

# Stable Learning via Sparse Variable Independence

Han Yu<sup>1</sup>, Peng Cui<sup>1\*</sup>, Yue He<sup>1</sup>, Zheyang Shen<sup>1</sup>, Yong Lin<sup>2</sup>, Renzhe Xu<sup>1</sup>, Xingxuan Zhang<sup>1</sup>

<sup>1</sup>Tsinghua University, <sup>2</sup>Hong Kong University of Science and Technology  
yuh21@mails.tsinghua.edu.cn, cuip@tsinghua.edu.cn, heyue18@mails.tsinghua.edu.cn  
shenzy13@qq.com, ylindf@connect.ust.hk, xrz199721@gmail.com, xingxuanzhang@hotmail.com

## Abstract

The problem of covariate-shift generalization has attracted intensive research attention. Previous stable learning algorithms employ sample reweighting schemes to decorrelate the covariates when there is no explicit domain information about training data. However, with finite samples, it is difficult to achieve the desirable weights that ensure perfect independence to get rid of the unstable variables. Besides, decorrelating within stable variables may bring about high variance of learned models because of the over-reduced effective sample size. A tremendous sample size is required for these algorithms to work. In this paper, with theoretical justification, we propose SVI (Sparse Variable Independence) for the covariate-shift generalization problem. We introduce sparsity constraint to compensate for the imperfection of sample reweighting under the finite-sample setting in previous methods. Furthermore, we organically combine independence-based sample reweighting and sparsity-based variable selection in an iterative way to avoid decorrelating within stable variables, increasing the effective sample size to alleviate variance inflation. Experiments on both synthetic and real-world datasets demonstrate the improvement of covariate-shift generalization performance brought by SVI.

## 1 Introduction

Most of the current machine learning techniques rely on the IID assumption, that the test and training data are independent and identically distributed, which is too strong to stand in wild environments (Koh et al. 2021). Test distribution often differs from training distribution especially if there is data selection bias when collecting the training data (Heckman 1979; Young et al. 2009). Covariate shift is a common type of distribution shift (Ioffe and Szegedy 2015; Tripuraneni, Adlam, and Pennington 2021), which assumes that the marginal distribution of covariates (i.e.  $P(\mathbf{X})$ ) may shift between training and test data while the generation mechanism of the outcome variable (i.e.  $P(Y|\mathbf{X})$ ) remains unchanged (Shen et al. 2021). To address this problem, there are some strands of works (Santurkar et al. 2018; Wilson and Cook 2020; Wang et al. 2021). When some information about the test distribution is known apriori (Peng et al. 2019;

Yang and Soatto 2020), domain adaptation methods are proposed based on feature space transformation or distribution matching (Ben-David et al. 2010; Weiss, Khoshgoftaar, and Wang 2016; Tzeng et al. 2017; Ganin and Lempitsky 2015; Saito et al. 2018). If there exists explicit heterogeneity in the training data, e.g. it is composed of multiple subpopulations corresponding to different source domains (Blanchard et al. 2021; Gideon, McInnis, and Provost 2021), domain generalization methods are proposed to learn a domain-agnostic model or invariant representation (Muandet, Balduzzi, and Schölkopf 2013; Li et al. 2017; Ganin et al. 2016; Li et al. 2018; Sun and Saenko 2016; He, Shen, and Cui 2021; Zhang et al. 2022a,b). In many real applications, however, neither the knowledge about test data nor explicit domain information in training data is available.

Recently, stable learning algorithms (Shen et al. 2018, 2020a,b; Kuang et al. 2018, 2020; Zhang et al. 2021; Liu et al. 2021a,b) are proposed to address a more realistic and challenging setting, where the training data consists of latent heterogeneity (without explicit domain information), and the goal is to achieve a model with good generalization ability under agnostic covariate shift. They make a structural assumption of covariates by splitting them into  $\mathcal{S}$  (i.e. stable variables) and  $\mathcal{V}$  (i.e. unstable variables), and suppose  $P(Y|\mathcal{S})$  remains unchanged while  $P(Y|\mathcal{V})$  may change under covariate shift. They aim to learn a group of sample weights to remove the correlations among covariates in observational data, and then optimize in the weighted distribution to capture stable variables. It is theoretically proved that, under the scenario of infinite samples, these models only utilize the stable variables for prediction (i.e. the coefficients on unstable variables will be perfectly zero) if the learned sample weights can strictly ensure the mutual independence among all covariates (Xu et al. 2022). However, with finite samples, it is almost impossible to learn weights that satisfy complete independence. As a result, the predictor cannot always get rid of the unstable variables (i.e. the unstable variables may have significantly non-zero coefficients). In addition, (Shen et al. 2020a) pointed out that it is unnecessary to remove the inner correlations of stable variables. The correlations inside stable variables could be very strong, thus decorrelating them could sharply decrease the effective sample size, even leading to the variance inflation of learned models. Taking these two factors together, the requirement

\*Corresponding author

of a tremendous effective sample size severely restricts the range of applications for these algorithms.

In this paper, we propose a novel algorithm named Sparse Variable Independence (SVI) to help alleviate the rigorous requirement of sample size. We integrate the sparsity constraint for variable selection and the sample reweighting process for variable independence into a linear/nonlinear predictive model. We theoretically prove that even when the data generation is nonlinear, the stable variables can surely be selected with a sparsity constraint like  $\ell_1$  penalty, if the correlations between stable variables and unstable variables are weak to some extent. Thus we do not require complete independence among covariates. To further reduce the requirement on sample size, we design an iterative procedure between sparse variable selection and sample reweighting to prevent attempts to decorrelate within stable variables. The experiments on both synthetic and real-world datasets clearly demonstrate the improvement of covariate-shift generalization performance brought by SVI.

The main contributions in this paper are listed below:

- We introduce the sparsity constraint to attain a more pragmatic independence-based sample reweighting algorithm which improves covariate-shift generalization ability with finite training samples. We theoretically prove the benefit of this.
- We design an iterative procedure to avoid decorrelating within stable variables, mitigating the problem of over-reduced effective sample size.
- We conduct extensive experiments on various synthetic and real-world datasets to verify the advantages of our proposed method.

