

Learning to Learn with Variational Information Bottleneck for Domain Generalization

Yingjun Du¹[0000-0001-7537-6457], Jun Xu², Huan Xiong⁴, Qiang Qiu⁵,
Xiantong Zhen^{1,3}, Cees G. M. Snoek¹, and Ling Shao^{3,4}

¹ AIM Lab, University of Amsterdam, The Netherlands

² College of Computer Science, Nankai University, China

³ Inception Institute of Artificial Intelligence, Abu Dhabi, UAE

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

⁵ Electrical and Computer Engineering, Duke University, USA

{y.du, x.zhen, cgmsnoek}@uva.nl, nankaimathxujun@gmail.com,
huan.xiong@mbzuai.ac.ae, qiang.qiu@duke.edu, ling.shao@ieee.org

Abstract. Domain generalization models learn to generalize to previously unseen domains, but suffer from prediction uncertainty and domain shift. In this paper, we address both problems. We introduce a probabilistic meta-learning model for domain generalization, in which classifier parameters shared across domains are modeled as distributions. This enables better handling of prediction uncertainty on unseen domains. To deal with domain shift, we learn domain-invariant representations by the proposed principle of meta variational information bottleneck, we call MetaVIB. MetaVIB is derived from novel variational bounds of mutual information, by leveraging the meta-learning setting of domain generalization. Through episodic training, MetaVIB learns to gradually narrow domain gaps to establish domain-invariant representations, while simultaneously maximizing prediction accuracy. We conduct experiments on three benchmarks for cross-domain visual recognition. Comprehensive ablation studies validate the benefits of MetaVIB for domain generalization. The comparison results demonstrate our method outperforms previous approaches consistently.

Keywords: Meta Learning, Domain Generalization, Variational Inference, Information Bottleneck

1 Introduction

This paper strives for domain generalization in image classification [24, 28, 31, 49]. The general challenge is to exploit the data variations of seen image domains with the aim to generalize well to unseen image domains. For example, by generalizing a *chair* classifier trained on PASCAL VOC to LabelMe [49], or by generalizing an *elephant* classifier trained on photo’s to sketches [24]. Domain generalization models typically suffer from two problems. First, since data from unseen domains is inaccessible during the learning stage, we do not know their statistical data distribution. This causes uncertainty in the predictions made on

the unseen domains. Second, data from different domains usually follows distinct distributions with great discrepancy, resulting in domain shift from seen to unseen domains. Domain shift has been extensively researched in domain generalization, mostly by learning feature representations that are invariant across domains [12, 17, 25, 27, 29, 31, 52]. Meta-learning [41, 46] that learns to generalize across tasks has been introduced to domain generalization by Li et al. [25] showing its great effectiveness in learning to generalize across domains [4, 25, 29]. To the best of our knowledge, none of these existing meta-learning methods deal with the prediction uncertainty on unseen domains.

In this paper, we address the two major domain generalization challenges jointly by one single probabilistic model under the meta-learning framework. We model parameters of classifiers shared across domains as probabilistic distributions that we infer from the data of the seen domains. The probabilistic modeling enables us to better handle the prediction uncertainty on previously unseen domains [15, 18]. To handle domain shift, we take inspiration from the information bottleneck (IB) theory [1, 2, 47] which learns robust representations to enhance generalization. IB encodes the input into compressed intermediate representations that maximize target prediction. It offers a promising technique to learn domain-invariant representations, but to the best of our knowledge has not yet been explored for domain generalization under the meta-learning framework. We propose the principle of meta variational information bottleneck (MetaVIB) for the optimization of the model. We derive MetaVIB from the variational bounds of mutual information by leveraging the meta-learning setting, and incorporate it as a data-driven regularizer into the optimization objective. The parameters of all classifiers and the network are jointly optimized during the meta-training stage and applied to the unseen domain in the meta-test stage. By episodic training, MetaVIB enables the network to learn to gradually close the gaps between domains to achieve domain-invariant representations that alleviate domain shift, while simultaneously being able to obtain accurate predictions.

We conduct extensive experiments on three benchmarks for cross-domain visual recognition. The ablation studies demonstrate the benefits of MetaVIB in the probabilistic framework for domain generalization. The comparison with state-of-the-art methods, shows that our method consistently delivers the best performance on all tasks, surpassing previous methods based on both regular learning and meta-learning.

2 Related Work

In this section, we review related work on domain generalization, information bottleneck and meta-learning.

Domain generalization has been a longstanding challenge in computer vision [24, 25, 27] and machine learning [6, 31], but recently regained increased research interest [4, 8, 10, 29, 42]. Learning domain-variant feature representation has been one of the main topics of focus in domain generalization [12, 17, 27, 28, 31, 52]. The core idea is to learn a model that generates invariant representations

for the source domains, without over-fitting, which generalizes to unseen target domains. We explore the domain discrepancy to learn invariant representations through the lens of mutual information [47].

Information bottleneck (IB) [47] provides an information-theoretic principle of encoding the input data into a compressed representation that maximizes target prediction. This is achieved by minimizing the mutual information $I(Z; X)$ between the input variable X and its latent representation Z , while maximizing the mutual information $I(Z; Y)$ between the output variable Y and the latent representation Z . To be more precise, the IB principle is to maximize the objective function:

$$\mathcal{L}_{\text{IB}}(\theta) = I(Z; Y|\theta) - \beta I(Z; X|\theta), \quad (1)$$

where $\beta \in [0, 1]$ is the hyperparameter that controls the size of the information bottleneck, and θ are the corresponding model parameters.

The IB principle has recently been introduced for theoretical understanding and analysis of deep neural networks [2, 21, 34, 43, 48]. The authors optimize the networks with an iterative Blahut-Arimoto algorithm, which is infeasible in practical systems. Alemi et al. [1] developed a variational approximation to the IB objective by leveraging variational inference, which allows the IB model to be parameterized with neural networks. Amjad et al. [2] investigated training deep neural networks (DNN) for classification based on minimization of the IB functional.

It is shown that for deterministic DNNs, the optimization can be ill-posed. This is because the IB functional can be infinite or not admitting gradient descent since it is piece-wise constant. The possible remedy indicated in their work is to train stochastic DNNs with the IB principle.

