



UvA-DARE (Digital Academic Repository)

MetaKernel: Learning Variational Random Features With Limited Labels

Du, Y.; Sun, H.; Zhen, X.; Xu, J.; Yin, Y.; Shao, L.; Snoek, C.G.M.

DOI

[10.1109/TPAMI.2022.3154930](https://doi.org/10.1109/TPAMI.2022.3154930)

Publication date

2024

Document Version

Final published version

Published in

IEEE Transactions on Pattern Analysis and Machine Intelligence

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Du, Y., Sun, H., Zhen, X., Xu, J., Yin, Y., Shao, L., & Snoek, C. G. M. (2024). MetaKernel: Learning Variational Random Features With Limited Labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(3), 1464-1478.
<https://doi.org/10.1109/TPAMI.2022.3154930>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

MetaKernel: Learning Variational Random Features With Limited Labels

Yingjun Du [✉], Haoliang Sun [✉], Xiantong Zhen [✉], Jun Xu [✉], Yilong Yin [✉],
Ling Shao [✉], *Fellow, IEEE*, and Cees G. M. Snoek [✉], *Senior Member, IEEE*

Abstract—Few-shot learning deals with the fundamental and challenging problem of learning from a few annotated samples, while being able to generalize well on new tasks. The crux of 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. In this paper, we propose meta-learning kernels with random Fourier features for few-shot learning, we call MetaKernel. Specifically, we propose learning variational random features in a data-driven manner to obtain task-specific kernels by leveraging the shared knowledge provided by related tasks in a meta-learning setting. We treat the random feature basis as the latent variable, which is estimated by variational inference. The shared knowledge from related tasks is incorporated into a context inference of the posterior, which we achieve via a long-short term memory module. To establish more expressive kernels, we deploy conditional normalizing flows based on coupling layers to achieve a richer posterior distribution over random Fourier bases. The resultant kernels are more informative and discriminative, which further improves the few-shot learning. To evaluate our method, we conduct extensive experiments on both few-shot image classification and regression tasks. A thorough ablation study demonstrates that the effectiveness of each introduced component in our method. The benchmark results on fourteen datasets demonstrate MetaKernel consistently delivers at least comparable and often better performance than state-of-the-art alternatives.

Index Terms—Meta learning, few-shot learning, normalizing flow, variational inference, random features

1 INTRODUCTION

HUMANS have the amazing ability to learn new concepts from only a few examples, and then effortlessly generalize this knowledge to new samples. In contrast, despite considerable progress, existing image classification models based on deep neural networks e.g., [2], [3], are still highly dependent on large amounts of annotated training data [4] to achieve satisfactory performance. This learnability gap between human intelligence and existing neural networks has motivated many to study learning from a few samples,

e.g., [5], [6], [7], [8]. Meta-learning, *a.k.a.* learning to learn [9], [10], emerged as a promising direction for few-shot learning [1], [7], [8], [11].

The working mechanism of meta-learning involves a meta-learner that exploits the common knowledge from various tasks to improve the performance of each individual task. Remarkable success has been achieved in learning good parameter initializations [8], [12], efficient optimization update rules [7], [11], and powerful common metrics [13], [14] from related tasks, which enables fast adaptation to new tasks with few training samples. Meta-learning has also proven to be effective in learning amortized networks shared by related tasks, which generate specific parameters [15] or normalization statistics [16] for individual few-shot learning tasks. However, how to properly define and exploit the prior knowledge from experienced tasks remains an open problem for few-shot learning, and is the one we address in this paper.

An effective base-learner should be powerful enough to solve individual tasks, while being able to absorb the information provided by the meta-learner for overall benefit. Kernels [17], [18], [19] have proven to be a powerful technique in the machine learning toolbox, e.g., [20], [21], [22], [23], [24], as they are able to produce strong performance without relying on a large amount of labelled data. Moreover, task-adaptive kernels with random features, leveraging data-driven sampling strategies [23], achieve improved performance over universal ones, at low sampling rates [25], [26], [27], [28]. This makes kernels with data-driven random features well-suited tools for learning tasks with limited data. Hence, we introduce kernels as base-learners into the meta-learning framework for few-shot learning. However, due to the

- Yingjun Du and Cees G. M. Snoek are with the University of Amsterdam, 1012, WX, Amsterdam, The Netherlands. E-mail: {y.du, cgmsnoek}@uva.nl.
- Haoliang Sun and Yilong Yin are with the School of Software, Shandong University, Jinan, Shandong 264209, China. E-mail: haolsun.cn@gmail.com, ylyin@sdu.edu.cn.
- Xiantong Zhen and Ling Shao are with the Inception Institute of Artificial Intelligence, Abu Dhabi 51133, UAE. E-mail: zhenxt@gmail.com, ling.shao@ieee.org.
- Jun Xu is with Nankai University, Tianjin 300071, China. E-mail: nankaimathxujun@gmail.com.

Manuscript received 30 Apr. 2021; revised 8 Nov. 2021; accepted 13 Feb. 2022. Date of publication 28 Feb. 2022; date of current version 5 Feb. 2024.

This work was 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. This work was also supported in part by the Natural Science Foundation of China under Grants 62176068 and 62106129, in part by the Major Basic Research Project of Natural Science Foundation of Shandong Province under Grant ZR2021ZD15, and in part by China Postdoctoral Science Foundation under Grant 2021TQ0195.

(Corresponding author: Yingjun Du.)

Recommended for acceptance by L. Liu, T. Hospedales, Y. Le Cun, M. Long, J. Luo, W. Ouyang, M. Pietikinen, and T. Tuytelaars.

Digital Object Identifier no. 10.1109/TPAMI.2022.3154930

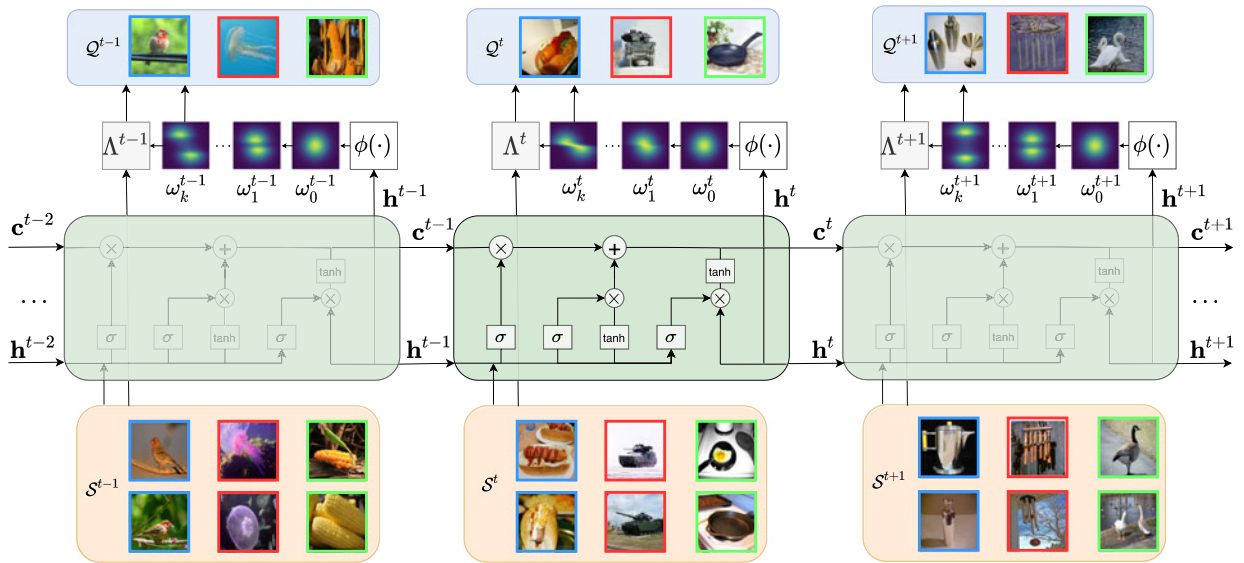


Fig. 1. MetaKernel learning framework. The meta-learner employs an LSTM-based context inference network $\phi(\cdot)$ to infer the spectral distribution over ω'_0 , the kernel from the support set S^t of the current task t , and the outputs h^{t-1} and c^{t-1} of the previous task. The enriched random bases ω'_k are obtained via conditional normalizing flows with a flow of length k . During the learning process, the cell state in the LSTM is deployed to accumulate shared knowledge by 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 the support set is combined with distilled information from the previous tasks to infer the adaptive spectral distribution of the kernels.

limited availability of samples, it is challenging to learn informative random features for few-shot tasks by solely relying on a task's own data. Therefore, exploring the shared prior knowledge from different but related tasks is essential for obtaining richer random features and few-shot learning.

We propose learning task-specific kernels in a data-driven way with variational random features by leveraging the shared knowledge provided by related tasks. To do so, we develop a latent variable model that treats the random Fourier basis of translation-invariant kernels as the latent variable. The posterior over the random feature basis corresponds to the spectral distribution associated with the kernel. The optimization of the model is formulated as a variational inference problem. Kernel learning with random Fourier features for few-shot learning allows us to leverage the universal approximation property of kernels to capture shared knowledge from related tasks. This probabilistic modelling framework provides a principled way of learning data-driven kernels with random Fourier features and, more importantly, fits well into the meta-learning framework for few-shot learning, providing us the flexibility to customize the variational posterior and leverage meta-knowledge to enhance individual tasks.

To incorporate the prior knowledge from experienced tasks, we further propose a context inference scheme to integrate the inference of random feature bases of the current task into the context of previous related tasks. The context inference provides a generalized way to integrate shared knowledge from the related tasks with task-specific information for the inference of random feature bases. To do so, we adopt a long short-term memory (LSTM) based inference network [29], leveraging its capability of learning long-term dependencies to collect and refine the shared meta-knowledge from a set of previously experienced tasks.

A preliminary conference version of this work, which also covers variational random features and task context

inference was published previously [1]. In this extended work, we further propose conditional normalizing flows to infer richer posteriors over the random bases, which allows us to obtain more informative random features. The normalizing flows (NFs) [30], [31], [32], [33], [34] model complicated high dimensional marginal distributions by transforming a simple base distribution (e.g., a standard normal) or priors through a learnable, invertible mapping and then applying the change of variables formula. Normalizing flows, which have not yet been explored in few-shot learning, provide a well-suited technique for learning more expressive random features by transforming a random basis into a richer distribution. While the technical contribution of conditional normalizing flows may appear modest, their conceptual and practical impact on multiple-kernel learning is considerable. Conditional normalizing flows allow us to obtain more complex distributions, which offer much more flexibility compared to the single kernel alternative. The overall learning framework of our MetaKernel is illustrated in Fig. 1.