## 2 Problem Definition

*Notations:* Generally in this paper, a bold-type letter represents a matrix or a vector, while a normal letter represents a scalar. Unless otherwise stated,  $\mathbf{X} \in \mathbb{R}^p$  denotes the covariates with dimension of  $p$ ,  $X_d$  denotes the  $d^{\text{th}}$  variable in  $\mathbf{X}$ , and  $Y \in \mathbb{R}$  denotes the outcome variable. Training data is drawn from the distribution  $P^{\text{tr}}(\mathbf{X}, Y)$  while test data is drawn from the unknown distribution  $P^{\text{te}}(\mathbf{X}, Y)$ . Let  $\mathcal{X}, \mathcal{X}_d, \mathcal{Y}$  denote the support of  $\mathbf{X}, X_d, Y$  respectively.

$\mathbf{S} \subseteq \mathbf{X}$  implies  $\mathbf{S}$  is a subset variables of  $\mathbf{X}$  with dimension  $p_s$ . We use  $\mathbf{A} \perp \mathbf{B}$  to denote that  $\mathbf{A}$  and  $\mathbf{B}$  are statistically independent of each other. We use  $\mathbb{E}_{Q(\cdot)}[\cdot]$  to denote expectation and  $\mathbb{E}_{Q(\cdot)}[\cdot|\cdot]$  to denote conditional expectation under distribution  $Q$ , which can be chosen as  $P^{\text{tr}}, P^{\text{te}}$  or any other proper distributions.

In this work, we focus on the problem of covariate shift. It is a typical and most common kind of distribution shift considered in OOD literature.

**Problem 1** (Covariate-Shift Generalization). *Given the samples  $\{(\mathbf{X}_i, Y_i)\}_{i=1}^N$  drawn from training distribution  $P^{\text{tr}}$ , the goal of covariate-shift generalization is to learn a prediction model so that it performs stably on predicting  $Y$  in agnostic test distribution where  $P^{\text{te}}(\mathbf{X}, Y) = P^{\text{te}}(\mathbf{X})P^{\text{tr}}(Y|\mathbf{X})$ .*

To address the covariate-shift generalization problem, we define minimal stable variable set (Xu et al. 2022).

**Definition 2.1** (Minimal stable variable set). *A minimal stable variable set of predicting  $Y$  under training distribution  $P^{\text{tr}}$  is any subset  $\mathbf{S}$  of  $\mathbf{X}$  satisfying the following equation, and none of its proper subsets satisfies it.*

$$\mathbb{E}_{P^{\text{tr}}}[Y|\mathbf{S}] = \mathbb{E}_{P^{\text{tr}}}[Y|\mathbf{X}]. \quad (1)$$

Under strictly positive density assumption, the minimal stable variable set  $\mathbf{S}$  is unique. In the setting of covariate shift where  $P^{\text{tr}}(\mathbf{X}) \neq P^{\text{te}}(\mathbf{X})$ , relationships between  $\mathbf{S}$  and  $\mathbf{X} \setminus \mathbf{S}$  can arbitrarily change, resulting in the unstable correlations between  $Y$  and  $\mathbf{X} \setminus \mathbf{S}$ . Demonstrably, according to (Xu et al. 2022),  $\mathbf{S}$  is a minimal and optimal predictor of  $Y$  under test distribution  $P^{\text{te}}$  if and only if it is a minimal stable variable set under  $P^{\text{tr}}$ . Hence, in this paper, we intend to capture the minimal stable variable set  $\mathbf{S}$  for stable prediction under covariate shift. Without ambiguity, we refer to  $\mathbf{S}$  as stable variables and  $\mathbf{V} = \mathbf{X} \setminus \mathbf{S}$  as unstable variables in the rest of this paper.

## 3 Method

### 3.1 Independence-based sample reweighting

First, we define the weighting function and the target distribution that we want the training distribution to be reweighted to.

**Definition 3.1** (Weighting function). *Let  $\mathcal{W}$  be the set of weighting functions that satisfy*

$$\mathcal{W} = \{w : \mathcal{X} \rightarrow \mathbb{R}^+ \mid \mathbb{E}_{P^{\text{tr}}(\mathbf{X})}[w(\mathbf{X})] = 1\}. \quad (2)$$

*Then  $\forall w \in \mathcal{W}$ , the corresponding weighted distribution is  $\tilde{P}_w(\mathbf{X}, Y) = w(\mathbf{X})P^{\text{tr}}(\mathbf{X}, Y)$ .  $\tilde{P}_w$  is well defined with the same support of  $P^{\text{tr}}$ .*

Since we expect that variables are decorrelated in the weighted distribution, we denote  $\mathcal{W}_\perp$  as the subset of  $\mathcal{W}$  in which  $\mathbf{X}$  are mutually independent in the weighted distribution  $\tilde{P}_w$ .

Under the infinite-sample setting, it is proved that if conducting weighted least squares using the weighting function in  $\mathcal{W}_\perp$ , almost surely there will be non-zero coefficients only on stable variables, no matter whether the data generation function is linear or nonlinear (Xu et al. 2022). However, the condition for this to hold is too strict and ideal. Under the finite-sample setting, we can hardly learn the sample weights corresponding to the weighting function in  $\mathcal{W}_\perp$ .

We now take a look at two specific techniques of sample reweighting, which will be incorporated into our algorithms.

**DWR** (Kuang et al. 2020) aims to remove linear correlations between each two variables, *i.e.*,

$$\hat{w}(\mathbf{X}) = \arg \min_{w(\mathbf{X})} \sum_{1 \leq i, j \leq p, i \neq j} (\text{Cov}(X_i, X_j; w))^2, \quad (3)$$

where  $\text{Cov}(X_i, X_j; w)$  represents the covariance of  $X_i$  and  $X_j$  after being reweighted by  $\tilde{P}_w$ . DWR is well fitted for the case where the data generation process is dominated by a linear function, since it focuses on linear decorrelation only.

