

Article

# Counterfactual Supervision-Based Information Bottleneck for Out-of-Distribution Generalization

Bin Deng \*  and Kui Jia \* 

School of Electronic and Information Engineering, South China University of Technology, Guangzhou 510641, China

\* Correspondence: eebindeng@mail.scut.edu.cn (B.D.); kuijia@scut.edu.cn (K.J.)

**Abstract:** Learning invariant (causal) features for out-of-distribution (OOD) generalization have attracted extensive attention recently, and among the proposals, invariant risk minimization (IRM) is a notable solution. In spite of its theoretical promise for linear regression, the challenges of using IRM in linear classification problems remain. By introducing the information bottleneck (IB) principle into the learning of IRM, the IB-IRM approach has demonstrated its power to solve these challenges. In this paper, we further improve IB-IRM from two aspects. First, we show that the key assumption of support overlap of invariant features used in IB-IRM guarantees OOD generalization, and it is still possible to achieve the optimal solution without this assumption. Second, we illustrate two failure modes where IB-IRM (and IRM) could fail in learning the invariant features, and to address such failures, we propose a *Counterfactual Supervision-based Information Bottleneck (CSIB)* learning algorithm that recovers the invariant features. By requiring counterfactual inference, CSIB works even when accessing data from a single environment. Empirical experiments on several datasets verify our theoretical results.

**Keywords:** out-of-distribution generalization; information bottleneck; causal learning



**Citation:** Deng, B.; Jia, K.

Counterfactual Supervision-Based Information Bottleneck for Out-of-Distribution Generalization. *Entropy* **2023**, *25*, 193. <https://doi.org/10.3390/e25020193>

Academic Editors: Sotiris Kotsiantis, Gerhard Bauch and Jan Lewandowski

Received: 10 October 2022  
Revised: 14 December 2022  
Accepted: 16 January 2023  
Published: 18 January 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Modern machine learning models are prone to catastrophic performance loss during deployment when the test distribution is different from the training distribution. This phenomenon has been repeatedly witnessed and intentionally exposed in many examples [1–5]. Among the explanations, shortcut learning [6] is considered as a main factor causing this phenomenon. A good example is the classification of images of cows and camels—a trained convolutional network tends to recognize cows or camels by learning spurious features from image backgrounds (e.g., green pastures for cows and deserts for camels), rather than learning the causal shape features of the animals [7]; decisions based on the spurious features would make the learned models fail when cows or camels appear in unusual or different environments. Machine learning models are expected to have the capability of out-of-distribution (OOD) generalization and avoid shortcut learning.

To achieve OOD generalization, recent theories [8–12] are motivated by causality literature [13,14] and resort to extraction of the invariant, causal features and establishing the relevant conditions under which machine learning models have guaranteed generalization. Among these works, invariant risk minimization (IRM) [8] is a notable learning paradigm that incorporates the invariance principle [15] into practice. In spite of the theoretical promise of IRM, it is only applicable to problems of linear regression. For other problems, such as linear classification, Ahuja et al. [12] first show that for OOD generalization, linear classification is more difficult (see Theorem 1) and propose a new learning method of information bottleneck-based invariant risk minimization (IB-IRM) based on the support overlap assumption (Assumption 7). In this work, we closely investigate the conditions identified in [12] and propose improved results for OOD generalization of linear classification.

Our technical contributions are as follows. In [12], a notion of support overlap of invariant features is assumed in order to make the OOD generalization of linear classification successful. In this work, we first show that this assumption is strong, but it is still possible to achieve such goal without this assumption. Then, we examine whether the IB-IRM proposed in [12] is sufficient to learn invariant features for linear classification and find that IB-IRM (and IRM) could fail in two modes. We then analyze two failure modes of IB-IRM and IRM, in particular when the spurious features in training environments capture sufficient information for the task of interest but have less information than the invariant features. Based on the above analyses, we propose a new method, termed counterfactual supervision-based information bottleneck (CSIB), to address such failures. We prove that without the need of the support overlap assumption, CSIB is theoretically guaranteed for the success of OOD generalization in linear classification. Notably, CSIB works even when accessing data from a single environment. Finally, we design three synthetic datasets and a colored MINST dataset based on our examples; experiments demonstrate the effectiveness of CSIB empirically.

The rest of this article is organized as follows. The learning problem of out-of-distribution (OOD) generalization is formulated in Section 2. In Section 3, we study the learnability of the OOD generalization with different assumptions to the training and test environments. Using these assumptions, two failure modes of previous methods (IRM and IB-IRM) are analysed in Section 4. Based on the above analysis, our method is then proposed in Section 5. The experiments are reported in Section 6. Finally, we discuss the related works in Section 7 and provide some conclusions and limitations of our work in Section 8. All the proofs and details of experiments are given in the Appendices A and B.

## 2. OOD Generalization: Background and Formulations

### 2.1. Background on Structural Equation Models

Before introducing our formulations of OOD generalization, we provide a detailed background on structural equation models (SEMs) [8,13].

**Definition 1** (Structural Equation Model (SEM)). *A structural equation model (SEM)  $\mathcal{C} := (\mathcal{S}, \mathcal{N})$  governing the random vector  $X = (X_1, \dots, X_d)$  is a set of structural equations:*

$$\mathcal{S}_i : X_i \leftarrow f_i(\text{Pa}(X_i), N_i),$$

where  $\text{Pa}(X_i) \subseteq \{X_1, \dots, X_d\} \setminus \{X_i\}$  are called the parents of  $X_i$ , and  $N_i$  are independent noise random variables. For every SEM, we yield a directed acyclic graph (DAG)  $\mathcal{G}$  by adding one vertex for each  $X_i$  and directed edges from each parent in  $\text{Pa}(X_i)$  (the causes) to child  $X_i$  (the effect).

**Definition 2** (Intervention). *Consider an SEM  $\mathcal{C} = (\mathcal{S}, \mathcal{N})$ . An intervention  $e$  on  $\mathcal{C}$  consists of replacing one or several of its structural equations to obtain an intervened SEM  $\mathcal{C}^e = (\mathcal{S}^e, \mathcal{N}^e)$ , with structural equations:*

$$\mathcal{S}_i^e : X_i^e \leftarrow f_i^e(\text{Pa}^e(X_i^e), N_i^e),$$

The variable  $X^e$  is intervened if  $\mathcal{S}_i \neq \mathcal{S}_i^e$  or  $N_i \neq N_i^e$ .

In an SEM  $\mathcal{C}$ , we can draw samples from the observational distribution  $\mathbb{P}(X)$  according to the topological ordering of its DAG  $\mathcal{G}$ . We can also manipulate (intervene) a unique SEM  $\mathcal{C}$  in different ways, indexed by  $e$ , to different but related SEMs  $\mathcal{C}^e$ , which results in different interventional distributions  $\mathbb{P}(X^e)$ . Such family of interventions are used to model the environments.

### 2.2. Formulations of OOD Generalization

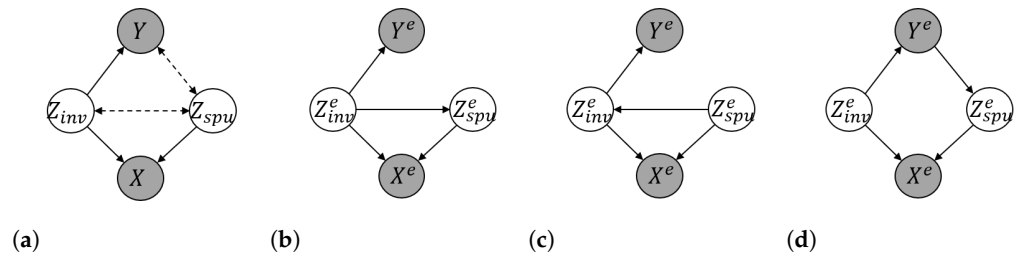
In this paper, we study the OOD generalization problem by following the linear classification structural equation model below [12].

**Assumption 1** (Linear classification SEM  $C_{ood}$ ).

$$\begin{aligned}
 Y &\leftarrow \mathbf{1}(w_{inv}^* \cdot Z_{inv}) \oplus N, \quad N \sim \text{Bernoulli}(q), \quad q < \frac{1}{2}; \\
 X &\leftarrow S(Z_{inv}, Z_{spu}),
 \end{aligned}
 \tag{1}$$

where  $w_{inv}^* \in \mathbb{R}^m$  is the labeling hyperplane,  $Z_{inv} \in \mathbb{R}^m$ ,  $Z_{spu} \in \mathbb{R}^o$ ,  $X \in \mathbb{R}^d$ ,  $\oplus$  is the XOR operator,  $S \in \mathbb{R}^{d \times (m+o)}$  is invertible ( $d = m + o$ ),  $\cdot$  is the dot product function, and  $\mathbf{1}(a) = 1$  if  $a \geq 0$  otherwise 0.

The SEM  $C_{ood}$  governs four random variables  $\{X, Y, Z_{inv}, Z_{spu}\}$ , and its directed acyclic graph (DAG) is illustrated in Figure 1a, where the exogenous noise variable  $N$  is omitted. Following Definition 2, each intervention  $e$  generates a new environment  $e$  with interventional distribution  $\mathbb{P}(X^e, Y^e, Z_{inv}^e, Z_{spu}^e)$ . We assume only the variables of  $X^e$  and  $Y^e$  are observable. In OOD generalization, we are interested in a set of environments  $\mathcal{E}_{all}$  defined as below.



**Figure 1.** (a) DAG of the SEM  $C_{ood}$  (Assumption 1); (b–d) DAGs of the interventional SEM  $C_{ood}^e$  in the training environments  $\mathcal{E}_{tr}$  with respect to different correlations between  $Z_{inv}$  and  $Z_{spu}$ . Grey nodes denote observed variables, and white nodes represent unobserved variables. Dashed lines denote the edges which might vary across the interventional environments and even be absent in some scenarios, whilst solid lines indicate that they are invariant across all the environments. All exogenous noise variables are omitted in the DAGs.

**Definition 3** ( $\mathcal{E}_{all}$ ). Consider the SEM  $C_{ood}$  (Assumption 1) and the learning goal of predicting  $Y$  from  $X$ . Then, the set of all environments  $\mathcal{E}_{all}(C_{ood})$  indexes all the interventional distributions  $\mathbb{P}(X^e, Y^e)$  obtainable by valid interventions  $e$ . An intervention  $e \in \mathcal{E}_{all}(C_{ood})$  is valid as long as (i) the DAG remains acyclic, (ii)  $\mathbb{P}(Y^e|Z_{inv}^e) = \mathbb{P}(Y|Z_{inv})$ , and (iii)  $\mathbb{P}(X^e|Z_{inv}^e, Z_{spu}^e) = \mathbb{P}(X|Z_{inv}, Z_{spu})$ .

Assumption 1 shows that  $Z_{inv}$  is the cause of the response  $Y$ . We name  $Z_{inv}$  the invariant features or causal features because  $\mathbb{P}(Y^e|Z_{inv}^e) = \mathbb{P}(Y|Z_{inv})$  always holds among all valid interventional SEMs  $C_{ood}^e$ , as defined in Definition 3. The  $Z_{spu}$  is called spurious features because  $\mathbb{P}(Y^e|Z_{spu}^e)$  may vary in different environments of  $\mathcal{E}_{all}$ .

Let  $D = \{D^e\}_{e \in \mathcal{E}_{tr}}$  be the training data gathered from a set of training environments  $\mathcal{E}_{tr} \subset \mathcal{E}_{all}$ , where  $D^e = \{(x_i^e, y_i^e)\}_{i=1}^{n_e}$  is the dataset from environment  $e$  with each instance  $(x_i^e, y_i^e)$  i.i.d. drawn from  $\mathbb{P}(X^e, Y^e)$ . Let  $\mathcal{X}^e \subseteq \mathbb{R}^d$  and  $\mathcal{Y} \subseteq \{0, 1\}$  be the support sets of  $X^e$  and  $Y$ , respectively. Given observed data  $D$ , the goal of OOD generalization is to find a predictor  $f : \mathbb{R}^d \rightarrow \mathcal{Y}$  such that it can perform well across a set of OOD environments (test environments)  $\mathcal{E}_{ood}$  of interest, where  $\mathcal{E}_{ood} \subseteq \mathcal{E}_{all}$ . Formally, it is expected to minimize

$$\max_{e \in \mathcal{E}_{ood}} R^e(f),
 \tag{2}$$

where  $R^e(f) := \mathbb{E}_{X^e, Y^e} [l(f(X^e), Y^e)]$  is the risk under the environment  $e$  with  $l(\cdot, \cdot)$  the 0-1 loss function. Since  $\mathcal{E}_{ood}$  may be different from  $\mathcal{E}_{tr}$ , this learning problem is called OOD generalization. We assume the predictor  $f = w \circ \Phi$  includes a feature extractor  $\Phi : \mathcal{X} \rightarrow \mathcal{H}$  and a classifier  $w : \mathcal{H} \rightarrow \mathcal{Y}$ . With a slight abuse of notation, we also let the classifier  $w$

and feature extractor  $\Phi$  be parameterized by themselves, respectively, as  $w \in \mathbb{R}^{c+1}$  and  $\Phi \in \mathbb{R}^{c \times d}$  with  $c$  the number of feature dimension.