Meta-learning, or learning to learn, endows models with the capacity to efficiently learn new tasks by acquiring common knowledge through experiencing a set of related tasks. It has been explored in several directions, e.g., by learning a meta learner on diverse tasks to adapt the parameters of the base learner on a specific task [14, 38, 44, 45, 51, 53], learning to optimize the parameters of deep neural networks [5, 32, 40], and learning to learn the gradient optimization process by recurrent neural networks [3, 35], *etc.* A representative meta-learning algorithm is the model agnostic meta-learning (MAML), which learns the models to be able to adapt to similar tasks with only a few gradient descent updates. Li et al. [25] introduced the idea of MAML [14] to domain generalization. They train models with generalization ability to unseen domains by leveraging the meta-learning setting. MetaReg [4] addresses the domain shifts by leveraging the insights from meta-learning [50]. They learn a meta regularizer to achieve the generalization from source to unseen target domains. Li et al. [29] proposed a meta-learning approach based on a feature-critic network, in which an auxiliary loss is introduced to improve generalization ability. Dou et al. [10] adopt a gradient-based model-agnostic learning algorithm to deal with domain shift for domain generalization. Two complementary losses are introduced for regularization of semantic features. The success of those works has indicated the effectiveness of meta-learning in

domain generalization. Probabilistic meta-learning has also been developed in few-shot learning to handle uncertainty [16, 18], which has not been explored for domain generalization.

3 Method

We describe the meta-learning setting for domain generalization. Following the setting in recent domain generalization by meta-learning [4, 25, 29], we divide a dataset into the Source domains \mathcal{S} used for train and the Target domains \mathcal{T} held-out for test. In the train phase, data in the source domains \mathcal{S} is episodically divided into sets of meta-train \mathcal{D}^s and meta-test \mathcal{D}^t domains. We train the model by optimizing over the prediction errors on meta-test \mathcal{D}^t domains. In the test phase, the learned model is applied to the target domains \mathcal{T} for performance evaluation. The training phase incorporates the idea of meta-learning which induces a higher level of learning by the split of meta-train and meta-test domains, rather than training on all source domains [26]. This episodic meta-learning process mimics the generalization from seen to previously unseen domains.

3.1 Probabilistic Modeling

We start with the probabilistic formulation of the domain generalization, based on which we develop the probabilistic model under the meta-learning framework. We consider the general estimation problem of conditionally predictive likelihood in the meta-test domain \mathcal{D}^t :

$$\max_{p(\mathbf{x}^t, \mathbf{y}^t)} \mathbb{E} [\log \int p(\mathbf{y}^t | \psi, \mathbf{x}^t) p(\psi | \mathbf{x}^t) d\psi], \quad (2)$$

where $(\mathbf{x}^t, \mathbf{y}^t)$ is the sample of paired input and label drawn from data distribution $p(\mathbf{x}^t, \mathbf{y}^t)$ in meta-test domain, $p(\mathbf{y}^t | \psi, \mathbf{x}^t)$ is the conditionally predictive distribution, ψ is the parameter set of the classifier. Note that we treat ψ as a stochastic variable that depends on the input \mathbf{x}^t and the optimization of (2) is with respect to the parameters of probabilities.

In this work, we parameterize the model by deep neural networks. From the information-theoretic point of view [1], we regard the feature representation from the neural network as a stochastic variable \mathbf{z}^t , which is the latent encoding of the input \mathbf{x}^t . In domain generalization, it is commonly assumed that the label space is shared across the source and target domains. By leveraging the meta-learning setting, we propose to use data D^s from the meta-train domains to estimate the parameters of the classifier by replacing $p(\psi | \mathbf{x}^t)$ with $q(\psi | D^s)$, which is applied to the meta-test domain. By incorporating the latent variable \mathbf{z} into (2), we obtain the following maximum conditionally predictive likelihood estimation,

$$\max_{p(\mathbf{x}^t, \mathbf{y}^t)} \mathbb{E} [\log \int p(\mathbf{y}^t | \psi, \mathbf{z}^t) p(\mathbf{z}^t | \mathbf{x}^t) q(\psi | D^s) d\mathbf{z} d\psi]. \quad (3)$$




Fig. 1. Computational graph of the probabilistic meta-learning model for domain generalization. Θ encloses the global model parameters and ψ contains the parameters of classifiers shared across domains. Θ and ψ are jointly optimized in the train phase on the source domains. In each episode, the source domain is divided into a meta-train (\mathcal{D}^s) and meta-test (\mathcal{D}^t) domain. ψ is produced by \mathcal{D}^s and applied to \mathcal{D}^t . In the test phase, the model (Θ) generates representations of data in the target domains and the classifier (ψ) predicts of data from the source domain.

This establishes a probabilistic latent model which can be represented in a computational graph as shown in Fig. 1, and the corresponding conditional joint distribution is defined as:

$$p(Y^t, Z^t, \psi | X^t, D^s, \Theta) = p(\psi | D^s; \Theta) \prod_{n=1}^N p(\mathbf{y}_n^t | \psi, \mathbf{z}_n^t) p(\mathbf{z}_n^t | \mathbf{x}_n^t; \Theta), \quad (4)$$

where Θ denotes the model parameters, $D^t = \{X^t, Y^t\} = \{\mathbf{x}_n^t, \mathbf{y}_n^t\}_{n=1}^N$, $D^s = \{X^s, Y^s\} = \{\mathbf{x}_m^s, \mathbf{y}_m^s\}_{m=1}^M$, and N (M) are the number of samples in the meta-test (meta-train) domains. It is possible to directly employ (3) as the optimization objective using the techniques of amortized inference [20, 36]. However, the learned representations \mathbf{z} would not be domain invariant, which is desired for domain generalization. To achieve domain-invariant representations, we resort to the information bottleneck (IB) principle [1, 48], which will be incorporated into the objective as a regularizer for joint optimization.

3.2 Meta Variational Information Bottleneck

We introduce the IB principle to learn domain-invariant representations under the meta-learning framework. We impose the information bottleneck on the feature representations to control the information flow in deep neural networks. This should largely remove domain related information while letting through the information that maximizes prediction of labels on the meta-test domain.