**SRDO** (Shen et al. 2020b) conducts sample reweighting by density ratio estimation. It simulates the target distribution  $\tilde{P}$  by means of random resampling on each covariate so that  $\tilde{P}(X_1, X_2, \dots, X_p) = \prod_{i=1}^p P^{tr}(X_i)$ . Then the weighting function can be estimated by

$$\hat{w}(\mathbf{X}) = \frac{\tilde{P}(\mathbf{X})}{P^{tr}(\mathbf{X})} = \frac{P^{tr}(X_1)P^{tr}(X_2)\dots P^{tr}(X_p)}{P^{tr}(X_1, X_2, \dots, X_p)}. \quad (4)$$

To estimate such density ratio, SRDO learns an MLP classifier to differentiate whether a sample belongs to the original distribution  $P^{tr}$  or the mutually independent target distribution  $\tilde{P}$ . Unlike DWR, this method can not only decrease linear correlations among covariates, but it can weaken non-linear dependence among them.

Under the finite-sample setting, for DWR, if the scale of sample size is not significantly greater than that of covariates dimension, it is difficult for Equation 3 to be optimized close to zero. For SRDO,  $\tilde{P}$  is generated by a rough process of resampling, further resulting in the inaccurate estimation of density ratio. In addition, both methods suffer from the over-reduced effective sample size when strong correlations exist inside stable variables, since they decorrelate variables globally. Therefore, they both require an enormous sample size to work.

### 3.2 Sample Reweighting with Sparsity Constraint

**Motivation and general idea** By introducing sparsity constraint, we can loose the demand for the perfectness of independence achieved by sample reweighting under the finite-sample setting, reducing the requirement on sample size. Inspired by (Zhao and Yu 2006), we theoretically prove the benefit of this.

**Theorem 3.1.** Set  $\mathbf{X}^{(n)} = [\mathbf{S}^{(n)}, \mathbf{V}^{(n)}], \beta = \begin{pmatrix} \beta_s \\ \beta_v \end{pmatrix}$ . For data generated by:  $\mathbf{Y}^{(n)} = \mathbf{X}^{(n)}\beta + g^{(n)}(\mathbf{S}) + \epsilon^{(n)} = \mathbf{S}^{(n)}\beta_s + \mathbf{V}^{(n)}\beta_v + g^{(n)}(\mathbf{S}) + \epsilon^{(n)}$ , where  $\epsilon^{(n)}$  is a vector of i.i.d. random variables with mean 0 and variance  $\sigma^2$ ,  $\mathbf{Y}^{(n)}$  is  $n \times 1$  outcome,  $\mathbf{S}^{(n)}$  and  $\mathbf{V}^{(n)}$  are  $n \times p_s$  and  $n \times p_v$  data matrix respectively,  $g^{(n)}(\mathbf{S})$  is  $n \times 1$  nonlinear term, which is linearly uncorrelated with the linear part.  $\beta_s$  is regression coefficients for  $p_s$  stable variables whose elements are all non-zero, while  $\beta_v = \mathbf{0}$ .

Assume  $\mathbf{S}$  and  $\mathbf{V}$  are normalized to zero mean with finite second-order moment, and both the covariates and the nonlinear term are bounded, i.e. almost surely  $\|\mathbf{X}\|_2 \leq B, g(\mathbf{S}) \leq \delta$ . If there exists a positive constant vector  $\eta$  such that:

$$|\text{Cov}^{(n)}(\mathbf{V}, \mathbf{S})\text{Cov}^{(n)}(\mathbf{S})^{-1}\text{sign}(\beta_s)| \leq 1 - \eta \quad (5)$$

where  $\text{Cov}^{(n)}$  represents sample covariance,  $\text{sign}$  represents element-wise sign function.

Then there exists a positive constant  $M$  which can be written as  $h(\delta, B, \text{Cov}(\mathbf{X}))$  unrelated to  $n$  for the following inequality to stand:  $\forall \lambda_n$  satisfying  $\lambda_n = o(n)$  and  $\lambda_n = \omega\left(n^{\frac{1+c}{2}}\right)$  with  $0 \leq c \leq 1$ :

$$P(\hat{\beta}(\lambda_n) =_s \beta) \geq 1 - o\left(n^{-\frac{c}{2}} e^{-\frac{\min_{i=1}^{p_v} \{\eta_i\}^2}{8M^2} n^c}\right) \quad (6)$$




Figure 1: Diagram of SVI.

where “ $=_s$ ” means equal after applying sign function,  $\hat{\beta}(\lambda_n)$  is the optimal solution of Lasso with regularizer coefficient  $\lambda_n$ .

The term  $\text{Cov}^{(n)}(\mathbf{V}, \mathbf{S})$  measures the dependency between stable variables and unstable variables. Thus Theorem 3.1 implies that if correlations between  $\mathbf{S}$  and  $\mathbf{V}$  are weakened to some extent, the probability of perfectly selecting stable variables  $\mathbf{S}$  approaches 1 at an exponential rate. This confirms that adding sparsity constraint to independence-based sample reweighting may loose the requirement for independence among covariates, even when the data generation is nonlinear.

Motivated by Theorem 3.1, we propose the novel Sparse Variable Independence (SVI) method. The algorithm mainly consists of two modules: the frontend for sample reweighting, and the backend for sparse learning under reweighted distribution. We then combine the two modules in an iterative way as shown in Figure 1. Details are provided as follows.

**Frontend implementation** We employ different techniques under different settings. For the case where data generation is dominated by a linear function, we employ DWR, namely Equation 3 as loss function to be optimized using gradient descent. For the case where data generation is totally governed by nonlinear function, we employ SRDO, namely Equation 4 to conduct density ratio estimation for reweighting.

**Backend implementation** As for the backend sparse learning module, in order to take the nonlinear setting into account, we implement the sparsity constraint following (Yamada et al. 2020) instead of Lasso. Typical for variable selection, we begin with an  $\ell_0$  constraint which is equivalent to multiplying the covariates  $\mathbf{X}$  with a hard mask  $\mathbf{M} = [m_1, m_2, \dots, m_p]^T$  whose elements are either one or zero. We approximate the elements in  $\mathbf{M}$  using clipped Gaussian random variables parametrized by  $\mu = [\mu_1, \mu_2, \dots, \mu_p]^T$ :

$$m_d = \max\{0, \min\{1, \mu_d + \epsilon_d\}\} \quad (7)$$

where  $\epsilon_d$  is drawn from zero-mean Gaussian distribution  $\mathcal{N}(0, \sigma^2)$ .