### 2.3. Background on IRM and IB-IRM

To minimize Equation (2), two notable solutions of IRM [8] and IB-IRM [12] are listed as follows:

$$\text{IRM: } \min_{w, \Phi} \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} R^e(w \circ \Phi), \text{ s.t. } w \in \arg \min_{\tilde{w}} R^e(\tilde{w} \circ \Phi), \forall e \in \mathcal{E}_{tr}, \tag{3}$$

$$\text{IB-IRM: } \min_{w, \Phi} \sum_{e \in \mathcal{E}_{tr}} h^e(\Phi), \text{ s.t. } \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} R^e(w \circ \Phi) \leq rth, w \in \arg \min_{\tilde{w}} R^e(\tilde{w} \circ \Phi), \forall e \in \mathcal{E}_{tr}, \tag{4}$$

where  $R^e(w \circ \Phi) = \mathbb{E}_{X^e, Y^e} [l(w \circ \Phi(X^e), Y^e)]$ , and  $h^e(\Phi) = H(\Phi(X^e))$  with  $H$  the Shannon entropy (or a lower bounded differential entropy), and  $rth$  is the threshold on the average risk. If we drop the invariance constraint from IRM and IB-IRM, we obtain standard empirical risk minimization (ERM) and information bottleneck-based empirical risk minimization (IB-ERM), respectively. The use of an entropy constraint in IB-IRM is inspired from the information bottleneck principle [16] where mutual information  $I(X; \Phi(X))$  is used for information compression. Since the representation  $\Phi(X)$  is a deterministic mapping of  $X$ , we have

$$I(X; \Phi(X)) = H(\Phi(X)) - H(\Phi(X)|X) = H(\Phi(X)), \tag{5}$$

thus minimizing the entropy of  $\Phi(X)$  is equivalent to minimizing the mutual information  $I(X; \Phi(X))$ . In brief, the optimization goal of IB-IRM is to select the one that has the least entropy among all highly predictive invariant predictors.

### 3. OOD Generalization: Assumptions and Learnability

To study the learnability of OOD generalization, we make following definition.

**Definition 4.** Given  $\mathcal{E}_{tr} \subset \mathcal{E}_{all}$  and  $\mathcal{E}_{ood} \subseteq \mathcal{E}_{all}$ . We say an algorithm succeeds to solve OOD generalization with respect to  $(\mathcal{E}_{tr}, \mathcal{E}_{ood})$  if the predictor  $f^* \in \mathcal{F}$  returned by this algorithm satisfies the following equation:

$$\max_{e \in \mathcal{E}_{ood}} R^e(f^*) = \min_{f \in \mathcal{F}} \max_{e \in \mathcal{E}_{ood}} R^e(f), \tag{6}$$

where  $\mathcal{F}$  is the learning hypothesis (a function set including all possible linear classifier). Otherwise we say it fails to solve OOD generalization.

So far, we have omitted how different environments of  $\mathcal{E}_{tr}$  and  $\mathcal{E}_{ood}$  exactly are to enable OOD generalization. Different assumptions about  $\mathcal{E}_{tr}$  and  $\mathcal{E}_{ood}$  make the OOD generalization problem different.

#### 3.1. Assumptions about the Training Environments $\mathcal{E}_{tr}$

Define the support set of the invariant (resp., spurious) features  $Z_{inv}^e$  (resp.,  $Z_{spu}^e$ ) in environment  $e$  as  $Z_{inv}^e$  (resp.,  $Z_{spu}^e$ ). In general, we make following assumptions to the invariant features  $Z_{inv}^e$  in the training environments  $\mathcal{E}_{tr}$ .

**Assumption 2** (Bounded invariant features).  $\cup_{e \in \mathcal{E}_{tr}} Z_{inv}^e$  is a bounded set. (A set  $\mathcal{Z}$  is bounded if  $\exists M < \infty$  such that  $\forall z \in \mathcal{Z}, \|z\| \leq M$ ).

**Assumption 3** (Strictly separable invariant features).  $\forall z \in \cup_{e \in \mathcal{E}_{tr}} Z_{inv}^e, w_{inv}^* \cdot z \neq 0$ .

The difficulties of OOD generalization are due to the spurious correlations between  $Z_{inv}$  and  $Z_{spu}$  in the training environments  $\mathcal{E}_{tr}$ . In this paper, we consider three modes induced by different correlations between  $Z_{inv}$  and  $Z_{spu}$  as shown below.

**Assumption 4** (Spurious correlation 1). Assume each  $e \in \mathcal{E}_{tr}$ ,

$$Z_{spu}^e \leftarrow AZ_{inv}^e + W^e; \tag{7}$$

where  $A \in \mathbb{R}^{o \times m}$ , and  $W^e \in \mathbb{R}^o$  is a continuous (or discrete with each component supported on at least two distinct values), bounded, and zero mean noise variable.

**Assumption 5** (Spurious correlation 2). Assume each  $e \in \mathcal{E}_{tr}$ ,

$$Z_{inv}^e \leftarrow AZ_{spu}^e + W^e; \tag{8}$$

where  $A \in \mathbb{R}^{m \times o}$ , and  $W^e \in \mathbb{R}^m$  is a continuous (or discrete with each component supported on at least two distinct values), bounded, and zero mean noise variable.

**Assumption 6** (Spurious correlation 3). Assume each  $e \in \mathcal{E}_{tr}$ ,

$$Z_{spu}^e \leftarrow W_1^e Y^e + W_0^e (1 - Y^e); \tag{9}$$

where  $W_0^e \in \mathbb{R}^o$  and  $W_1^e \in \mathbb{R}^o$  are independent noise variables.

For each  $e \in \mathcal{E}_{tr}$ , the DAGs of its corresponding interventional SEMs  $\mathcal{C}_{ood}^e$  with respect to Assumptions 4–6 are illustrated in Figure 1b–d, respectively. It is worth noting that although the DAGs are identical across all training environments in each mode of Assumptions 4–6, the interventional SEMs  $\mathcal{C}_{ood}^e$  among different training environments are different due to the interventions on the exogenous noise variables.

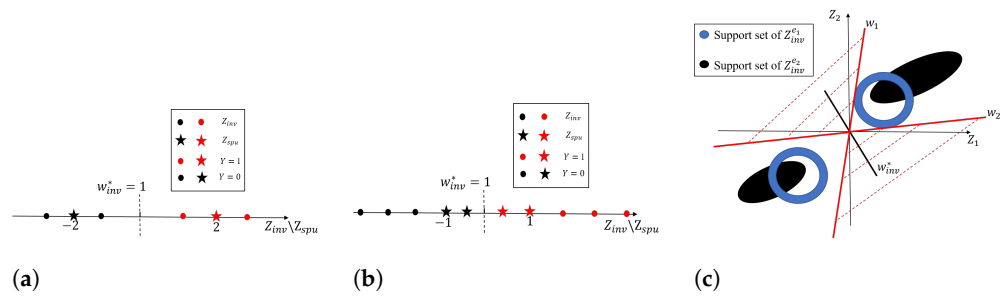
### 3.2. Assumptions about the OOD Environments $\mathcal{E}_{ood}$

**Theorem 1** (Impossibility of guaranteed OOD generalization for linear classification [12]). Suppose  $\mathcal{E}_{ood} = \mathcal{E}_{all}$ . If for all the training environments  $\mathcal{E}_{tr}$ , the latent invariant features are bounded and strictly separable, i.e., Assumptions 2 and 3 hold, then every deterministic algorithm fails to solve the OOD generalization.

The above theorem shows that it is impossible to solve OOD generalization if  $\mathcal{E}_{ood} = \mathcal{E}_{all}$ . To make it learnable, Ahuja et al. [12] propose the support overlap assumption (Assumption 7) to the invariant features.

**Assumption 7** (Invariant feature support overlap).  $\forall e \in \mathcal{E}_{ood}, \mathcal{Z}_{inv}^e \subseteq \cup_{e' \in \mathcal{E}_{tr}} \mathcal{Z}_{inv}^{e'}$ .

However, Assumption 7 is strong, and we would show that it is still possible to solve OOD generalization without this assumption. For better illustration, consider an OOD generalization task from  $\mathbb{P}(X^{e_1}, Y^{e_2})$  to  $\mathbb{P}(X^{e_2}, Y^{e_2})$  with  $\mathcal{E}_{tr} = \{e_1\}$  and  $\mathcal{E}_{ood} = \{e_2\}$ , and the support sets of the corresponding invariant features  $Z_{inv}^{e_1}$  and  $Z_{inv}^{e_2}$  are intuitively illustrated in Figure 2c (assume  $dim(Z_{inv}) = 2$  in this example). From Figure 2c, it is clear that although the support sets of invariant features between the two environments are different, it is still possible to solve OOD generalization if the learned feature extractor  $\Phi$  only captures the invariant features, e.g.,  $\Phi(X) = Z_{inv}$ .



**Figure 2.** (a) Example 1; (b) example 2; (c) example illustration. Here,  $\dim(Z_{inv}) = 2$  and  $Z_{inv} = (Z_1, Z_2)$ . The blue and black regions represent the support sets of  $Z_{inv}^{e_1}$  and  $Z_{inv}^{e_2}$ , corresponding to the environments  $e_1$  and  $e_2$ , respectively.  $\mathcal{E}_{tr} = \{e_1\}$  is the training environment and  $\mathcal{E}_{ood} = \{e_2\}$  is the OOD environment. Although Assumption 7 does not hold in this example, any zero-error classifier with  $\Phi(X) = Z_{inv}$  on the  $e_1$  environment data would clearly make the classification error zero in  $e_2$ , thus succeeding to solve OOD generalization.

To make Assumption 7 weaker, we propose the following assumption.

**Assumption 8.** Let  $\mathbb{P}(Z_{inv}^{tr}, Y^{tr}) = \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} \mathbb{P}(Z_{inv}^e, Y^e)$  be the mixture distribution of invariant features in the training environments. Denote  $\mathcal{A}$  be a hypothesis set including all linear classifiers mapping from  $\mathbb{R}^m$  to  $\mathcal{Y}$ .  $\forall e \in \mathcal{E}_{ood}$ , assume  $F_l(\mathbb{P}(Z_{inv}^{tr}, Y^{tr})) \subseteq F_l(\mathbb{P}(Z_{inv}^e, Y^e))$ , where  $l$  is the 0-1 loss function and  $F_l(\mathbb{P}(Z, Y)) = \arg \min_{f \in \mathcal{A}} \mathbb{E}_{Z, Y}[l(f(Z), Y)]$ .

Clearly, under the assumption of separable invariant features (Assumption 3), for any  $e \in \mathcal{E}_{ood}$ , Assumption 7 holds  $\Rightarrow \mathcal{Z}_{inv}^e \subseteq \mathcal{Z}_{inv}^{tr} \Rightarrow F_l(\mathbb{P}(Z_{inv}^{tr}, Y^{tr})) \subseteq F_l(\mathbb{P}(Z_{inv}^e, Y^e)) \Rightarrow$  Assumption 8 holds, but not vice versa. Therefore, Assumption 8 is weaker than Assumption 7. We show that Assumption 8 could be substituted for Assumption 7 for the success of OOD generalization in our proposed method in Section 5.

**4. Failures of IRM and IB-IRM**

Under Spurious Correlation 1 (Assumption 4), the IB-IRM algorithm has been shown to enable OOD generalization, while IRM fails [12]. In this section, we would show that both IRM and IB-IRM could fail under Spurious Correlations 2 and 3 (Assumptions 5 and 6).

**4.1. Failure under Spurious Correlation 2**

**Example 1 (Counter-Example 1).** Under Assumption 5, let  $Z_{inv}^e \leftarrow Z_{spu}^e + W^e$  with  $\dim(Z_{inv}^e) = \dim(Z_{spu}^e) = \dim(W^e) = 1$  and  $w_{inv}^* = 1$  be the generated classifier in Assumption 1. We assume two training environments and a OOD environment as:

$$\begin{aligned} \mathcal{E}_{tr} &= \{e_1, e_2\}; & \mathcal{E}_{ood} &= \{e_3\}; \\ e_1 &: \mathbb{P}(Z_{spu}^{e_1} = -2) = 1, \mathbb{P}(W^{e_1} = -1) = 0.5, \mathbb{P}(W^{e_1} = 1) = 0.5; \\ e_2 &: \mathbb{P}(Z_{spu}^{e_2} = 2) = 1, \mathbb{P}(W^{e_2} = -1) = 0.5, \mathbb{P}(W^{e_2} = 1) = 0.5; \\ e_3 &: \mathbb{P}(Z_{spu}^{e_3} = 1) = 1, \mathbb{P}(W^{e_3} = -2) = 0.5, \mathbb{P}(W^{e_3} = 2) = 0.5. \end{aligned}$$

Figure 2a shows the support points of these features in the training environments. Then, by applying any algorithm to solve the above example with  $rth = q$ , we would obtain a predictor of  $f^* = w^* \circ \Phi^*$ . Consider the prediction made by this model as (we ignore the classifier bias for convenience)

$$f^*(X^e) = f^*(S(Z_{inv}^e, Z_{spu}^e)) = \mathbf{1}(\Phi_{inv}^* Z_{inv}^e + \Phi_{spu}^* Z_{spu}^e). \tag{10}$$

It is trivial to show that the  $f^*$  of  $\Phi_{inv}^* = 0$  and  $\Phi_{spu}^* = 1$  is an invariant predictor across training environments with classification error  $R^{e_1} = R^{e_2} = q$ , and it achieves the least entropy of  $h^e(\Phi^*) = 0$  for each training environment  $e$ . Therefore, it is a solution of IB-IRM



and IRM. However, the predictor of  $f^*$  relies on spurious features and has the test error  $R^{e_3} = 0.5$ ; thus, it fails to solve the OOD generalization.

