)	98.6	99.7	94.7	98.9
IMP (Allen et al., 2019)	98.4 0.3	99.5 0.1	95.0 0.1	98.6 0.1
RBF KERNEL	95.5 0.2	99.1 0.2	92.8 0.3	97.8 0.2
RFFs (2048d)	99.5 0.2	99.5 0.2	97.2 0.3	98.3 0.2
META VRF (w/o LSTM, 780d)	99.6 0.2	99.6 0.2	97.0 0.3	98.4 0.2
META VRF (vanilla LSTM, 780d)	99.7 0.2	99.8 0.1	97.5 0.3	99.0 0.2
META VRF (bi-LSTM, 780d)	99.8 0.1	99.9 0.1	97.8 0.3	99.2 0.2

Table 3. Performance (%) on *miniImageNet* (5-way)

Method	1-shot	5-shot
META-SGD (Li et al., 2017)	54.24 0.03	70.86 0.04
(Gidaris & Komodakis, 2018)	56.20 0.86	73.00 0.64
(Bauer et al., 2017)	56.30 0.40	73.90 0.30
(Munkhdalai et al., 2017)	57.10 0.70	70.04 0.63
(Qiao et al., 2018)	59.60 0.41	73.54 0.19
LEO (Rusu et al., 2019)	61.76 0.08	77.59 0.12
SNAIL (Mishra et al., 2018)	55.71 0.99	68.88 0.92
TADAM (Oreshkin et al., 2018)	58.50 0.30	76.70 0.30
META VRF (w/o LSTM, 780d)	62.12 0.07	77.05 0.28
META VRF (vanilla LSTM, 780d)	63.21 0.06	77.83 0.28
META VRF (bi-LSTM, 780d)	63.80 0.05	77.97 0.28

ion to (Rusu et al., 2019; Bauer et al., 2017; Qiao et al., 2018). We choose activations in the 21-st layer, with average pooling over spatial dimensions, as feature embeddings. The dimension of pre-trained embeddings is 640. We show the comparison results on the *miniImageNet* dataset for 5-way 1-shot and 5-shot settings in Table 3. Our MetaVRF with bi-LSTM achieves the best performance under both settings and largely surpasses LEO, a recently proposed meta-learning method, especially on the challenging 5-way 1-shot setting. Note that the MetaVRF with vanilla LSTM and without LSTM also produce competitive performance.

Efficiency. Regular random Fourier features (RFFs) usually require high sampling rates to achieve satisfactory performance. However, our MetaVRF can achieve high performance with a relatively low sampling rate compared, which guarantees its high efficiency. In Figure 4, we compare with regular RFFs using different sampling rates. We show the performance change of fully trained models using RFFs and our MetaVRF with bi-LSTM under a different number D of bases. We show the comparison results for the 5-way 5-shot setting in Figure 4. MetaVRF with bi-LSTM consistently




Figure 4. Performance with different numbers D of bases. Our MetaVRF consistently achieves better performance than regular RFFs, especially with relatively low sampling rates.




Figure 5. Performance with varied ways and shots on Omniglot.

tently yields higher performance than regular RFFs with the same number of sampled bases. The results verify the efficiency of our MetaVRF in learning adaptive kernels and the effectiveness in improving performance by exploring dependencies of related tasks.

Versatility. In contrast to most existing meta-learning methods, our MetaVRF can be used for versatile settings. We evaluate the performance of MetaVRF on more challenging scenarios where the number of ways C and shots k between training and testing are inconsistent. Specifically, we test the performance of MetaVRF on tasks with varied C and k , when it is trained on one particular C -way- k -shot task. As shown in Figure 5, the results demonstrate that the

trained model can still produce good performance, even on the challenging condition with a far higher number of ways. In particular, the model trained on the 20-way-5-shot task can retain a high accuracy of 94% on the 100-way setting, as shown in Figure 5(a). The results also indicate that our MetaVRF exhibits considerable robustness and flexibility to a great variety of testing conditions.