To validate our method, we conduct extensive experiments on fourteen benchmark datasets for a variety of few-shot learning tasks including image classification and regression. Unlike our prior work [1], we also experiment on the large-scale Meta-Dataset by Triantafillou *et al.* [35] and the challenging few-shot domain generalization setting suggested by Du *et al.* [16]. MetaKernel consistently delivers at least comparable and often better performance than state-of-the-art alternatives on all datasets, and the ablation analysis demonstrates the effectiveness of each MetaKernel component for few-shot learning.

The rest of this paper is organized as follows: Section 2 summarizes related work. Section 3 presents the proposed MetaKernel framework. Section 4 summarizes experimental details, state-of-the-art comparisons and detailed ablation studies. Section 5 closes with concluding remarks.

2 RELATED WORK

2.1 Meta-Learning

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, e.g., [8], [12], [15], [36], [37]. Existing methods can be roughly categorized into four groups.

Models in the first group are based on distance metrics and generally learn a shared or adaptive embedding space in which query images are accurately matched to support images for classification. They rely on the assumption that a common metric space is shared across related tasks and usually do not employ an explicit base-learner for each task. By extending the matching network [13] to few-shot scenarios, Snell *et al.* [14] constructed a prototype for each class by averaging the feature representations of samples from the class in the metric space. The classification is conducted by matching the query samples to prototypes by computing their distances. To enhance the prototype representation, Allen *et al.* [38] proposed an infinite mixture of prototypes (IMP) to adaptively represent data distributions for each class, using multiple clusters instead of a single vector. Oreshkin *et al.* [39] proposed a task-dependent adaptive metric for few-shot learning and established prototypes of classes conditioned on a task representation encoded by a task embedding network. Yoon *et al.* [40] proposed a few-shot learning algorithm aided by a linear transformer that performs task-specific null-space projection of the network output. Graphical neural network based models generalize the matching methods by learning the message propagation from the support set and transferring it to the query set [41]. Prototype based methods have recently been improved in a variety of ways [35], [42], [43]. In this work, we design an explicit base-learner based on kernels for each individual task.

Algorithms in the second group learn an optimization that is shared across tasks, while being adaptable to new tasks. Finn *et al.* [8] proposed model-agnostic meta-learning (MAML) to learn an appropriate initialization of model parameters and adapt it to new tasks with only a few gradient steps. To make MAML less prone to meta-overfitting, easier to parallelize and more interpretable, Zintgraf *et al.* [44] proposed fast context adaptation via meta-learning (CAVIA), a single model that adapts to a new task via gradient descent by updating only a set of input parameters at test time, instead of the entire network. Ravi and Larochelle [7] proposed an LSTM-based meta-learner that is trained to optimize a neural network classifier. It captures both the short-term knowledge in individual tasks and the long-term knowledge common to all tasks. Learning a shared optimization algorithm has also been explored to quickly learn new tasks [11], [45]. Bayesian meta-learning methods [15], [46], [47], [48] usually rely on hierarchical Bayesian models to learn the shared statistical information from different tasks and to infer the uncertainty of the models. Rusu *et al.* [12] proposed to learn a low-dimensional latent embedding of model parameters and perform optimization-based meta-learning in this space, which allows for a task-specific parameter initialization and achieves adaptation more effectively. Our method is orthogonal to

optimization based methods and learns a specific base-learner for each task.

The third group of explicitly learned base-learners incorporate what meta-learners have learned and effectively addresses individual tasks [1], [15], [49]. Gordon *et al.* [15] avoided the need for gradient based optimization at test time by amortizing the posterior inference of task-specific parameters in their VERSA. It amortizes the cost of inference and alleviates the need for second derivatives during training by replacing test-time optimization with a forward pass through the inference network. To enable efficient adaptation to unseen learning problems, Bertinetto *et al.* [49] incorporated fast solvers with closed-form solutions as the base learning component of their meta-learning framework. These teach the deep network to use ridge regression as part of its own internal model, enabling it to quickly adapt to novel data. In our method, we also deploy an explicit base-learner but, differently, we leverage a memory mechanism based on an LSTM to collect shared knowledge from related tasks and enhance the base-learners for individual tasks.

In the fourth group, a memory mechanism is part of the solution, where an external memory module is deployed to store and leverage key knowledge for quick adaptation [50], [51], [52]. Santoro *et al.* [50] introduced neural Turing machines into meta-learning by augmenting their neural network with an external memory module, which is used to rapidly assimilate new data to help make accurate predictions with only a few samples. Munkhdalai *et al.* [51] proposed a Meta Network (MetaNet) to learn meta-level knowledge across tasks and shifting the inductive biases via fast parameterization for rapid generalization. Munkhdalai *et al.* [52] designed conditionally shifted neurons within the framework of meta-learning, which modify their activation values with task-specific shifts retrieved from a memory module. In this work, we also leverage a memory mechanism, but, differently, we deploy an LSTM module to collect shared knowledge from related tasks experienced previously to help solve individual tasks.

2.2 Kernel Learning

Kernel learning with random Fourier features is a versatile and powerful tool in machine learning [19], [53], [54]. Pioneering works [24], [55], [56] learn to combine predefined kernels in a multi-kernel learning manner. Kernel approximation by random Fourier features (RFFs) [22] is an effective technique for efficient kernel learning [57], which has recently become increasingly popular [23], [26]. RFFs [22] are derived from Bochner's theorem [58].

Theorem 1 (Bochner's theorem [58]). *A continuous, real valued, symmetric and shift-invariant function $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$ on \mathbb{R}^d is a positive definite kernel if and only if it is the Fourier transform $p(\omega)$ of a positive finite measure such that*

$$k(\mathbf{x}, \mathbf{x}') = \int_{\mathbb{R}^d} e^{i\omega^\top(\mathbf{x}-\mathbf{x}')} dp(\omega) = \mathbb{E}_\omega[\zeta_\omega(\mathbf{x})\zeta_\omega(\mathbf{x}')^*], \quad (1)$$

where $\zeta_\omega(\mathbf{x}) = e^{i\omega^\top \mathbf{x}}$.

It is guaranteed that $\zeta_\omega(\mathbf{x})\zeta_\omega(\mathbf{x}')^*$ is an unbiased estimation of $k(\mathbf{x}, \mathbf{x}')$ with sufficient RFF bases $\{\omega\}$ drawn from $p(\omega)$ [22]. For a predefined kernel, e.g., radial basis function (RBF), we sample from its spectral distribution using the

Monte Carlo method, and obtain the explicit feature map

$$\mathbf{z}(\mathbf{x}) = \frac{1}{\sqrt{D}} [\cos(\boldsymbol{\omega}_1^\top \mathbf{x} + b_1), \dots, \cos(\boldsymbol{\omega}_D^\top \mathbf{x} + b_D)], \quad (2)$$

where $\{\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_D\}$ are the random bases sampled from $p(\boldsymbol{\omega})$, and $[b_1, \dots, b_D]$ are D biases sampled from a uniform distribution with a range of $[0, 2\pi]$. Finally, the kernel value $\mathbf{k}(\mathbf{x}, \mathbf{x}') = \mathbf{z}(\mathbf{x})\mathbf{z}(\mathbf{x}')^\top$ in K is computed as the dot product of their random feature maps with the same bases.

Wilson and Adams [59] learn kernels in the frequency domain by modelling the spectral distribution as a mixture of Gaussians and computing its optimal linear combination. Instead of modelling the spectral distribution with explicit density functions, other works focus on optimizing the random base sampling strategy [23], [60]. Nonetheless, it has been shown that accurate approximation of kernels does not necessarily result in high classification performance [61], [62]. This suggests that learning adaptive kernels with random features by data-driven sampling strategies [23] can improve the performance, even with a low sampling rate, compared to using universal random features [61], [62].

Our work introduces kernels into few-shot meta-learning. We propose to learn kernels with random features in a data-driven way by formulating it as a variational inference problem. This allows us to generate task-specific kernels as well as to leverage shared knowledge from related tasks.

2.3 Normalizing Flows

Normalizing flows (NFs) [30], [32], [63] are promising methods for expressive probability density estimation with tractable distributions. Unlike variational methods, sampling and density evaluation can be efficient and exact for NFs with neat architectures. Generally, NFs are categorized into five types based on how they construct a flow: 1) Autoregressive flows were one of the first classes of flows with invertible autoregressive functions. Examples of such flows include inverse autoregressive flow [64] and masked autoregressive flow [65]. 2) Linear flows generalize the idea of permutating of input variables via an invertible linear transformation [33]. 3) Residual flows [66] are designed as residual networks. The invertible property can be preserved under appropriate constraints; 4) volume-preserved flows with effective invertible architecture, such as coupling layers [31], are typically used in generative tasks. 5) Infinitesimal flows provide another alternative strategy for constructing flows in continuous time by parameterizing its infinitesimal dynamics [32]. Normalizing flows are known to be effective in applications with probabilistic models, including probabilistic modelling [33], [67], [68], [69], inference [32], [64] and representation learning [70].

In this work, we introduce conditional normalizing flows into our kernel learning framework to infer richer posteriors over the random bases, which yields more informative random features. To our knowledge, this is the first work that introduces conditional normalizing flows into the meta-learning framework for few-shot learning.

3 METHODOLOGY

In this section, we present our methodology for learning kernels with random Fourier features under the meta-

learning framework with limited labels. In Section 3.1, we describe the base-learner based on kernel ridge regression. We introduce kernel learning with random features by formulating it as a variational inference problem in Section 3.2. We describe the context inference to leverage the shared knowledge provided by related tasks in Section 3.3. We further enrich the variational random features by conditional normalizing flows in Section 3.4.

3.1 Meta-Learning With Kernels

We adopt the episodic training strategy [7] commonly used for few-shot meta-learning, which involves *meta-training* and *meta-test* 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 to as an episode [7]. In the *meta-test* stage, the base-learner is evaluated on a *meta-test* set with different classes of data samples from the *meta-training* set.

