

Invariant Information Bottleneck for Domain Generalization

Bo Li¹, Yifei Shen², Yezhen Wang¹,
Wenzhen Zhu³, Colorado Reed⁴, Kurt Keuzter⁴,
Dongsheng Li¹, and Han Zhao⁵

¹Microsoft Research Asia, China

²HKUST, China

³WUSTL, USA

⁴UC Berkeley, USA

⁵UIUC, USA

November 5, 2021

Abstract

Invariant risk minimization (IRM) has recently emerged as a promising alternative for domain generalization. Nevertheless, the loss function is difficult to optimize for nonlinear classifiers and the original optimization objective could fail when pseudo-invariant features and geometric skews exist. Inspired by IRM, in this paper we propose a novel formulation for domain generalization, dubbed invariant information bottleneck (IIB). IIB aims at minimizing invariant risks for nonlinear classifiers and simultaneously mitigating the impact of pseudo-invariant features and geometric skews. Specifically, we first present a novel formulation for invariant causal prediction via mutual information. Then we adopt the variational formulation of the mutual information to develop a tractable loss function for nonlinear classifiers. To overcome the failure modes of IRM, we propose to minimize the mutual information between the inputs and the corresponding representations. IIB significantly outperforms IRM on synthetic datasets, where the pseudo-invariant features and geometric skews occur, showing the effectiveness of proposed formulation in overcoming failure modes of IRM. Furthermore, experiments on DomainBed show that IIB outperforms 13 baselines by 0.7% on average across 7 real datasets.

Introduction

In most statistical machine learning algorithms, a fundamental assumption is that the training data and test data are *independently and identically distributed* (i.i.d.). However, the data we have in many real-world applications are not i.i.d. Distributional shifts are ubiquitous. Under such circumstances, classic statistical learning paradigms with strong generalization guarantees, e.g., Empirical Risk Minimization (ERM) [48], often fail to generalize due to the violation of the i.i.d. assumption. It has been widely observed that the performance of a model often deteriorates dramatically when it is faced with samples from a different domain, even under a mild distributional shift [5]. On the other hand, collecting training samples from all possible future scenarios is essentially infeasible. Hence, understanding and improving the generalization of models on *out-of-distribution* data is crucial.

Domain generalization (DG), which aims to learn a model from several different domains so that it generalizes to *unseen* related domains, has recently received much attention. From the perspective of representation learning, there are several paradigms towards this goal, including invariant representation learning [33, 56], causality [5, 21], meta-learning [7, 12], and feature disentanglement [12, 39]. Of particular interest is the invariant learning methods. Some early works, e.g., DANN [16], CDANN [30], aim at finding representations that are invariant across domains. Nevertheless, learning invariant representations fails for domain adaptation or generalization when the marginal label distributions change between source and target domains [55]. Recently, Invariant Causal Prediction (ICP), and its follow-up Invariant Risk Minimization (IRM), have attracted much interest. ICP assumes that the data are generated according to a structural causal model (SCM) [37]. The causal mechanism for the data generating process is the same across domains, while the *non-causal* mechanisms can vary among different domains. Under such data generative assumptions, IRM [5] attempts to learn an optimal classifier that is invariant across domains. Theoretical analysis in ICP reveals that under the SCM, such a classifier can generalize across domains.

Despite the intuitive motivations, IRM falls short in several aspects. First, the proposed loss function in [5] is difficult to optimize when the classifier is nonlinear. Furthermore, it has been shown that IRM fails when the pseudo-invariant features [42] or geometric skews exist [34]. Under such circumstances, the classifier will utilize both the causal and spurious features, leading to a violation of invariant causal prediction. To address the first issue, we propose an information-theoretical formulation of invariant causal prediction and adopt a variational approximation to ease the optimization procedure. To tackle the second issue, we emphasize that the use of pseudo-invariant features or geometric skews will inevitably increase the mutual information between the inputs and the representations. Thus, to miti-

gate the impact of pseudo-invariant features and geometric skews, we propose to constrain this mutual information, which naturally leads to a formulation of information bottleneck. Our empirical results show that the proposed approach can effectively improve the accuracy when the pseudo-invariant features and geometric skews exist.

Contributions: In this paper, we propose a novel information-theoretic formulation for domain generalization, termed as invariant information bottleneck (IIB). IIB aims at minimizing invariant risks for non-linear classifiers while mitigating the impact of pseudo-invariant features and geometric skews. Specifically, our contributions can be summarized as follows:

(1) We propose a novel formulation for invariant causal prediction via mutual information. We further adopt variational approximation to develop tractable loss functions for nonlinear classifiers.

(2) To mitigate the impact of pseudo-invariant features and geometric skews, inspired by the information bottleneck principle, we propose to constrain the mutual information between the inputs and the representations. The effectiveness is verified by the synthetic experiments of failure modes [1, 34], where IIB significantly improves the performance of IRM.

(3) Empirically, we analyze IIB’s performance with extensive experiments on both synthetic and large-scale benchmarks. Compared to existing DG algorithms, IIB is able to eliminate the spurious information better than other methods. and achieves consistent improvements on 7 datasets by 0.7% on DomainBed [18].

Related Works

Domain Generalization

Existing methods of DG can be divided into three categories: (1) **Data Manipulation:** Machine learning models typically rely on diverse training data to enhance the generalization ability. Data manipulation/augmentation methods [36, 41] aim to increase the diversity of existing training data with operations including flipping, rotation, scaling, cropping, adding noise, etc. Domain randomization [10, 51, 52] provides more complex operations for image data, such as altering the location/texture of objects, replicating and resizing objects, modifying the illumination, contrastness and angle of view. In addition, there are some methods [41, 40, 28, 45, 4, 46, 59] that exploits generated data samples to enhance the model generalization ability. (2) **Ensemble Learning** methods [32, 44] assume that any sample in the test domain can be regarded as an integrated sample of the multiple-source do-




Figure 1: Illustrations of features in OOD generalization. For all the bears in training domains, the predictions $P(Y | Z)$ conditioning on the invariant features (e.g. *outline*) will be correct and invariant, and also the most generalizable. While the predictions conditioning on pseudo invariant features (possibly *fur color* in this example) are misleading and may affect the generalization ability on test domains. Geometric skews are the spurious features used as a short-cut for max-margin classifiers. In this example, ERM will use all 5 features as they are informative to labels. IRM, with the invariance constraint, will utilize the first 3 features. IIB, by selecting the minimal sufficient features, only includes the **shape** or **outline**.

mains, so the overall prediction should be inferred by a combination of the models trained on different domains. (3) **Meta-Learning** aims at learning a general model from multiple domains. In terms of domain generalization, MLDG [26] divides data from the multiple domains into meta-train and meta-test to simulate the domain shift situation to learn the general representations. In particular, Meta-Reg [7] learns a meta-regularizer for the classifier, and Meta-VIB [12] learns to generate the weights in the meta-learning paradigm by regularizing the KL divergence between marginal distributions of representations of the same category but from different domains.