The standard  $\ell_0$  constraint for a general function  $f$  parametrized by  $\theta$  can be written as:

$$\mathcal{L}(\theta, \mu) = \mathbb{E}_{P(\mathbf{X}, Y)} \mathbb{E}_{\mathbf{M}} [l(f(\mathbf{M} \odot \mathbf{X}; \theta), Y) + \lambda \|\mathbf{M}\|_0]. \quad (8)$$

In Equation 8, with the help of continuous probabilistic approximation, we can derive the  $\ell_0$  norm of the mask as:

$\mathbb{E}_{\mathbf{M}} \|\mathbf{M}\|_0 = \sum_{d=1}^p P(m_d > 0) = \sum_{d=1}^p \Phi(\frac{\mu_d}{\sigma})$ , where  $\Phi$  is the cumulative distribution function of standard Gaussian. After learning sample weights corresponding to  $\hat{w}(\mathbf{X})$ , and combining them with the sparsity constraint, we rewrite Equation 8 as below:

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\mu}) = \mathbb{E}_{P(\mathbf{X}, Y)} [\hat{w}(\mathbf{X}) \mathbb{E}_{\mathbf{M}} [l(f(\mathbf{M} \odot \mathbf{X}; \boldsymbol{\theta}), Y) + \lambda \sum_{d=1}^p \Phi(\frac{\mu_d}{\sigma})]] \quad (9)$$

As a result, the optimization of Equation 9 outputs model parameters  $\boldsymbol{\theta}$ , and soft masks  $\boldsymbol{\mu}$  which are continuous variables in the range  $[0, 1]$ . Therefore,  $\mu_d$  can be viewed as the probability of selecting  $X_d$  as a stable variable. We may set a threshold to conduct variable selection, and then retrain a model only using selected variables for better covariate-shift generalization.

**Iterative procedure** As mentioned before, global decorrelation between each pair of covariates could be too aggressive to accomplish. In reality, correlations inside stable variables could be strong. Global decorrelation may give rise to the shrinkage of effective sample size, causing inflated variance. It is worth noting that outputs of the backend module can be interpreted as  $P(X_d \in \mathcal{S})$ , namely the probability of each variable belonging to stable variables. They contain information on the covariate structure. Therefore, when using DWR as the frontend, we propose an approach to taking advantage of this information as feedback for the frontend module to mitigate the decrease in effective sample size.

We first denote  $\mathbf{A} \in [0, 1]^{p \times p}$  as the covariance matrix mask, where  $A_{ij}$  represents the strength of decorrelation for  $X_i$  and  $X_j$ . Obviously, since we hope to preserve correlation inside stable variables  $\mathcal{S}$ , when the pair of variables is more likely to belong to  $\mathcal{S}$ , they should be less likely to be decorrelated. Thus the elements in  $\mathbf{A}$  can be calculated as:  $A_{ij} = 1 - P(X_i \in \mathcal{S})P(X_j \in \mathcal{S}) = 1 - \mu_i \mu_j$ . We incorporate this term into the loss function of DWR in Equation 3, revising it as:

$$\mathcal{L}(\mathbf{w}) = \sum_{1 \leq i, j \leq p, i \neq j} (A_{ij} \text{Cov}(X_i, X_j; \mathbf{w}))^2. \quad (10)$$

Through Equation 10, we implement SVI in an iterative way by combining sample reweighting and sparse learning. The details of this algorithm are described in Algorithm 1. We also present a diagram to illustrate it in Figure 1.

As we can see, when initialized, the frontend module  $\mathcal{M}_{reweight}$  learns a group of sample weights corresponding to the weighting function  $\hat{w}(\mathbf{X})$ . Given such sample weights, the backend module  $\mathcal{M}_{sparse}$  conducts sparse learning in a way of soft variable selection under the reweighted distribution, outputting the probability  $P(X_d \in \mathcal{S})$  for each variable  $X_d$  to be in the stable variable set. Such structural information can be utilized by  $\mathcal{M}_{reweight}$  to learn better sample weights, since some of the correlations inside stable variables will be preserved. Therefore, sample reweighting and sparse learning modules benefit each other through such a loop of iteration and feedback. The iterative procedure and its convergence are hard to be analyzed theoretically, like in previous works (Liu et al. 2021a; Zhou et al.

---

#### Algorithm 1: Sparse Variable Independence (SVI)

---

**Input:** Dataset  $[\mathbf{X}, \mathbf{Y}]$ , sparse learning regularizer coefficient  $\lambda$ , number of loops  $T$ , period of moving average for soft mask  $T_m$ , maximum number of iterations for learning weights  $T_w$ , maximum number of iterations for sparse learning  $T_\theta$ , selection threshold  $\gamma$ .

**Output:** Selected variable set  $\mathcal{S}'$ .

Initialize  $\mathbf{A}$  with  $A_{ij} = 1, \forall 1 \leq i, j \leq p$

Initialize  $\mathbf{U}$  as empty list.

**for**  $t = 1$  to  $T$  **do**

**while** not convergence or reach  $T_w$  **do**

    Update sample weights  $\mathbf{w}$  to minimize  $\mathcal{L}(\mathbf{w})$  in Equation 10.

**end while**

**while** not convergence or reach  $T_\theta$  **do**

    Update model parameters  $\boldsymbol{\theta}$  and soft mask probability  $\boldsymbol{\mu}$  to minimize  $\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\mu})$  in Equation 9.

**end while**

  Delete first element of  $\mathbf{U}$  if its length is greater than  $T_m$ .

  Append current  $\boldsymbol{\mu}$  to the end of  $\mathbf{U}$ .

  Calculate the moving average  $\boldsymbol{\mu}'$  from  $\mathbf{U}$ .

  Calculate  $A_{ij} = 1 - \mu'_i \mu'_j, \forall 1 \leq i, j \leq p$ .

**end for**

Calculate  $\mathcal{S}' = \{X_d \mid \mu_d \geq \gamma\}$

**return:**  $\mathcal{S}'$

---

2022), so we illustrate them through empirical experiments in Figure 3(c) and in appendix.