#### 4.2. Failure under Spurious Correlation 3

**Example 2** (Counter-Example 2). Under Assumption 6, let  $Z_{spu}^e \leftarrow W_1^e Y^e + W_0^e (1 - Y^e)$  with  $\dim(Z_{inv}) = \dim(Z_{spu}) = \dim(W_0^e) = \dim(W_1^e) = 1$ ,  $Z_{inv}^e$  be a discrete variable supported uniformly on six points  $\{-4, -3, -2, 2, 3, 4\}$  among all environments, and  $w_{inv}^* = 1$  be the generated classifier in Assumption 1. We assume two training environments and a OOD environment as:

$$\begin{aligned} \mathcal{E}_{tr} &= \{e_1, e_2\}; \quad \mathcal{E}_{ood} = \{e_3\} \\ e_1 &: \mathbb{P}(W_0^{e_1} = -1) = 1, \mathbb{P}(W_1^{e_1} = 1) = 1; \\ e_2 &: \mathbb{P}(W_0^{e_2} = -0.5) = 1, \mathbb{P}(W_1^{e_2} = 0.5) = 1; \\ e_3 &: \mathbb{P}(W_0^{e_3} = 1) = 1, \mathbb{P}(W_1^{e_3} = -1) = 1; \end{aligned}$$

Figure 2b shows the support points of these features in the training environments. Then, by applying any algorithm to solve the above example with  $rth = q$ , we would obtain a predictor of  $f^* = w^* \circ \Phi^*$ . Consider the prediction made by this model as (we ignore the classifier bias for convenience):

$$f^*(X^e) = f^*(S(Z_{inv}^e, Z_{spu}^e)) = \mathbf{1}(\Phi_{inv}^* Z_{inv}^e + \Phi_{spu}^* Z_{spu}^e). \tag{11}$$

It is trivial to show that the  $f^*$  of  $\Phi_{inv}^* = 0$  and  $\Phi_{spu}^* = 1$  is an invariant predictor across training environments with classification error  $R^{e_1} = R^{e_2} = 0$ , and it achieves the least entropy of  $h^e(\Phi^*) = 1$  among all highly predictive predictors for each training environment  $e$ . and Therefore, it is a solution of IB-IRM and IRM. However, the predictor of  $f^*$  relies on spurious features and has the test error  $R^{e_3} = 1$ ; thus, it fails to solve the OOD generalization.

#### 4.3. Understanding the Failures

From the illustrations of the above simple examples, we can conclude that the failure of the invariance constraint for removing the spurious features is because the spurious features among all training environments are strictly linearly separable by their corresponding labels. This would make the predictor rely only on spurious features to achieve minimum training error and also be the invariant predictor across training environments. Since the label set is finite (with only two values in binary classification) in classification problems, such a phenomenon may exist. We state such failure mode formally as below.

**Theorem 2.** Given any  $\mathcal{E}_{tr} \subset \mathcal{E}_{all}$  and  $\mathcal{E}_{ood} \subseteq \mathcal{E}_{all}$  satisfying Assumptions 2, 3, and 7, if two sets  $\cup_{e \in \mathcal{E}_{tr}} Z_{spu}^e (Y^e = 1)$  and  $\cup_{e \in \mathcal{E}_{tr}} Z_{spu}^e (Y^e = 0)$  are linearly separable and  $H(Z_{inv}^e) > H(Z_{spu}^e)$  on each training environment  $e$ , then IB-IRM (and IRM, ERM, or IB-ERM) with any  $rth \in \mathbb{R}$  fails to solve the OOD generalization.

The understanding of Theorem 2 is intuitive since when the spurious features in the training environments with respect to different labels are linearly separable, there is no algorithm that can distinguish spurious features from invariant features. Although the assumption of linear separation of the spurious features seems strong for this failure, it is easy to hold in high-dimensional space when  $\dim(Z_{spu})$  is large (common cases in practice such as image data). We show one case in Appendix A.3 that if the number of environments is  $|\mathcal{E}_{tr}| < \dim(Z_{spu})/2$  under Assumption 6, the spurious features in the training environments are probably separable by their labels. This is because in  $o$ -dimensional space there is a high probability that  $o$  randomly drawn distinct points are linearly separable for any two subsets.

### 5. Counterfactual Supervision-Based Information Bottleneck

In the above analyses, we have shown two failure modes of IB-IRM and IRM for OOD generalization in the linear classification problem. The key reason for the failure is due to the learned features  $\Phi(X)$  that rely on spurious features. To prevent such failure, we present the counterfactual supervision-based information bottleneck (CSIB) learning algorithm for removing the spurious features progressively.

In general, the IB-ERM method is applied to extract features from the beginning of each iteration:

$$\min_{w, \Phi} \sum_{e \in \mathcal{E}_{tr}} h^e(\Phi) \quad \text{s.t.} \quad \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} R^e(w \circ \Phi) \leq rth \tag{12}$$

Due to the information bottleneck, only a part of the information of the input  $X$  are exploited in  $\Phi(X)$ . If the information of spurious features  $Z_{spu}$  exists in the learned features  $\Phi(X)$ , the idea of CSIB is to drop such information and meanwhile maintain the causal information (represented by invariant features  $Z_{inv}$ ) as well. However, achieving such a goal faces two challenges: (1) how to determine whether  $\Phi(X)$  contains spurious information of  $Z_{spu}$ ? and (2) how to remove the information of  $Z_{spu}$ ?

Fortunately, due to the orthogonality in the linear space, it is possible to disentangle the features that are exploited by  $\Phi(X)$  (denoted as  $X_1$ ) and the features that are not exploited by  $\Phi(X)$  (denoted as  $X_2$ ) via Singular Value Decomposition (SVD). Based on that, we could construct an SEM  $\mathcal{C}_{new}$  governing three variables of  $X_1$ ,  $X_2$ , and  $X$ . Therefore, by conducting counterfactual interventions on  $X_1$  and  $X_2$  in  $\mathcal{C}_{new}$ , we could solve the first challenge by requiring a single supervision on the counterfactual examples  $X'$ . For example, if we intervene on  $X_1$  and find that the causal information remains in the resulting  $X'$ , then the extracted features  $\Phi(X)$  are definitely the spurious features. To address the second challenge, we replace the input by  $X_2$  by filtering out the information of  $X_1$  and conduct the same learning procedure from the beginning.

The learning algorithm of CSIB is illustrated in Algorithm 1, and Figure 3 shows the framework of CSIB. We show in Theorem 3 that CSIB is theoretically guaranteed to succeed to solve OOD generalization.

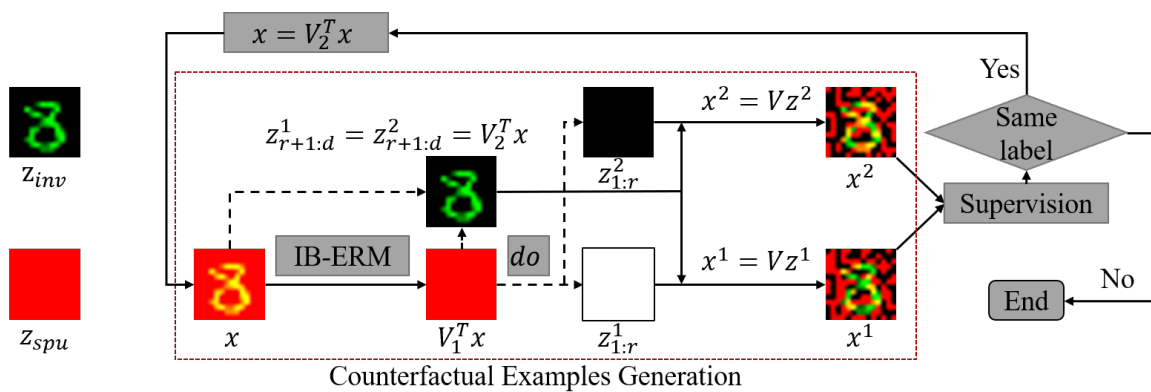


Figure 3. A simplified framework for the illustration of the proposed CSIB method.

**Theorem 3** (Guarantee of CSIB). *Given any  $\mathcal{E}_{tr} \subset \mathcal{E}_{all}$  and  $\mathcal{E}_{ood} \subseteq \mathcal{E}_{all}$  satisfying Assumptions 2, 3, and 8, then for every spurious correlation of Assumptions 4, 5, and 6 (in this correlation mode, assume the spurious features are linearly separable in the training environments), the CSIB algorithm with  $rth = q$  succeeds in solving the OOD generalization.*



**Algorithm 1** Counterfactual Supervision-based Information Bottleneck (CSIB)

**Input:**  $\mathbb{P}(X^e, Y^e)$ ,  $e \in \mathcal{E}_{tr}$ ,  $rth > 0$ ,  $c \geq \dim(Z_{inv})$ ,  $M \gg 0$ , and  $(x, y)$  is an example randomly drawn from  $\mathbb{P}(X^e, Y^e)$ .

**Output:** classifier  $w \in \mathbb{R}^{c+1}$ , feature extractor  $\Phi = \mathbb{R}^{c \times d}$ .

**Begin:**

```

1:  $Lv \leftarrow []$ ;  $Lr \leftarrow []$ ;  $\Phi' \leftarrow \mathbb{I}^{d \times d}$ 
2:  $d' \leftarrow \dim(X^e)$ 
3: Apply IB-ERM method (Equation (12)) to  $\mathbb{P}(X^e, Y^e)$  and obtain  $w^* \in \mathbb{R}^{c+1}$  and  $\Phi^* \in \mathbb{R}^{c \times d'}$ 
4: Apply SVD to  $\Phi^*$  as  $\Phi^* = U\Lambda V^T = [U_1, U_2][\Lambda_1, \mathbf{0}; \mathbf{0}, \mathbf{0}][V_1^T; V_2^T]$ 
5:  $r \leftarrow \text{rank}(\Phi^*)$ 
6:  $z_{1:r}^1 \leftarrow [-M, \dots, -M]$ ;  $z_{r+1:d'}^1 \leftarrow V_2^T \Phi' x$ 
7:  $z_{1:r}^2 \leftarrow [M, \dots, M]$ ;  $z_{r+1:d'}^2 \leftarrow V_2^T \Phi' x$ 
8:  $x^1 \leftarrow Vz^1$ ;  $x^2 \leftarrow Vz^2$ 
9: if  $Lv$  is not empty then
10:    $z_{old} \leftarrow []$ ;  $i \leftarrow 0$ ;  $x' \leftarrow x$ 
11:   while  $i < \text{len}(Lv)$  do
12:      $z \leftarrow Lv[i]x'$ 
13:      $z_{old}.\text{append}(z)$ 
14:      $x' \leftarrow z_{Lr[i]}$ 
15:      $i \leftarrow i + 1$ 
16:   end while
17:    $i \leftarrow 0$ 
18:   while  $i < \text{len}(Lv)$  do
19:      $j \leftarrow \text{len}(Lv) - i$ 
20:      $z^1 \leftarrow z_{old}[j]$ ;  $z^2 \leftarrow z_{old}[j]$ 
21:      $z_{Lr[j]}^1 \leftarrow x^1$ ;  $z_{Lr[j]}^2 \leftarrow x^2$ 
22:      $x^1 \leftarrow Lv[j]^T z^1$ ;  $x^2 \leftarrow Lv[j]^T z^2$ 
23:      $i \leftarrow i + 1$ 
24:   end while
25: end if
26: if  $\text{label}(x^1) = \text{label}(x^2)$  then
27:    $Lr.\text{append}(r)$ ;  $Lv.\text{append}(V^T)$ 
28:    $X^e \leftarrow V_2^T X^e$ ;  $\Phi' \leftarrow V_2^T \Phi'$ 
29:   Goto Step 2
30: end if
31:  $w \leftarrow w^*$ ;  $\Phi \leftarrow \Phi^*$ 

```

**End**

**Remark 1.** CSIB succeeds to solve OOD generalization without assuming the support overlap to invariant features and could apply to multiple spurious modes where IB-IRM (as well as ERM, IRM, and IB-ERM) may fail. By introducing counterfactual inference and further supervision (usually conducted by a human) with several steps, CSIB works even when accessing data from a single environment, which is significant especially in the cases where multiple environments' data are not available.

## 6. Experiments

### 6.1. Toy Experiments on Synthetic Datasets

We perform experiments on three synthetic datasets from different spurious correlations modes to verify our method—counterfactual, supervision-based, and information bottleneck (CSIB)—and compare them to ERM, IB-ERM, IRM, and IB-IRM. We follow the same protocol for tuning hyperparameters from [8,12,17] and report the classification error for all experiments. In the following, we first briefly describe the designed datasets and then report the main results. More experimental details can be found in the Appendix.

6.1.1. Datasets

**Example 1/1S.** The example is a modified one from the linear unit tests introduced in [17], which generalizes the cow/camel classification task with relevant backgrounds.

$$\begin{aligned} \theta_{cow} &= \mathbf{1}_m, & \theta_{camel} &= -\theta_{cow}, & v_{animal} &= 10^{-2} \\ \theta_{grass} &= \mathbf{1}_o, & \theta_{sand} &= -\theta_{grass}, & v_{background} &= 1. \end{aligned}$$

The dataset  $D_e$  of each environment  $e \in \mathcal{E}_{tr}$  is sampled from the following distribution

$$\begin{aligned} U^e &\sim \text{Categorical}(p^e s^e, (1-p^e)s^e, p^e(1-s^e), (1-p^e)(1-s^e)), \\ Z_{inv}^e &\sim \begin{cases} (\mathcal{N}_m(0, 0.1) + \theta_{cow})v_{animal} & \text{if } U^e \in \{1, 2\}, \\ (\mathcal{N}_m(0, 0.1) + \theta_{camel})v_{animal} & \text{if } U^e \in \{3, 4\}, \end{cases} \\ Z_{spu}^e &\sim \begin{cases} (\mathcal{N}_o(0, 0.1) + \theta_{grass})v_{background} & \text{if } U^e \in \{1, 4\}, \\ (\mathcal{N}_o(0, 0.1) + \theta_{sand})v_{background} & \text{if } U^e \in \{2, 3\}, \end{cases} \\ Z^e &\leftarrow (Z_{inv}^e, Z_{spu}^e), \quad X^e \leftarrow S(Z^e), \quad N \sim \text{Bernoulli}(q), \quad q < 0.5, \quad Y^e \leftarrow \mathbf{1}(\mathbf{1}_m^T Z_{inv}^e) \oplus N \end{aligned}$$

We set  $s^{e_0} = 0.5, s^{e_1} = 0.7$ , and  $s^{e_2} = 0.3$  for the first three environments, and  $s^{e_j} \sim \text{Uniform}(0.3, 0.7)$  for  $j > 3$ . The scrambling matrix  $S$  is an identical matrix in Example 1 and a random unitary matrix in Example 1S. Here, we set  $p^e = 1$  and  $q = 0$  for all environments to make the spurious features and the invariant features both linearly separable to confuse each other. The experiments on different values of  $q$  and  $p^e$  are presented in the Appendix, where we have found very interesting observations related to the inductive bias of neural networks.