Mutual Information-based Domain Adaptation

Domain Adaptation is an important topic in the direction of transfer learning [29, 15, 47, 30, 58, 60, 57, 23]. The mutual information-based approaches have been widely applied in this area. The key idea is to learn a domain-invariant representation that are informative to the label, which can be formulated as [54, 24]

$$\max I(Z, Y) - \lambda I(Z, A) \tag{1}$$

where A is the identity of domains, Z denotes the representation, and Y denotes the labels. Commonly adopted implementations of (1) are DANN [16] and CDANN

[30]. These implementations are also often adopted in domain generalization as baselines [18].

Invariant Risk Minimization

The above approaches enforces the invariance of the learned representation. Invariant Risk Minimization (IRM) suggest the invariance of feature-conditioned label distribution. Specifically, IRM seeks for an invariant causal prediction such that $\mathbb{E}[Y^e | \Phi(X^e)] = \mathbb{E}[Y^{e'} | \Phi(X^{e'})]$, for all $e, e' \in \mathcal{E}$. The objective of IRM is given by

$$\begin{aligned} \min_{\mathbf{w}, \Phi} \quad & \sum_{e \in \mathcal{E}_{\text{train}}} R^e(\mathbf{w} \circ \Phi), \\ \text{s.t. } \quad & \mathbf{w} \in \underset{\hat{\mathbf{w}}}{\text{argmin}} R^e(\hat{\mathbf{w}} \circ \Phi), \end{aligned}$$

where R^e is the cross-entropy loss for environment e , Φ is the feature extractor and \mathbf{w} is a linear classifier. Note that the above objective is a bilevel optimization and difficult to optimize. Thus, in [5], first-order approximation is adopted and the loss function is given by

$$\min_{\Phi} \sum_{e \in \mathcal{E}_{\text{train}}} R^e(\Phi) + \lambda \cdot \|\nabla_{\mathbf{w}|_{\mathbf{w}=1.0}} R^e(\mathbf{w} \circ \Phi)\|, \quad (2)$$

where $\mathbf{w} \in \mathbb{R}$ is a dummy classifier.

Preliminaries

Recently, some works have focused on characterizing the failure modes of OOD generalization [42, 34], as discussed in the sequel.

Pseudo-invariant Features

[42] The first problem is that the invariant features in the training dataset may not be the invariant features in the test dataset. Specifically, we denote the causal feature and spurious feature as z_c and z_s respectively. According to the analysis in [42], there exists a transformation Φ such that $[z_c, \Phi z_s]$ are invariant features across the training dataset. Furthermore, the classifier will utilize $[z_c, \Phi z_s]$ instead of z_c to achieve a lower training error. The OOD generalization may fail due to the inclusion of z_s , which can be arbitrary in the test dataset. An illustration of pseudo-invariant features is shown in Fig. 1.

Geometric Skews

[34] The OOD generalization can fail even if we assume the invariant features in the training dataset are invariant features in the test dataset due to the *geometric skews*. It is observed in [34] that as the number of training points increase, the ℓ_2 -norm of this max-margin classifier grows. For imbalanced data, the spurious feature can be used as a short-cut for classification. Specifically, we consider an invariant feature z_{inv} is concatenated with a spurious feature z_{sp} such that $\mathbb{P}[z_{\text{sp}} \cdot y > 0] > 0.5$. The dataset consists of a majority group S_{maj} where $z_{\text{sp}} \cdot y > 0$ (e.g., cows/camels with green/yellow backgrounds) and a minority group S_{min} where $z_{\text{sp}} \cdot y < 0$ (e.g., cows/camels with yellow/green backgrounds). Let w_{all} denote the least-norm classifier using invariant features to classify all samples and w_{min} denote the least-norm classifier using invariant features to classify the samples in S_{min} , and we have $\|w_{\text{min}}\| \ll \|w_{\text{all}}\|$. We can use the spurious feature as a short-cut to classify S_{maj} and S_{min} , and then adopt w_{min} to classify the remaining S_{min} . This classifier using spurious feature will have a smaller norm than the invariant classifier, which leads to the failure of OOD generalization.




Figure 2: A structural causal model explaining that different parts of an input X have different probabilistic relationships with the model output Y . Observed variables are shaded, while others are with dotted outlines.

Methodology Design

In this section, we propose a novel information-theoretic objective of finding invariant causal relationship and overcoming the issues in OOD generalization.

Invariant Causal Prediction via Mutual Information

Like other casual related works [11, 31], we begin with a structural causal model, shown in Figure 2. For simplicity, we leave out all the unnecessary elements. In general, we can see that an input X can be divided into two variables, the *causal* feature Z_c and *environmental* feature Z_a . In Figure 2, we can readout that both features are highly correlated with Y , but only Z_c is regarded as a plausible explanation. Through the concept of d -separation [37], we can readout the conditional independence conditions that all data distributions $\mathcal{P}(D, X, Y)$ should satisfy:

1. $Y \not\perp\!\!\!\perp D$ means the marginal distribution of class label Y can change across domains.
2. $Y \perp\!\!\!\perp D \mid Z_c$ means the class label Y is independent of domain D conditioned on causal feature Z_c . The underlying causal mechanism determines that the value of Y comes from its unique causal parent Z_c , which does not change across domains.
3. $Y \not\perp\!\!\!\perp D \mid Z_c, Z_a$ means the conditional independence won't hold true if conditioned on both *causal* and *environmental* features since Z_a is a collider between D and Y .

The conditional independence tells us that only the real causal relation is stable and remains invariant across domains. In other words, we should eliminate the spurious environmental feature Z_a by seeking the causal feature Z_c that is independent of D from $\Phi(X)$. Particularly, the representation $Z = \Phi(X)$ should have the following two merits: (1) Z does not change with different domains for the same class label Y , hence achieving the conditional invariance of $Y \perp\!\!\!\perp D \mid Z$; (2) Z should be informative of the class label Y (otherwise even a constant $\phi(\cdot)$ would minimize the above loss). The above two conditions coincide with the objective of IRM [5], and also suggest the following learning objective:

$$\max_{\Phi} I(\Phi(X), Y) - \lambda I(Y, D \mid \Phi(X)), \quad (3)$$

where Φ is the feature extractor.