We derive new variational bounds of mutual information by leveraging the setting of meta-learning for domain generalization. This gives rise to a meta version of variational information bottleneck, which we call MetaVIB in contrast to its original form in a standard learning framework [1]. To avoid confusion, we omit the superscript t for the meta-test domain in this subsection.

Let the random variables X , Y , and Z denote the input, output, and the intermediate feature representation in the deep neural network, which encodes

X . The mutual information $I(Z; Y)$ between the latent encoding Z of data X and its output label Y is defined as follows:

$$I(Z; Y) = \int p(\mathbf{y}, \mathbf{z}) \log \frac{p(\mathbf{y}, \mathbf{z})}{p(\mathbf{y})p(\mathbf{z})} d\mathbf{y}d\mathbf{z} = \int p(\mathbf{y}, \mathbf{z}) \log \frac{p(\mathbf{y}|\mathbf{z})}{p(\mathbf{y})} d\mathbf{y}d\mathbf{z}. \quad (5)$$

Since $p(\mathbf{y}|\mathbf{z})$ is intractable, we introduce $q(\mathbf{y}|\mathbf{z}, \psi)$ to be a variational approximation of $p(\mathbf{y}|\mathbf{z})$, where conditioning on the classifier parameter ψ is indicated by (4), and the prior distribution of ψ is denoted as $p(\psi)$. Then we have:

$$D_{\text{KL}}[p(\mathbf{y}|\mathbf{z})||q(\mathbf{y}|\mathbf{z}, \psi)] = \int p(\mathbf{y}|\mathbf{z}) \log \frac{p(\mathbf{y}|\mathbf{z})}{q(\mathbf{y}|\mathbf{z}, \psi)} d\mathbf{y} \geq 0, \quad (6)$$

which leads to

$$I(Z; Y) \geq \int p(\mathbf{y}, \mathbf{z}) \log q(\mathbf{y}|\mathbf{z}, \psi) d\mathbf{y}d\mathbf{z} + H(Y), \quad (7)$$

where $H(Y) = -\int p(\mathbf{y}) \log p(\mathbf{y}) d\mathbf{y}$ is the entropy of Y . Taking expectation values of both sides with respect to $\psi \sim p(\psi)$, we have

$$\begin{aligned} I(Z; Y) - H(Y) &\geq \mathbb{E}_{\psi \sim p(\psi)} \int p(\mathbf{y}, \mathbf{z}) \log q(\mathbf{y}|\mathbf{z}, \psi) d\mathbf{y}d\mathbf{z} \\ &= \int p(\psi) p(\mathbf{y}, \mathbf{z}) \log q(\mathbf{y}|\mathbf{z}, \psi) d\mathbf{y}d\mathbf{z}d\psi. \end{aligned} \quad (8)$$

Note that the entropy $H(Y)$ is independent of our optimization procedure and can thus be ignored. By replacing the prior $p(\psi)$ with a meta prior $q(\psi|D^s)$ conditioned on data D^s from the meta-train domains, leveraging the fact that $p(\mathbf{y}, \mathbf{z}) = \int p(\mathbf{y}, \mathbf{z}|\mathbf{x})p(\mathbf{x})d\mathbf{x} = \int p(\mathbf{y}|\mathbf{x})p(\mathbf{z}|\mathbf{x})p(\mathbf{x})d\mathbf{x}$, and ignoring the $H(Y)$ term, we obtain a new variational lower bound:

$$I(Z; Y) \geq \int p(\mathbf{x})p(\mathbf{y}|\mathbf{x})p(\mathbf{z}|\mathbf{x})q(\psi|D^s) \log q(\mathbf{y}|\psi, \mathbf{z}) d\mathbf{x}d\mathbf{y}d\mathbf{z}d\psi, \quad (9)$$

which is tractable in general by approximation [1].

Now we consider the second term $I(Z; X)$, which can be written as follows:

$$I(Z; X) = \int p(\mathbf{x}, \mathbf{z}) \log \frac{p(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} d\mathbf{x}d\mathbf{z}. \quad (10)$$

Instead of simply using an uninformative prior $p(\mathbf{z})$, we leverage the meta setting and introduce a meta prior $q(\mathbf{z}|D^s)$ as a variational approximation to $p(\mathbf{z})$. Due to the fact that $D_{\text{KL}}[p(Z)||q(Z|D^s)] > 0$, we obtain the following upper bound:

$$I(Z; X) \leq \int p(\mathbf{x})p(\mathbf{z}|\mathbf{x}) \log \frac{p(\mathbf{z}|\mathbf{x})}{q(\mathbf{z}|D^s)} d\mathbf{x}d\mathbf{z}. \quad (11)$$

By combining the two bounds (9) and (11), we establish the meta variational information bottleneck (MetaVIB)

$$\begin{aligned} \mathcal{L}_{\text{IB}} \geq & \int p(\mathbf{x})p(\mathbf{y}|\mathbf{x})p(\mathbf{z}|\mathbf{x})p(\psi|D^s) \log q(\mathbf{y}|\mathbf{z}, \psi) d\mathbf{x} d\mathbf{y} d\mathbf{z} d\psi \\ & - \beta \int p(\mathbf{x})p(\mathbf{z}|\mathbf{x}) \log \frac{p(\mathbf{z}|\mathbf{x})}{q(\mathbf{z}|D^s)} d\mathbf{x} d\mathbf{z} = \mathcal{L}_{\text{MetaVIB}} \end{aligned} \quad (12)$$

which extends the IB theory [47] into the meta-learning scenario, offering a new principle of learning domain-invariant representations for domain generalization.

We follow [1] to approximate $p(\mathbf{x}, \mathbf{y}) = p(\mathbf{x})p(\mathbf{y}|\mathbf{x})$ and $p(\mathbf{x})$ with empirical data distribution $p(\mathbf{x}, \mathbf{y}) = \frac{1}{N} \sum_{n=1}^N \delta_{\mathbf{x}_n}(\mathbf{x})\delta_{\mathbf{y}_n}(\mathbf{y})$ and $p(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N \delta_{\mathbf{x}_n}(\mathbf{x})$, where N is the number of samples in the meta-test domain. This essentially regards the data points $(\mathbf{x}_n, \mathbf{y}_n)$ and \mathbf{x}_n as the samples drawn from the data distributions $p(\mathbf{x}, \mathbf{y})$ and $p(\mathbf{x})$, respectively.