For the few-shot classification problem, we sample N -way k -shot classification tasks from the *meta-training* set, where k is the number of labelled examples for each of the N classes. Given the t th task with a support set $\mathcal{S}^t = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N \times k}$ and query set $\mathcal{Q}^t = \{(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)\}_{i=1}^m$ ($\mathcal{S}^t, \mathcal{Q}^t \subseteq \mathcal{X}$), we learn the parameters α^t of the predictor f_{α^t} using a standard learning algorithm with a kernel trick $\alpha^t = \Lambda(\Phi(X), Y)$, where $\mathcal{S}^t = \{X, Y\}$. Here, Λ is the base-learner and $\Phi: \mathcal{X} \rightarrow \mathbb{R}^{\mathcal{X}}$ is a mapping function from \mathcal{X} to a dot product space \mathcal{H} . The similarity measure $\mathbf{k}(\mathbf{x}, \mathbf{x}') = \langle \Phi(\mathbf{x}), \Phi(\mathbf{x}') \rangle$ is called a kernel [19].

In traditional supervised learning, the base-learner for the t th single task usually relies on a universal kernel to map the input into a dot product space for efficient learning. Once the base-learner is trained on the support set, its performance is evaluated on the query set using the following loss function

$$\sum_{(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \mathcal{Q}^t} L(f_{\alpha^t}(\Phi(\tilde{\mathbf{x}})), \tilde{\mathbf{y}}), \quad (3)$$

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

$$\sum_t^T \sum_{(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \mathcal{Q}^t} L(f_{\alpha^t}(\Phi^t(\tilde{\mathbf{x}})), \tilde{\mathbf{y}}), \text{ s.t. } \alpha^t = \Lambda(\Phi^t(X), Y), \quad (4)$$

where Φ^t is the feature mapping function which can be obtained by learning a task-specific kernel \mathbf{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 Λ for few-shot learning. The kernel value in the Gram matrix $K \in \mathbb{R}^{C_k \times C_k}$ is computed as $\mathbf{k}(\mathbf{x}, \mathbf{x}') = \Phi(\mathbf{x})\Phi(\mathbf{x}')^\top$, where “ \top ” is the transpose operation. The base-learner Λ for a single task is obtained by solving the following objective *w.r.t.* the support set of this task

$$\Lambda = \arg \min_{\alpha} \text{Tr}[(Y - \alpha K)(Y - \alpha K)^\top] + \lambda \text{Tr}[\alpha K \alpha^\top], \quad (5)$$

which admits a closed-form solution

$$\alpha = Y(\lambda I + K)^{-1}. \quad (6)$$

The learned predictor is then applied to samples in the query set \tilde{X}

$$\hat{Y} = f_\alpha(\tilde{X}) = \alpha \tilde{K}, \quad (7)$$

Here, $\tilde{K} = \Phi(X)\Phi(\tilde{X})^\top \in \mathbb{R}^{Ck \times m}$, with each element as $\kappa(\mathbf{x}, \tilde{\mathbf{x}})$ between the samples from the support and query sets. Note that we also treat λ in (5) as a trainable parameter by leveraging the meta-learning setting, and all these parameters are learned by the meta-learner.

In order to obtain task-specific kernels, we consider learning adaptive kernels with random Fourier features in a data-driven way. This also enables shared knowledge of different tasks to be captured by exploring their dependencies in the meta-learning framework.

3.2 Variational Random Features

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}

$$\max_p \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}} \log p(\mathbf{y}|\mathbf{x}, S) \quad (8)$$

$$= \max_p \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}} \log \int p(\mathbf{y}|\mathbf{x}, S, \omega) p(\omega|\mathbf{x}, S) d\omega. \quad (9)$$

We adopt a conditional prior distribution $p(\omega|\mathbf{x}, S)$ over the base ω , as in the conditional variational autoencoder [71], rather than an uninformative prior [72], [73].

In order to infer the posterior $p(\omega|\mathbf{y}, \mathbf{x}, S)$ over ω , which is generally intractable, we use a variational distribution $q_\phi(\omega|S)$ to approximate it, where the base is conditioned on the support set S . We obtain the variational distribution by minimizing the Kullback-Leibler (KL) divergence: $D_{\text{KL}}[q_\phi(\omega|S)||p(\omega|\mathbf{y}, \mathbf{x}, S)]$. By applying Bayes' rule to the posterior $p(\omega|\mathbf{y}, \mathbf{x}, S)$, we derive the evidence lower bound (ELBO) as

$$\log p(\mathbf{y}|\mathbf{x}, S) \geq \mathbb{E}_{q_\phi(\omega|S)} \log p(\mathbf{y}|\mathbf{x}, S, \omega) - D_{\text{KL}}[q_\phi(\omega|S)||p(\omega|\mathbf{x}, S)]. \quad (10)$$

The first term of the ELBO is the predictive log-likelihood conditioned on the observation \mathbf{x} , S and the inferred RFF bases ω . 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 $q_\phi(\omega|S)$ and the meta prior $p(\omega|\mathbf{x}, S)$, 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, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2022.3154930>.

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




Fig. 2. Graphical illustration of variational inference of the random Fourier basis under the meta-learning framework for few-shot learning, where (\mathbf{x}, \mathbf{y}) is a sample in the query set \mathcal{Q}^t . The base ω^t of the t th task is dependent on the support set S^t of the current task and the context C of related tasks. The dashed lines indicate variational inference.

$$\mathcal{L} = \frac{1}{T} \sum_{t=1}^T \left[\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t} \mathbb{E}_{q_\phi(\omega^t|S^t)} \log p(\mathbf{y}|\mathbf{x}, S^t, \omega^t) - D_{\text{KL}}[q_\phi(\omega^t|S^t)||p(\omega^t|\mathbf{x}, S^t)] \right], \quad (11)$$

where S^t is the support set of the t th task associated with its specific bases $\{\omega_d^t\}_{d=1}^D$ and $(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t$ is the sample from the query set of the t th task.

3.3 Task 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 with a conditional distribution $q_\phi(\omega^t|S^t, C)$, where we use C to contain the shared knowledge provided by related tasks. This makes the bases $\{\omega_d^t\}_{d=1}^D$ of the current t th task conditioned also on the context C of related tasks, which gives rise to a new ELBO, as follows:

$$\log p(\mathbf{y}|\mathbf{x}, S^t) \geq \mathbb{E}_{q_\phi(\omega|S^t, C)} \log p(\mathbf{y}|\mathbf{x}, S^t, \omega) - D_{\text{KL}}[q_\phi(\omega|S^t, C)||p(\omega|\mathbf{x}, S^t)]. \quad (12)$$

This can be represented in a directed graphical model, as shown in Fig. 2. In a practical sense, the KL term in (12) 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 information absorbed from experienced tasks. The base-learner built on the inferred kernel with the informative random features effectively solves the current task.

However, since there is usually a large number of related tasks, it is non-trivial to model them all 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, leveraging its innate capability of remembering long-term information [74]. The LSTM offers a well-suited structure to implement the context inference. The cell state c stores and accrues the meta

knowledge shared among related tasks. It can also be updated when experiencing a new task in each episode over the course of learning, where the output \mathbf{h} is used to adapt the model to each specific task.

To be more specific, we model the variational posterior $q_\phi(\omega^t | \mathcal{S}^t, \mathcal{C})$ through $q_\phi(\omega | \mathbf{h}^t)$, which is parameterized as a multi-layer perceptron (MLP) $\phi(\mathbf{h}^t)$. Note that \mathbf{h}^t is the output from an LSTM that takes \mathcal{S}^t and \mathcal{C} as inputs. To address the impact of the order of tasks We implement the inference network with bidirectional LSTMs [75], [76], which can incorporate more context information, which absorbs knowledge about the previous tasks as well as information about the later tasks. For the LSTM, we have

$$[\mathbf{h}^t, \mathbf{c}^t] = g_{\text{LSTM}}(\mathcal{S}^t, \mathbf{h}^{t-1}, \mathbf{c}^{t-1}), \quad (13)$$

where $g_{\text{LSTM}}(\cdot)$ is a LSTM network that takes the current support set, the output \mathbf{h}^{t-1} and the cell state \mathbf{c}^{t-1} as input. \mathcal{S}^t is the average over the feature representation vectors of samples in the support set [77]. The feature representation is obtained by a shared convolutional network $\psi(\cdot)$. To incorporate more context information, we also implement the inference with a bidirectional LSTM. We thus have $\mathbf{h}^t = [\mathbf{h}^t \mathbf{h}^t]$, where \mathbf{h}^t and \mathbf{h}^t are the outputs from the forward and backward LSTMs, respectively, and $[\cdot, \cdot]$ indicates a concatenation operation.

Therefore, the optimization objective with the context inference is:

$$\begin{aligned} \mathcal{L} = & \frac{1}{T} \sum_{t=1}^T \left[\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t} \mathbb{E}_{q_\phi(\omega^t | \mathbf{h}^t)} \log p(\mathbf{y} | \mathbf{x}, \mathcal{S}^t, \omega^t) \right. \\ & \left. - D_{\text{KL}}[q_\phi(\omega^t | \mathbf{h}^t) \| p(\omega^t | \mathbf{x}, \mathcal{S}^t)] \right], \quad (14) \end{aligned}$$

where the variational approximate posterior $q_\phi(\omega^t | \mathbf{h}^t)$ is taken as a multivariate Gaussian with a diagonal covariance. Given the support set as input, the mean ω_μ and standard deviation ω_σ are output from the inference network $\phi(\cdot)$. The conditional prior $p(\omega^t | \mathbf{x}, \mathcal{S}^t)$ is implemented with a prior network which takes an aggregated representation using the cross attention [78] between \mathbf{x} and \mathcal{S}^t . The details of the prior network are provided in the supplementary material, available online. To enable back propagation with the sampling operation during training, we adopt the reparametrization trick [72], [73] as $\omega = \omega_\mu + \omega_\sigma \odot \epsilon$, where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$.

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 information across all the tasks in the meta-training set, the output and the cell state of the LSTMs are passed down across batches. As a result, the final the cell state contains the distilled prior knowledge from all the tasks experienced in the meta-training set.

3.4 Enriching Random Features by Normalizing Flows

The posterior distribution $q_\phi(\omega | \mathbf{h}^t)$ is assumed to be a fully factorized Gaussian corresponding to a single kernel, resulting in

limited expressive ability to approximate the true posterior over random Fourier bases. Motivated by the empirical success of normalizing flows [32] and conditional normalizing flows [34], we propose the conditional normalizing flows that provide a principled way to learn richer posteriors. Conditional normalizing flows allow us to obtain more complex distributions corresponding to the multiple-kernel, which results in better expressiveness.