**Example 2/2S.** This example is extended from Example 1 to show one of the failure modes of IB-IRM (as well as ERM, IRM, and IB-ERM) and how our method can be improved by intervention (counterfactual supervision). Given  $w^e \in \mathbb{R}$ , each instance in the environment data  $D^e$  is sampled by

$$\begin{aligned} \theta_{spu} &= 5 \cdot \mathbf{1}_o, & \theta_w &= w^e \cdot \mathbf{1}_m, & v_{spu} &= 10^{-2}, & v_w &= 1, & p, q &\sim \text{Bernoulli}(0.5), \\ Z_{spu}^e &= \mathcal{N}_o(0, 1)v_{spu} + (2p - 1) \cdot \theta_{spu}, & W^e &= \mathcal{N}_m(0, 1)v_w + (2q - 1) \cdot \theta_w \\ Z_{inv}^e &= AZ_{spu}^e + W^e, & Z^e &\leftarrow (Z_{inv}^e, Z_{spu}^e), & X^e &\leftarrow S(Z^e), & Y^e &= \mathbf{1}(\mathbf{1}_m^T Z_{inv}^e), \end{aligned}$$

where we set  $m = o = 5$ , and  $A \in \mathbb{R}^{m \times o}$  is the identical matrix in our experiments. We set  $w^{e_0} = 3, w^{e_1} = 2, w^{e_2} = 1$ , and  $w^{e_j} = \text{Uniform}(0, 3)$  if  $j > 3$  for different training environments. This example shows clearer smaller entropy of spurious features than that of invariant features, which is opposite Example 1/1S.

**Example 3/3S.** This example extends from Example 2 and is similar to the construction of Example 2/2S. Let  $w^e \sim \text{Uniform}(0, 1)$  for different training environments. Each instance in the environments  $e$  is sampled by

$$\begin{aligned} \theta_{inv} &= \cdot 10 \cdot \mathbf{1}_m, & v_{inv} &= 10, & v_{spu} &= 1, & p, q &\sim \text{Bernoulli}(0.5), \\ Z_{inv}^e &= \mathcal{N}_m(0, 1)v_{inv} + (2p - 1) \cdot \theta_{inv}, & Y^e &= \mathbf{1}(\mathbf{1}_m^T Z_{inv}^e), \\ Z_{spu}^e &= 2(Y^e - 1) \cdot v_{spu} + (2q - 1) \cdot w^e \cdot \mathbf{1}_o, & Z^e &\leftarrow (Z_{inv}^e, Z_{spu}^e), & X^e &\leftarrow S(Z^e), \end{aligned}$$

where we set  $m = o = 5$  in our experiments. The spurious features have smaller entropy than the invariant features in this example, which is similar to Example 2/2S, but the invariant features significantly enjoy much larger margin than the spurious features, which is very different from the above two examples. We show a summary of the properties of these three datasets in Table 1 for a general view.

**Table 1.** Summary of three synthetic datasets. Note that for linearly separable features, their margin levels significantly influence the final learning classifier due to the implicit bias of the gradient descent [18]. Such bias would push the standard learning (such as cross-entropy loss) to focus more on the large-margin features. The margin with respect to a dataset (or features)  $\mathcal{Z}$  (each instance has a label 0 or 1) is the minimum distance between a point in  $\mathcal{Z}$  and the max-margin hyperplane, which separates  $\mathcal{Z}$  by its labels.

Datasets	Margin Relationship	Entropy Relationship	Dim <sub>inv</sub>	Dim <sub>spu</sub>
Example 1/1S	Margin <sub>inv</sub> $\ll$ Margin <sub>spu</sub>	Entropy <sub>inv</sub> $<$ Entropy <sub>spu</sub>	5	5
Example 2/2S	Margin <sub>inv</sub> $\approx$ Margin <sub>spu</sub>	Entropy <sub>inv</sub> $>$ Entropy <sub>spu</sub>	5	5
Example 3/3S	Margin <sub>inv</sub> $\gg$ Margin <sub>spu</sub>	Entropy <sub>inv</sub> $>$ Entropy <sub>spu</sub>	5	5

### 6.1.2. Summary of Results

Table 2 shows the classification errors of different methods when training data comes from single, three, and six environments. We can see that ERM and IRM fail to recognize the invariant features in the experiment of Example 1/1S, where invariant features have smaller margin than spurious features do, while information bottleneck-based methods (IB-ERM, IB-IRM, and CSIB) show improved results due to the smaller entropy of the invariant features. Our method CSIB shows results consistent with IB-IRM in Example 1/1S when invariant features are extracted in the first run, which verifies the effectiveness of using the information bottleneck for OOD generalization. In another dataset of Example 2/2S, where the invariant features have larger entropy than spurious features do, we can see that only CSIB can remove the spurious features compared with the other method, although the information bottleneck-based method IB-ERM would degrade the performance of ERM by focusing more on the spurious features. In the third experiment of Example 3/3S, we can see that although ERM shows not-bad results due to the significantly larger margin of invariant features, our CSIB method still shows improvements by removing more spurious features. Notably, comparing the IB-ERM and IB-IRM when only spurious features are extracted (Example 2/2S, Example 3/3S), our CSIB method could effectively remove them by counterfactual supervision and then refocus on the invariant features. Note that the reason of non-zero average error and the fluctuant results of CSIB in some experiments is that the entropy minimization in the training process is less accurate, where entropy is substituted by variance for the ease of the optimization. Nevertheless, there always exists a case where the entropy is indeed truly minimized, and the error reaches zero (see (min) in the table) in Example 2/2S and Example 3/3S. In summary, CSIB consistently performs better in different spurious correlations modes and is especially more effective than IB-ERM and IB-IRM when the spurious features enjoy much smaller entropy than the invariant features do.

### 6.2. Experiments on Color MNIST Dataset

In this experiment, we set up a binary classification task for digit recognition and identify whether the digit is less than five or more than five. We use real-world dataset, the MNIST database of handwritten digits (<http://yann.lecun.com/exdb/mnist/>), for the construction. Following our learning setting, we use color information as the spurious features that correlates strongly with the class label. By construction, the label is more strongly correlated with the color than with the digit in the training environments, but this correlation is broken in the test environment. Specifically, the three designed environments (two training environments and one test environment containing 10,000 points each) of the color MNIST are as follows: first, we define a preliminary binary label  $\hat{y}$  to the image base on the digit:  $\hat{y} = 0$  for digits 0–4 and  $\hat{y} = 1$  for 5–9. Second, we obtain the final label  $y$  by flipping  $\hat{y}$  with probability 0.25. Then, we flip the final labels to obtain the color id, where the flipping probabilities with respect to two training environments and one test environment are 0.2 and 0.1, and 0.9. For better understanding, we randomly draw 20 examples for each label from each environment and visualize them in Figure 4.

**Table 2.** Main results: #Envs means the number of training environments, and (min) reports the minimal test classification error across different running seeds.

	#Envs	ERM (min)	IRM (min)	IB-ERM (min)	IB-IRM (min)	CSIB (min)
Example 1	1	0.50 ± 0.01 (0.49)	0.50 ± 0.01 (0.49)	<b>0.23 ± 0.02 (0.22)</b>	0.31 ± 0.10 (0.25)	<b>0.23 ± 0.02 (0.22)</b>
Example 1S	1	0.50 ± 0.00 (0.49)	0.50 ± 0.00 (0.50)	0.46 ± 0.04 (0.39)	<b>0.30 ± 0.10 (0.25)</b>	0.46 ± 0.04 (0.39)
Example 2	1	0.40 ± 0.20 (0.00)	0.50 ± 0.00 (0.49)	0.50 ± 0.00 (0.49)	0.46 ± 0.02 (0.45)	<b>0.00 ± 0.00 (0.00)</b>
Example 2S	1	0.50 ± 0.00 (0.50)	0.31 ± 0.23 (0.00)	0.50 ± 0.00 (0.50)	0.45 ± 0.01 (0.43)	<b>0.10 ± 0.20 (0.00)</b>
Example 3	1	0.16 ± 0.06 (0.09)	0.18 ± 0.03 (0.14)	0.50 ± 0.01 (0.49)	0.40 ± 0.20 (0.01)	<b>0.11 ± 0.20 (0.00)</b>
Example 3S	1	0.17 ± 0.07 (0.10)	<b>0.09 ± 0.02 (0.07)</b>	0.50 ± 0.00 (0.50)	0.50 ± 0.00 (0.50)	0.21 ± 0.24 (0.00)
Example 1	3	0.45 ± 0.01 (0.45)	0.45 ± 0.01 (0.45)	<b>0.22 ± 0.01 (0.21)</b>	0.23 ± 0.13 (0.02)	<b>0.22 ± 0.01 (0.21)</b>
Example 1S	3	0.45 ± 0.00 (0.45)	0.45 ± 0.00 (0.45)	0.41 ± 0.04 (0.34)	<b>0.27 ± 0.11 (0.11)</b>	0.41 ± 0.04 (0.34)
Example 2	3	0.40 ± 0.20 (0.00)	0.50 ± 0.00 (0.50)	0.50 ± 0.00 (0.50)	0.33 ± 0.04 (0.25)	<b>0.00 ± 0.00 (0.00)</b>
Example 2S	3	0.50 ± 0.00 (0.50)	0.37 ± 0.15 (0.15)	0.50 ± 0.00 (0.50)	0.34 ± 0.01 (0.33)	<b>0.10 ± 0.20 (0.00)</b>
Example 3	3	0.18 ± 0.04 (0.15)	0.21 ± 0.02 (0.20)	0.50 ± 0.01 (0.49)	0.50 ± 0.01 (0.49)	<b>0.11 ± 0.20 (0.00)</b>
Example 3S	3	0.18 ± 0.04 (0.15)	0.08 ± 0.03 (0.03)	0.50 ± 0.00 (0.50)	0.43 ± 0.09 (0.31)	<b>0.01 ± 0.00 (0.00)</b>
Example 1	6	0.46 ± 0.01 (0.44)	0.46 ± 0.09 (0.41)	<b>0.22 ± 0.01 (0.20)</b>	0.37 ± 0.14 (0.17)	<b>0.22 ± 0.01 (0.20)</b>
Example 1S	6	0.46 ± 0.02 (0.44)	0.46 ± 0.02 (0.44)	<b>0.35 ± 0.10 (0.23)</b>	0.42 ± 0.12 (0.28)	<b>0.35 ± 0.10 (0.23)</b>
Example 2	6	0.49 ± 0.01 (0.48)	0.50 ± 0.01 (0.48)	0.50 ± 0.00 (0.50)	0.30 ± 0.01 (0.28)	<b>0.00 ± 0.00 (0.00)</b>
Example 2S	6	0.50 ± 0.00 (0.50)	0.35 ± 0.12 (0.25)	0.50 ± 0.00 (0.50)	0.30 ± 0.01 (0.29)	<b>0.20 ± 0.24 (0.00)</b>
Example 3	6	0.18 ± 0.04 (0.15)	0.20 ± 0.01 (0.19)	0.50 ± 0.00 (0.49)	0.37 ± 0.16 (0.16)	<b>0.01 ± 0.01 (0.00)</b>
Example 3S	6	0.18 ± 0.04 (0.14)	<b>0.05 ± 0.04 (0.01)</b>	0.50 ± 0.00 (0.50)	0.50 ± 0.00 (0.50)	0.11 ± 0.20 (0.00)

**Figure 4.** Visualization of the color mnist dataset.

The classification results on the color MNIST dataset are shown in Table 3. From the results, we can see that both ERM and IB-ERM methods almost surely use the color features to achieve the task. Although IRM and IB-IRM methods have shown some improvements over ERM, only our method can perform better than a random prediction, which demonstrates the effectiveness of CSIB.

**Table 3.** Classification accuracy (%) on color MNIST dataset. “Oracle” in the table means that the training and test data are in the same environment.

Methods	ERM	IRM	IB-ERM	IB-IRM	CSIB	Oracle
Accuracy	9.94 ± 0.28	20.39 ± 2.76	9.94 ± 0.28	43.84 ± 12.48	<b>60.03 ± 1.28</b>	84.72 ± 0.65

## 7. Related Works

We divide the works related to OOD generalization into two categories: theory and methods, though some of them belong to both.

### 7.1. Theory of OOD Generalization

Based on different definitions to the distributional changes, we review the corresponding theory by the following three categories.

**Based on causality.** Due to the close connection between the distributional changes and the interventions discussed in the theory of causality [13,14], the problem of OOD generalization is usually built in the framework of causal learning. The theory states that a response  $Y$  is directly caused only by its parents variables  $X_{Pa(Y)}$ , and all interventions other than those on  $Y$  do not change the conditional distribution of  $\mathbb{P}(Y|X_{Pa(Y)})$ . Such theory inspires a popular learning principle—the invariance principle—that aims to discover a set of variables such that they remain invariant to the response  $Y$  in all observed environments [15,19,20]. Invariant risk minimization (IRM) [8] is then proposed to learn a feature extractor  $\Phi$  in an end-to-end way such that the optimal classifier based on the extracted features  $\Phi(X)$  remains unchanged in each environment. The theory in [8] shows the guarantee of IRM for OOD generalization under some general assumptions but only focuses on the linear regression tasks. Different from the failure analyses of IRM for the classification tasks in [21,22], where the response  $Y$  is the cause of the spurious feature, Ahuja et al. [12] analyse another scenario when the invariant feature is the cause of the spurious feature and show that in this case, linear classification is more difficult than linear regression, where the invariance principle itself is insufficient to ensure the success of OOD generalization. They also claim that the assumption of support overlap of invariant features is necessarily needed. They then propose a learning principle of information bottleneck-based invariant risk minimization (IB-IRM) for linear classification, which shows how to address the failures of IRM by adding information bottleneck [16] into the learning. In this work, we closely investigate the conditions identified in [12] and first show that support overlap of invariant features is not necessarily needed for the success of OOD generalization. We further show several failure cases of IB-IRM and propose improved results for it.