Thus, the approximated lower bound $\tilde{\mathcal{L}}_{\text{MetaVIB}}$ in practice can be written as:

$$\begin{aligned} \tilde{\mathcal{L}}_{\text{MetaVIB}} = & \frac{1}{N} \sum_{n=1}^N \int [p(\mathbf{z}_n|\mathbf{x}_n)p(\psi|D^s) \log q(\mathbf{y}_n|\mathbf{z}_n, \psi) \\ & - \beta p(\mathbf{z}_n|\mathbf{x}_n) \log \frac{p(\mathbf{z}_n|\mathbf{x}_n)}{q(\mathbf{z}_n|D^s)}] d\mathbf{z}_n d\psi. \end{aligned} \quad (13)$$

We use Monte Carlo sampling to draw samples from $p(\psi|D^s)$ for ψ and from $p(\mathbf{z}_n|\mathbf{x}_n)$ for \mathbf{z}_n in the lower bound of MetaVIB in (13). We attain the following objective function:

$$\begin{aligned} \mathcal{L} = & - \frac{1}{NC} \sum_{c=1}^C \sum_{n=1}^{N_c} \left(\frac{1}{L_{\mathbf{z}} L_{\psi}} \sum_{\ell_{\mathbf{z}}=1}^{L_{\mathbf{z}}} \sum_{\ell_{\psi}=1}^{L_{\psi}} \log q(\mathbf{y}_n|\mathbf{z}^{(\ell_{\mathbf{z}})}, \psi_c^{(\ell_{\psi})}) \right. \\ & \left. + \beta D_{\text{KL}} [p(\mathbf{z}|\mathbf{x}_n)||q(\mathbf{z}|D_c^s)] \right). \end{aligned} \quad (14)$$

where C is the number of classes and D_c^s contains the samples from the c -th category in the meta-train domains. We amortize the posterior distribution $q(\psi|D_c^s)$ and the meta prior $q(\mathbf{z}_n|D_c^s)$ across classes, that is, the variational distribution of each class is inferred individually by the samples from its corresponding class D_c^s , which further alleviates the computational overhead. In addition, the KL term can be calculated in a closed form. Here, to enable back-propagation, we adopt the re-parameterization trick [20], that is,

$$\mathbf{z}_n^{(\ell_{\mathbf{z}})} = f(\mathbf{x}_n, \epsilon^{(\ell_{\mathbf{z}})}), \quad \epsilon^{(\ell_{\mathbf{z}})} \sim \mathcal{N}(0, I) \quad (15)$$

and

$$\psi_c^{(\ell_{\psi})} = f(D_c^s, \epsilon^{(\ell_{\psi})}), \quad \epsilon^{(\ell_{\psi})} \sim \mathcal{N}(0, I) \quad (16)$$

where $f(\cdot)$ is a deterministic function which is usually parameterized by a multiple layer perception (MLP) and $L_{\mathbf{z}}$ and L_{ψ} are the number of samples for \mathbf{z}_n and ψ_c , respectively.

Taking a closer look at the objective (14), we observe that the first term is the negative log predictive likelihood in the meta-test domain, where the label \mathbf{y}_n of \mathbf{x}_n is predicted from its latent encoding \mathbf{z}_n and the classifier parameter ψ . Minimizing the first term guarantees maximal prediction accuracy. The second term is the KL divergence between distributions of latent encoding of the sample in the target domain and that estimated by the samples from the same category in the meta-train domains. It is the minimization of the KL term in (14) that enables the model to learn domain-invariant representations. This is in contrast to the regular IB principle [1, 48] which is to compress the input and does not necessarily result in domain-invariant representations.

3.3 Learning with Stochastic Neural Networks

We implement the proposed model by end-to-end learning with stochastic neural networks that are comprised of convolutional layers and fully-connected layers. The inference is parameterized by a feed-forward multiple layer perception (MLP). During the training phase, given K domains, we randomly sample one domain as the meta-test domain, the remaining $K - 1$ domains are used as the meta-train domains. Then we choose a batch of M samples $\{(\mathbf{x}_m^s, \mathbf{y}_m^s)\}_{m=1}^M$ from the meta-train domain \mathcal{D}^s , and a batch of N samples $\{(\mathbf{x}_n^t, \mathbf{y}_n^t)\}_{n=1}^N$ from the meta-test domain \mathcal{D}^t . Note that M samples from meta-train domains cover all the C classes. For each sample $\mathbf{x}_{m,c}^s$ of the c -th class, we first extract its features via $h_\theta(\mathbf{x}_{m,c}^s)$, where $h_\theta(\cdot)$ is the feature extraction network and we use permutation-invariant instance-pooling operations to get the mean feature $\bar{\mathbf{h}}_c^s$ of samples in the c -th class. The mean feature $\bar{\mathbf{h}}_c^s$ will be fed into a small MLP network $g_{\phi_1}(\cdot)$ to calculate the mean $\boldsymbol{\mu}_c^\psi$ and variance $\boldsymbol{\sigma}_c^\psi$ of the weight vector distribution ψ_c for c -th class, which is then used to sample the weight vector ψ_c of this class by $\psi_c \sim \mathcal{N}(\boldsymbol{\mu}_c^\psi, \text{diag}((\boldsymbol{\sigma}_c^\psi)^2))$. The weight vectors $\{\psi_c\}_{c=1}^C$ of all C classes are combined column by column to form a weight matrix $\psi = [\psi_1, \psi_2, \dots, \psi_C]$.

We calculate the parameters of the latent distribution, i.e., the mean $\boldsymbol{\mu}_c^s$ and variance $\boldsymbol{\sigma}_c^s$ of the c -th class in the meta-train domain by another small MLP network $g_{\phi_2}(\cdot)$. Then the parameter \mathbf{z}_c is sampled from the distribution $\mathbf{z}_c \sim \mathcal{N}(\boldsymbol{\mu}_c^s, \text{diag}((\boldsymbol{\sigma}_c^s)^2))$. For each sample $\mathbf{x}_{n,c}^t$ in the meta-test domain, we also calculate the mean $\boldsymbol{\mu}_{n,c}^t$ and variance $\boldsymbol{\sigma}_{n,c}^t$ of the distribution. Thus its latent coding vector $\mathbf{z}_{n,c}$ can be naturally sampled from $\mathbf{z}_{n,c} \sim \mathcal{N}(\boldsymbol{\mu}_{n,c}^t, \text{diag}((\boldsymbol{\sigma}_{n,c}^t)^2))$. Denote $\bar{\mathbf{h}}_c^s$ as the mean feature of all the samples of the c -th class from the meta-train domains, i.e., $\bar{\mathbf{h}}_c^s = \frac{1}{M_c} \sum_{m=1}^{M_c} \mathbf{x}_{m,c}^s$. We provide the detailed step-by-step algorithm of the proposed MetaVIB for training in the supplemental material.