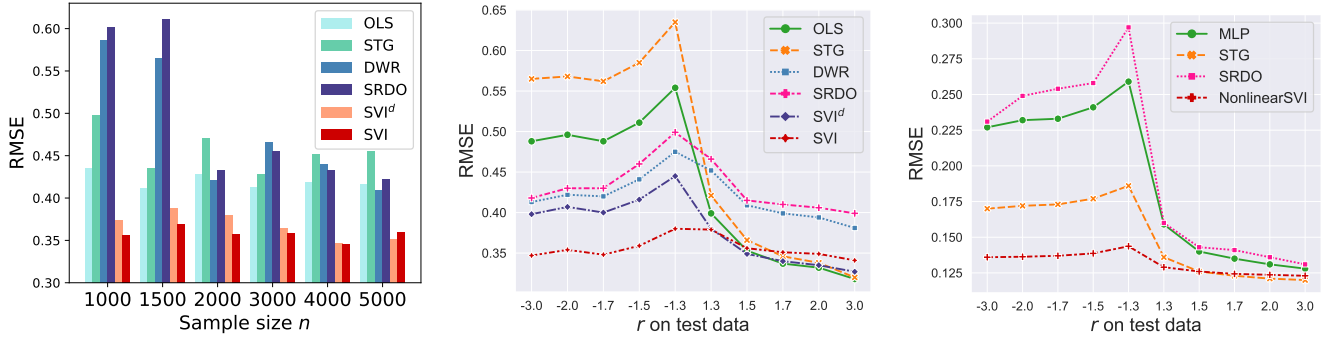
We denote Algorithm 1 as **SVI** which applies to the linear setting since its frontend is DWR. For nonlinear settings where DWR cannot address, we employ SRDO as the frontend module, denoted as **NonlinearSVI**. We do not implement it iteratively because it is hard for the resampling process of SRDO to incorporate the feedback of structural information from the backend module, which we leave for future extensions. For both settings, unlike previous methods that directly learn a prediction model by optimizing the weighted loss, we first apply our algorithm to select the expected stable variables, then we retrain a model using these variables for prediction under covariate shift.

## 4 Experiments

### 4.1 Baselines

We compare SVI with the following methods. We tune the hyperparameters by grid search and validation on data from the environment corresponding to  $r_{train}$ .

- OLS (Ordinary Least Squares): For linear settings.
- MLP (Multi-Layer Perceptron): For nonlinear settings.
- STG (Stochastic Gates) (Yamada et al. 2020): Directly optimizing Equation 9 without sample reweighting.
- DWR (Kuang et al. 2020): Optimizing Equation 3 and conducting Weighted Least Squares (for linear settings).
- SRDO (Shen et al. 2020b): Conducting density ratio estimation through Equation 4.



(a) RMSE of linear settings when fixing  $r_{train} = 2.5$  with varying  $n$ . (b) RMSE of the linear setting with bias rate  $r_{train} = 2.5$  and sample size  $n = 2000$ . (c) RMSE of the nonlinear setting with bias rate  $r_{train} = 1.8$  and sample size  $n = 25000$ .

Figure 2: Detailed results of experiments on synthetic data under linear and nonlinear settings.

Table 1: Results of the linear setting under varying sample size  $n$  and train data bias rate  $r$ .

Scenario 1: Varying sample size $n(r = 2.5)$									
$n$	$n = 1000$			$n = 1500$			$n = 2000$		
Methods	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error
OLS	0.435	0.084	0.572	0.411	0.072	0.531	0.428	0.084	0.554
STG	0.498	0.133	0.694	0.435	0.090	0.579	0.471	0.116	0.635
DWR	0.586	0.043	0.676	0.565	0.044	0.666	0.421	0.027	0.475
SRDO	0.601	0.040	0.698	0.611	0.045	0.690	0.433	0.035	0.499
$SVI^d$	0.374	0.023	<b>0.413</b>	0.388	0.043	0.465	0.380	0.038	0.445
SVI	<b>0.356</b>	<b>0.020</b>	0.415	<b>0.369</b>	<b>0.030</b>	<b>0.430</b>	<b>0.357</b>	<b>0.013</b>	<b>0.380</b>
Scenario 2: Varying bias rate $r(n = 1500)$									
$r$	$r = 2.0$			$r = 2.5$			$r = 3.0$		
Methods	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error
OLS	0.374	0.039	0.453	0.411	0.072	0.531	0.443	0.100	0.589
STG	0.395	0.047	0.475	0.435	0.090	0.579	0.526	0.165	0.749
DWR	0.431	0.031	0.499	0.565	0.044	0.666	0.576	0.038	0.655
SRDO	0.457	0.058	0.511	0.611	0.045	0.690	0.592	0.037	0.643
$SVI^d$	0.356	0.017	0.395	0.388	0.043	0.465	0.364	0.015	0.393
SVI	<b>0.345</b>	<b>0.015</b>	<b>0.380</b>	<b>0.369</b>	<b>0.030</b>	<b>0.430</b>	<b>0.330</b>	<b>0.014</b>	<b>0.377</b>

Table 2: Results of the nonlinear setting under varying sample size  $n$  and train data bias rate  $r$ .