Normalizing flows map a complex distribution $p_{\mathbf{x}}(\mathbf{X})$ to a simpler distribution $p_{\mathbf{z}}(\mathbf{Z})$ through a chain of transformations. Let $\mathbf{x} \in X$ denote data sampled from an unknown distribution $\mathbf{x} \sim p_X(\mathbf{x})$. The key idea in normalizing flows is to represent $p_X(\mathbf{x})$ as a transformation $\mathbf{x} = g(\mathbf{z})$ of a single Gaussian distribution $\mathbf{z} \sim p_Z = \mathcal{N}(0, \mathbf{I})$. Moreover, we assume that the mapping is bijective: $\mathbf{x} = g(\mathbf{z}) = f^{-1}(\mathbf{z})$. Therefore, the log-likelihood of the data is given by the change of variable formula

$$\log(p_X(\mathbf{x})) = \log(p_Z(f(\mathbf{x}))) + \log \left(\left| \det \left(\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}^T} \right) \right| \right), \quad (15)$$

where $\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}^T}$ is the Jacobian of the map $f(\mathbf{x})$ at \mathbf{x} . The functions f can be learned by maximum likelihood (15), where the bijectivity assumption allows expressive mappings to be trained by gradient backpropagation.

To make the Jacobian tractable for the map $f(\mathbf{x})$, NICE [30] and RealNVP [31] proposed to stack a sequence of simple bijective transformations, such that their Jacobian is a triangular matrix. In this way, the log-determinant depends only on the sum of its diagonal elements. Dinh *et al.* [30], [31] proposed the additive coupling layer for each transformation. In each affine coupling transformation, the input vector $\mathbf{x} \in \mathbb{R}^d$ is split into upper and lower halves, $\mathbf{x}_{I_1}, \mathbf{x}_{I_2} \in \mathbb{R}^{d/2}$. These are plugged into the following transformation, referred to as a single flow-block f_i

$$\mathbf{z}_1 = \mathbf{x}_{I_1}, \quad \mathbf{z}_2 = \mathbf{x}_{I_2} \circ \exp(s_i(\mathbf{x}_{I_1})) + t_i(\mathbf{x}_{I_1}), \quad (16)$$

where \circ denotes element-wise multiplication. It is important to note that the mappings s_i and t_i can be arbitrarily complicated functions of \mathbf{x}_i and need not be invertible themselves. In practice, s_i and t_i are achieved via neural networks.

Given the outputs \mathbf{z}_1 and \mathbf{z}_2 , this affine transformation is invertible by

$$\mathbf{x}_{I_1} = \mathbf{z}_1, \quad \mathbf{x}_{I_2} = (\mathbf{z}_2 - t_i(\mathbf{z}_1)) \circ \exp(-s_i(\mathbf{z}_1)). \quad (17)$$

The RealNVP [31] flow comprises k reversible flow-blocks interleaved with switch-permutations

$$f_{\text{RealNVP}} = f_k \cdot r \dots f_2 \cdot r \cdot f_1, \quad (18)$$

where r denotes a switch-permutation, which permutes the order of \mathbf{x}_1 and \mathbf{x}_2 . According to the chain rule, the log-determinant of the Jacobian of the whole transformation f is computed by summing the log-determinants of the Jacobian of each f_i , making the likelihood calculation tractable.

Conditional normalizing flows [34] learn conditional likelihoods for complicated target distributions in multivariate prediction tasks. Take an input $\mathbf{x} \in \mathcal{X}$ and a regression target $\mathbf{y} \in \mathcal{Y}$. CNFs learn a complicated distribution $p_{\mathcal{Y}|\mathcal{X}}(\mathbf{y}|\mathbf{x})$ using a conditional prior $p_{\mathcal{Z}|\mathcal{X}}(\mathbf{z}|\mathbf{x})$ and a mapping

$f_\phi : \mathcal{Y} \times \mathcal{X} \rightarrow \mathcal{Z}$, which is bijective in \mathcal{Y} and \mathcal{Z} . The log-likelihood of CNFs is

$$\begin{aligned} \log(p_{\mathcal{Y}|\mathcal{X}}(\mathbf{y}|\mathbf{x})) &= \log(p_{\mathcal{Z}|\mathcal{X}}(\mathbf{z}|\mathbf{x})) + \log\left(\left|\frac{\partial \mathbf{z}}{\partial \mathbf{y}}\right|\right) \\ &= \log(p_{\mathcal{Z}|\mathcal{X}}(f_\phi(\mathbf{y}, \mathbf{x})|\mathbf{x})) + \log\left(\left|\frac{\partial f_\phi(\mathbf{y}, \mathbf{x})}{\partial \mathbf{y}}\right|\right). \end{aligned} \quad (19)$$

Different from NFs, in the log-likelihood of CNFs, all distributions are conditional and the flow has a conditioning argument for \mathbf{x} .

We parameterize the approximate posterior distribution $q_\phi(\omega|\mathbf{h}^t)$ with a flow of length K , $q_\phi(\omega|\mathbf{h}^t) := q_K(\omega_K)$. The ELBO (10) is thus written as an expectation over the initial distribution $q_0(\omega)$

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x}, \mathcal{S}) &\geq -\mathbb{E}_{q_\phi(\omega|\mathbf{h}^t)}[\log q_\phi(\omega|\mathbf{h}^t) + \log p(\mathbf{y}, \omega|\mathcal{S}, \mathbf{x})] \\ &= -\mathbb{E}_{q_0(\omega_0)}[\ln q_K(\omega_K) + \log p(\mathbf{y}, \omega_K|\mathcal{S}, \mathbf{x})] \\ &= -\mathbb{E}_{q_0(\omega_0)}\left[\ln q_0(\omega_0) - \sum_{k=1}^K \ln \left|\det \frac{\partial f}{\partial \omega_k}\right|\right] \\ &\quad + \mathbb{E}_{q_0(\omega_0)}[\log p(\mathbf{y}, \omega_K|\mathcal{S}, \mathbf{x})], \end{aligned} \quad (20)$$

where $q_0(\omega_0)$ is obtained from the approximate posterior distribution $q_\phi(\omega|\mathbf{h}^t)$ without transformation.

We then obtain the objective by maximizing the log-likelihood $\log p(\mathbf{y}|\mathbf{x}, \mathcal{S})$ with respect to a batch of T tasks

$$\begin{aligned} \mathcal{L} &= \frac{1}{T} \sum_{t=1}^T \left[\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}^t} \mathbb{E}_{q_0(\omega_0^t)} \left[-\ln q_0(\omega_0^t) + \sum_{k=1}^K \ln \left|\det \frac{\partial f}{\partial \omega_k^t}\right|\right] \right. \\ &\quad \left. + \mathbb{E}_{q_0(\omega_0^t)}[\log p(\mathbf{y}, \omega_K^t|\mathcal{S}^t, \mathbf{x})] \right], \end{aligned} \quad (21)$$

where ω_k^t is the random base after k transformations.

We rely on the conditional coupling layer from [34] to transform the random base distribution. This layer is an extension of the affine coupling layer from RealNVP [31] to make the computation of the Jacobian for the map $f(x)$ tractable. The input $\omega_{k-1} = [\omega_{k-1}^{I_0}, \omega_{k-1}^{I_1}]$ of an affine coupling layer is split into two parts, which are transformed individually

$$\begin{aligned} \omega_k^{I_i} &= \omega_{k-1}^{I_i} \odot \exp(s_{i+1}(\omega_{k-(1-i)}^{I_{(1-i)}}, \mathbf{h}^t)) \\ &\quad + t_{(i+1)}(\omega_{k-(1-i)}^{I_{(1-i)}}, \mathbf{h}^t) \end{aligned} \quad (22)$$

where $i \in \{0, 1\}$. Note that the transformations s_{i+1}, t_{i+1} do not need to be invertible and are modelled as convolutional neural networks. The inverse of an affine coupling layer is

$$\begin{aligned} \omega_{k-1}^{I_i} &= (\omega_k^{I_i} - t_{(i+1)}(\omega_{k-(1-i)}^{I_{(1-i)}}, \mathbf{h}^t)) \\ &\quad \odot \exp(-s_{(i+1)}(\omega_{k-i}^{I_i}, \mathbf{h}^t)). \end{aligned} \quad (23)$$

The log-determinant of the Jacobian for one affine coupling layer is calculated as the sum over s_i , i.e., $\sum_j s_1(\omega_{k-1}^{I_1}, \mathbf{h}^t)_j + \sum_j s_2(\omega_k^{I_0}, \mathbf{h}^t)_j$. A deep invertible network is built as a sequence of multiple such layers, with a permutation of the dimensions after each layer. The conditional input \mathbf{h}^t is added as an extra input to each transformation in the coupling layer.




Fig. 3. Effect of conditional normalizing flows on the random bases. They transform the single Gaussian distribution of the random bases into a more complex distribution, which yields more informative random features.

We refer to the kernel constructed based on the random bases by conditional normalizing flows as MetaKernel.

We visualize the distribution of the random bases produced by the CNFs in Fig. 3. ω_k indicates the distribution of the random bases after k transformations. This visualization shows that we can transform a single Gaussian distribution of random bases into a more complex distribution, which achieves more informative random features, resulting in improved performance, as we will demonstrate in our experiments.

4 EXPERIMENTS

In this section, we report our experiments to demonstrate the effectiveness of the proposed MetaKernel for both regression and classification with limited labels. We also provide thorough ablation studies to gain insight into our method by showing the efficacy of each introduced component.

4.1 Few-Shot Classification

The few-shot classification experiments are conducted on four commonly used benchmarks, i.e., Omniglot [6], *miniImageNet* [13], CIFAR-FS [79] and Meta-Dataset [35]. We also perform experiments on DomainNet [80] for few-shot domain generalization. Sample images from each dataset are provided in Fig. 4.