4 Experiments

We conduct our experiments on three benchmarks commonly used in domain generalization [4, 25, 29, 42]. We first provide ablation studies to gain insights into the properties and benefits of MetaVIB. Then we compare with previous

Table 1. Benefit of MetaVIB under the probabilistic framework on VLCS

	VOC2007	LabelMe	Caltech-101	SUN09	Mean
AlexNet	68.41	62.11	93.40	64.16	72.02
Baseline	69.87 \pm 0.63	61.32 \pm 0.27	95.97 \pm 0.43	66.32 \pm 0.25	73.37
VIB	70.02 \pm 0.52	62.17 \pm 0.29	95.93 \pm 0.32	67.93 \pm 0.41	74.01
MetaVIB	70.28 \pm 0.71	62.66 \pm 0.35	97.37 \pm 0.63	67.85 \pm 0.17	74.54

methods based on both regular learning and meta-learning for domain generalization. We put more results in the supplementary material due to space limit.

4.1 Datasets

VLCS [49] is a real-world dataset that contains four domains collected from **VOC2007** [13], **LabelMe** [37], **Caltech-101** [19], and **SUN09** [9]. Images are from 5 classes, i.e., *bird*, *car*, *chair*, *dog*, *person*. The domain shift across those datasets makes VLCS a suitable benchmark for domain generalization.

PACS [24] contains 9991 images from 4 domains, i.e., **Photo**, **Art painting**, **Cartoon**, and **Sketch**, which cover huge domain gaps. Images are from 7 object classes, i.e., *dog*, *elephant*, *giraffe*, *guitar*, *horse*, *house*, and *person*.

Rotated MNIST [42] is a synthetic dataset consisting of 6 domains, each containing 1000 images of the 10 digits (i.e., $\{0, 1, \dots, 9\}$, 100 for each) randomly selected from the training set of MNIST [23], with 6 rotation degrees: 0° , 15° , 30° , 45° , 60° , and 75° .

4.2 Implementation Details

Splits, Metrics and Backbone On all datasets, we follow the train-test splits suggested by [4, 24, 25], and perform experiments with the “leave-one-domain-out” strategy: we take the samples from one domain as the target domain for testing, and the samples from the remaining domains as the source domain for training. We use the AlexNet [22] pre-trained on ImageNet and fine-tuned on the source domains of each dataset to perform testing on the target domain of that dataset. We use the average accuracy of all classes as the evaluation metric [17]. To benchmark previous methods, we employ the pre-trained AlexNet [22] on ImageNet as the backbone on VLCS and PACS. For Rotated MNIST we use a backbone network with two convolutions and one fully-connected layer. Even more implementation details about training stage, the feature extraction network and inference networks for different datasets are provided in the supplemental materials.

4.3 Ablation Study

To study the benefit of the MetaVIB under the probabilistic framework for domain generalization, we compare with several alternative models on VLCS and PACS in Tables 1 and 2.

Table 2. Benefit of MetaVIB under the probabilistic framework on PACS

	Photo	Art painting	Cartoon	Sketch	Mean
AlexNet	88.47	67.21	66.12	55.32	69.28
Baseline	90.32 \pm 0.35	68.12 \pm 0.51	70.25 \pm 0.17	61.81 \pm 0.26	72.63
VIB	90.17 \pm 0.28	69.93 \pm 0.34	71.01 \pm 0.27	62.37 \pm 0.42	73.37
MetaVIB	91.93 \pm 0.23	71.94 \pm 0.34	73.17 \pm 0.21	65.94 \pm 0.24	75.74

Benefit of probabilistic modeling To show the benefit of probabilistic modeling, we first consider AlexNet [22] which is pre-trained on ImageNet, fine-tuned on the source domains and applied to the target domains. We define our *Baseline* model as the probabilistic model that predicts parameter distributions of the classifiers, without regular VIB or MetaVIB. The probabilistic model outperforms the pre-trained AlexNet by 1.35% and 3.35% on the VLCS and PACS benchmarks. The results indicate that the classifiers learned by probabilistic modeling better generalize to the target domains.

Benefit of MetaVIB We show the benefit of MetaVIB by comparing with the regular VIB [1], which is applied to the baseline model as a regularization in the optimization, and the Baseline model. We first establish the probabilistic model with the regular *VIB* which performs better than the baseline (74.01% - up 0.64%) on VLCS and (73.37% - up 0.74%) on PACS. The VIB regularization term maximizes the mutual information between Z and the target Y , which will encourage better prediction performance compared to the *Baseline* model. However, our MetaVIB learns an even better domain-invariant representation, as it consistently outperforms VIB by up to 2.37% on **PACS** [24]. As indicated in the optimization objective in (14) minimizing the KL term makes the representations of samples in the meta-target domain to be close to the representations obtained by the samples of the same class from the meta-source domains. As a result, the learned model acquires the ability to generate domain-invariant representations by the episodic training. In contrast, the regular VIB is to simply compress the input with no explicit mechanism to narrow the gaps across domains. The obtained representations with regular VIB are not necessarily domain-invariant. Actually, there is no evident causal relation between compression and generalization as indicated in [39].