Proposition 1. Assume $I(Y, D \mid Z) = 0$, we achieve invariant causal prediction in the sense that $\mathbb{E}[Y \mid \Phi(X) = x]$ is fixed.

Proof. Note that $I(Y, D \mid Z) = 0$ implies Y and D are independent conditioned on $\Phi(X)$. The conditional independence indicates that $\mathbb{P}(Y \mid \Phi(X) = x, D) = \mathbb{P}(Y \mid \Phi(X) = x)$, which shows that $\mathbb{P}(Y \mid \Phi(X) = x)$ is fixed across domain D . Thus $\mathbb{E}[Y \mid \Phi(X) = x]$ is fixed and we can achieve invariant causal prediction. \square

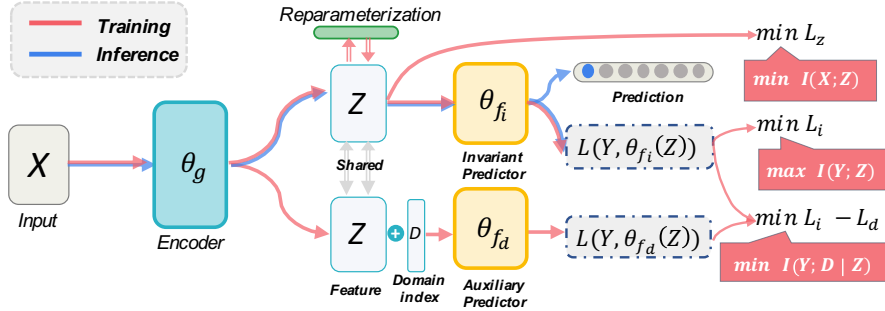


Figure 3: IIB optimizes a model consisting of three parts: (1) an invariant predictor $f_i(Z)$; (2) a domain-dependent predictor $f_d(Z, D)$; (3) an encoder $g(X)$. The three loss terms on the right hand side respectively correspond to the optimization of the three mutual information terms.

On the Failures of OOD Generalization

In this subsection, we first scrutinize the failure conditions of OOD generalization, i.e., pseudo-invariant features and geometric skews. Then we propose to utilize the least informative feature, as well as minimizing $I(X, Z)$ as a solution.

With pseudo-invariant features and geometric skews, the failure of OOD generalization is due to the inclusion of (transformations of) spurious features. We first give an example when the features are one-dimensional and the classifier is linear [34]. Denote the invariant feature, pseudo-invariant feature, feature causing geometric skews, spurious feature as $Z_i, Z_p, Z_{sk},$ and Z_{sp} . The overall features are $\mathbf{Z} = [Z_i, Z_p, Z_{sk}, Z_{sp}]$. In the ERM model, all the features will be adopted and OOD generalization fails. We consider the following optimization problem

$$\begin{aligned} \min_w \sum_{e \in \mathcal{E}_{\text{train}}} R^e(\mathbf{w} \cdot \mathbf{Z}), \\ \text{s.t. } \|\mathbf{w}\|_0 \leq 1, \mathbf{w} \in \underset{\hat{\mathbf{w}}}{\text{argmin}} R^e(\hat{\mathbf{w}} \cdot \mathbf{Z}), \end{aligned} \quad (4)$$

where $\|\mathbf{w}\|_0 \leq 1$ is the sparsity constraint, and $\mathbf{w} \in \underset{\hat{\mathbf{w}}}{\text{argmin}} R^e(\hat{\mathbf{w}} \cdot \mathbf{Z})$ is the invariant risk constraint of IRM. Due to the sparsity constraint, there are only four choices. Choosing Z_{sp} cannot satisfy the invariant constraint while choosing Z_p or Z_{sk} cannot minimize the empirical risk. Thus, the only optimal solution is $\mathbf{w} = [w_1^*, 0, 0, 0]$. Without the sparsity constraint, the optimization problem becomes IRM and Z_i, Z_p, Z_{sk} will be used for classification. Without invariance constraint, Z_{sp} might be chosen as the spurious feature for a lower empirical risk.

We then extend this intuition into the loss function design of deep neural net-

works in the view of mutual information. Suppose Z_1, Z_2 are features extracted from X , we have $I(X, [Z_1, Z_2]) \geq I(X, Z_1)$ as Z_1 is a subset of $[Z_1, Z_2]$. Thus, we expect that the sparse feature selection can be done by penalizing a larger $I(X, Z)$. To this end, we formulate our objective as

$$\max_{\Phi} I(\Phi(X), Y) - \lambda I(Y, D | \Phi(X)) - \beta I(X, \Phi(X)). \quad (5)$$

The term $I(Z, Y) - \beta I(X, Z)$ is the information bottleneck and $I(Y, D | Z)$ is an invariant constraint. As a result, we refer (5) as *invariant information bottleneck (IIB)*.

Loss Function Design

The objective in (5) is still not a tractable loss function as the mutual information of high dimensional vectors is hard to estimate. We further leverage variational approximation to solve this issue. According to variational IB [3], the loss of information bottleneck can be written as

$$\begin{aligned} & I(Z, Y) - \beta I(Z, X) \\ & \geq \mathbb{E}_{p_{x,y,z}} \left[\log q(y | z) \right] - \beta \mathbb{E}_{p_{x,z}} \left[\log \frac{p(z | x)}{r(z)} \right]. \end{aligned} \quad (6)$$

Optimizing (6) is still a difficult task. As a result, we follow the reparametrization operation in VIB [3], we use an encoder of the form $p(z|x; g) = \mathcal{N}(z|g^\mu(x), g^\Sigma(x))$. g outputs a K -dimensional mean μ of z and a $K \times K$ covariance matrix Σ . Then we have $q(z|x)d(z) = q(\varepsilon)d\varepsilon$, where $z = g(x, \varepsilon)$, $\varepsilon \sim \mathcal{N}(0, 1)$, so we can optimize (6) by optimizing

$$\mathcal{L}_i(g, f_i) + \beta \mathcal{L}_z(g), \quad (7)$$

where $\mathcal{L}_i = \min_{g, f_i} \mathbb{E}_{x,y} [L(y, f_i(g(x)))]$ and $\mathcal{L}_z = \min_g \mathbb{E}_x [KL[q(z|x; g)||r(z)]]$. such that KL is the KL-divergence, $g(x)$ is the feature extractor, f_i is the classifier, and L is the cross-entropy loss.