Scenario 1: Varying sample size $n(r = 2.0)$									
$n$	$n=15000$			$n = 20000$			$n = 25000$		
Methods	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error
MLP	0.221	0.080	0.331	0.262	0.113	0.416	0.249	0.104	0.389
STG	0.177	0.049	0.243	0.176	0.048	0.241	0.176	0.048	0.243
SRDO	0.244	0.123	0.380	0.288	0.133	0.469	0.231	0.090	0.373
NonlinearSVI	<b>0.130</b>	<b>0.002</b>	<b>0.133</b>	<b>0.125</b>	<b>0.001</b>	<b>0.128</b>	<b>0.126</b>	<b>0.002</b>	<b>0.129</b>
Scenario 2: Varying bias rate $r(n = 25000)$									
$r$	$r = 1.8$			$r = 2.0$			$r = 2.2$		
Methods	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error	Mean_Error	Std_Error	Max_Error
MLP	0.188	0.051	0.259	0.249	0.104	0.389	0.498	0.312	0.901
STG	0.150	0.026	0.186	0.176	0.048	0.243	0.208	0.076	0.308
SRDO	0.200	0.071	0.297	0.236	0.089	0.353	0.469	0.203	0.717
NonlinearSVI	<b>0.132</b>	<b>0.007</b>	<b>0.144</b>	<b>0.126</b>	<b>0.002</b>	<b>0.129</b>	<b>0.126</b>	<b>0.002</b>	<b>0.129</b>

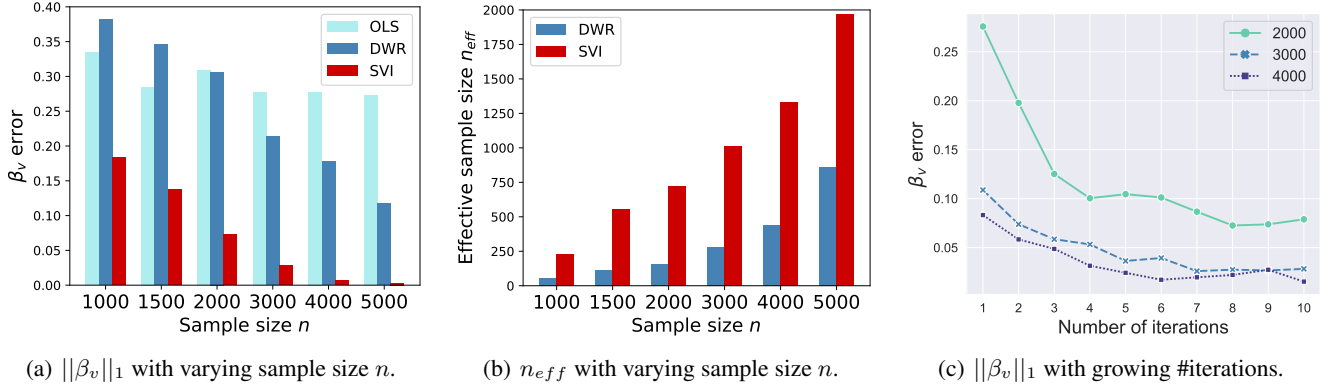


Figure 3: Additional results for illustrating effectiveness of SVI, when fixing  $r_{train} = 2.5$ .

- $SVI^d$ : A degenerated version of SVI by running only one iteration for ablation study to demonstrate the benefit of iterative procedure under linear settings.

## 4.2 Evaluation Metrics

To evaluate the covariate-shift generalization performance and stability, we assess the algorithms on test data from multiple different environments, thus we adopt the following metrics:  $Mean\_Error = \frac{1}{|\varepsilon_{te}|} \sum_{e \in \varepsilon_{te}} \mathcal{L}^e$ ;  $Std\_Error = \sqrt{\frac{1}{|\varepsilon_{te}|-1} \sum_{e \in \varepsilon_{te}} (\mathcal{L}^e - Mean\_Error)^2}$ ;  $Max\_Error = \max_{e \in \varepsilon_{te}} \mathcal{L}^e$ ;  $\varepsilon_{te}$  denotes the testing environments,  $\mathcal{L}^e$  denotes the empirical loss in the test environment  $e$ .

## 4.3 Experiments on Synthetic Data

**Dataset** We generate  $\mathbf{X} = \{\mathbf{S}, \mathbf{V}\}$  from a multivariate Gaussian distribution  $\mathbf{X} \sim N(\mathbf{0}, \Sigma)$ . In this way, we can simulate different correlation structures of  $\mathbf{X}$  through the control of covariance matrix  $\Sigma$ . In our experiments, we tend to make the correlations inside stable variables strong.

For linear and nonlinear settings, we employ different data generation functions. It is worth noting that a certain degree of model misspecification is needed, otherwise the model will be able to directly learn the true stable variables simply using OLS or MLP.

For the linear setting, we introduce the model misspecification error by adding an extra polynomial term to the dominated linear term, and later use a linear model to fit the data. The generation function is as follows:

$$Y = f(\mathbf{S}) + \epsilon = [\mathbf{S}, \mathbf{V}] \cdot [\beta_s, \beta_v]^T + \mathbf{S}_{\cdot,1} \mathbf{S}_{\cdot,2} \mathbf{S}_{\cdot,3} + \epsilon \quad (11)$$

For the nonlinear setting, we generate data in a totally nonlinear fashion. We employ random initialized MLP as the data generation function.

$$Y = f(\mathbf{S}) + \epsilon = MLP(\mathbf{S}) + \epsilon \quad (12)$$

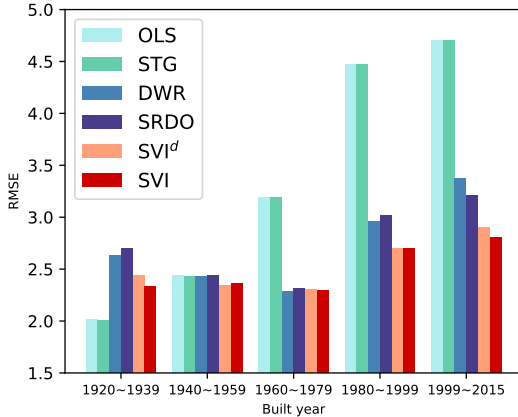
Later we use MLP with a smaller capacity to fit the data. More details are included in appendix.

**Generating Various Environments** To simulate the scenario of covariate shift and test not only the prediction accuracy but prediction stability, we generate a set of environments each with a distinct distribution. Specifically, following (Shen et al. 2020a), we generate different environments in our experiments through changing  $P(\mathbf{V}|\mathbf{S})$ , further leading to the change of  $P(Y|\mathbf{V})$ . Among all the unstable variables  $\mathbf{V}$ , we simulate unstable correlation  $P(\mathbf{V}_b|\mathbf{S})$  on a subset  $\mathbf{V}_b \in \mathbf{V}$ . We vary  $P(\mathbf{V}_b|\mathbf{S})$  through different strengths of selection bias with a bias rate  $r \in [-3, -1) \cup (1, 3]$ . For each sample, we select it with probability  $Pr = \prod_{V_i \in \mathbf{V}_b} |r|^{-5D_i}$ , where  $D_i = |f(\mathbf{S}) - \text{sign}(r)V_i|$ ,  $\text{sign}$  denotes sign function. In our experiments, we set  $p_{v_b} = 0.1 * p$ .