Influence of information bottleneck size β The bottleneck size β controls the amount of information flow that goes through the bottleneck of the networks. To measure its influence on the performance, we plot the information plane dynamics of different network layers with varying β in Fig. 2. We observe that MetaVIB with $\beta = 0.01$ achieves the highest $I(Z; Y)$ while at the same time $I(Z; X)$ is minimal. We also report the influence of β in Table 3. MetaVIB achieves the best performance when $\beta = 0.01$, which is consistent with the information dynamic in Fig. 2. We observe in Fig. 2 (c) that with $\beta = 0.01$, the




Fig. 2. Influence of information bottleneck size β on domain generalization for “Sketch” as the test domain on PACS. X , Y , and T denote input image, output target, and outputs per layer of the inference network that generates the latent encoding Z . The horizontal (vertical) axis plots mutual information between the features of each layer and the input (output). Each of the three layers of the inference network produces a curve in the information plane with the input layer at the far right and output layer at the far left. The color-scale denotes training iterations from 0 to 15,000. The mutual information of different layers in the same iteration are connected by fine lines. Compared to other values of β , for $\beta = 0.01$, $I(T; Y)$ reaches the highest value, which explains the best performance.

Table 3. Influence of information bottleneck size β on domain generalization for PACS. MetaVIB obtains best results for $\beta = 0.01$. We obtain similar results on other datasets, see supplemental material.

	Photo	Art painting	Cartoon	Sketch	Mean
$\beta = 1$	89.05 \pm 0.45	69.02 \pm 0.41	71.13 \pm 0.17	58.87 \pm 0.43	72.02
$\beta = 0.1$	90.51 \pm 0.14	70.71 \pm 0.28	70.78 \pm 0.11	62.05 \pm 0.26	73.51
$\beta = 0.01$	91.93 \pm 0.23	71.94 \pm 0.34	73.17 \pm 0.21	65.94 \pm 0.24	75.74
$\beta = 0.001$	90.17 \pm 0.25	70.07 \pm 0.32	71.75 \pm 0.17	63.90 \pm 0.38	73.89

$I(X; T)$ is lowest and $I(T; Y)$ is the highest, compared to those with other values of β . A larger $I(Z; Y)$ indicates that we can make more accurate predictions Y from Z , while a smaller $I(Z; X)$ indicates Z contains the minimal information from X that is required for prediction, suggesting a domain-invariant representation Z . This explains why $\beta = 0.01$ produces the best prediction results compared to other values of β . In our experiments, the optimal value of β is obtained by using a validation set for each dataset and we found $\beta = 0.01$ produces the best performance on all datasets.

Analyzing domain-invariance We visualize the features learned by the pre-trained AlexNet, VIB and MetaVIB in Fig. 3. For better illustration, we use t-SNE [30] to reduce the feature dimension into a two-dimensional subspace. We observe that the features of the same category learned by pre-trained Alexnet (Fig. 3 (a)) show large discrepancy among the four domains. The regular VIB reduces this discrepancy to some extent, but still suffers from considerable gaps

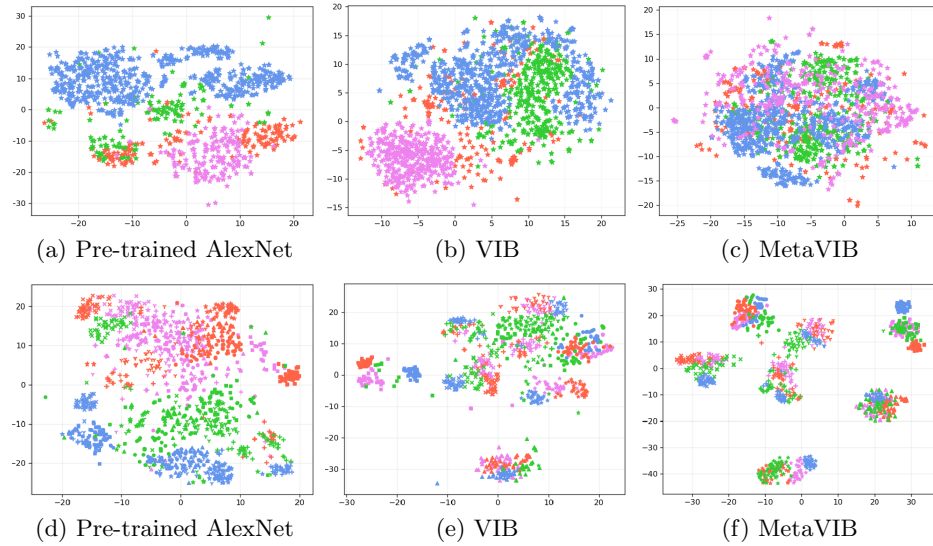


Fig. 3. Analyzing domain-invariance. Visualization of feature representations from pre-trained AlexNet, VIB, and MetaVIB on PACS. The top row shows different features for the *horse* category from four different domains, where the violet shapes denotes the unseen domain *cartoon*. Bottom row shows the distributions of feature representations from all seven PACS classes for four domains, where the unseen domain (green) is *art painting*. MetaVIB reduces the domain gap to achieve domain-invariant yet discriminative representations, which enables accurate predictions. MetaVIB fills the gap between domains (c), while maximally separating samples of different classes (f).

between the unseen domain (violet shapes) (Fig. 3 (b)). MetaVIB largely reduces the discrepancy of different domains including the unseen domains as shown in Fig. 3 (c). In Fig. 3 (d), we observe again that the gaps of features among 4 domains by the pre-trained AlexNet are larger than those between the 7 classes in each domain. Fig. 3 (e) shows that the VIB reduces the domain gaps to certain extent. From Fig. 3 (f), we observe MetaVIB reduces domain gaps considerably while at the same time scatters the samples of 7 classes in each domain. Overall, the proposed MetaVIB principle demonstrates effectiveness in learning domain-invariant representations to tackle domain shift.

Success and failure cases We show some success and failure cases in Fig. 4. MetaVIB successfully predicts the labels for ambiguous images. The dog in the second image in Fig. 4 (a) wears human clothes, showing strong characteristics of a person. Yet, MetaVIB correctly predicts it with a high confidence probability of 0.732. The sketch of the horse looks like a dog in the fourth image, but MetaVIB predicts it correctly with a high probability of 0.636. In the failure cases (b), MetaVIB fails to make the correct prediction, but provides reasonable




Fig. 4. Success and failure cases of MetaVIB. The numbers associated with each image are the top two prediction probabilities of MetaVIB, with ground truth labels in red. MetaVIB successfully distinguishes ambiguous cases in (a). For more challenging cases in (b), MetaVIB provides a high probability for the true label, but fails to make the correct prediction.

Table 4. State-of-the-Art comparison on VLCS, in classification accuracy (%).