We next proceed to deal with $I(Y, D|Z)$. Following the rules of variational approximation [14], we have

$$I(Y, D | Z) = H(Y | Z) - H(Y | D, Z), \quad (8)$$

where $H(Y | Z) = -\sup_q \mathbb{E}_{p_{y,z}} [\log q(y|z)]$ and $H(Y | D, Z) = -\sup_h \mathbb{E}_{p_{y,z,d}} [\log h(y | z, d)]$. Thanks to the universal approximation ability of neural networks, (8) can be written as the subtraction of two classification loss [14]:

Table 1: Accuracy on CS-CMNIST experiment. We split 20% from train set as validation set.

Methods	Validation Acc. (%) \uparrow	Test Acc.(%) \uparrow
ERM [48]	95.38 \pm 0.03	11.16 \pm 0.31
IRM [6]	97.59 \pm 1.39	57.98 \pm 0.86
IB-ERM [1]	97.64 \pm 0.04	58.47 \pm 0.86
IB-IRM [1]	97.51 \pm 1.09	71.79 \pm 0.70
IIB ($\lambda = 0$)	92.95 \pm 0.50	69.52 \pm 0.80
IIB ($\beta = 0$)	92.39 \pm 0.50	66.93 \pm 0.33
IIB	98.11 \pm 0.84	74.23 \pm 4.80

$$\begin{aligned}
 I(Y, D | Z) = & \min_{f_i, g} \underbrace{\mathbb{E}_{x, y} [L(y, f_i(g(x)))]}_{\mathcal{L}_i} \\
 & - \min_{f_d, g} \underbrace{\mathbb{E}_{x, y, d} [L(y, f_d(g(x), d))]}_{\mathcal{L}_d}, \tag{9}
 \end{aligned}$$

where f_i takes feature z as the input, and $f_d, d = 1, \dots, D$ takes both feature z and domain index d as the input. Overall, we can maximize R_{IIB} by optimizing its tractable lower bound:

$$\min_{g, f_i} \max_{f_d} \mathcal{L}_i(g, f_i) + \lambda |\mathcal{L}_i(g, f_i) - \mathcal{L}_d(g, f_d)| + \beta \mathcal{L}_z(g).$$

Guided by this objective, as illustrated in Figure 3, IIB optimizes a model consisting of three parts: (1) an invariant predictor $f_i(Z)$; (2) an domain-dependent predictor $f_d(Z, D)$; (3) an encoder $g(X)$.

Synthetic Experiments

Experimental Setup

To validate IIB’s efficacy of mitigating the impact of pseudo-invariant features and geometric skews, we adopt two types of synthetic experiments. Both pseudo-invariant features and geometric skews exist in the two experiments.

CS-CMNIST [1] CS-CMNIST is a ten-way classification task. The images are all drawn from MNIST. There are three environments, two training environments contain each 20,000 images, one test environment contains 20,000 images. There are




Figure 4: (a) The above figure represents the image examples in majority/minority group in train set in **Cross Lines** experiment, while colored lines are not included in the test data. The below figure represents the image examples in train/test sets with different spurious feature scales in **Vertical Line** experiment. (b) Accuracy at test domains with different spurious feature scales \mathcal{B} . The upward-pointing red triangle denotes different \mathcal{B} at training domains (we set them to 4 and 0 respectively).

Table 2: Accuracy on Cross-Lines experiment. We split 20% from train set as validation set.

Methods	Validation Acc. (%) \uparrow	Test Acc.(%) \uparrow
ERM [48]	90.12 \pm 0.12	65.60 \pm 0.27
IRM [6]	63.82 \pm 0.25	42.68 \pm 0.32
IB-ERM [1]	83.93 \pm 0.10	69.70 \pm 0.42
IB-IRM [1]	81.61 \pm 0.69	65.82 \pm 0.77
IIB ($\lambda = 0$)	79.97 \pm 0.50	69.52 \pm 0.80
IIB ($\beta = 0$)	78.47 \pm 0.50	66.93 \pm 0.33
IIB	92.86 \pm 0.29	71.04 \pm 0.37

ten colors associated with ten digit class correspondingly. The probability p_e denotes that the image is colored with associated color. In two training environments, p_e is set to 1 and 0.9, which means the images with certain class are colored with associated color with probability p_e and are colored with random color with probability $1 - p_e$. In test environment, p_e is set to 0, which means all images are colored at random. Overall, the color of images in training domains can be fully predictive to label with spurious features, i.e. using the associated color, but the information disappear at test domain.

In CS-CMNIST, if the accuracy drops more at test time, it reflects that relying more on spurious features during training. We will give results of IIB on AC-CMNIST (in DomainBed it’s known as CMNIST) in next section.

Geometric Skew CIFAR10 [34] There are two types of tasks (as shown in Figure 4 (a)). For the first type, we name it *Cross Lines* experiment, we create ten-valued spurious feature and add a vertical line passing through the middle of each channel, and also a horizontal line through the first channel. For these four lines added, we take the value of $(0.5 \pm 0.5\mathcal{B})$ where $\mathcal{B} \in [-1, 1]$. Four lines, each with 2 choices, then we have a total of $2^4 = 16$ configurations. Among them, we choose the first 10 and denote the 10 configurations to each class in CIFAR10. For i -th configuration, corresponding to i -th class, we add this line with a probability of $p_{ii} = 0.5$; for other j -th class, we set $p_{ij} = (1 - p_{ii})/10 = 0.05$. Taking the probability means 50% data (the majority group) are correlated with spurious features (the specific colored line corresponding to each class), while other 5% data (the minority group) are correlated with other 9 configurations at random. For the second type, we name it *Vertical Line*, we add a colored line to the last channel of CIFAR10, regardless of the label during training, and vary its brightness during testing. In detail, we add a line with value choose from $\mathcal{B} \in [-4, 4]$. To avoid negative values, all pixels in last channel are added by 4, and then added by \mathcal{B} , and then divided by 9 to make sure the values lie in the range of $[0, 1]$. Such an experiment would artificially create non-orthogonal components, where each data-point is represented on the plane of $(x_{\text{inv}}, x_{\text{inv}} + x_{\text{env}})$, rather than a more easy-to-disentangle representation under $(x_{\text{inv}}, x_{\text{env}})$. As discussed in [34], the model would be more susceptible to spurious features that may shift during testing.