4.1.1 Datasets

Omniglot [6] is a few-shot classification benchmark that contains 1623 handwritten characters (each with 20 examples). All characters are grouped into one of 50 alphabets. For fair comparison against the state of the art, we follow the same data split and pre-processing used by Vinyals *et al.* [13]. Specifically, the training, validation, and test sets are composed of a random split of [1100, 200, 423]. The dataset is augmented with rotations of 90 degrees, which results in 4000 classes for training, 400 for validation, and 1292 for testing. The number of examples is fixed to 20. All images are resized to 28×28 . For a N -way, k -shot task at training time, we randomly sample N classes from the 4000 classes, each with $(k + 15)$ examples. Thus, there are $C \times k$ examples in the support set and $C \times 15$ examples in the query set. The same sampling strategy is followed for validation and testing.

miniImageNet [13] is a challenging dataset constructed from ImageNet [81], which comprises a total of 100 different classes (each with 600 instances). All images are down-sampled to 84×84 . We use the same splits as Ravi and Larochelle [7], with [64, 16, 20] classes for training, validation

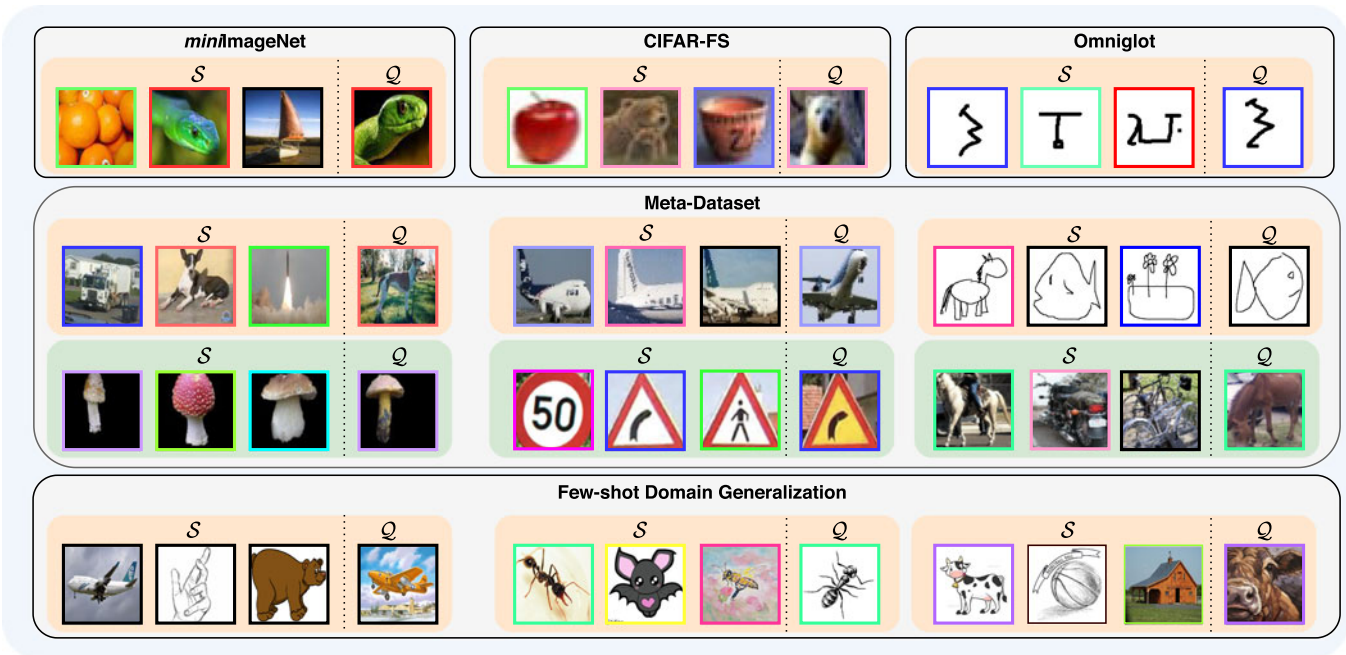


Fig. 4. Examples from each dataset. S and Q indicate the support and query sets for each task. For Meta-Dataset, we only show examples from *ImageNet* [81], *Aircraft* [82], *Quick Draw* [83], *Fungi* [84], *Traffic Signs* [85] and *MS-COCO* [86]. For the few-shot domain generalization, we only show the examples from *DomainNet* using *Quick Draw* as the target domain during the meta-test stage.

and testing. We use the same episodic sampling strategy as for Omniglot.

CIFAR-FS [49] is adapted from *CIFAR-100* [79] for few-shot learning. In the many-shot image classification benchmark *CIFAR-100*, there are 100 classes grouped into 20 super-classes (each with 600 instances). *CIFAR-FS* uses the same split criteria (64, 16, 20) with which *miniImageNet* has been generated. The resolution of all images is 32×32 .

Meta-Dataset [35] is composed of ten existing image classification datasets (eight for training, two for testing). These are: *ILSVRC-2012* (*ImageNet*, [81]), *Omniglot* [6], *Aircraft* [82], *CUB-200-2011* (*Birds*, [87]), *Describable Textures* [88], *Quick Draw* [83], *Fungi* [84], *VGG Flowr* [89], *Traffic Signs* [85] and *MS-COCO* [86]. Each episode generated in *Meta-Dataset* uses classes from a single dataset. Two of these datasets, *Traffic Signs* and *MSCOCO*, are fully reserved for evaluation, which means that no classes from these sets are participated in the training set. Apart from for *Traffic Signs* and *MS-COCO*, the remaining datasets contribute some classes to the training, validation and test splits. There are about 14 million images in total in *Meta-Dataset*.

DomainNet. [80]. Du *et al.* [16] introduced the setting of few-shot domain generalization, which combines the challenges of both few-shot classification and domain generalization. It is based on the *DomainNet* dataset by Peng *et al.* [80], which contains six distinct domains, i.e., *clipart*, *infograph*, *painting*, *quickdraw*, *real*, and *sketch*, for 345 categories. The categories are from 24 divisions.

4.1.2 Implementation Details

We extract image features using two different convolutional neural networks *Conv-4* and *ResNet-12* for *miniImageNet*, and *CIFAR-FS*. We do not use any fully connected layers in this CNNs. For the *Meta-Dataset* experiments, we use a *ResNet-18* [3] as our base learner to be consistent with [35]. The

dimension of all feature vectors is 256. We also evaluate the random Fourier features (RFFs) and the radial basis function (RBF) kernel, where we take the bandwidth σ as the mean of the pair-wise distances between samples in the support set of each task. The inference network $\phi(\cdot)$ is a three-layer MLP with 256 units in the hidden layers and rectifier non-linearity, where the input sizes is 512 for the bidirectional LSTMs. Note that the parameters of feature extractor $\Psi(\cdot)$, inference network $\phi(\cdot)$, and LSTM are optimized by meta-learning among different episodes. The base-learner λ , the kernel K could be directly computed inside each episode. We use an SGD optimizer with a momentum of 0.9 in all experiments.

The key hyperparameter for the number of bases D in (2) is set to $D=780$ for *MetaKernel* in all experiments, while we use RFFs with $D=2048$ as this produces the best performance. The sampling rate in *MetaKernel* 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 [22], [103]. We adopt a similar meta-testing protocol as [8], [15], but we test on 3000 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. Our code will be publicly released.¹

4.1.3 Comparison to the State of the Art

Few-Shot Image Classification. We first evaluate *MetaKernel* on the *miniImageNet*, *CIFAR-FS* and *Omniglot* datasets under various way (the number of classes used in each task) and shot (the number of support set examples used per class) configurations. The results are reported in Table 1. We report the results of two experiments using MAML [8]. To

1. <https://github.com/YDU-uva/MetaKernel>

TABLE 1
Few-Shot Image Classification Performance (%) on *miniImageNet*, *CIFAR-FS* Under Different Backbones

Method	Backbone	<i>miniImageNet</i> , 5-way		<i>CIFAR-FS</i> , 5-way	
		1-shot	5-shot	1-shot	5-shot
MatchingNet [13]	Conv-4	44.2	57	—	—
Meta-LSTM [7]	Conv-4	43.4±0.8	60.6±0.7	—	—
MAML [8]	Conv-4	48.7±1.8	63.1±0.9	58.9±1.9	71.5±1.0
MAML (64C)	Conv-4	46.7±1.7	61.1±0.1	58.9±1.8	71.5±1.1
ProtoNets [14]	Conv-4	47.4±0.6	65.4±0.5	55.5±0.7	72.0±0.6
RelationNet [90]	Conv-4	50.4±0.8	65.3±0.7	55.0±1.0	69.3±0.8
IMPS [91]	Conv-4	49.60±0.8	68.10±0.80	—	—
SNAIL [49]	Conv-4	45.1	55.2	—	—
GNN [41]	Conv-4	50.3	66.4	61.9	75.3
PLATIPUS [47]	Conv-4	50.1±1.9	—	—	—
VERSA [15]	Conv-5	53.3±1.8	67.3±0.9	62.5±1.7	75.1±0.9
R2-D2 (64C) [49]	Conv-4	49.5±0.2	65.4±0.2	62.3±0.2	77.4±0.2
R2-D2 [92]	Conv-4	51.7±1.8	63.3±0.9	60.2±1.8	70.9±0.9
CAVIA [44]	Conv-4	51.8±0.7	65.6±0.6	—	—
iMAML [36]	Conv-4	49.3±1.9	—	—	—
TPN [93]	Conv-4	51.4±0.8	67.1±0.7	—	—
SAML [94]	Conv-4	52.2	66.49	—	—
VSM [43]	Conv-4	54.7±1.6	68.0±0.9	63.5±1.7	76.2±0.8
MetaVRF [1]*	Conv-4	54.2±0.8	67.8±0.7	63.1±0.7	76.5±0.9
MetaKernel	Conv-4	55.5±0.9	68.5±0.8	64.3±0.8	77.5±0.9
ProtoNets [14]	ResNet-12	59.3 ±0.6	75.6±0.5	72.2±0.7	83.5±0.5
SNAIL [49]	ResNet-12	55.7±0.9	68.8±0.9	—	—
Dynamic FS [95]	ResNet-12	55.5±0.8	70.1±0.6	—	—
TADAM [96]	ResNet-12	58.5±0.3	76.7±0.3	—	—
MTL [97]	ResNet-12	61.2±1.8	75.5±0.8	—	—
VariationalFSL [98]	ResNet-12	61.2±0.3	77.6±0.1	—	—
TapNet [99]	ResNet-12	61.6±0.2	76.4±0.1	—	—
MetaOptNet [100]	ResNet-12	62.6±0.6	78.6±0.4	72.8±0.7	84.3±0.5
DSN-MR [101]	ResNet-12	64.6±0.7	79.5±0.5	75.6±0.9	86.2±0.6
VSM [†] [43]	ResNet-12	64.2±0.3	79.7±0.3	75.3±0.7	86.5±0.4
MetaVRF [1] [†]	ResNet-12	64.1±0.2	79.5±0.4	75.1±0.5	85.8±0.5
MetaKernel	ResNet-12	66.9±0.3	81.3±0.2	76.1±0.4	86.9±0.5

Best performing methods and any other runs within the 95% confidence margin are denoted in bold. MetaKernel consistently achieves the top performance.