Recently, some works tackle the challenge of OOD generalization in the nonlinear regime [23,24]. Commonly, both of them use variational autoencoder (VAE)-based models [25,26] to identify the latent variables from observations in the first stage. Then, these inferring latent variables are separated into two distinct parts of invariant (causal) and spurious (non-causal) features based on different assumptions. Specifically, Lu et al. [23,27] assume that the latent variables conditioned on some accessible side information such as the environment index or class label follow the exponential family distributions, and Liu et al. [24] directly disentangle the latent variables to two different parts during the inferring stage and assume that the marginal distributions of them are independent of each other. These assumptions, however, are rather strong in general. Nevertheless, these solutions aim to capture the latent variables such that the response given these variables is invariant for different environments, which could still fail because the invariance principle itself is insufficient for OOD generalization in the classification tasks, as shown in [12]. In this work, we focus on the linear classification only and show a new theory of a new method that addresses several OOD generalization failures in the linear settings. Our method could extend to the nonlinear regime by combining with the disentangled representation learning [28] or causal representation learning [29]. Specifically, once the latent representations are well disentangled, i.e., the latent features are represented by a linear transform of the causal features and spurious features, we then could apply our method to filter out the spurious features in the latent space such that only causal features remain.

**Based on robustness.** Different from those based on the causality, where different distributions are generated by intervention on a same SEM and the goal is to discover causal features, the robustness-based methods aim to protect the model against the potential distributional shifts within the uncertainty set, which is usually constrained by  $f$ -divergence [30] or Wasserstein distance [31]. This series of works is theoretically addressed by distributionally robust optimization (DRO) under a minimax framework [32,33]. Recently, some works tend to discover the connections between causality and robustness [34]. Although these

works show less relevance to us, it is possible that a well-defined measure of distribution divergence could help to effectively extract causal features under the robustness framework. This would be an interesting avenue for future research.

**Others.** Some other works assume that the distributions (domains) are generated from a hyper-distribution and aim to minimize the average risk estimation error bound [35–37]. These works are often built based on the generalization theory under the independent and identically distributed (IID) assumption. The authors in [38] do not make any assumption on the distributional changes and only study the learnability of OOD generalization in a general way. All of these theories do not cover the OOD generalization problem under a single training environment or domain.

## 7.2. Methods of OOD Generalization

**Based on the invariance principle.** Inspired from the invariance principle [15,19], many methods are proposed by designing various loss to extract features to better satisfy the principle itself. IRMv1 [8] is the first objective to address this in an end-to-end way by adding a gradient penalty to the classifier. Following this work, Krueger et al. [9] suggest penalizing the variance of the risks, while Xie et al. [39] give the same objective but take the square root of the variance, and many other alternatives can also be found [40–42]. It is clear that all of these methods aim to find an invariant predictor. Recently, Ahuja et al. [12] found that for the classification problem, finding the invariant predictor is not enough to extract causal features since the features could include spurious information to make the predictor invariant across training environments, and they propose IB-IRM to address such a failure. Similar ideas to IB-IRM can also be found in the work [43,44], where different loss functions are proposed to achieve the same purpose. Specifically, Alesiani et al. [44] also use the information bottleneck (IB) for the help in dropping spurious correlations, but their analyses only focus on the scenario when spurious features are independent from the causal features, which could be considered as a special case of ours. More recently, Wang et al. [45] propose similar ideas to ours but only tackle the situation when the invariant features have the same distribution among all environments. In this work, we further show that IB-IRM could still fail in two cases due to the model only relying on spurious features to meet the task of interest. We then propose a counterfactual supervision-based information bottleneck (CSIB) method to address such failures and show improving results to prior works.

**Based on distribution matching.** It is worth noting that there are many works focused on learning domain invariant features representations [46–48]. Most of these works are inspired by the seminal theory of domain adaptation [49,50]. The goal of these methods is to learn a feature extractor  $\Phi$  such that the marginal distribution of  $\mathbb{P}(\Phi(X))$  or the conditional distribution of  $\mathbb{P}(\Phi(X)|Y)$  is invariant across different domains. This is different from the invariance principle, where the goal is to make  $\mathbb{P}(Y|\Phi(X))$  (or  $\mathbb{E}(Y|\Phi(X))$ ) invariant. We refer readers to the papers of [8,51] for better understanding the details of why these distribution-matching-based methods often fail to address OOD generalization.

**Others.** Other related methods are varied, including by using data augmentation in both image level [52] or feature level [53], by removing spurious correlations through stable learning [54], and by utilizing the inductive bias of neural networks [3,55], etc. Most of these methods are empirically inspired from experiments and are verified on some specific datasets. Recently, empirical studies in [56,57] notice that the real effects of many OOD generalization (domain generalization) methods are weak, which indicates that the benchmark-based evaluation criteria may be inadequate to validate the OOD generalization algorithms.

## 8. Conclusions, Limitations and Future Work

In this paper, we focus on the OOD generalization problem of linear classification. We first revisit the fundamental assumptions and results of prior works and show that the condition of invariant features supporting overlap is not necessarily needed for the success of OOD generalization and thus propose a weaker counterpart. Then, we show two failure



cases of IB-IRM (as well as ERM, IB-ERM, and IRM) and illustrate its intrinsic causes by theoretical analysis. We further propose a new method—counterfactual supervision-based information bottleneck (CSIB)—and theoretically prove its effectiveness under some weaker assumptions. CSIB works even when accessing data from a single environment and can easily extend to the multi-class problems. Finally, we design several synthetic datasets with our examples for experimental verification. Empirical observations among all comparing methods illustrate the effectiveness of the CSIB.

Since we only take the linear problem into account, including linear representation and linear classifier, any nonlinear case would not be guaranteed by our theoretical results, and thus CSIB may fail. Therefore, the same as prior works (IRM [8] and IB-IRM [12]), the nonlinear challenge is still an unsolved problem [21,22]. We believe this is of great value for investigating in future work since widely used data in the wild are nonlinearly generated. Another fruitful direction is to design a powerful algorithm for entropy minimization during the learning process of CSIB. Currently, we use the variance of features to replace the entropy of the features during optimization. However, variance and entropy are essentially different. A truly effective entropy minimization is the key to the success of CSIB. Another limitation of our method is that we have to require further supervision to the counterfactual examples during the learning process, although it only takes one time for a single step.

**Author Contributions:** Conceptualization, B.D. and K.J.; methodology, B.D.; software, B.D.; validation, B.D. and K.J.; formal analysis, B.D.; investigation, B.D.; resources, B.D.; data curation, B.D.; writing—original draft preparation, B.D.; writing—review and editing, B.D. and K.J.; visualization, B.D.; supervision, K.J. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research (including the APC) was funded by Guangdong R&D key project of China (No.: 2019B010155001).

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Data is contained within the article or supplementary material.

**Conflicts of Interest:** The authors declare no conflict of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

OOD	Out-of-distribution
ERM	Empirical risk minimization
IRM	Invariant risk minimization
IB-ERM	Information bottleneck-based empirical risk minimization
IB-IRM	Information bottleneck-based invariant risk minimization
CSIB	Counterfactual supervision-based information bottleneck
DAG	Directed acyclic graph
SEM	Structure equation model
SVD	Singular value decomposition

## Appendix A. Experiments Details

In this section, we provide more details on the experiments. The code to reproduce the experiments can be found at <https://github.com/szubing/CSIB>.

### Appendix A.1. Optimization Loss of IB-ERM

The objective function of IB-ERM is as follows:

$$\min_{w, \Phi} \sum_{e \in \mathcal{E}_{tr}} h^e(\Phi) \quad \text{s.t.} \quad \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} R^e(w \circ \Phi) \leq rth. \quad (\text{A1})$$

Since the entropy of  $h^e(\Phi) = H(\Phi(X^e))$  is hard to estimate by a differential variable that can be optimized by using gradient descent, we follow [12] by using the variance instead of the entropy for optimization. The total loss function is given by

$$\text{loss}(w, \Phi) = \sum_{e \in \mathcal{E}_{tr}} (R^e(w \circ \Phi) + \lambda \text{Var}(\Phi)) \quad (\text{A2})$$

with a hyperparameter  $\lambda$  onto it.

#### Appendix A.2. Experiments Setup

**Model, hyperparameters, loss, and evaluation.** In all experiments, we follow the same protocol as prescribed by [12,17] for the model / hyperparameter selection, training, and evaluation. Except those specified, for all experiments across three examples and five comparing methods, the model is the same with a linear feature extractor  $\Phi \in \mathbb{R}^{d \times d}$  followed by a linear classifier  $w \in \mathbb{R}^{d+1}$ . We use binary cross-entropy loss for classification. All hyperparameters, including the learning rate, the penalty term in IRM, or the  $\lambda$  associated with the  $\text{Var}(\Phi)$  in Equation (A2), etc., are randomly searched and selected by using 20 test samples for validation. The results reported in the main manuscript use three hyperparameter queries of each and average over five data seeds. The results when searching over more hyperparameter values are reported in the supplementary experiments. The search spaces of all the hyperparameters follow the same as in [12,17]. The classification test errors between 0 and 1 are reported.

**Compute description.** Our computing resource is one GPU of NVIDIA GeForce GTX 1080 Ti with 6 CPU cores of Intel(R) Core(TM) i7-8700 CPU @ 3.20GHz.

**Existing codes and datasets used.** In our experiments, we mainly rely on the following two github repositories: InvarianceUnitTests (<https://github.com/facebookresearch/InvarianceUnitTests>) and IB-IRM (<https://github.com/ahujak/IB-IRM>).

#### Appendix A.3. Supplementary Experiments

The purpose of the first supplementary experiment is to illustrate what the result would be when we increase the number of running seeds in the hyperparameters selection. These results are shown in Table A1, where we increase the number of hyperparameter queries to 10 of each. It is clear that overall, the results of the CSIB in Table A1 are much better and have less fluctuations than those in Table 2, and the conclusions remain almost the same as we have summarized in Section 6.1.2. This further verifies the effectiveness of the CSIB method.

**Observation on different settings in Example 1/1S.** In our main experiments of Example 1/1S, we set  $p^e = 1$  and  $q = 0$  to make the spurious features and the invariant features both linearly separable to confuse each other. Here, we analyse what the result would be if we vary their values. Following [17], we set  $p^{e_0} = 0.95$ ,  $p^{e_1} = 0.97$ ,  $p^{e_2} = 0.99$ , and  $p^{e_j} \sim \text{Uniform}(0.9, 1)$  to make spurious features linearly inseparable, and  $q$  is set to 0/0.05 to make invariant features linearly separable/inseparable. Table A2 shows the corresponding results. Interestingly, we find that all methods except for IB-IRM have an ideal error rate (the same as the Oracle) when the spurious features are linearly inseparable ( $p^e \neq 1$ ), even when the invariant features are linearly inseparable too ( $q = 0.05$ ). Why would this happen? We then remove the linear embedding  $\Phi$ . The results are presented in Table A3. Comparing the results between Tables A2 and A3, we found there is a significant inductive bias of the neural network, though the model is linear. Further analysis to such observation is out of the scope of this paper, but this would be an interesting avenue for future research.

**Table A1.** Supplementary results when using 10 hyperparameter queries. #Envs means the number of training environments, and (min) reports the minimal test classification error across different running data seeds.

	#Envs	ERM (min)	IRM (min)	IB-ERM (min)	IB-IRM (min)	CSIB (min)	Oracle (min)
Example 1	1	0.50 ± 0.01 (0.49)	0.50 ± 0.01 (0.49)	0.23 ± 0.02 (0.22)	0.31 ± 0.10 (0.25)	0.23 ± 0.02 (0.22)	0.00 ± 0.00 (0.00)
Example 1S	1	0.50 ± 0.00 (0.49)	0.50 ± 0.00 (0.49)	0.09 ± 0.04 (0.04)	0.30 ± 0.10 (0.25)	0.08 ± 0.04 (0.04)	0.00 ± 0.00 (0.00)
Example 2	1	0.40 ± 0.20 (0.00)	0.00 ± 0.00 (0.00)	0.50 ± 0.00 (0.49)	0.48 ± 0.03 (0.43)	0.00 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 2S	1	0.50 ± 0.00 (0.50)	0.30 ± 0.25 (0.00)	0.50 ± 0.00 (0.50)	0.50 ± 0.01 (0.48)	0.00 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 3	1	0.16 ± 0.06 (0.09)	0.03 ± 0.00 (0.03)	0.50 ± 0.01 (0.49)	0.41 ± 0.09 (0.25)	0.02 ± 0.01 (0.00)	0.00 ± 0.00 (0.00)
Example 3S	1	0.16 ± 0.06 (0.10)	0.04 ± 0.01 (0.02)	0.50 ± 0.00 (0.50)	0.41 ± 0.12 (0.26)	0.01 ± 0.01 (0.00)	0.00 ± 0.00 (0.00)
Example 1	3	0.44 ± 0.01 (0.44)	0.44 ± 0.01 (0.44)	0.21 ± 0.00 (0.21)	0.21 ± 0.10 (0.06)	0.21 ± 0.00 (0.21)	0.00 ± 0.00 (0.00)
Example 1S	3	0.45 ± 0.00 (0.44)	0.45 ± 0.00 (0.44)	0.09 ± 0.03 (0.05)	0.23 ± 0.13 (0.01)	0.09 ± 0.03 (0.05)	0.00 ± 0.00 (0.00)
Example 2	3	0.13 ± 0.07 (0.00)	0.00 ± 0.00 (0.00)	0.50 ± 0.00 (0.50)	0.33 ± 0.04 (0.25)	0.00 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 2S	3	0.50 ± 0.00 (0.50)	0.14 ± 0.20 (0.00)	0.50 ± 0.00 (0.50)	0.34 ± 0.01 (0.33)	0.00 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 3	3	0.17 ± 0.04 (0.14)	0.02 ± 0.00 (0.02)	0.50 ± 0.01 (0.49)	0.43 ± 0.08 (0.29)	0.01 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 3S	3	0.17 ± 0.04 (0.13)	0.02 ± 0.00 (0.02)	0.50 ± 0.00 (0.50)	0.36 ± 0.18 (0.07)	0.01 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 1	6	0.46 ± 0.01 (0.44)	0.46 ± 0.09 (0.41)	0.22 ± 0.01 (0.21)	0.41 ± 0.11 (0.26)	0.22 ± 0.01 (0.21)	0.00 ± 0.00 (0.00)
Example 1S	6	0.46 ± 0.02 (0.44)	0.46 ± 0.02 (0.44)	0.06 ± 0.04 (0.02)	0.45 ± 0.07 (0.41)	0.06 ± 0.04 (0.02)	0.00 ± 0.00 (0.00)
Example 2	6	0.21 ± 0.03 (0.17)	0.00 ± 0.00 (0.00)	0.50 ± 0.00 (0.50)	0.36 ± 0.03 (0.31)	0.00 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 2S	6	0.50 ± 0.00 (0.50)	0.10 ± 0.20 (0.00)	0.50 ± 0.00 (0.50)	0.19 ± 0.16 (0.01)	0.00 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 3	6	0.17 ± 0.03 (0.14)	0.02 ± 0.00 (0.02)	0.50 ± 0.00 (0.49)	0.37 ± 0.16 (0.16)	0.01 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)
Example 3S	6	0.17 ± 0.03 (0.14)	0.02 ± 0.00 (0.02)	0.50 ± 0.00 (0.50)	0.46 ± 0.09 (0.28)	0.01 ± 0.00 (0.00)	0.00 ± 0.00 (0.00)