Observation for results on synthetic experiments

In CS-CMNIST, we compare IIB with several methods, including ERM [48], IRM [6], IB-IRM [1]. Particularly, IB-IRM [1] is from a concurrent work, which propose to combine information bottleneck and IRM to eliminate geometric skews. Among them (see Table 1), IIB has observable improvements over two synthetic datasets compared with other algorithms. Compare to IB-IRM, which is a direct combination of IB and IRM, our approach took a different approach to optimize the learning objective, which led to further enhancements. In the *Cross Lines* experiment (see Table 2), we train the network on images with colored cross lines (each color corresponds to a specific class in CIFAR10), and test on normal images. From the improvements of IB over IRM, we observe that the information bottleneck structure can help mitigate the failure of IRM in geometric skews. In the *Vertical Line* experiment (see Figure 4 (b)), we train the network on $\mathcal{B} = 4$ or 0, and test on domains with different spurious feature scale \mathcal{B} . The results show that as the offset of spurious feature scale increases, the accuracy of training and testing environments decreases a lot. However, IIB still keeps good results even with large offset, indicating that it’s effectiveness in alleviating the dependence on spurious feature. We have similar observations that information bottleneck (IB) structure would over-

come the geometric skews which fails IRM.

Table 3: Performance comparison (Acc. %) between the proposed IIB method and the state-of-the-art domain generalization methods with *leave one domain out* model selection strategy. The best accuracy in each dataset is presented in boldface. The average accuracy over all the datasets is also reported.

Methods	Colored-MNIST	Rotated-MNIST	VLCS	PACS	OfficeHome	TerraIncognita	DomainNet	Average
ERM [48]	36.7 ± 0.1	97.7 ± 0.0	77.2 ± 0.4	83.0 ± 0.7	65.7 ± 0.5	41.4 ± 1.4	40.6 ± 0.2	63.2
DANN [16]	40.7 ± 2.3	97.6 ± 0.2	76.9 ± 0.4	81.0 ± 1.1	64.9 ± 1.2	44.4 ± 1.1	38.2 ± 0.2	63.4
CDANN [27]	39.1 ± 4.4	97.5 ± 0.2	77.5 ± 0.2	78.8 ± 2.2	64.3 ± 1.7	39.9 ± 3.2	38.0 ± 0.1	62.2
MLDG [26]	36.7 ± 0.2	97.6 ± 0.0	77.2 ± 0.9	82.9 ± 1.7	66.1 ± 0.5	46.2 ± 0.9	41.0 ± 0.2	64.0
IRM [5]	40.3 ± 4.2	97.0 ± 0.2	76.3 ± 0.6	81.5 ± 0.8	64.3 ± 1.5	41.2 ± 3.6	33.5 ± 3.0	62.0
GroupDRO [43]	36.8 ± 0.1	97.6 ± 0.1	77.9 ± 0.5	83.5 ± 0.2	65.2 ± 0.2	44.9 ± 1.4	33.0 ± 0.3	62.7
MMD [2]	36.8 ± 0.1	97.8 ± 0.1	77.3 ± 0.5	83.2 ± 0.2	60.2 ± 5.2	46.5 ± 1.5	23.4 ± 9.5	60.7
VREx [22]	36.9 ± 0.3	93.6 ± 3.4	76.7 ± 1.0	81.3 ± 0.9	64.9 ± 1.3	37.3 ± 3.0	33.4 ± 3.1	60.6
ARM [53]	36.8 ± 0.0	98.1 ± 0.1	76.6 ± 0.5	81.7 ± 0.2	64.4 ± 0.2	42.6 ± 2.7	35.2 ± 0.1	62.2
Mixup [50]	33.4 ± 4.7	97.8 ± 0.0	77.7 ± 0.6	83.2 ± 0.4	67.0 ± 0.2	48.7 ± 0.4	38.5 ± 0.3	63.8
RSC [20]	36.5 ± 0.2	97.6 ± 0.1	77.5 ± 0.5	82.6 ± 0.7	65.8 ± 0.7	40.0 ± 0.8	38.9 ± 0.5	62.7
MTL [9]	35.0 ± 1.7	97.8 ± 0.1	76.6 ± 0.5	83.7 ± 0.4	65.7 ± 0.5	44.9 ± 1.2	40.6 ± 0.1	63.5
SagNet [35]	36.5 ± 0.1	94.0 ± 3.0	77.5 ± 0.3	82.3 ± 0.1	67.6 ± 0.3	47.2 ± 0.9	40.2 ± 0.2	63.6
IIB(Ours)	39.9 ± 1.2	97.2 ± 0.2	77.2 ± 1.6	83.9 ± 0.2	68.6 ± 0.1	45.8 ± 1.4	41.5 ± 2.3	64.9

DomainBed Experiments

To empirically corroborate the effectiveness of IIB, we conduct experiments on DomainBed [18]. From small to large, we experiment on 7 different datasets of domain generalization task, including Colored-MINIST [5], Rotated-MNIST [17], PACS [25], VLCS [13], Office-Home [49], Terra Incognita [8], DomainNet [38].

Model Selection Strategy We choose two types of model selection strategies out of three in DomainBed. We don’t test on test-domain validation set, it’s too optimistic to estimate an algorithm’s performance since it allows access to test domain while training. In training-domain validation set, the validation set is subset of training set, we choose the model that performs best on the overall validation set for each domain. This strategy characterizes the in-distribution generalization capability of the model. In leave-one-domain-out cross validation, the training domains are separated from the test domain. This strategy characterizes the out-of-domain distribution generalization capacity of the model. Due to the space limit, we give results on leave-one-domain-out cross validation in Table 3, and put the results on training-domain validation set in supplementary materials.

Hyper-parameters and Implementation Details In both selection strategies, for default hyper-parameters (e.g. learning rate, weight decay), we use default settings in DomainBed (e.g. learning rate is set to $1e-3$ for small images and with a selection range of $lr \in [10^{-4.5}, 10^{-2.5}]$). For IIB specific hyper-parameters, we set $\lambda \in [1, 10^2]$, and $\beta \in [10^{-3}, 10^{-4}]$. For backbone feature extractor, in Rotated/Colored-MNIST, we use 4-layers 3x3 ConvNet. For VLCS and PACS, we use ResNet-18 [19]. For larger datasets, we opt to ResNet-50. For classifier, we both test linear and non-linear invariant (environment) classifiers. Specifically, in linear classifier, it has only one layer, otherwise it has three MLP layers with two RELU activation layers. For the increased number of parameters in the non-linear classifier, we correspondingly reduce the number of conv-layers in the backbone network to achieve a balance. We test the hyper-parameters and different model implementations on RotatedMNIST, the network is trained for 5000 iterations with batch size set to 128. We report the results in Table 4. We observe that the overall network parameters under non-linear classifier setting are not increased too much. The best results are usually obtained with lambda=100 and beta=1e-4.