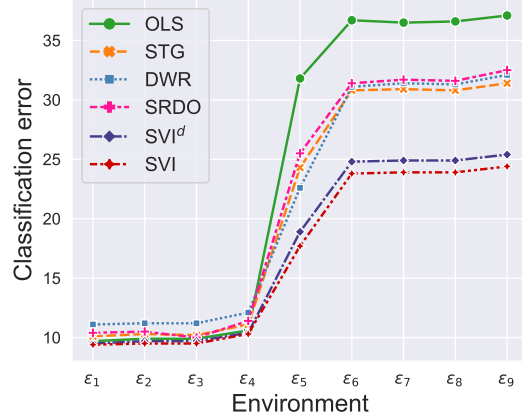
**Experimental settings** We train our models on data from one single environment generated with a bias rate  $r_{train}$  and test on data from multiple environments with bias rates  $r_{test}$  ranging in  $[-3, -1) \cup (1, 3]$ . Each model is trained 10 times independently with different training datasets from the same bias rate  $r_{train}$ . Similarly, for each  $r_{test}$ , we generate 10 different test datasets. The metrics we report are the mean results of these 10 times.

**Results** Results are shown in Table 1 and 2 when varying sample size  $n$  and training data bias rate  $r_{train}$ . Detailed results for two specific settings are illustrated in Figure 2(b) and 2(c). In addition to prediction performance, we illustrate the effectiveness of our algorithm in weakening residual correlations and increasing effective sample size in Figure 3(a) and 3(b). Analysis of the results is as follows:

- From Table 1 and 2, for almost every setting, SVI consistently outperforms other baselines in Mean\_Error, Std\_Error, and Max\_Error, indicating its superior covariate-shift generalization ability and stability. From Figure 2(b) and 2(c), when  $r_{test} < -1$ , i.e. the correlation between  $\mathbf{V}$  and  $Y$  reverses in test data compared with that in training data, other baselines fail significantly while SVI remains stable against such challenging distribution shift.
- From Figure 2(a), we find that there is a sharp rise of prediction error for DWR and SRDO with the decrease of



(a) Prediction error for house price



(b) Classification error for people income prediction

Figure 4: Detailed results of experiments on real-world data.

sample size, confirming the severity of the over-reduced effective sample size. Meanwhile, SVI maintains great performance and generally outperforms  $SVI^d$ , demonstrating the superiority of the iteration procedure which avoids aggressive decorrelation within  $\mathcal{S}$ .

- In Figure 3(a), we calculate  $\|\beta_v\|_1$  to measure the residual correlations between unstable variables  $\mathbf{V}$  and the outcome  $Y$ . DWR always preserves significant non-zero coefficients on unstable variables especially when the sample size gets smaller, while SVI truncates the scale of residual correlations by one or two orders of magnitude sometimes. This strongly illustrates that SVI indeed helps alleviate the imperfectness of variable decorrelation in previous independence-based sample reweighting methods.
- In Figure 3(b), we calculate  $n_{eff} = \frac{(\sum_{i=1}^n w_i)^2}{\sum_{i=1}^n w_i^2}$  as effective sample size following (Kish 2011). It is evident that SVI greatly boosts the effective sample size compared with global decorrelation schemes like DWR, since there are strong correlations inside stable variables. The shrinkage of  $n_{eff}$  compared with the original sample size  $n$  becomes rather severe for DWR when  $n$  is relatively small, even reaching 1/10.
- In Figure 3(c) we plot the change of  $\|\beta_v\|_1$  with the increasing number of iterations of SVI. We can observe that the residual correlations decline as the algorithm evolves. This proves the benefit brought by our iterative procedure. Also, this demonstrates the numerical convergence of SVI since coefficients of unstable variables  $\mathbf{V}$  gradually approach zero with the iterative process. More experiments for analyzing the iterative process and convergence are included in appendix.

#### 4.4 Experiments on Real-World Data

**House Price Prediction** It is a regression task for predicting house price, splitting data into 6 environments, 1 for training and 5 for testing, according to the time period that

the house was built.<sup>1</sup> In 4 out of 5 test environments, SVI and  $SVI^d$  outperform other baselines, especially in the last 3 environments. The gap significantly increases along the time axis, which represents a longer time span between test data and training data. This implies that a severer distribution shift may embody the superiority of our algorithm. Besides, overall SVI performs slightly better than  $SVI^d$ , further demonstrating the benefit of iterative procedure on real-world data.

**People Income Prediction** It is a binary classification task for income prediction where 10 environments are generated by the combination of race and sex. We train the models on the first environment (*White, Female*) and test on the other 9 environments to simulate distribution shift.

For the first 4 environments where people’s gender remains the same as that in training data, i.e. both are female, these methods make little difference in prediction. However, for the last 5 environments, when the sex category is male, performance drops for every method with varying degrees. We can see that SVI methods are the most stable ones, whose performance is affected much less than other baselines in the presence of distribution shift. Moreover, SVI still outperforms  $SVI^d$  slightly, indicating the practical use of the iteration procedure again.

## 5 Conclusion

In this paper, we combined sample reweighting and sparsity constraint to compensate for the deficiency of independence-based sample reweighting when there is a residual dependency between stable and unstable variables, and the variance inflation caused by removing strong correlations inside stable variables. Experiments on both synthetic datasets and real-world datasets demonstrated the effectiveness of our algorithm.

<sup>1</sup>Due to the limited space, a detailed description of the two real-world datasets is presented in appendix.