	VOC2007	LabelMe	Caltech-101	SUN09	Mean
D’Innocente & Caputo [11]	66.06	57.45	94.95	65.87	71.08
Li et al. [24]	69.99	63.49	93.63	61.32	72.11
Li et al. [25]	67.70	62.60	94.40	64.40	72.28
Li et al. [26]	67.10	64.30	94.10	65.90	72.90
Carlucci et al. [8]	70.62	60.90	96.93	64.30	73.19
Dou et al. [10]	69.14	64.90	94.78	67.64	74.11
MetaVIB	70.28 \pm 0.21	62.66 \pm 0.35	97.37 \pm 0.23	67.85 \pm 0.17	74.54

probabilities for both a person and a dog, which shows the effectiveness in handling uncertainty. It is hard to distinguish which object needs to be predicted in these images, as shown in the first image in Fig. 4 (b).

4.4 State-of-the-Art Comparison

We compare with regular and meta-learning methods for domain generalization. The results on the three datasets are reported in Tables 4-6. On the **VLCS** dataset [49], our MetaVIB achieves high recognition accuracy, surpassing the second best method, i.e., MASF [10], by a margin of 0.43%. Note that on all domains, our MetaVIB consistently outperforms MLDG [25], which is a gradient-based meta-learning algorithm. On the **PACS** dataset [24], our MetaVIB again achieves the best overall performance. It outperforms most of the previous methods, showing clear performance advantages over JiGen [8]. Again, our MetaVIB performs better than other meta-learning based methods, e.g., MetaReg [4], Reptile [33], MLDG [25], Feature-Critic [29], and MASF [10]. It is worth highlighting that our MetaVIB exceeds those meta-learning methods on the “Cartoon” domain by phenomenal margins. On the **Rotated MNIST** dataset [42], the proposed MetaVIB achieves consistently high performance on the test domains, exceeding the alternative methods. It is worthwhile to mention that our MetaVIB outperforms the meta-learning algorithms MetaReg [4], and Reptile [33]. showing its effectiveness as a meta-learning method for domain generalization. To

Table 5. State-of-the-Art comparison on PACS, in classification accuracy (%).

	Photo	Art painting	Cartoon	Sketch	Mean
Ghifary et al. [17]	91.12	60.27	58.65	47.68	64.48
Bousmalis et al. [7]	83.25	61.13	66.54	58.58	67.37
Li et al. [24]	89.50	62.86	66.97	57.51	69.21
Shankar et al. [42]	89.48	64.84	67.69	57.52	69.63
Li et al. [25]	88.00	66.23	66.88	58.96	70.01
Nichol et al. [33]	88.78	64.35	70.09	59.91	70.78
Li et al. [26]	86.10	64.70	72.30	65.00	72.00
Li et al. [29]	89.94	64.89	71.72	61.85	72.10
Balaji et al. [4]	91.70	69.82	70.35	59.26	72.62
Carlucci et al. [8]	89.00	67.63	71.71	65.18	73.38
Dou et al. [10]	90.68	70.35	72.46	67.33	75.21
MetaVIB	91.93 ± 0.23	71.94 ± 0.34	73.17 ± 0.21	65.94 ± 0.24	75.74

Table 6. State-of-the-Art comparison on Rotated MNIST, in averaged classification accuracy (%) of different methods over 10 runs. MetaVIB consistently achieves the best performance on different domains with different rotation angles.

	M _{0°}	M _{15°}	M _{30°}	M _{45°}	M _{60°}	M _{75°}	Mean
Shankar et al. [42]	86.03 ± 0.69	98.92 ± 0.53	98.60 ± 0.51	98.38 ± 0.29	98.68 ± 0.28	88.94 ± 0.47	94.93
Balaji et al. [4]	85.70 ± 0.31	98.87 ± 0.41	98.32 ± 0.44	98.58 ± 0.28	98.93 ± 0.32	89.44 ± 0.37	94.97
Li et al. [27]	86.42 ± 0.24	98.61 ± 0.27	99.19 ± 0.19	98.22 ± 0.24	99.48 ± 0.19	88.92 ± 0.43	95.15
Nichol et al. [33]	87.78 ± 0.30	99.44 ± 0.22	98.42 ± 0.24	98.80 ± 0.20	99.03 ± 0.28	87.42 ± 0.33	95.15
Li et al. [29]	89.23 ± 0.25	99.68 ± 0.24	99.20 ± 0.20	99.24 ± 0.18	99.53 ± 0.23	91.44 ± 0.34	96.39
MetaVIB	91.28 ± 0.21	99.90 ± 0.02	99.29 ± 0.11	99.78 ± 0.10	99.57 ± 0.13	92.75 ± 0.31	97.08

conclude, on all datasets, our MetaVIB accomplishes better performance than previous methods based on both regular learning and meta-learning. The best results on all benchmarks validate the effectiveness of our method for domain generalization.

5 Conclusion

In this work, we propose a new probabilistic model for domain generalization under the meta-learning framework. To address prediction uncertainty, we model the parameters of the classifiers shared across domains by a probabilistic distribution, which is inferred from the source domain and directly used for the target domains. To reduce domain shift, our method learns domain-invariant representations by a new Meta Variational Information Bottleneck principle, derived from a variational bound of mutual information. MetaVIB integrates the strengths of meta-learning, variational inference and probabilistic modeling for domain generalization. Our MetaVIB has been evaluated by extensive experiments on three benchmark datasets for cross-domain visual recognition. Ablation studies validate the benefits of our contributions. MetaVIB consistently achieves high performance and advances the state of the art on all three benchmarks.