6. Conclusion

In this paper, we introduce kernel approximation based on random Fourier features into the meta-learning framework for few-shot learning. We propose meta variational random features (MetaVRF), which leverage variational inference and meta-learning to infer the spectral distribution of random Fourier features in a data-driven way. MetaVRF generates random Fourier features of high representational power with a relatively low spectral sampling rate by using an LSTM based inference network to explore the shared knowledge. In practice, our LSTM-based inference network demonstrates a great ability to collect and absorb common knowledge from experience tasks, which enables quick adaptation to specific tasks for improved performance. Experimental results on both regression and classification tasks demonstrate the effectiveness for few-shot learning.

Acknowledgements

This research was supported in part by Natural Science Foundation of China (No. 61976060, 61871016, 61876098).

References

- Allen, K. R., Shelhamer, E., Shin, H., and Tenenbaum, J. B. Infinite mixture prototypes for few-shot learning. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 232–241, 2019.
- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., Shillingford, B., and de Freitas, N. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, 2016.
- Aravind Rajeswaran, Chelsea Finn, S. K. S. L. Meta-learning with implicit gradients. *arXiv preprint arXiv:1909.04630*, 2019.
- Avron, H., Sindhvani, V., Yang, J., and Mahoney, M. W. Quasi-monte carlo feature maps for shift-invariant kernels. *The Journal of Machine Learning Research*, 17(1):4096–4133, 2016.
- Bach, F. R., Lanckriet, G. R., and Jordan, M. I. Multiple kernel learning, conic duality, and the smo algorithm. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 6, 2004.
- Bauer, M., Rojas-Carulla, M., Świątkowski, J. B., Schölkopf, B., and Turner, R. E. Discriminative k-shot learning using probabilistic models. *arXiv preprint arXiv:1706.00326*, 2017.
- Bertinetto, L., Henriques, J. F., Torr, P. H., and Vedaldi, A. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019.
- Bishop, C. M. *Pattern recognition and machine learning*. springer, 2006.
- Bullins, B., Zhang, C., and Zhang, Y. Not-so-random features. In *International Conference on Learning Representations*, 2018.
- Carratino, L., Rudi, A., and Rosasco, L. Learning with sgd and random features. In *Advances in Neural Information Processing Systems*, pp. 10192–10203, 2018.
- Chang, W.-C., Li, C.-L., Yang, Y., and Poczos, B. Data-driven random fourier features using stein effect. *arXiv preprint arXiv:1705.08525*, 2017.
- Chen, Y., Hoffman, M. W., Colmenarejo, S. G., Denil, M., Lillicrap, T. P., Botvinick, M., and De Freitas, N. Learning to learn without gradient descent by gradient descent. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 748–756. JMLR. org, 2017.
- Devos, A., Chatel, S., and Grossglauser, M. Reproducing meta-learning with differentiable closed-form solvers. In *ICLR Workshop*, 2019.
- Duvenaud, D., Lloyd, J. R., Grosse, R., Tenenbaum, J. B., and Ghahramani, Z. Structure discovery in nonparametric regression through compositional kernel search. *arXiv preprint arXiv:1302.4922*, 2013.
- Finn, C. and Levine, S. Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm. In *International Conference on Learning Representations*, 2018.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pp. 1126–1135. JMLR. org, 2017.
- Finn, C., Xu, K., and Levine, S. Probabilistic model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, pp. 9516–9527, 2018.