[†]Results based on our re-implementation.

keep MAML [8] consistent with our backbone for *miniImageNet* and *CIFAR-FS*, in addition to its original results, we also implement MAML (64C) with 64 channels in each convolutional layer for fair comparisons. While it obtains modest performance, we believe the increased model size leads to overfitting. As the original SNAIL uses a very deep ResNet-12 network for embedding, we cite the results of SNAIL reported in [49] using a similar shallow network as ours. For fair comparison, we also cite the original results of R2-D2 [49] using 64 channels. On all benchmark datasets, MetaKernel delivers the best performance. It is worth noting that MetaKernel achieves an accuracy of 55.5% under the 5-way 1-shot setting by ResNet-12 on the *miniImageNet* dataset, surpassing the second-best model by 2.3%. This is a good improvement considering the challenge of this setting. On *CIFAR-FS*, our model surpasses the second-best method, i.e., VERSA [15] and has a smaller margin of error bar under the 5-way 1-shot setting using the same backbone.

Few-Shot Meta-Dataset Classification. Next, we evaluate MetaKernel on the most challenging few-shot classification benchmark i.e., Meta-Dataset [35], which is composed of 10 image classification datasets. For Meta-Dataset, we train our model on the ILSVRC [81] training split and test on the 10 diverse datasets. As shown in Table 2, MetaKernel

outperforms fo-Proto-MAML [35] across all 10 datasets. MetaKernel also surpasses the second-best method, RFS [102], on 7 out of 10 datasets. Overall, we perform well against previous methods, achieving new state-of-the-art results on the challenging Meta-Dataset.

Few-Shot Domain Generalization. We also evaluate our method on few-shot domain generalization [16], which combines the challenges of both few-shot classification and domain generalization. For few-shot domain generalization, each task has only a few samples in the support set for training and we test the model on tasks in a query set, which come from a different domain than the support set. The results are reported in Table 3. MetaKernel obtains the best performance, surpassing the MetaNorm [16] by a margin of up to 2.0% on the 5-way 1-shot and 1.8% on the 5-way 5-shot setting. Its performance on the few-shot domain generalization task demonstrates that MetaKernel is not only able to handle the problem of few-shot learning, but also thrives under domain-shifts.

4.2 Few-Shot Regression

We also consider regression tasks with a varying number of shots k , and compare MetaKernel with MAML [8], a

TABLE 2
Few-Shot Meta-Dataset Classification Accuracy (%) With Variable Number of Ways and Shots, Following the Setup in [35]

Dataset	Matching Net [13]	ProtoNet [14]	fo-MAML [8]	Relation Net [90]	fo-Proto-MAML [35]	RFS [102]		MetaKernel
						LR-Simple	LR-Distill	
ILSVRC	45.00	50.50	45.51	34.69	49.53	60.14	61.48	61.71
Omniglot	52.27	59.98	55.55	45.35	63.37	64.92	64.31	65.43
Aircraft	48.97	53.10	56.24	40.73	55.95	63.12	62.32	65.37
Birds	62.21	68.79	63.61	49.51	68.66	77.69	79.47	77.13
Textures	64.15	66.56	68.04	52.97	66.49	78.59	79.28	82.01
Quick Draw	42.87	48.96	43.96	43.30	51.52	62.48	60.83	58.46
Fungi	33.97	39.71	32.10	30.55	39.96	47.12	48.53	49.73
VGG Flower	80.13	85.27	81.74	68.76	87.15	91.60	91.00	93.16
Traffic Signs	47.80	47.12	50.93	33.67	48.83	77.51	76.33	77.91
MSCOCO	34.99	41.00	35.30	29.15	43.74	57.00	59.28	56.97
<i>Mean Accuracy</i>	51.24	56.10	53.30	42.87	57.52	68.02	68.28	68.79

1000 tasks are sampled for evaluation. MetaKernel is a consistent top-performer.

representative meta-learning algorithm. We follow MAML [8] and fit a target sine function $y=A \sin (w x+b)$, 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, which follow a uniform distribution within the corresponding interval. The goal is to estimate the target sine function given only n randomly sampled data points. Here, we consider inputs within the range of $x \in [-5, 5]$, and conduct three tests under the conditions of $k=3, 5, 10$. For fair comparison, we compute the feature embedding using a small MLP with two hidden layers of size 40, following the same settings used in MAML.

The results in Fig. 5 show that MetaKernel fits the function well with only three shots, even when we do not use the full model. It performs better with an increasing number of shots, almost entirely fitting the target function with ten shots. We observe all MetaKernel variants perform better than MAML [8] for all three settings with varying numbers of shots, both visually and in terms of MSE. Best results are obtained with our full model.

4.3 Ablation Studies

To study how our proposed components bring performance gains to MetaKernel on few-shot learning, our ablations consider: (1) the benefit of random Fourier features; (2) the benefit of task context inference; (3) the benefit of enriching random features by normalizing flows; (4) the effect of

deeper embeddings; (5) the efficiency of the model; (6) the versatility of the model.

Benefit of Random Fourier Features. We first show the benefit of random Fourier features (RFFs) by comparing them with the regular RBF kernel. As can be seen from the first two rows in Table 4, RFFs perform 10.7% better than an RBF kernel on the 5-way 1-shot setting of *miniImageNet*, and 14.9% better on the 5-way 5-shot setting of *CIFAR-FS*. The considerable performance gain over RBF kernels on both datasets indicates the benefit of adaptive kernels based on random Fourier features for few-shot image classification. The modest performance obtained by RBF kernels is due to the mean of pair-wise distances of support samples being unable to provide a proper estimate of the kernel bandwidth. Note that the performance of RFFs is better than the variational RFFs on the 5-way 1-shot setting of *miniImageNet*. This may be due to the fact that the support samples are too small, resulting in the random bases generated from the samples not accurately representing the current task, while the parameters in the random bases of RFFs are sampled from a standard Gaussian distribution. Therefore, the context information among previous related tasks should be integrated into the variational RFFs. In addition, RFFs cannot use the context information directly since it generates random base parameters sampled from a deterministic distribution.

Benefit of Task Context Inference. We investigate the benefit of adding task context inference to the MetaKernel. Specifically, we leverage a bi-LSTM cell state c to store and accrue the meta-knowledge shared among related tasks. The experimental results are reported in Table 4. Adding task context inference on top of the MetaKernel with variational random features leads to a consistent gain under all settings, for both datasets. This demonstrates the effectiveness of using an LSTM to explore task dependency.

Benefit of Enriching Features by Normalizing Flows. We show the benefit of enriching the variational random features by normalizing flows in the last two row of Table 4. we find that MetaKernel performs better than MetaVRF (55.5% -up 1.3%) under the 5-way 1-shot setting on *miniImageNet* and (64.3% -up 1.2%) under the 5-way 1-shot setting

TABLE 3
Few-Shot Domain Generalization Performance (5-Way %)

Method	1-shot	5-shot
ProtoNets [14]	28.4±1.8	47.9±0.8
MAML [8]	28.7±1.8	49.3±0.8
VERSA [15]	30.9±1.7	51.7±0.8
MetaNorm [16]	32.7±1.7	51.9±0.8
METAKERNEL	34.7±1.7	53.7±0.8

The best performing methods and any other runs within 95% confidence margin are denoted in bold. MetaKernel is on par with MetaNorm for 1-shot and outperforms all previous methods for the 5-shot setting.




Fig. 5. Few-shot regression performance comparison (MSE). MetaKernel fits the target function well, even with variational random features only using three shots, and consistently outperforms MAML for all settings. Legend: --- MAML; - - - MetaKernel (variational RFFs only); - - - MetaKernel (variational RFFs & task context); - - - MetaKernel (full model); — Ground Truth; ▲ Support Samples.

TABLE 4
Ablation Studies on *miniImageNet* and *CIFAR-FS* Demonstrating Benefit of Random Fourier Features, Task Context Inference and Feature Enriching by Normalizing Flows

Method	<i>miniImageNet</i> , 5-way		<i>CIFAR-FS</i> , 5-way	
	1-shot	5-shot	1-shot	5-shot
RBF Kernel	42.1±1.2	54.9±1.1	46.0±1.2	59.8±1.0
RFFs	52.8±0.9	65.4±0.9	61.1±0.8	74.7±0.9
Variational RFFs	51.3±0.8	66.1±0.7	61.1±0.7	74.3±0.9
Variational RFFs & task context inference	54.2±0.8	67.8±0.7	63.1±0.7	76.5±0.9
METAKERNEL	55.5±0.9	68.5±0.8	64.3±0.8	77.5±0.9

Best settings within the 95% confidence margin denoted in bold.

on *CIFAR-FS*. These results indicate that the CNFs provide more informative kernels for the new task, which allows the learned distribution of random bases to more closely approximate the real random bases distribution and therefore improves few-shot classification performance.

Deep Embeddings. MetaKernel is independent of the convolutional architecture for feature extraction and works 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. Specifically, we adopt the pre-trained embeddings from a

28-layer wide residual network (WRN-28-10) [109], in a similar fashion to [12], [107], [108]. We choose activations in the 21-st layer, with average pooling over spatial dimensions, as feature embeddings. The dimension of the 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 5. MetaKernel achieves the best performance under both settings and surpasses LEO [12], a recently proposed meta-learning method, especially on the challenging 5-way 1-shot setting. Compared with our conference paper, MetaVRF [1], MetaKernel performs 1.23% better on the 5-way 1-shot setting of *miniImageNet*, which also validates the effectiveness of the CNFs. The consistent state-of-the-art results on all benchmarks using both shallow and deep feature extraction networks validate the effectiveness of MetaKernel for few-shot learning.

Efficiency. Regular RFFs usually require high sampling rates to achieve satisfactory performance. However, our MetaKernel achieves high performance with a relatively low sampling rate, which guarantees its high efficiency. In Fig. 6, we compare with regular RFFs using different sampling rates. We provide the performance change of fully trained models using RFFs and MetaKernel under a varying number of bases D . We show the comparison results for the 5-way 5-shot setting on *miniImageNet* and *CIFAR-FS* in Fig. 6. MetaKernel consistently yields higher performance than regular RFFs with the same number of sampled bases. The results verify the efficiency of MetaKernel in learning adaptive kernels and its effectiveness in improving performance by exploring the dependencies of related tasks.