**Observation on linearly separable properties of high-dimensional data.** Here, we empirically show that for  $o$ -dimensional data, we have high probability that  $o$  randomly drawn points are linearly separable for any two subsets. To verify that, we design a random experiment as follows: (1) Let  $o \in [100, 10,000]$ , and we randomly draw  $o$  points from  $[-1, 1]^o$ , and give random labels to these  $o$  points of 0 or 1. (2) We train a linear classifier to fit these  $o$  points and report the final training error. (3) We perform (1) and (2) 100 times for different seeds. Our results show that for 100 runs, all training errors reach 0 for every  $o$ , which proves our conjecture.

**Table A2.** Results in Example 1/1S, where the learning model is a linear embedding  $\Phi \in \mathbb{R}^{d \times d}$  followed by a linear classifier  $w \in \mathbb{R}^{d+1}$ .

	#Envs	$p^e = 1?$	$q$	ERM	IB-ERM	IB-IRM	CSIB	IRM	Oracle
Example 1	1	Yes	0	0.50 ± 0.01	0.23 ± 0.02	0.31 ± 0.10	0.23 ± 0.02	0.50 ± 0.01	0.00 ± 0.00
Example 1S	1	Yes	0	0.50 ± 0.00	0.46 ± 0.04	0.30 ± 0.10	0.46 ± 0.04	0.50 ± 0.00	0.00 ± 0.00
Example 1	3	Yes	0	0.45 ± 0.01	0.22 ± 0.01	0.23 ± 0.13	0.22 ± 0.01	0.45 ± 0.01	0.00 ± 0.00
Example 1S	3	Yes	0	0.45 ± 0.00	0.41 ± 0.04	0.27 ± 0.11	0.41 ± 0.04	0.45 ± 0.00	0.00 ± 0.00
Example 1	6	Yes	0	0.46 ± 0.01	0.22 ± 0.01	0.37 ± 0.14	0.22 ± 0.01	0.46 ± 0.09	0.00 ± 0.00
Example 1S	6	Yes	0	0.46 ± 0.02	0.35 ± 0.10	0.42 ± 0.12	0.35 ± 0.10	0.46 ± 0.02	0.00 ± 0.00
Example 1	1	No	0	0.00 ± 0.00	0.00 ± 0.00	0.15 ± 0.20	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Example 1S	1	No	0	0.00 ± 0.00	0.00 ± 0.00	0.12 ± 0.19	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Example 1	3	No	0	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Example 1S	3	No	0	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.01	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Example 1	6	No	0	0.00 ± 0.00	0.00 ± 0.00	0.30 ± 0.20	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Example 1S	6	No	0	0.00 ± 0.00	0.00 ± 0.00	0.31 ± 0.20	0.00 ± 0.00	0.04 ± 0.06	0.00 ± 0.00
Example 1	1	No	0.05	0.05 ± 0.00	0.05 ± 0.00	0.32 ± 0.22	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00
Example 1S	1	No	0.05	0.05 ± 0.00	0.05 ± 0.00	0.19 ± 0.17	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00
Example 1	3	No	0.05	0.05 ± 0.00	0.05 ± 0.00	0.07 ± 0.03	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00
Example 1S	3	No	0.05	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00
Example 1	6	No	0.05	0.05 ± 0.00	0.05 ± 0.00	0.30 ± 0.21	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00
Example 1S	6	No	0.05	0.05 ± 0.00	0.05 ± 0.00	0.32 ± 0.19	0.05 ± 0.00	0.05 ± 0.00	0.05 ± 0.00

**Table A3.** Results in Example 1/1S, where the learning model is a linear classifier  $w \in \mathbb{R}^{d+1}$  without linear embedding  $\Phi$ . The CSIB must require a feature extractor, so there are not results related to the CSIB.

	#Envs	$p^e = 1?$	$q$	ERM	IB-ERM	IB-IRM	IRM	Oracle
Example 1	1	Yes	0	0.50 ± 0.01	0.25 ± 0.01	0.31 ± 0.10	0.50 ± 0.01	0.00 ± 0.00
Example 1S	1	Yes	0	0.50 ± 0.00	0.49 ± 0.01	0.30 ± 0.10	0.50 ± 0.00	0.00 ± 0.00
Example 1	3	Yes	0	0.44 ± 0.01	0.23 ± 0.01	0.21 ± 0.10	0.44 ± 0.01	0.00 ± 0.00
Example 1S	3	Yes	0	0.45 ± 0.00	0.44 ± 0.01	0.42 ± 0.04	0.45 ± 0.00	0.00 ± 0.00
Example 1	6	Yes	0	0.46 ± 0.01	0.27 ± 0.07	0.41 ± 0.11	0.46 ± 0.01	0.01 ± 0.01
Example 1S	6	Yes	0	0.46 ± 0.02	0.42 ± 0.08	0.46 ± 0.09	0.46 ± 0.02	0.01 ± 0.02
Example 1	1	No	0	0.50 ± 0.01	0.00 ± 0.00	0.15 ± 0.20	0.50 ± 0.01	0.00 ± 0.00
Example 1S	1	No	0	0.50 ± 0.00	0.00 ± 0.00	0.13 ± 0.19	0.50 ± 0.00	0.00 ± 0.00
Example 1	3	No	0	0.45 ± 0.01	0.00 ± 0.00	0.00 ± 0.00	0.45 ± 0.01	0.00 ± 0.00
Example 1S	3	No	0	0.45 ± 0.00	0.01 ± 0.02	0.08 ± 0.14	0.46 ± 0.02	0.00 ± 0.00
Example 1	6	No	0	0.46 ± 0.01	0.10 ± 0.16	0.30 ± 0.20	0.46 ± 0.01	0.01 ± 0.01
Example 1S	6	No	0	0.46 ± 0.01	0.24 ± 0.19	0.41 ± 0.12	0.47 ± 0.03	0.01 ± 0.02
Example 1	1	No	0.05	0.50 ± 0.01	0.05 ± 0.00	0.32 ± 0.22	0.50 ± 0.01	0.05 ± 0.00
Example 1S	1	No	0.05	0.50 ± 0.01	0.05 ± 0.01	0.20 ± 0.17	0.50 ± 0.00	0.05 ± 0.00
Example 1	3	No	0.05	0.45 ± 0.01	0.05 ± 0.00	0.07 ± 0.03	0.47 ± 0.01	0.05 ± 0.00
Example 1S	3	No	0.05	0.45 ± 0.01	0.07 ± 0.03	0.11 ± 0.11	0.46 ± 0.01	0.05 ± 0.00
Example 1	6	No	0.05	0.47 ± 0.01	0.14 ± 0.14	0.30 ± 0.21	0.47 ± 0.01	0.05 ± 0.00
Example 1S	6	No	0.05	0.47 ± 0.01	0.27 ± 0.18	0.42 ± 0.11	0.47 ± 0.01	0.05 ± 0.01

Then, we look back to Theorem 2. For real data, such as an image, the dimension of spurious features  $o$  is often high. Assume different environments enjoy different spurious points randomly; then, from the above observation, there is a high probability that the following events will occur: For any labeling data in the  $n$  training environments with  $n < o/2$  (2 is due to the binary label), models could achieve zero training error by relying on spurious features only. This illustrates why prior methods easily fail to address OOD generalization under Assumption 6.

## Appendix B. Proofs

### Appendix B.1. Preliminary

Before our proofs, we first review some useful properties related to the entropy [12,58].

**Entropy.** For discrete random variable  $X \sim \mathbb{P}_X$  with support  $\mathcal{X}$ , its entropy (Shannon entropy) is defined as

$$H(X) = - \sum_{x \in \mathcal{X}} \mathbb{P}_X(X = x) \log(\mathbb{P}_X(X = x)) \quad (\text{A3})$$

The differential entropy of the continuous random variable  $X \sim \mathbb{P}_X$  with support  $\mathcal{X}$  is given by

$$h(X) = - \int_{x \in \mathcal{X}} p_X(x) \log(p_X(x)) dx, \quad (\text{A4})$$

where  $p_X(x)$  is the probability density function of the distribution  $\mathbb{P}_X$ . Sometimes, we may confuse using  $H(X)$  or  $h(X)$  to represent its entropy no matter whether  $X$  is discrete or continuous.

**Lemma A1.** If  $X$  and  $Y$  are discrete random variables that are independent, then

$$H(X + Y) \geq \max\{H(X), H(Y)\}. \quad (\text{A5})$$

**Proof.** Define  $Z = X + Y$ . Since  $X \perp Y$ , we have

$$\begin{aligned} H(Z|X) &= - \sum_{x \in \mathcal{X}} \mathbb{P}_X(x) \sum_{z \in \mathcal{Z}} \mathbb{P}_{Z|X}(Z = z|X = x) \log(\mathbb{P}_{Z|X}(Z = z|X = x)) \\ &= - \sum_{x \in \mathcal{X}} \mathbb{P}_X(x) \sum_{z \in \mathcal{Z}} \mathbb{P}_{Y|X}(Y = z - x|X = x) \log(\mathbb{P}_{Y|X}(Y = z - x|X = x)) \\ &= - \sum_{x \in \mathcal{X}} \mathbb{P}_X(x) \sum_{z \in \mathcal{Z}} \mathbb{P}_Y(Y = z - x) \log(\mathbb{P}_Y(Y = z - x)) \\ &= - \sum_{x \in \mathcal{X}} \mathbb{P}_X(x) \sum_{y \in \mathcal{Y}} \mathbb{P}_Y(Y = y) \log(\mathbb{P}_Y(Y = y)) \\ &= H(Y), \end{aligned}$$

and similar we have  $H(Z|Y) = H(X)$ . Therefore,

$$H(X + Y) = I(Z, X) + H(Z|X) = I(Z, X) + H(Y) \geq H(Y) \tag{A6}$$

$$H(X + Y) = I(Z, Y) + H(Z|Y) = I(Z, Y) + H(X) \geq H(X). \tag{A7}$$

This completes the proof.  $\square$

**Lemma A2.** If  $X$  and  $Y$  are continuous random variables that are independent, then

$$h(X + Y) \geq \max\{h(X), h(Y)\}. \tag{A8}$$

**Proof.** Define  $Z = X + Y$ . Since  $X \perp Y$ , we have

$$\begin{aligned} h(Z|X) &= - \int_{x \in \mathcal{X}} p_X(x) \int_{z \in \mathcal{Z}} p_{Z|X}(Z = z|X = x) \log(p_{Z|X}(Z = z|X = x)) dx dz \\ &= - \int_{x \in \mathcal{X}} p_X(x) \int_{z \in \mathcal{Z}} p_{Y|X}(Y = z - x|X = x) \log(p_{Y|X}(Y = z - x|X = x)) dx dz \\ &= - \int_{x \in \mathcal{X}} p_X(x) \int_{z \in \mathcal{Z}} p_Y(Y = z - x) \log(p_Y(Y = z - x)) dx dz \\ &= - \int_{x \in \mathcal{X}} p_X(x) dx \int_{y \in \mathcal{Y}} p_Y(Y = y) \log(p_Y(Y = y)) dy \\ &= h(Y), \end{aligned}$$

and similar, we have  $h(Z|Y) = h(X)$ . Therefore,

$$h(X + Y) = I(Z, X) + h(Z|X) = I(Z, X) + h(Y) \geq h(Y) \tag{A9}$$

$$h(X + Y) = I(Z, Y) + h(Z|Y) = I(Z, Y) + h(X) \geq h(X). \tag{A10}$$

This completes the proof.  $\square$

**Lemma A3.** If  $X$  and  $Y$  are discrete random variables that are independent with the supports satisfying  $2 \leq |\mathcal{X}| < \infty, 2 \leq |\mathcal{Y}| < \infty$ , then

$$H(X + Y) > \max\{H(X), H(Y)\}. \tag{A11}$$

**Proof.** From Lemma A1 and due to the symmetry of  $X$  and  $Y$ , we only need to prove  $H(X + Y) \neq H(X)$ . The proof is by contradiction. Suppose  $H(X + Y) = H(X)$ , then from Equation (A7) it follows that  $I(X + Y, Y) = 0$ , thus  $X + Y \perp Y$ . However,  $\mathbb{P}(Y = y_{max}|X + Y = x_{max} + y_{max}) = 1$ , which is different from  $\mathbb{P}(Y = y_{max}) < 1$  (due to  $|Y| \geq 2$ ). This contradicts  $X + Y \perp Y$ .  $\square$