- Garcia, V. and Bruna, J. Few-shot learning with graph neural networks. In *International Conference on Learning Representations*, 2018.
- Gärtner, T., Flach, P. A., Kowalczyk, A., and Smola, A. J. Multi-instance kernels. In *International Conference on Machine Learning*, 2002.
- Gers, F. A. and Schmidhuber, J. Recurrent nets that time and count. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks*, volume 3, pp. 189–194. IEEE, 2000.
- Gidaris, S. and Komodakis, N. Dynamic few-shot visual learning without forgetting. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4367–4375, 2018.
- Gönen, M. and Alpaydm, E. Multiple kernel learning algorithms. *Journal of machine learning research*, 12(Jul): 2211–2268, 2011.
- Gordon, J., Bronskill, J., Bauer, M., Nowozin, S., and Turner, R. E. Meta-learning probabilistic inference for prediction. In *International Conference on Learning Representations*, 2019.
- Graves, A. and Schmidhuber, J. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural networks*, 18(5-6):602–610, 2005.
- Hensman, J., Durrande, N., and Solin, A. Variational fourier features for gaussian processes. *The Journal of Machine Learning Research*, 18(1):5537–5588, 2017.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hofmann, T., Schölkopf, B., and Smola, A. J. Kernel methods in machine learning. *The annals of statistics*, pp. 1171–1220, 2008.
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. In *International Conference on Learning Representations*, 2019.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Koch, G. Siamese neural networks for one-shot image recognition. In *ICML Workshop*, 2015.
- Krizhevsky, A. et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Li, C.-L., Chang, W.-C., Mroueh, Y., Yang, Y., and Poczos, B. Implicit kernel learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2007–2016, 2019.
- Li, Z., Zhou, F., Chen, F., and Li, H. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- Mishra, N., Rohaninejad, M., Chen, X., and Abbeel, P. A simple neural attentive meta-learner. In *International Conference on Learning Representations*, 2018.
- Munkhdalai, T., Yuan, X., Mehri, S., and Trischler, A. Rapid adaptation with conditionally shifted neurons. *arXiv preprint arXiv:1712.09926*, 2017.
- Oreshkin, B., López, P. R., and Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 721–731, 2018.
- Qiao, S., Liu, C., Shen, W., and Yuille, A. L. Few-shot image recognition by predicting parameters from activations. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7229–7238, 2018.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pp. 1177–1184, 2007.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Rudin, W. *Fourier analysis on groups*, volume 121967. Wiley Online Library, 1962.
- Rusu, A. A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., and Hadsell, R. Meta-learning with latent embedding optimization. In *International Conference on Learning Representations*, 2019.
- Satorras, V. G. and Estrach, J. B. Few-shot learning with graph neural networks. In *International Conference on Learning Representations*, 2018.
- Schmidhuber, J. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992.
- Schuster, M. and Paliwal, K. K. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681, 1997.

- Shervashidze, N., Schweitzer, P., Leeuwen, E. J. v., Mehlhorn, K., and Borgwardt, K. M. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(Sep):2539–2561, 2011.
- Sinha, A. and Duchi, J. C. Learning kernels with random features. In *Advances in Neural Information Processing Systems*, pp. 1298–1306, 2016.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 4077–4087, 2017.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pp. 3483–3491, 2015.
- Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P. H., and Hospedales, T. M. Learning to compare: Relation network for few-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1199–1208, 2018.
- Thrun, S. and Pratt, L. *Learning to learn*. Springer Science & Business Media, 2012.
- Titsias, M. K., Schwarz, J., Matthews, A. G. d. G., Pascanu, R., and Teh, Y. W. Functional regularisation for continual learning using gaussian processes. *arXiv preprint arXiv:1901.11356*, 2019.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pp. 3630–3638, 2016.
- Wilson, A. and Adams, R. Gaussian process kernels for pattern discovery and extrapolation. In *International Conference on Machine Learning*, pp. 1067–1075, 2013.
- Yang, Z., Wilson, A., Smola, A., and Song, L. A la carte-learning fast kernels. In *Artificial Intelligence and Statistics*, pp. 1098–1106, 2015.
- Yu, F. X. X., Suresh, A. T., Choromanski, K. M., Holtmann-Rice, D. N., and Kumar, S. Orthogonal random features. In *Advances in Neural Information Processing Systems*, pp. 1975–1983, 2016.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in Neural Information Processing Systems*, pp. 3391–3401, 2017.
- Zintgraf, L., Shiarli, K., Kurin, V., Hofmann, K., and Whiteson, S. Fast context adaptation via meta-learning. In *International Conference on Machine Learning*, pp. 7693–7702, 2019.