Observation for results on DomainBed

From Table 3, we can see that IIB ends up working better than existing algorithms in *average* performance on 7 datasets. Since DomainBed is a publicly tested benchmark that was originally proposed to demonstrate that, none of the currently available Domain Generalization algorithms can significantly outperform each other. The results in Table 3 also indirectly illustrate this matter, as it does not show big advantage in small datasets (Colored-MNSIT, Rotated-MNIST), but performs better than others in larger datasets (PACS, Office-Home, DomainNet). We opine that the Information Bottleneck structure is able to better eliminate the noise brought by spurious features in large datasets, while when the data set is too small, this noise may still be useful as the short-cut in test domain for prediction, thus achieving better results.

Conclusion

In this paper, we developed a novel information-theoretical approach for domain generalization, namely, IIB. IIB aims at tackling the optimization difficulty and the failure modes of IRM. The superior performance is demonstrated on both synthetic and real datasets through extensive experiments. As for future directions, it is interesting to investigate the theoretical foundations of adopting IB in nonlinear invariant causal prediction and the effectiveness of IIB on regression tasks.

Table 4: We test the different hyper-parameters’ impact to the proposed IIB method on RotatedMNIST with leave-one-domain-out strategy. The results of multiply-add cumulation (MAC) operations and network parameters (Params) are reported.

Classifier Type	MACs	Params	β	λ	Acc. (%) \uparrow
linear	5.83G	370.95K	1e-3	100	61.1
			1e-4	1	94.7
				10	95.3
				100	95.1
non-linear	5.83G	375.33K	1e-3	100	63.2
			1e-4	1	96.8
				10	97.2
				100	97.3

References

- [1] Kartik Ahuja, Ethan Caballero, Dinghui Zhang, Yoshua Bengio, Ioannis Mitliagkas, and Irina Rish. Invariance principle meets information bottleneck for out-of-distribution generalization. *CoRR*, abs/2106.06607, 2021.
- [2] Kei Akuzawa, Yusuke Iwasawa, and Yutaka Matsuo. Adversarial invariant feature learning with accuracy constraint for domain generalization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 315–331. Springer, 2019.
- [3] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [4] Asha Anoosheh, Eirikur Agustsson, Radu Timofte, and Luc Van Gool. Com-bogan: Unrestrained scalability for image domain translation. In *2018 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 783–790. IEEE, 2018.
- [5] Martín Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *CoRR*, abs/1907.02893, 2019.
- [6] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- [7] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In *Proceedings*

of the 32nd International Conference on Neural Information Processing Systems, pages 1006–1016, 2018.

- [8] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *The 15th European Conference on Computer Vision (ECCV)*, volume 11220 of *Lecture Notes in Computer Science*, pages 472–489. Springer, 2018.
- [9] Gilles Blanchard, Aniket Anand Deshmukh, Ürün Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *J. Mach. Learn. Res.*, 22:2:1–2:55, 2021.
- [10] João Borrego, Atabak Dehban, Rui Figueiredo, Plinio Moreno, Alexandre Bernardino, and José Santos-Victor. Applying domain randomization to synthetic data for object category detection. *CoRR*, abs/1807.09834, 2018.
- [11] Shiyu Chang, Yang Zhang, Mo Yu, and Tommi S. Jaakkola. Invariant rationalization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020*, pages 1448–1458. PMLR, 2020.
- [12] Ying-Jun Du, Jun Xu, Huan Xiong, Qiang Qiu, Xiantong Zhen, Cees G. M. Snoek, and Ling Shao. Learning to learn with variational information bottleneck for domain generalization. In *The 16th European Conference on Computer Vision*, volume 12355 of *Lecture Notes in Computer Science*, pages 200–216. Springer, 2020.
- [13] Chen Fang, Ye Xu, and Daniel N. Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, pages 1657–1664. IEEE Computer Society, 2013.
- [14] Farzan Farnia and David Tse. A minimax approach to supervised learning. In *Advances in Neural Information Processing Systems*, pages 4233–4241, 2016.
- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- [16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. Domain-adversarial training of neural networks. In *Domain Adaptation in Computer Vision Applications*, Advances in Computer Vision and Pattern Recognition, pages 189–209. Springer, 2017.

- [17] Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *2015 IEEE International Conference on Computer Vision*, pages 2551–2559. IEEE Computer Society, 2015.
- [18] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *CoRR*, abs/2007.01434, 2020.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *The 14th European Conference on Computer Vision*, volume 9908 of *Lecture Notes in Computer Science*, pages 630–645. Springer, 2016.
- [20] Zeyi Huang, Haohan Wang, Eric P. Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, volume 12347 of *Lecture Notes in Computer Science*, pages 124–140. Springer, 2020.
- [21] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). *arXiv preprint arXiv:2003.00688*, 2020.
- [22] David Krueger, Ethan Caballero, Jörn-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Rémi Le Priol, and Aaron C. Courville. Out-of-distribution generalization via risk extrapolation (rex). *CoRR*, abs/2003.00688, 2020.
- [23] Bo Li, Yezhen Wang, Tong Che, Shanghang Zhang, Sicheng Zhao, Pengfei Xu, Wei Zhou, Yoshua Bengio, and Kurt Keutzer. Rethinking distributional matching based domain adaptation. *arXiv preprint arXiv:2006.13352*, 2020.
- [24] Bo Li, Yezhen Wang, Shanghang Zhang, Dongsheng Li, Trevor Darrell, Kurt Keutzer, and Han Zhao. Learning invariant representations and risks for semi-supervised domain adaptation. *CoRR*, abs/2010.04647, 2020.
- [25] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. In *IEEE International Conference on Computer Vision*, pages 5543–5551. IEEE Computer Society, 2017.
- [26] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the*

Thirty-Second AAAI Conference on Artificial Intelligence, pages 3490–3497. AAAI Press, 2018.