**Lemma A4.** If  $X$  and  $Y$  are continuous random variables that are independent and have a bounded support, then

$$h(X + Y) > \max\{h(X), h(Y)\}. \tag{A12}$$

**Proof.** From Lemma A2 and due to the symmetry of  $X$  and  $Y$ , we only need to prove  $h(X + Y) \neq h(X)$ . The proof is by contradiction. Suppose  $h(X + Y) = h(X)$ , then from Equation (A10) it follows that  $I(X + Y, Y) = 0$ , thus  $X + Y \perp Y$ . For any  $\delta > 0$ , define an event  $\mathcal{M} : x_{max} + y_{max} - \delta \leq X + Y \leq x_{max} + y_{max}$ . If  $\mathcal{M}$  occurs, then  $Y \geq y_{max} - \delta$  and  $X \geq x_{max} - \delta$ . Thus,  $\mathbb{P}_Y(Y \leq y_{max} - \delta | \mathcal{M}) = 0$ . However, we can always choose a  $\delta > 0$  that is small enough to make  $\mathbb{P}_Y(Y \leq y_{max} - \delta) > 0$ . This contradicts  $X + Y \perp Y$ .  $\square$

### Appendix B.2. Proof of Theorem 2

**Proof.** The proof is trivial. Since two sets  $\cup_{e \in \mathcal{E}_{tr}} Z_{spu}^e (Y^e = 1)$  and  $\cup_{e \in \mathcal{E}_{tr}} Z_{spu}^e (Y^e = 0)$  are linearly separable, there exists a linear classifier  $w$  that only relies on spurious features and can achieve zero classification error on each environment. Therefore,  $w$  is an invariant predictor across different training environments. In addition,  $H(Z_{inv}^e) > H(Z_{spu}^e)$  would make IB-IRM prefer to choose these spurious features. Therefore,  $w$  would be an optimal solution of IB-IRM, ERM, IRM, and IB-ERM. However, since  $w$  relies on spurious features which may change arbitrary in unseen environments, it thus fails to solve OOD generalization.  $\square$

### Appendix B.3. Proof of Theorem 3

**Proof.** Assume  $\Phi^* \in \mathbb{R}^{c \times d}$  and  $w^*$  are the feature extractor and classifier learned by IB-ERM. Consider the feature variable extracted by  $\Phi^*$  as

$$\Phi^* X^e = \Phi^* S(Z_{inv}^e, Z_{spu}^e) = \Phi_{inv} Z_{inv}^e + \Phi_{spu} Z_{spu}^e. \quad (\text{A13})$$

We first show that  $\Phi_{inv} = \mathbf{0}$  or  $\Phi_{spu} = \mathbf{0}$ . We prove this by contradiction. Assume  $\Phi_{inv} \neq \mathbf{0}$  and  $\Phi_{spu} \neq \mathbf{0}$ . By observing that a solution of  $\Phi_{inv} = \mathbf{1}, \Phi_{spu} = \mathbf{0}, w^* = w_{inv}^*$  could make the average training error to  $q$ ; therefore any solution returned by IB-ERM should also achieve the error no larger than  $q$  (because  $rth = q$  in the constraint of Equation (12)). Therefore  $w^* \neq \mathbf{0}$ .

1. In the case when each  $e \in \mathcal{E}_{tr}$  follows Assumption 4 of  $Z_{spu}^e \leftarrow AZ_{inv}^e + W^e$ , we have

$$\begin{aligned} w^* \cdot (\Phi_{inv} Z_{inv}^e + \Phi_{spu} Z_{spu}^e) &= w^* \cdot \Phi_{inv} Z_{inv}^e + w^* \cdot \Phi_{spu} (AZ_{inv}^e + W^e) \\ &= w^* \cdot (\Phi_{inv} + \Phi_{spu} A) Z_{inv}^e + w^* \cdot \Phi_{spu} W^e. \end{aligned}$$

Then, for any  $z = (z_{inv}^e, z_{spu}^e)$  of  $\mathbf{1}(w_{inv}^* \cdot z_{inv}^e) = 1$ , we must have  $w^* \cdot (\Phi_{inv} + \Phi_{spu} A) z_{inv}^e + w^* \cdot \Phi_{spu} w^e \geq 0$  for any  $w^e$  to make error no larger than  $q$ . Since  $W^e$  is zero mean with at least two distinct points in each component, we can conclude that  $w^* \cdot (\Phi_{inv} + \Phi_{spu} A) z_{inv}^e \geq 0$ . Similarly, for any  $z = (z_{inv}^e, z_{spu}^e)$  of  $\mathbf{1}(w_{inv}^* \cdot z_{inv}^e) = 0$ , we have  $w^* \cdot (\Phi_{inv} + \Phi_{spu} A) z_{inv}^e < 0$ . From Lemma A3 or Lemma A4, we obtain  $H((\Phi_{inv} + \Phi_{spu} A) Z_{inv}^e + \Phi_{spu} W^e) > H((\Phi_{inv} + \Phi_{spu} A) Z_{inv}^e)$ . Therefore, there exists a more optimal solution to IB-ERM with zero weight to  $Z_{spu}^e$ , which contradicts the assumption.

2. In the case when each  $e \in \mathcal{E}_{tr}$  follows Assumption 5 of  $Z_{inv}^e \leftarrow AZ_{spu}^e + W^e$ , we have

$$\begin{aligned} w^* \cdot (\Phi_{inv} Z_{inv}^e + \Phi_{spu} Z_{spu}^e) &= w^* \cdot \Phi_{inv} (AZ_{spu}^e + W^e) + w^* \cdot \Phi_{spu} Z_{spu}^e \\ &= w^* \cdot (\Phi_{spu} + \Phi_{inv} A) Z_{spu}^e + w^* \cdot \Phi_{inv} W^e. \end{aligned}$$

From Lemma A3 or Lemma A4, we obtain  $H((\Phi_{spu} + \Phi_{inv} A) Z_{spu}^e + \Phi_{inv} W^e) > H((\Phi_{spu} + \Phi_{inv} A) Z_{spu}^e)$ . In addition, the spurious features are assumed to be linearly separable. Therefore, there exists a more optimal solution to IB-ERM with zero weight to  $Z_{inv}^e$ , which contradicts the assumption.



3. In the case when each  $e \in \mathcal{E}_{tr}$  follows Assumption 6 of  $Z_{spu}^e \leftarrow W_1^e Y^e + W_0^e (1 - Y^e)$ , we have

$$\begin{aligned} w^* \cdot (\Phi_{inv} Z_{inv}^e + \Phi_{spu} Z_{spu}^e) &= w^* \cdot \Phi_{inv} Z_{inv}^e + w^* \cdot \Phi_{spu} (W_1^e Y^e + W_0^e (1 - Y^e)) \\ &= w^* \cdot \Phi_{inv} Z_{inv}^e + w^* \cdot \Phi_{spu} W_1^e Y^e + w^* \cdot \Phi_{spu} W_0^e (1 - Y^e). \end{aligned}$$

Then, for any  $z = (z_{inv}^e, z_{spu}^e)$  of  $\mathbf{1}(w_{inv}^* \cdot z_{inv}^e) = 1$ , we must have  $w^* \cdot \Phi_{inv} z_{inv}^e + w^* \cdot \Phi_{spu} w_1^e y^e + w^* \cdot \Phi_{spu} w_0^e (1 - y^e) \geq 0$  for any  $w_1^e$  and  $w_0^e$  to make error no larger than  $q$ . Since  $W_1^e$  and  $W_0^e$  are both zero mean variables with at least two distinct points in each component, we can conclude that  $w^* \cdot \Phi_{inv} z_{inv}^e \geq 0$ ; Similarly, for any  $z = (z_{inv}^e, z_{spu}^e)$  of  $\mathbf{1}(w_{inv}^* \cdot z_{inv}^e) = 0$ , we have  $w^* \cdot \Phi_{inv} z_{inv}^e < 0$ . From Lemma A3 or Lemma A4, we obtain  $H(\Phi_{inv} Z_{inv}^e + \Phi_{spu} W_1^e Y^e + \Phi_{spu} W_0^e (1 - Y^e)) > H(\Phi_{inv} Z_{inv}^e)$ . Therefore, there exists a more optimal solution to IB-ERM with zero weight to  $Z_{spu}^e$ , which contradicts the assumption.

So far, we have proved that the feature extractor  $\Phi^*$  learned by IB-ERM would never extract both spurious features and invariant features together. Then, we perform singular value decomposition (SVD) to the  $\Phi^*$  as

$$\Phi^* = U \Lambda V^T = [U_1, U_2] [\Lambda_1, \mathbf{0}; \mathbf{0}, \mathbf{0}] [V_1^T; V_2^T] = U_1 \Lambda_1 V_1^T \tag{A14}$$

Let  $S \in \mathbb{R}^{d \times d}$  be the orthogonal matrix. Set  $r$  as the rank of the matrix  $\Phi^*$ , i.e.,  $r = Rank(\Phi^*)$ , and let  $V_1^T S = [V'_1, V'_2]$  with  $V'_1 \in \mathbb{R}^{r \times m}$  and  $V'_2 \in \mathbb{R}^{r \times o}$ , and  $V_2^T S = [V''_1, V''_2]$  with  $V''_1 \in \mathbb{R}^{(d-r) \times m}$  and  $V''_2 \in \mathbb{R}^{(d-r) \times o}$ , then

$$\Phi^* X^e = U_1 \Lambda_1 V_1^T S [Z_{inv}^e; Z_{spu}^e] = U_1 \Lambda_1 (V'_1 Z_{inv}^e + V'_2 Z_{spu}^e). \tag{A15}$$

Since  $\Phi^* X^e$  contains the information either from spurious features or from invariant features, we must have  $U_1 \Lambda_1 V'_1 = \mathbf{0}$  or  $U_1 \Lambda_1 V'_2 = \mathbf{0}$ , and thus,  $V'_1 = \mathbf{0}$  or  $V'_2 = \mathbf{0}$  due to  $Rank(U_1 \Lambda_1) = r$ . If  $V'_2 = \mathbf{0}$ , then  $\Phi^*$  extract invariant features only. Otherwise when  $V'_1 = \mathbf{0}$ , we decompose the  $V^T S$  by

$$V^T S = [V_1^T; V_2^T] S = [V_1^T S; V_2^T S] = [V'_1, V'_2; V''_1, V''_2]. \tag{A16}$$

Since  $V^T$  and  $S$  are both the orthogonal matrix,  $V^T S$  is also orthogonal; thus  $V'_1 = \mathbf{0} \Rightarrow V_2^T V''_2 = \mathbf{0}$ , and then  $Rank(V''_2) = Rank([V'_2; V''_2]) - Rank(V'_2) = o - r$  (note that  $r \leq \min\{m, o\}$ ). Then,

$$V_2^T X^e = V_2^T S [Z_{inv}^e; Z_{spu}^e] = [V''_1, V''_2] [Z_{inv}^e; Z_{spu}^e] = V''_1 Z_{inv}^e + V''_2 Z_{spu}^e. \tag{A17}$$

Therefore, by running the CSIB for one iteration, the rank of spurious features would be decreased by  $r > 0$ . This would result in zero weight to spurious features by finite runs of CSIB.

Then, we intend to show why the counterfactual supervision step could help to distinguish whether  $V'_1$  is  $\mathbf{0}$  or not. For a specific instance  $x = S[z_{inv}; z_{spu}]$ , let two new features be  $z^1$  and  $z^2$ , then  $do(z^1_{1:r}) = [-M, \dots, -M]$  and  $do(z^1_{r+1:d}) = V_2^T x$ ;  $do(z^2_{1:r}) = [M, \dots, M]$  and  $do(z^2_{r+1:d}) = V_2^T x$ . Back the new features  $z^1$  and  $z^2$  to the input space as  $x^1 = Vz^1$  and  $x^2 = Vz^2$ . If  $V'_1 = \mathbf{0}$ , then

$$\begin{aligned} S^{-1} x^1 &= S^{-1} Vz^1 = S^{-1} V [z^1_{1:r}; V''_1 z_{inv} + V''_2 z_{spu}] \\ &= (V^T S)^T [z^1_{1:r}; V''_1 z_{inv} + V''_2 z_{spu}] \\ &= [V_1^T, V_1^{''T}; V_2^T, V_2^{''T}] [z^1_{1:r}; V''_1 z_{inv} + V''_2 z_{spu}] \\ &= [V_1^T z^1_{1:r} + V_1^{''T} (V''_1 z_{inv} + V''_2 z_{spu}); V_2^T z^1_{1:r} + V_2^{''T} (V''_1 z_{inv} + V''_2 z_{spu})] \\ &= [z_{inv}; V_2^T z^1_{1:r} + V_2^{''T} V''_2 z_{spu}], \end{aligned}$$

and similarly we have  $S^{-1}x^2 = [z_{inv}; V_2^T z_{1:r}^2 + V_2^{''T} V_2'' z_{spu}]$ . Therefore, the ground truths of  $x^1$  and  $x^2$  are the same. On other hand, if  $V_1' \neq \mathbf{0}$ , then  $V_2' = \mathbf{0}$ , and

$$\begin{aligned} S^{-1}x^1 &= S^{-1}Vz^1 = S^{-1}V[z_{1:r}^1; V_1'' z_{inv} + V_2'' z_{spu}] \\ &= (V^T S)^T [z_{1:r}^1; V_1'' z_{inv} + V_2'' z_{spu}] \\ &= [V_1'^T, V_1''^T; V_2'^T, V_2''^T] [z_{1:r}^1; V_1'' z_{inv} + V_2'' z_{spu}] \\ &= [V_1'^T z_{1:r}^1 + V_1''^T (V_1'' z_{inv} + V_2'' z_{spu}); V_2'^T z_{1:r}^1 + V_2''^T (V_1'' z_{inv} + V_2'' z_{spu})] \\ &= [V_1'^T z_{1:r}^1 + V_1''^T V_1'' z_{inv}; z_{spu}], \end{aligned}$$

and similarly we have  $S^{-1}x^2 = [V_1'^T z_{1:r}^2 + V_1''^T V_1'' z_{inv}; z_{spu}]$ . Since  $z_{1:r}^1 = -z_{1:r}^2$  and their magnitudes are large enough to make  $\text{sgn}(w_{inv}^* \cdot (V_1'^T z_{1:r}^1 + V_1''^T V_1'' z_{inv})) \neq \text{sgn}(w_{inv}^* \cdot (V_1'^T z_{1:r}^2 + V_1''^T V_1'' z_{inv}))$ ; thus the ground truths of  $x^1$  and  $x^2$  would be different. Therefore, the counterfactual supervision step could help to detect whether invariant features or spurious features are extracted by using a single sample only.