References

1. Alemi, A.A., Fischer, I., Dillon, J.V., Murphy, K.: Deep variational information bottleneck. In: International Conference on Learning Representations (2017)
2. Amjad, R.A., Geiger, B.C.: Learning representations for neural network-based classification using the information bottleneck principle. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019)
3. Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., de Freitas, N.: Learning to learn by gradient descent by gradient descent. In: *Advances in Neural Information Processing Systems* (2016)
4. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: *Advances in Neural Information Processing Systems*. pp. 998–1008 (2018)
5. Bertinetto, L., Henriques, J.F., Valmadre, J., Torr, P.H.S., Vedaldi, A.: Learning feed-forward one-shot learners. In: *Advances in Neural Information Processing Systems* (2016)
6. Blanchard, G., Lee, G., Scott, C.: Generalizing from several related classification tasks to a new unlabeled sample. In: *Advances in neural information processing systems*. pp. 2178–2186 (2011)
7. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: *Advances in Neural Information Processing Systems*. pp. 343–351 (2016)
8. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2019)
9. Choi, M.J., Lim, J.J., Torralba, A., Willsky, A.S.: Exploiting hierarchical context on a large database of object categories. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 129–136. IEEE (2010)
10. Dou, Q., de Castro, D.C., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. In: *Advances in Neural Information Processing Systems*. pp. 6447–6458 (2019)
11. D’Innocente, A., Caputo, B.: Domain generalization with domain-specific aggregation modules. In: *German Conference on Pattern Recognition*. pp. 187–198. Springer (2018)
12. Erfani, S., Baktashmotlagh, M., Moshtaghi, M., Nguyen, V., Leckie, C., Bailey, J., Kotagiri, R.: Robust domain generalisation by enforcing distribution invariance. In: *Proceedings of the International Joint Conference on Artificial Intelligence*. pp. 1455–1461 (2016)
13. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International journal of computer vision* **88**(2), 303–338 (2010)
14. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *Proceedings of International Conference on Machine Learning*. pp. 1126–1135 (2017)
15. Finn, C., Levine, S.: Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm. In: *International Conference on Learning Representations* (2018)
16. Finn, C., Xu, K., Levine, S.: Probabilistic model-agnostic meta-learning. In: *Advances in Neural Information Processing Systems*. pp. 9516–9527 (2018)

17. Ghifary, M., Bastiaan Kleijn, W., Zhang, M., Balduzzi, D.: Domain generalization for object recognition with multi-task autoencoders. In: The IEEE International Conference on Computer Vision. pp. 2551–2559 (2015)
18. Gordon, J., Bronskill, J., Bauer, M., Nowozin, S., Turner, R.E.: Meta-learning probabilistic inference for prediction. arXiv preprint arXiv:1805.09921 (2018)
19. Griffin, G., Holub, A., Perona, P.: Caltech-256 object category dataset (2007)
20. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: International Conference on Learning Representations (2014)
21. Kolchinsky, A., Tracey, B.D., Van Kuyk, S.: Caveats for information bottleneck in deterministic scenarios. arXiv preprint arXiv:1808.07593 (2018)
22. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems. pp. 1097–1105 (2012)
23. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
24. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5542–5550 (2017)
25. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Learning to generalize: Meta-learning for domain generalization. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
26. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: *IEEE International Conference on Computer Vision* (2019)
27. Li, H., Jialin Pan, S., Wang, S., Kot, A.C.: Domain generalization with adversarial feature learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5400–5409 (2018)
28. Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., Tao, D.: Deep domain generalization via conditional invariant adversarial networks. In: *Proceedings of the European Conference on Computer Vision*. pp. 624–639 (2018)
29. Li, Y., Yang, Y., Zhou, W., Hospedales, T.M.: Feature-critic networks for heterogeneous domain generalization. In: *Proceedings of International Conference on Machine Learning* (2019)
30. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
31. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: *Proceedings of Proceedings of International Conference on Machine Learning*. pp. 10–18 (2013)
32. Munkhdalai, T., Yu, H.: Meta networks. In: *Proceedings of International Conference on Machine Learning* (2017)
33. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. arXiv preprint arXiv:1803.02999 (2018)
34. Peng, X.B., Kanazawa, A., Toyer, S., Abbeel, P., Levine, S.: Variational discriminator bottleneck: Improving imitation learning, inverse rl, and gans by constraining information flow. arXiv preprint arXiv:1810.00821 (2018)
35. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: *International Conference on Learning Representations* (2017)
36. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. arXiv preprint arXiv:1401.4082 (2014)

37. Russell, B.C., Torralba, A., Murphy, K.P., Freeman, W.T.: Labelme: a database and web-based tool for image annotation. *International journal of computer vision* **77**(1-3), 157–173 (2008)
38. Satorras, V.G., Estrach, J.B.: Few-shot learning with graph neural networks. In: *International Conference on Learning Representations* (2018)
39. Saxe, A.M., Bansal, Y., Dapello, J., Advani, M., Kolchinsky, A., Tracey, B.D., Cox, D.D.: On the information bottleneck theory of deep learning. In: *International Conference on Learning Representations* (2018)
40. Schmidhuber, J.: Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation* **4**(1), 131–139 (1992)
41. Schmidhuber, J., Zhao, J., Wiering, M.: Shifting inductive bias with success-story algorithm, adaptive levin search, and incremental self-improvement. *Machine Learning* **28**(1), 105–130 (1997)
42. Shankar, S., Piratla, V., Chakrabarti, S., Chaudhuri, S., Jyothi, P., Sarawagi, S.: Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745* (2018)
43. Shwartz-Ziv, R., Tishby, N.: Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810* (2017)
44. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: *Advances in Neural Information Processing Systems*. pp. 4077–4087 (2017)
45. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., 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)
46. Thrun, S., Pratt, L.: *Learning to learn*. Springer Science & Business Media (2012)
47. Tishby, N., Pereira, F.C., Bialek, W.: The information bottleneck method. *arXiv preprint physics/0004057* (2000)
48. Tishby, N., Zaslavsky, N.: Deep learning and the information bottleneck principle. 2015 *IEEE Information Theory Workshop (ITW)* (Apr 2015). <https://doi.org/10.1109/itw.2015.7133169>
49. Torralba, A., Efros, A.A., et al.: Unbiased look at dataset bias. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2011)
50. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. *Artificial Intelligence Review* **18**(2), 77–95 (2002)
51. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: *Advances in Neural Information Processing Systems*. pp. 3637–3645 (2016)
52. Xie, Q., Dai, Z., Du, Y., Hovy, E., Neubig, G.: Controllable invariance through adversarial feature learning. In: *Advances in Neural Information Processing Systems*. pp. 585–596 (2017)
53. Zhen, X., Sun, H., Du, Y., Xu, J., Yin, Y., Shao, L., Snoek, C.: Learning to learn kernels with variational random features. *International Conference on Machine Learning* (2020)