- [27] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the 15th European Conference on Computer Vision*, pages 647–663. Springer, 2018.
- [28] Alexander H. Liu, Yen-Cheng Liu, Yu-Ying Yeh, and Yu-Chiang Frank Wang. A unified feature disentangler for multi-domain image translation and manipulation. In *Advances in Neural Information Processing Systems*, pages 2595–2604, 2018.
- [29] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791*, 2015.
- [30] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1640–1650, 2018.
- [31] Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. *CoRR*, abs/2006.07500, 2020.
- [32] Massimiliano Mancini, Samuel Rota Bulò, Barbara Caputo, and Elisa Ricci. Best sources forward: Domain generalization through source-specific nets. In *2018 IEEE International Conference on Image Processing*, pages 1353–1357. IEEE, 2018.
- [33] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *Proceedings of the 30th International Conference on Machine Learning*, pages 10–18. JMLR.org, 2013.
- [34] Vaishnavh Nagarajan, Anders Andreassen, and Behnam Neyshabur. Understanding the failure modes of out-of-distribution generalization. *arXiv preprint arXiv:2010.15775*, 2020.
- [35] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias, 2021.
- [36] Narges Honarvar Nazari and Adriana Kovashka. Domain generalization using shape representation. In *Proceedings of European Conference on Computer Vision - ECCV 2020 Workshops*, pages 666–670. Springer, 2020.

- [37] Judea Pearl. Causal inference. In Isabelle Guyon, Dominik Janzing, and Bernhard Schölkopf, editors, *Causality: Objectives and Assessment (NIPS 2008 Workshop)*, pages 39–58. JMLR.org, 2010.
- [38] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *2019 IEEE/CVF International Conference on Computer Vision*, pages 1406–1415. IEEE, 2019.
- [39] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5102–5112. PMLR, 2019.
- [40] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 12553–12562. IEEE, 2020.
- [41] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauero. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *Proceedings of The 7th International Conference on Learning Representations*, 2019.
- [42] Elan Rosenfeld, Pradeep Ravikumar, and Andrej Risteski. The risks of invariant risk minimization. *CoRR*, abs/2010.05761, 2020.
- [43] Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *CoRR*, abs/1911.08731, 2019.
- [44] Mattia Segù, Alessio Tonioni, and Federico Tombari. Batch normalization embeddings for deep domain generalization. *CoRR*, abs/2011.12672, 2020.
- [45] Thanh-Dat Truong, Chi Nhan Duong, Khoa Luu, and Minh-Triet Tran. Recognition in unseen domains: Domain generalization via universal non-volume preserving models. *CoRR*, abs/1905.13040, 2019.
- [46] Thanh-Dat Truong, Khoa Luu, Chi Nhan Duong, Ngan Le, and Minh-Triet Tran. Image alignment in unseen domains via domain deep generalization. *CoRR*, abs/1905.12028, 2019.

- [47] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [48] Vladimir Vapnik. An overview of statistical learning theory. *IEEE Trans. Neural Networks*, 10(5):988–999, 1999.
- [49] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 5385–5394. IEEE Computer Society, 2017.
- [50] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. *CoRR*, abs/2001.00677, 2020.
- [51] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto L. Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *2019 IEEE/CVF International Conference on Computer Vision*, pages 2100–2110. IEEE, 2019.
- [52] Sergey Zakharov, Wadim Kehl, and Slobodan Ilic. Deceptionnet: Network-driven domain randomization. In *2019 IEEE/CVF International Conference on Computer Vision*, pages 532–541. IEEE, 2019.
- [53] Marvin Zhang, Henrik Marklund, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group shift. *CoRR*, abs/2007.02931, 2020.
- [54] Han Zhao, Chen Dan, Bryon Aragam, Tommi S. Jaakkola, Geoffrey J. Gordon, and Pradeep Ravikumar. Fundamental limits and tradeoffs in invariant representation learning. *CoRR*, abs/2012.10713, 2020.
- [55] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International Conference on Machine Learning*, pages 7523–7532, 2019.
- [56] Han Zhao, Shanghang Zhang, Guanhang Wu, José M. F. Moura, João Paulo Costeira, and Geoffrey J. Gordon. Adversarial multiple source domain adaptation. In *Advances in Neural Information Processing Systems*, pages 8568–8579, 2018.

- [57] Sicheng Zhao, Bo Li, Pengfei Xu, and Kurt Keutzer. Multi-source domain adaptation in the deep learning era: A systematic survey. *arXiv preprint arXiv:2002.12169*, 2020.
- [58] Sicheng Zhao, Bo Li, Pengfei Xu, Xiangyu Yue, Guiguang Ding, and Kurt Keutzer. Madan: multi-source adversarial domain aggregation network for domain adaptation. *International Journal of Computer Vision*, pages 1–26, 2021.
- [59] Sicheng Zhao, Bo Li, Xiangyu Yue, Yang Gu, Pengfei Xu, Runbo Hu, Hua Chai, and Kurt Keutzer. Multi-source domain adaptation for semantic segmentation. In *Advances in Neural Information Processing Systems*, pages 7285–7298, 2019.
- [60] Sicheng Zhao, Xiangyu Yue, Shanghang Zhang, Bo Li, Han Zhao, Bichen Wu, Ravi Krishna, Joseph E Gonzalez, Alberto L Sangiovanni-Vincentelli, Sanjit A Seshia, et al. A review of single-source deep unsupervised visual domain adaptation. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

Supplementary Materials

Construction Details of Synthetic Dataset

In this section, we will give more details on constructing the **Cross Line** and **Vertical Line** experiments.

Cross Lines Experiment

Based on CIFAR10 dataset, we create ten-valued spurious feature and add a vertical line passing through the middle of each channel, and a horizontal line passing through the first channel. For each line added to the channel, we implement by adding the value taken of $0.5 \pm 0.5\mathcal{B}$ where $\mathcal{B} \in [-1, 1]$. Four lines, each with two choices of $+\mathcal{B}$ or $-\mathcal{B}$. Then we have in total $2^4 = 16$ choices. We select 10 of the 16 configurations to map each configuration into one specific class. In detail, we select the images with specific class (e.g. bird), and add the the line with specific configuration (e.g. $0.5 + 0.5\mathcal{B}$, $\mathcal{B} = 0.8$). Similar with [34], we add the line of i -th configuration to corresponding class images with a probability of $p_{ii} = 0.5$; for other j -th class, we set $p_{ij} = (1 - p_{ii})/10 = 0.05$. We call the i -th class images added with i -th configuration line with p_{ii} the majority group. We call the i -th class images added with other j -th configuration line with p_{ij} the minority group. The specific configurations are in the Table 5.