Versatility. In contrast to most existing meta-learning methods, MetaKernel is applicable to versatile settings. We

TABLE 5
Deep Embedding Performance (%) on *miniImageNet* (5-way) Using a 28-Layer Wide Residual Feature Extractor

Method	1-shot	5-shot
META-SGD [104]	54.24±0.03	70.86±0.04
SNAIL [105]	55.71±0.99	68.88±0.92
GIDARISet al. [106]	56.20±0.86	73.00±0.64
BAUERet al. [107]	56.30±0.40	73.90±0.30
ADARESNET (DF) [52]	57.10±0.70	70.04±0.63
TADAM [39]	58.50±0.30	76.70±0.30
QIAOet al. [108]	59.60±0.41	73.54±0.19
LEO [12]	61.76±0.08	77.59±0.12
METAVRF [1]	63.80±0.05	77.97±0.28
METAKERNEL	65.03±0.03	79.01±0.28

The best performing methods and any other runs within the 95% confidence margin are denoted in bold. MetaKernel also outperforms previous methods for deeper networks.




Fig. 6. Efficiency with varying numbers D of bases. MetaKernel consistently achieves better performance than regular RFFS, especially with relatively low sampling rates.




Fig. 7. Versatility of MetaKernel with varied ways and shots on Omniglot.

evaluate the performance of MetaKernel on more challenging scenarios where the number of ways N and shots k between training and testing are inconsistent. Specifically, we test the performance of MetaKernel on Omniglot tasks with varied N and k , when it is trained on one particular N -way k -shot task. As shown in Fig. 7, the results demonstrate the trained model still produces good performance, even under the challenging conditions with a far higher number of ways. In particular, the model trained on the 20-way 5-shot task retains a high accuracy of 94% on the 100-way setting, as shown in Fig. 7a. The results also indicate that our model exhibits considerable robustness and flexibility to a variety of testing conditions.

5 CONCLUSION

In this paper, we introduced kernel approximation based on random Fourier features into the meta-learning framework for few-shot learning. We proposed to learn random features for each few-shot task in a data-driven way by formulating it as a variational inference problem, where the random Fourier basis is defined as the latent variable. We introduced an inference network based on an LSTM module, which enables the shared knowledge from related tasks to be incorporated into each individual task. To further enhance the kernels, we introduced conditional normalizing flows to generate richer posteriors over random bases, resulting in more informative random features. Experimental results on both regression and classification tasks demonstrated the effectiveness for few-shot learning. The extensive ablation study demonstrates the efficacy of each component in our MetaKernel.

In the literature, the few-shot learning setting typically assumes all tasks are sampled from the same task distribution and tasks are related to each other. Hence, there is no need for an explicit metric evaluating the relatedness between two tasks. Future work could investigate how to


collect more general shared knowledge in cross-domain few-shot learning [110] [111], where the the training tasks and the test task come from different distributions.

REFERENCES

- [1] X. Zhen *et al.*, "Learning to learn kernels with variational random features," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 11409–11419.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Informat. Process. Syst.*, 2012, pp. 1097–1105.
- [3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.
- [5] L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 4, pp. 594–611, Apr. 2006.
- [6] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, "Human-level concept learning through probabilistic program induction," *Science*, vol. 350, no. 6266, pp. 1332–1338, 2015.
- [7] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [8] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1126–1135.
- [9] J. Schmidhuber, "Learning to control fast-weight memories: An alternative to dynamic recurrent networks," *Neural Comput.*, vol. 4, no. 1, pp. 131–139, 1992.
- [10] S. Thrun and L. Pratt, *Learning to Learn*, Berlin, Germany: Springer, 2012.
- [11] M. Andrychowicz *et al.*, "Learning to learn by gradient descent by gradient descent," in *Proc. 30th Conf. Adv. Neural Informat. Process. Syst.*, 2016, pp. 3981–3989.
- [12] A. A. Rusu *et al.*, "Meta-learning with latent embedding optimization," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [13] O. Vinyals *et al.*, "Matching networks for one shot learning," in *Proc. 30th Int. Conf. Adv. Neural Informat. Process. Syst.*, 2016, pp. 3630–3638.
- [14] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. Adv. Neural Informat. Process. Syst.*, 2017, pp. 4077–4087.
- [15] J. Gordon, J. Bronskill, M. Bauer, S. Nowozin, and R. E. Turner, "Meta-learning probabilistic inference for prediction," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [16] Y. Du, X. Zhen, L. Shao, and C. G. M. Snoek, "MetaNorm: Learning to normalize few-shot batches across domains," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [17] B. Schölkopf and A. J. Smola, *Learning with Kernels*. Cambridge, MA, USA: MIT Press, 2002.
- [18] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA, USA: MIT Press, 2018.
- [19] T. Hofmann, B. Schölkopf, and A. J. Smola, "Kernel methods in machine learning," *The Ann. Statist.*, vol. 36, no. 3, pp. 1171–1220, 2008.
- [20] N. Cristianini *et al.*, *An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [21] A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statist. Comput.*, vol. 14, pp. 199–222, 2004.
- [22] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *Proc. 20th Int. Conf. Adv. Neural Informat. Process. Syst.*, 2007, pp. 1177–1184.
- [23] A. Sinha and J. C. Duchi, "Learning kernels with random features," in *Proc. 30th Int. Conf. Adv. Neural Informat. Process. Syst.*, 2016, pp. 1298–1306.
- [24] F. R. Bach, G. R. Lanckriet, and M. I. Jordan, "Multiple kernel learning, conic duality, and the smo algorithm," in *Proc. Int. Conf. Mach. Learn.*, 2004, Art. no. 6.
- [25] J. Hensman, N. Durrande, and A. Solin, "Variational fourier features for gaussian processes," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 5537–5588, 2017.

- [26] L. Carratino, A. Rudi, and L. Rosasco, "Learning with SGD and random features," in *Proc. 32nd Int. Conf. Adv. Neural Informat. Process. Syst.*, 2018, pp. 10192–10203.
- [27] B. Bullins, C. Zhang, and Y. Zhang, "Not-so-random features," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [28] C.-L. Li, W.-C. Chang, Y. Mroueh, Y. Yang, and B. Póczos, "Implicit kernel learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2019, pp. 2007–2016.
- [29] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [30] L. Dinh, D. Krueger, and Y. Bengio, "NICE: Non-linear independent components estimation," 2014, *arXiv:1410.8516*.
- [31] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using real NVP," *Proc. Int. Conf. Learn. Representations*, 2017.
- [32] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1530–1538.
- [33] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1×1 convolutions," in *Proc. Adv. Neural Informat. Process. Syst.*, 2018, pp. 10236–10245.
- [34] C. Winkler, D. Worrall, E. Hoogeboom, and M. Welling, "Learning likelihoods with conditional normalizing flows," 2019, *arXiv:1912.00042*.
- [35] E. Triantafyllou *et al.*, "Meta-dataset: A dataset of datasets for learning to learn from few examples," 2019, *arXiv:1903.03096*.
- [36] A. Rajeswaran, C. Finn, S. Kakade, and S. Levine, "Meta-learning with implicit gradients," 2019, *arXiv:1909.04630*.
- [37] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-learning in neural networks: A survey," 2020, *arXiv:2004.05439*.
- [38] K. R. Allen, E. Shelhamer, H. Shin, and J. B. Tenenbaum, "Infinite mixture prototypes for few-shot learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 232–241.
- [39] B. Oreshkin, P. R. López, and A. Lacoste, "Tadam: Task dependent adaptive metric for improved few-shot learning," in *Proc. Adv. Neural Informat. Process. Syst.*, 2018, pp. 721–731.
- [40] S. W. Yoon, J. Seo, and J. Moon, "Tapnet: Neural network augmented with task-adaptive projection for few-shot learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 7115–7123.
- [41] V. Garcia and J. Bruna, "Few-shot learning with graph neural networks," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [42] T. Cao, M. Law, and S. Fidler, "A theoretical analysis of the number of shots in few-shot learning," 2019, *arXiv:1909.11722*.
- [43] X. Zhen, Y. Du, H. Xiong, Q. Qiu, C. G. M. Snoek, and L. Shao, "Learning to learn variational semantic memory," in *Proc. 34th Int. Conf. Adv. Neural Informat. Process. Syst.*, 2020, pp. 9122–9134.
- [44] L. Zintgraf, K. Shiarli, V. Kurin, K. Hofmann, and S. Whiteson, "Fast context adaptation via meta-learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 7693–7702.
- [45] Y. Chen *et al.*, "Learning to learn without gradient descent by gradient descent," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 748–756.
- [46] H. Edwards and A. Storkey, "Towards a neural statistician," 2016, *arXiv:1606.02185*.
- [47] C. Finn, K. Xu, and S. Levine, "Probabilistic model-agnostic meta-learning," in *Proc. Adv. Neural Informat. Process. Syst.*, 2018, pp. 9516–9527.
- [48] S. Sæmundsson, K. Hofmann, and M. P. Deisenroth, "Meta reinforcement learning with latent variable gaussian processes," 2018, *arXiv:1803.07551*.
- [49] L. Bertinetto, J. F. Henriques, P. H. Torr, and A. Vedaldi, "Meta-learning with differentiable closed-form solvers," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [50] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap, "Meta-learning with memory-augmented neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1842–1850.
- [51] T. Munkhdalai and H. Yu, "Meta networks," in *Proc. 34th Int. Conf. Int. Conf. Mach. Learn.*, 2017, pp. 2554–2563.
- [52] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler, "Rapid adaptation with conditionally shifted neurons," 2017, *arXiv:1712.09926*.
- [53] C. M. Bishop, *Pattern Recognition and Machine Learning*. Berlin, Germany: Springer, 2006.
- [54] N. Shervashidze, P. Schweitzer, E. J. v. Leeuwen, K. Mehlhorn, and K. M. Borgwardt, "Weisfeiler-Lehman graph kernels," *J. Mach. Learn. Res.*, vol. 12, no. 9, pp. 2539–2561, 2011.
- [55] M. Gönen and E. Alpaydm, "Multiple kernel learning algorithms," *J. Mach. Learn. Res.*, vol. 12, pp. 2211–2268, 2011.
- [56] D. Duvenaud, J. R. Lloyd, R. Grosse, J. B. Tenenbaum, and Z. Ghahramani, "Structure discovery in nonparametric regression through compositional kernel search," 2013, *arXiv:1302.4922*.
- [57] T. Gärtner, P. A. Flach, A. Kowalczyk, and A. J. Smola, "Multi-instance kernels," in *Proc. 19th Int. Conf. Mach. Learn.*, 2002, pp. 179–186.
- [58] W. Rudin, *Fourier Analysis on Groups*. Hoboken, NJ, USA: Wiley, 1962.
- [59] A. Wilson and R. Adams, "Gaussian process kernels for pattern discovery and extrapolation," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 1067–1075.
- [60] Z. Yang, A. Wilson, A. Smola, and L. Song, "A la carte-learning fast kernels," in *Proc. 18th Int. Conf. Artif. Intell. Statist.*, 2015, pp. 1098–1106.
- [61] H. Avron, V. Sindhwani, J. Yang, and M. W. Mahoney, "Quasi-monte carlo feature maps for shift-invariant kernels," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 4096–4133, 2016.
- [62] W.-C. Chang, C.-L. Li, Y. Yang, and B. Póczos, "Data-driven random fourier features using stein effect," 2017, *arXiv:1705.08525*.
- [63] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, "Normalizing flows for probabilistic modeling and inference," *J. Mach. Learn. Res.*, vol. 22, no. 57, pp. 1–64, 2021.
- [64] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, "Improved variational inference with inverse autoregressive flow," *Adv. Neural Informat. Process. Syst.*, vol. 29, pp. 4743–4751, 2016.
- [65] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation," in *Proc. Adv. Neural Informat. Process. Syst.*, 2017, pp. 2335–2344.
- [66] R. T. Chen, J. Behrmann, D. Duvenaud, and J.-H. Jacobsen, "Residual flows for invertible generative modeling," in *Proc. 33rd Int. Conf. Adv. Neural Informat. Process. Syst.*, 2019, pp. 9916–9926.
- [67] J. Ho, X. Chen, A. Srinivas, Y. Duan, and P. Abbeel, "Flow++: Improving flow-based generative models with variational dequantization and architecture design," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 2722–2730.
- [68] P. Esling *et al.*, "Universal audio synthesizer control with normalizing flows," 2019, *arXiv:1907.00971*.
- [69] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2019, pp. 3617–3621.
- [70] J.-H. Jacobsen, A. Smeulders, and E. Oyallon, "i-RevNet: Deep invertible networks," in *Proc. Int. Conf. on Learn. Representations*, 2018.
- [71] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Proc. 28th Int. Conf. Adv. Neural Informat. Process. Syst.*, 2015, pp. 3483–3491.
- [72] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013, *arXiv:1312.6114*.
- [73] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic back-propagation and approximate inference in deep generative models," 2014, *arXiv:1401.4082*.
- [74] F. A. Gers and J. Schmidhuber, "Recurrent nets that time and count," in *Proc. IEEE-INNS-ENNS Int. Joint Conf. Neural Netw.*, 2000, pp. 189–194.
- [75] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Trans. Signal Process.*, vol. 45, no. 11, pp. 2673–2681, Nov. 1997.
- [76] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional LSTM and other neural network architectures," *Neural Netw.*, vol. 18, no. 5–6, pp. 602–610, 2005.
- [77] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. R. Salakhutdinov, and A. J. Smola, "Deep sets," in *Proc. 31st Int. Conf. Adv. Neural Informat. Process. Syst.*, 2017, pp. 3394–3404.
- [78] H. Kim *et al.*, "Attentive neural processes," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [79] A. Krizhevsky, "Learning multiple layers of features from tiny images," Univ. Toronto, Toronto, ON, Canada, Tech. Rep., 2009.
- [80] X. Peng, Q. Bai, X. Xia, Z. Huang, K. Saenko, and B. Wang, "Moment matching for multi-source domain adaptation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 1406–1415.
- [81] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *Proc. Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.