## 6 Acknowledgements

This work was supported in part by National Key R&D Program of China (No. 2018AAA0102004), National Natural Science Foundation of China (No. U1936219, 62141607), Beijing Academy of Artificial Intelligence (BAAI). Peng Cui is the corresponding author. All opinions in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

## References

- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79(1): 151–175.
- Blanchard, G.; Deshmukh, A. A.; Dogan, Ü.; Lee, G.; and Scott, C. D. 2021. Domain Generalization by Marginal Transfer Learning. *J. Mach. Learn. Res.*, 22: 2:1–2:55.
- Ganin, Y.; and Lempitsky, V. S. 2015. Unsupervised Domain Adaptation by Backpropagation. *ArXiv*, abs/1409.7495.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1): 2096–2030.
- Gideon, J.; McInnis, M. G.; and Provost, E. M. 2021. Improving Cross-Corpus Speech Emotion Recognition with Adversarial Discriminative Domain Generalization (AD-DoG). *IEEE Transactions on Affective Computing*, 12: 1055–1068.
- He, Y.; Shen, Z.; and Cui, P. 2021. Towards non-iid image classification: A dataset and baselines. *Pattern Recognition*, 110: 107383.
- Heckman, J. J. 1979. Sample selection bias as a specification error. *Applied Econometrics*, 31: 129–137.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *ArXiv*, abs/1502.03167.
- Kish, L. 2011. Survey sampling. In *Survey sampling*, 643–643.
- Koh, P. W.; Sagawa, S.; Marklund, H.; Xie, S. M.; Zhang, M.; Balsubramani, A.; Hu, W.; Yasunaga, M.; Phillips, R. L.; Gao, I.; et al. 2021. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, 5637–5664. PMLR.
- Kuang, K.; Cui, P.; Athey, S.; Xiong, R.; and Li, B. 2018. Stable prediction across unknown environments. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1617–1626.
- Kuang, K.; Xiong, R.; Cui, P.; Athey, S.; and Li, B. 2020. Stable prediction with model misspecification and agnostic distribution shift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 4485–4492.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, 5542–5550.
- Li, H.; Pan, S. J.; Wang, S.; and Kot, A. C. 2018. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5400–5409.
- Liu, J.; Hu, Z.; Cui, P.; Li, B.; and Shen, Z. 2021a. Heterogeneous Risk Minimization. *arXiv preprint arXiv:2105.03818*.
- Liu, J.; Hu, Z.; Cui, P.; Li, B.; and Shen, Z. 2021b. Kernelized heterogeneous risk minimization. *arXiv preprint arXiv:2110.12425*.
- Muandet, K.; Balduzzi, D.; and Schölkopf, B. 2013. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, 10–18. PMLR.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment Matching for Multi-Source Domain Adaptation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 1406–1415.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum Classifier Discrepancy for Unsupervised Domain Adaptation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3723–3732.
- Santurkar, S.; Tsipras, D.; Ilyas, A.; and Madry, A. 2018. How Does Batch Normalization Help Optimization? (No, It Is Not About Internal Covariate Shift). *ArXiv*, abs/1805.11604.
- Shen, Z.; Cui, P.; Kuang, K.; Li, B.; and Chen, P. 2018. Causally regularized learning with agnostic data selection bias. In *Proceedings of the 26th ACM international conference on Multimedia*, 411–419.
- Shen, Z.; Cui, P.; Liu, J.; Zhang, T.; Li, B.; and Chen, Z. 2020a. Stable learning via differentiated variable decorrelation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2185–2193.
- Shen, Z.; Cui, P.; Zhang, T.; and Kunag, K. 2020b. Stable learning via sample reweighting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5692–5699.
- Shen, Z.; Liu, J.; He, Y.; Zhang, X.; Xu, R.; Yu, H.; and Cui, P. 2021. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*.
- Sun, B.; and Saenko, K. 2016. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, 443–450. Springer.
- Tripuraneni, N.; Adlam, B.; and Pennington, J. 2021. Overparameterization Improves Robustness to Covariate Shift in High Dimensions. In *NeurIPS*.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial Discriminative Domain Adaptation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2962–2971.
- Wang, J.; Lan, C.; Liu, C.; Ouyang, Y.; and Qin, T. 2021. Generalizing to Unseen Domains: A Survey on Domain Generalization. *ArXiv*, abs/2103.03097.
- Weiss, K.; Khoshgoftaar, T. M.; and Wang, D. 2016. A survey of transfer learning. *Journal of Big data*, 3(1): 1–40.



- Wilson, G.; and Cook, D. J. 2020. A Survey of Unsupervised Deep Domain Adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11: 1 – 46.
- Xu, R.; Zhang, X.; Shen, Z.; Zhang, T.; and Cui, P. 2022. A Theoretical Analysis on Independence-driven Importance Weighting for Covariate-shift Generalization. In *International Conference on Machine Learning*, 24803–24829. PMLR.
- Yamada, Y.; Lindenbaum, O.; Negahban, S.; and Kluger, Y. 2020. Feature selection using stochastic gates. In *International Conference on Machine Learning*, 10648–10659. PMLR.
- Yang, Y.; and Soatto, S. 2020. FDA: Fourier Domain Adaptation for Semantic Segmentation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4084–4094.
- Young, M. D.; Wakefield, M. J.; Smyth, G. K.; and Oshlack, A. 2009. Gene ontology analysis for RNA-seq: accounting for selection bias. *Genome Biology*, 11: R14 – R14.
- Zhang, X.; Cui, P.; Xu, R.; Zhou, L.; He, Y.; and Shen, Z. 2021. Deep stable learning for out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5372–5382.
- Zhang, X.; Zhou, L.; Xu, R.; Cui, P.; Shen, Z.; and Liu, H. 2022a. NICO++: Towards Better Benchmarking for Domain Generalization. *ArXiv*, abs/2204.08040.
- Zhang, X.; Zhou, L.; Xu, R.; Cui, P.; Shen, Z.; and Liu, H. 2022b. Towards Unsupervised Domain Generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4910–4920.
- Zhao, P.; and Yu, B. 2006. On model selection consistency of Lasso. *The Journal of Machine Learning Research*, 7: 2541–2563.
- Zhou, X.; Lin, Y.; Pi, R.; Zhang, W.; Xu, R.; Cui, P.; and Zhang, T. 2022. Model Agnostic Sample Reweighting for Out-of-Distribution Learning. In *International Conference on Machine Learning*, 27203–27221. PMLR.