Finally, when only invariant features are extracted by  $\Phi$ , the training error is minimized, i.e.,  $w^* \Phi_{inv} \in \arg \min_f \mathbb{E}_{\mathbb{P}} [l(f(Z_{inv}^{tr}), Y^{tr})]$ . Then, based on our assumption to the OOD environments (Assumptions 8), i.e.,  $\forall e \in \mathcal{E}_{ood}, F_l(\mathbb{P}(Z_{inv}^{tr}, Y^{tr})) \subseteq F_l(\mathbb{P}(Z_{inv}^e, Y^e))$ , therefore, for any  $e \in \mathcal{E}_{ood}$ , we have  $\mathbb{E}_{\mathbb{P}} [l((X^e, Y^e), w^* \Phi)] = \mathbb{E}_{\mathbb{P}} [l((Z_{inv}^e, Y^e), w^* \Phi_{inv})] = \mathbb{E}_{\mathbb{P}} [l((Z_{inv}^{tr}, Y^{tr}), w^* \Phi_{inv})] = q$ .  $\square$

It is worth noting that the proof of Theorem 3 does not rely on how many labels there would be, so it is easily extended to the multi-class classification case as long as the corresponding assumptions and conditions are satisfied.

## References

1. Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; Fergus, R. Intriguing properties of neural networks. *arXiv* **2013**, arXiv:1312.6199.
2. Rosenfeld, A.; Zemel, R.; Tsotsos, J.K. The elephant in the room. *arXiv* **2018**, arXiv:1808.03305.
3. Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F.A.; Brendel, W. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
4. Nguyen, A.; Yosinski, J.; Clune, J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In Proceedings of the Computer Vision and Pattern Recognition Conference, Boston, MA, USA, 7–12 June 2015; pp. 427–436.
5. Gururangan, S.; Swayamdipta, S.; Levy, O.; Schwartz, R.; Bowman, S.R.; Smith, N.A. Annotation Artifacts in Natural Language Inference Data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*; Association for Computational Linguistics: New Orleans, LA, USA, 2018.
6. Geirhos, R.; Jacobsen, J.H.; Michaelis, C.; Zemel, R.; Brendel, W.; Bethge, M.; Wichmann, F.A. Shortcut learning in deep neural networks. *Nat. Mach. Intell.* **2020**, *2*, 665–673. [\[CrossRef\]](#)
7. Beery, S.; Van Horn, G.; Perona, P. Recognition in terra incognita. In Proceedings of the European Conference on Computer Vision, Munich, Germany, 8–14 September 2018; pp. 456–473.
8. Arjovsky, M.; Bottou, L.; Gulrajani, I.; Lopez-Paz, D. Invariant risk minimization. *arXiv* **2019**, arXiv:1907.02893.
9. Krueger, D.; Caballero, E.; Jacobsen, J.H.; Zhang, A.; Binas, J.; Zhang, D.; Le Priol, R.; Courville, A. Out-of-distribution generalization via risk extrapolation (rex). In Proceedings of the International Conference on Machine Learning, PMLR, Virtual, 18–24 July 2021; pp. 5815–5826.
10. Ahuja, K.; Shanmugam, K.; Varshney, K.; Dhurandhar, A. Invariant risk minimization games. In Proceedings of the International Conference on Machine Learning, PMLR, Virtual, 13–18 July 2020; pp. 145–155.
11. Pezeshki, M.; Kaba, O.; Bengio, Y.; Courville, A.C.; Precup, D.; Lajoie, G. Gradient starvation: A learning proclivity in neural networks. In Proceedings of the Neural Information Processing Systems, Virtual, 6–14 December 2021; Volume 34.
12. Ahuja, K.; Caballero, E.; Zhang, D.; Gagnon-Audet, J.C.; Bengio, Y.; Mitliagkas, I.; Rish, I. Invariance principle meets information bottleneck for out-of-distribution generalization. In Proceedings of the Neural Information Processing Systems, Virtual, 6–14 December 2021; Volume 34.
13. Pearl, J. *Causality*; Cambridge University Press: Cambridge, UK, 2009.
14. Peters, J.; Janzing, D.; Schölkopf, B. *Elements of Causal Inference: Foundations and Learning Algorithms*; The MIT Press: Cambridge, MA, USA, 2017.
15. Peters, J.; Bühlmann, P.; Meinshausen, N. Causal inference by using invariant prediction: Identification and confidence intervals. *J. R. Stat. Soc. Ser. B* **2016**, *78*, 947–1012. [\[CrossRef\]](#)

16. Tishby, N. The information bottleneck method. In Proceedings of the Annual Allerton Conference on Communications, Control and Computing, Monticello, IL, USA, 22–24 September 1999; pp. 368–377.
17. Aubin, B.; Słowiak, A.; Arjovsky, M.; Bottou, L.; Lopez-Paz, D. Linear unit-tests for invariance discovery. *arXiv* **2021**, arXiv:2102.10867.
18. Soudry, D.; Hoffer, E.; Nacson, M.S.; Gunasekar, S.; Srebro, N. The implicit bias of gradient descent on separable data. *J. Mach. Learn. Res.* **2018**, *19*, 2822–2878.
19. Heinze-Deml, C.; Peters, J.; Meinshausen, N. Invariant causal prediction for nonlinear models. *arXiv* **2018**, arXiv:1706.08576.
20. Rojas-Carulla, M.; Schölkopf, B.; Turner, R.; Peters, J. Invariant models for causal transfer learning. *J. Mach. Learn. Res.* **2018**, *19*, 1309–1342.
21. Rosenfeld, E.; Ravikumar, P.K.; Risteski, A. The Risks of Invariant Risk Minimization. In Proceedings of the International Conference on Learning Representations, Virtual, 3–7 May 2021.
22. Kamath, P.; Tangella, A.; Sutherland, D.; Srebro, N. Does invariant risk minimization capture invariance? In Proceedings of the International Conference on Artificial Intelligence and Statistics, PMLR, San Diego, CA, USA, 13–15 April 2021; pp. 4069–4077.
23. Lu, C.; Wu, Y.; Hernández-Lobato, J.M.; Schölkopf, B. Invariant Causal Representation Learning for Out-of-Distribution Generalization. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 December 2022.
24. Liu, C.; Sun, X.; Wang, J.; Tang, H.; Li, T.; Qin, T.; Chen, W.; Liu, T.Y. Learning causal semantic representation for out-of-distribution prediction. In Proceedings of the Neural Information Processing Systems, Virtual, 6–14 December 2021; Volume 34.
25. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. *arXiv* **2013**, arXiv:1312.6114.
26. Rezende, D.J.; Mohamed, S.; Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In Proceedings of the International Conference on Machine Learning, PMLR, Beijing, China, 21–26 June 2014; pp. 1278–1286.
27. Lu, C.; Wu, Y.; Hernández-Lobato, J.M.; Schölkopf, B. Nonlinear invariant risk minimization: A causal approach. *arXiv* **2021**, arXiv:2102.12353.
28. Bengio, Y.; Courville, A.; Vincent, P. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **2013**, *35*, 1798–1828. [[CrossRef](#)] [[PubMed](#)]
29. Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N.R.; Kalchbrenner, N.; Goyal, A.; Bengio, Y. Toward causal representation learning. *Proc. IEEE* **2021**, *109*, 612–634. [[CrossRef](#)]
30. Namkoong, H.; Duchi, J.C. Stochastic gradient methods for distributionally robust optimization with f-divergences. In Proceedings of the Neural Information processing Systems, Barcelona, Spain, 5–10 December 2016; Volume 29.
31. Sinha, A.; Namkoong, H.; Volpi, R.; Duchi, J. Certifying some distributional robustness with principled adversarial training. *arXiv* **2017**, arXiv:1710.10571.
32. Lee, J.; Raginsky, M. Minimax statistical learning with wasserstein distances. In Proceedings of the Neural Information Processing Systems, Montreal, Canada, 3–8 December 2018; Volume 31.
33. Duchi, J.C.; Namkoong, H. Learning models with uniform performance via distributionally robust optimization. *Ann. Stat.* **2021**, *49*, 1378–1406. [[CrossRef](#)]
34. Bühlmann, P. Invariance, causality and robustness. *Stat. Sci.* **2020**, *35*, 404–426. [[CrossRef](#)]
35. Blanchard, G.; Lee, G.; Scott, C. Generalizing from several related classification tasks to a new unlabeled sample. In Proceedings of the Neural Information Processing Systems, Granada, Spain, 12–15 December 2011; Volume 24.
36. Muandet, K.; Balduzzi, D.; Schölkopf, B. Domain generalization via invariant feature representation. In Proceedings of the International Conference on Machine Learning, PMLR, Atlanta, GA, USA, 16–21 June 2013; pp. 10–18.
37. Deshmukh, A.A.; Lei, Y.; Sharma, S.; Dogan, U.; Cutler, J.W.; Scott, C. A generalization error bound for multi-class domain generalization. *arXiv* **2019**, arXiv:1905.10392.
38. Ye, H.; Xie, C.; Cai, T.; Li, R.; Li, Z.; Wang, L. Towards a Theoretical Framework of Out-of-Distribution Generalization. In Proceedings of the Neural Information Processing Systems, Virtual, 6–14 December 2021.
39. Xie, C.; Chen, F.; Liu, Y.; Li, Z. Risk variance penalization: From distributional robustness to causality. *arXiv* **2020**, arXiv:2006.07544.
40. Jin, W.; Barzilay, R.; Jaakkola, T. Domain extrapolation via regret minimization. *arXiv* **2020**, arXiv:2006.03908.
41. Mahajan, D.; Tople, S.; Sharma, A. Domain generalization using causal matching. In Proceedings of the International Conference on Machine Learning, PMLR, Virtual, 18–24 July 2021; pp. 7313–7324.
42. Bellot, A.; van der Schaar, M. Generalization and invariances in the presence of unobserved confounding. *arXiv* **2020**, arXiv:2007.10653.
43. Li, B.; Shen, Y.; Wang, Y.; Zhu, W.; Reed, C.J.; Zhang, J.; Li, D.; Keutzer, K.; Zhao, H. Invariant information bottleneck for domain generalization. In Proceedings of the Association for the Advancement of Artificial Intelligence, Virtual, 22 February–1 March 2022.
44. Alesiani, F.; Yu, S.; Yu, X. Gated information bottleneck for generalization in sequential environments. *Knowl. Informat. Syst.* **2022**, 1–23, *in press*. [[CrossRef](#)]
45. Wang, H.; Si, H.; Li, B.; Zhao, H. Provable Domain Generalization via Invariant-Feature Subspace Recovery. In Proceedings of the International Conference on Machine Learning, Baltimore, MD, USA, 17–23 July 2022.
46. Ganin, Y.; Lempitsky, V. Unsupervised domain adaptation by backpropagation. In Proceedings of the International conference on machine learning, PMLR, Lille, France, 6–11 July 2015; pp. 1180–1189.
47. 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, Munich, Germany, 8–14 September 2018; pp. 624–639.

48. Zhao, S.; Gong, M.; Liu, T.; Fu, H.; Tao, D. Domain generalization via entropy regularization. In Proceedings of the Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 16096–16107.
49. Ben-David, S.; Blitzer, J.; Crammer, K.; Pereira, F. Analysis of representations for domain adaptation. In Proceedings of the Neural Information Processing Systems, Hong Kong, China, 3–6 October 2006; Volume 19.
50. Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; Vaughan, J.W. A theory of learning from different domains. *Mach. Learn.* **2010**, *79*, 151–175. [[CrossRef](#)]
51. Zhao, H.; Des Combes, R.T.; Zhang, K.; Gordon, G. On learning invariant representations for domain adaptation. In Proceedings of the International Conference on Machine Learning. PMLR, Long Beach, CA, USA, 10–15 June 2019; pp. 7523–7532.
52. Xu, Q.; Zhang, R.; Zhang, Y.; Wang, Y.; Tian, Q. A fourier-based framework for domain generalization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 14383–14392.
53. Zhou, K.; Yang, Y.; Qiao, Y.; Xiang, T. Domain Generalization with MixStyle. In Proceedings of the International Conference on Learning Representations, Vienna, Austria, 3–7 May 2021.
54. Zhang, X.; Cui, P.; Xu, R.; Zhou, L.; He, Y.; Shen, Z. Deep stable learning for out-of-distribution generalization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 21–25 June 2021; pp. 5372–5382.
55. Wang, H.; Ge, S.; Lipton, Z.; Xing, E.P. Learning robust global representations by penalizing local predictive power. In Proceedings of the Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; Volume 32.
56. Gulrajani, I.; Lopez-Paz, D. In Search of Lost Domain Generalization. In Proceedings of the International Conference on Learning Representations, Virtual, 2–4 December 2020.
57. Wiles, O.; Goyal, S.; Stimberg, F.; Rebuffi, S.A.; Ktena, I.; Dvijotham, K.D.; Cemgil, A.T. A Fine-Grained Analysis on Distribution Shift. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
58. Thomas, M.; Joy, A.T. *Elements of Information Theory*; Wiley-Interscience: Hoboken, NJ, USA, 2006.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.