- [82] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi, "Fine-grained visual classification of aircraft," 2013, *arXiv:1306.5151*.
- [83] J. Jongejan, H. Rowley, T. Kawashima, J. Kim, and N. Fox-Gieg, "The quick, draw! – a.i. experiment," 2016. [Online]. Available: quickdraw.withgoogle.com
- [84] B. Schroeder and Y. Cui, "FGVCx fungi classification challenge," 2018. [Online]. Available: github.com/visipedia/fgvcx_fungi_comp
- [85] S. Houben, J. Stallkamp, J. Salmen, M. Schlipsing, and C. Igel, "Detection of traffic signs in real-world images: The German traffic sign detection benchmark," in *Proc. Int. Joint Conf. Neural Netw.*, 2013, pp. 1–8.
- [86] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [87] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD Birds-200–2011 Dataset," California Institute of Technology, Pasadena, CA, USA, Tech. Rep. CNS-TR-2011-001, 2011.
- [88] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi, "Describing textures in the wild," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 3606–3613.
- [89] M.-E. Nilsback and A. Zisserman, "Automated flower classification over a large number of classes," in *Proc. Indian Conf. Comput. Vis., Graph. Image Process.*, 2008, pp. 722–729.
- [90] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 1199–1208.
- [91] K. Allen, E. Shelhamer, H. Shin, and J. Tenenbaum, "Infinite mixture prototypes for few-shot learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 232–241.
- [92] A. Devos, S. Chatel, and M. Grossglauser, "Reproducing meta-learning with differentiable closed-form solvers," in *Proc. Int. Conf. Learn. Representations.*, 2019.
- [93] Y. Liu *et al.*, "Learning to propagate labels: Transductive propagation network for few-shot learning," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [94] F. Hao, F. He, J. Cheng, L. Wang, J. Cao, and D. Tao, "Collect and select: Semantic alignment metric learning for few-shot learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 8459–8468.
- [95] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4367–4375.
- [96] B. Oreshkin, P. López, and A. Lacoste, "Tadam: Task dependent adaptive metric for improved few-shot learning," in *Proc. 32nd Int. Conf. Adv. Neural Informat. Process. Syst.*, 2018, pp. 719–729.
- [97] Q. Sun, Y. Liu, T. Chua, and B. Schiele, "Meta-transfer learning for few-shot learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 403–412.
- [98] J. Zhang, C. Zhao, B. Ni, M. Xu, and X. Yang, "Variational few-shot learning," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 1685–1694.
- [99] S. Yoon, J. Seo, and J. Moon, "TapNet: Neural network augmented with task-adaptive projection for few-shot learning," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 7115–7123.
- [100] K. Lee, S. Maji, A. Ravichandran, and S. Soatto, "Meta-learning with differentiable convex optimization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 10649–10657.
- [101] C. Simon, P. Koniusz, R. Nock, and M. Harandi, "Adaptive subspaces for few-shot learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 4135–4144.
- [102] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola, "Rethinking few-shot image classification: A good embedding is all you need?," 2020, *arXiv:2003.11539*.
- [103] F. X. Yu, A. T. Suresh, K. M. Choromanski, D. N. Holtmann-Rice, and S. Kumar, "Orthogonal random features," in *Proc. Adv. Neural Informat. Process. Syst.*, 2016, pp. 1975–1983.
- [104] Z. Li, F. Zhou, F. Chen, and H. Li, "Meta-sgd: Learning to learn quickly for few-shot learning," 2017, *arXiv:1707.09835*.
- [105] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel, "A simple neural attentive meta-learner," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [106] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4367–4375.
- [107] M. Bauer, M. Rojas-Carulla, J. B. Świątkowski, B. Schölkopf, and R. E. Turner, "Discriminative k-shot learning using probabilistic models," 2017, *arXiv:1706.00326*.
- [108] S. Qiao, C. Liu, W. Shen, and A. L. Yuille, "Few-shot image recognition by predicting parameters from activations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7229–7238.
- [109] S. Zagoruyko and N. Komodakis, "Wide residual networks," 2016, *arXiv:1605.07146*.
- [110] H.-Y. Tseng, H.-Y. Lee, J.-B. Huang, and M.-H. Yang, "Cross-domain few-shot classification via learned feature-wise transformation," 2020, *arXiv:2001.08735*.
- [111] Y. Guo *et al.*, "A broader study of cross-domain few-shot learning," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 124–141.



Yingjun Du received his master's degree from the College of Software, Beihang University, Beijing, China, in 2020. He is currently working toward the PhD degree in Informatics Institute, University of Amsterdam, The Netherlands. His research interests include meta learning, few-shot learning, variational inference, domain generalization, and neural memory, etc.



Haoliang Sun received the BE degree from Shandong University, Shandong, China, in 2014, and the PhD degree from Shandong University, Shandong, China, in 2020. He was supported by China Scholarship Council (CSC) and spent one and half years with the University of Wisconsin-Madison. He currently pursues a postdoctoral researcher with Shandong University. His research interests span both machine learning and computer vision, including meta-learning, flow-based generative model, few-shot learning, and kernels.




Xiantong Zhen received the BS and ME degrees from Lanzhou University, Lanzhou, China, in 2007 and 2010, respectively, and the PhD degree from the Department of Electronic and Electrical Engineering, the University of Sheffield, U.K., in 2013. He is with Inception Institute of Artificial Intelligence (IIAI) in the United Arab Emirates and College of Computer Science. He is currently an assistant professor with the University of Amsterdam. His research interests include machine learning and computer vision.



Jun Xu received the BSc degree in pure mathematics and the MSc degree in information and probability both from the School of Mathematics Science, Nankai University, China, in 2011 and 2014, respectively, and the PhD degree from the Department of Computing, The Hong Kong Polytechnic University, in 2018. He is currently an assistant professor in the School of Statistics and Data Science, Nankai University, Tianjin, China. He worked as a research scientist in Inception Institute of Artificial Intelligence (IIAI), Abu Dhabi, UAE. For more information, please visit <https://csjunxu.github.io/>.



Yilong Yin received the PhD degree from Jilin University, Changchun, China, in 2000. He is currently the director of the Machine Learning and Data Mining Group and a professor with Shandong University, Jinan, China. From 2000 to 2002, he was a postdoctoral fellow with the Department of Electronic Science and Engineering, Nanjing University, Nanjing, China. His research interests include machine learning, data mining, and computer vision.



Ling Shao (Fellow, IEEE) is currently the executive vice president and provost of the Mohamed bin Zayed University of Artificial Intelligence. He is also the CEO and chief scientist of the Inception Institute of Artificial Intelligence (IIAI), Abu Dhabi, United Arab Emirates. His research interests include computer vision, machine learning, and medical imaging. He is a fellow of IAPR, IET, and BCS.



Cees G. M. Snoek (Senior Member, IEEE) received the MSc degree in business information systems and the PhD degree in computer science both from the University of Amsterdam, The Netherlands, in 2000 and 2005, respectively. He is currently a full professor in computer science with the University of Amsterdam, where he heads the Video & Image Sense Lab. He is also a director of three public-private AI research labs: QUVA Lab with Qualcomm, Atlas Lab with TomTom and AIM Lab with the Inception Institute of Artificial Intelligence. He currently serves as an associate editor for *Computer Vision and Image Understanding* and *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.**