Vertical Line Experiment

Based on CIFAR10 dataset, we add a vertical line to the last channel of all images. In detail, we add the line with value from $\mathcal{B} \in [-4, 4]$. To avoid negative values, all pixels in last channel are added by 4, and then added by \mathcal{B} , and then divided by 9 to ensure the pixels lie in the range of $[0, 1]$. Such operations will add non-orthogonal components to images, where each data-point is represented on the plane of $(x_{\text{inv}}, x_{\text{env}} + x_{\text{inv}})$ because of the added line with constant value in last channel. In [34], they show the non-orthogonal versus orthogonal experiment, the results suggest that the non-orthogonal images $(x_{\text{inv}}, x_{\text{env}} + x_{\text{inv}})$ are more hard-to-disentangle than orthogonal ones $(x_{\text{inv}}, x_{\text{env}})$, and would cause geometric skews of a max-margin classifier. In our experiment, during training, we set the $\mathcal{B} = 4$ and 0, and test on domains with different $\mathcal{B} \in \{-4, -2, 0, 2, 4\}$.

More Results on DomainBed

We report the rest experiment results in this section. In training-domain validation set, the validation set is subset of training set, we choose the model that performs

Table 5: 10 configurations in **Cross Lines** experiments.

Configuration #	Channel	Line's Position	Sign
0	0	Vertical	+
	1	Vertical	+
	2	Vertical	+
	0	Horizontal	+
1	0	Vertical	-
	1	Vertical	+
	2	Vertical	+
	0	Horizontal	+
2	0	Vertical	+
	1	Vertical	-
	2	Vertical	+
	0	Horizontal	+
3	0	Vertical	+
	1	Vertical	+
	2	Vertical	+
	0	Horizontal	+
4	0	Vertical	+
	1	Vertical	+
	2	Vertical	-
	0	Horizontal	+
5	0	Vertical	+
	1	Vertical	+
	2	Vertical	+
	0	Horizontal	-
6	0	Vertical	-
	1	Vertical	-
	2	Vertical	+
	0	Horizontal	+
7	0	Vertical	+
	1	Vertical	-
	2	Vertical	-
	0	Horizontal	+
8	0	Vertical	+
	1	Vertical	+
	2	Vertical	-
	0	Horizontal	-
9	0	Vertical	-
	1	Vertical	+
	2	Vertical	+
	0	Horizontal	-
10	0	Vertical	-
	1	Vertical	+
	2	Vertical	-
	0	Horizontal	+

best on the overall validation set for each domain. This strategy characterizes the in-distribution generalization capability of the model.

The results are recorded in Table 6. From the Table, we can see that IIB achieves 67.5% across 7 datasets on average, which is comparable to the best algorithm CORAL on DomainBed. Also IIB shows better performance on larger datasets (e.g. OfficeHome, DomainNet). The results demonstrate IIB’s in-domains generalization ability.

Table 6: Performance comparison (Acc. %) between the proposed IIB method and the state-of-the-art domain generalization methods with *training-domain validation set* model selection strategy. The best accuracy in each dataset is presented in boldface. The average accuracy over all the datasets is also reported.

Algorithm	ColoredMNIST	RotatedMNIST	VLCS	PACS	OfficeHome	TerraIncognita	DomainNet	Avg
ERM	51.5 ± 0.1	98.0 ± 0.0	77.5 ± 0.4	85.5 ± 0.2	66.5 ± 0.3	46.1 ± 1.8	40.9 ± 0.1	66.6
IRM	52.0 ± 0.1	97.7 ± 0.1	78.5 ± 0.5	83.5 ± 0.8	64.3 ± 2.2	47.6 ± 0.8	33.9 ± 2.8	65.4
GroupDRO	52.1 ± 0.0	98.0 ± 0.0	76.7 ± 0.6	84.4 ± 0.8	66.0 ± 0.7	43.2 ± 1.1	33.3 ± 0.2	64.8
Mixup	52.1 ± 0.2	98.0 ± 0.1	77.4 ± 0.6	84.6 ± 0.6	68.1 ± 0.3	47.9 ± 0.8	39.2 ± 0.1	66.7
MLDG	51.5 ± 0.1	97.9 ± 0.0	77.2 ± 0.4	84.9 ± 1.0	66.8 ± 0.6	47.7 ± 0.9	41.2 ± 0.1	66.7
CORAL	51.5 ± 0.1	98.0 ± 0.1	78.8 ± 0.6	86.2 ± 0.3	68.7 ± 0.3	47.6 ± 1.0	41.5 ± 0.1	67.5
MMD	51.5 ± 0.2	97.9 ± 0.0	77.5 ± 0.9	84.6 ± 0.5	66.3 ± 0.1	42.2 ± 1.6	23.4 ± 9.5	63.3
DANN	51.5 ± 0.3	97.8 ± 0.1	78.6 ± 0.4	83.6 ± 0.4	65.9 ± 0.6	46.7 ± 0.5	38.3 ± 0.1	66.1
CDANN	51.7 ± 0.1	97.9 ± 0.1	77.5 ± 0.1	82.6 ± 0.9	65.8 ± 1.3	45.8 ± 1.6	38.3 ± 0.3	65.6
MTL	51.4 ± 0.1	97.9 ± 0.0	77.2 ± 0.4	84.6 ± 0.5	66.4 ± 0.5	45.6 ± 1.2	40.6 ± 0.1	66.2
SagNet	51.7 ± 0.0	98.0 ± 0.0	77.8 ± 0.5	86.3 ± 0.2	68.1 ± 0.1	48.6 ± 1.0	40.3 ± 0.1	67.2
ARM	56.2 ± 0.2	98.2 ± 0.1	77.6 ± 0.3	85.1 ± 0.4	64.8 ± 0.3	45.5 ± 0.3	35.5 ± 0.2	66.1
VREx	51.8 ± 0.1	97.9 ± 0.1	78.3 ± 0.2	84.9 ± 0.6	66.4 ± 0.6	46.4 ± 0.6	33.6 ± 2.9	65.6
RSC	51.7 ± 0.2	97.6 ± 0.1	77.1 ± 0.5	85.2 ± 0.9	65.5 ± 0.9	46.6 ± 1.0	38.9 ± 0.5	66.1
IIB(Ours)	52.0 ± 0.3	98.1 ± 0.2	77.6 ± 0.2	85.7 ± 0.6	69.0 ± 0.1	48.5 ± 0.4	41.6 ± 0.8	67.5