

Stable Adversarial Learning under Distributional Shifts

Jiashuo Liu¹, Zheyang Shen¹, Peng Cui¹, Linjun Zhou¹, Kun Kuang², Bo Li¹, Yishi Lin³

¹Tsinghua University

²Zhejiang University

³Tencent

liujiashuo77@gmail.com, shenzy17@mails.tsinghua.edu.cn, cuip@tsinghua.edu.cn, zhoulj16@mails.tsinghua.edu.cn, kunkuang@zju.edu.cn, libo@sem.tsinghua.edu.cn, yishilin14@gmail.com

Abstract

Machine learning algorithms with empirical risk minimization are vulnerable under distributional shifts due to the greedy adoption of all the correlations found in training data. Recently, there are robust learning methods aiming at this problem by minimizing the worst-case risk over an uncertainty set. However, they equally treat all covariates to form the decision sets regardless of the stability of their correlations with the target, resulting in the overwhelmingly large set and low confidence of the learner. In this paper, we propose Stable Adversarial Learning (SAL) algorithm that leverages heterogeneous data sources to construct a more practical uncertainty set and conduct differentiated robustness optimization, where covariates are differentiated according to the stability of their correlations with the target. We theoretically show that our method is tractable for stochastic gradient-based optimization and provide the performance guarantees for our method. Empirical studies on both simulation and real datasets validate the effectiveness of our method in terms of uniformly good performance across unknown distributional shifts.

Introduction

Traditional machine learning algorithms which optimize the average loss often suffer from the poor generalization performance under distributional shifts induced by latent heterogeneity, unobserved confounders or selection biases in training data (Daume and Marcu 2006; Torralba and Efros 2011; Kuang et al. 2018; Shen et al. 2019). However, in high-stake applications such as medical diagnosis (Kukar 2003), criminal justice (Berk et al. 2018; Rudin and Ustun 2018) and autonomous driving (Huval et al. 2015), it is critical for the learning algorithms to ensure the robustness against potential unseen data. Therefore, robust learning methods have recently aroused much attention due to its favorable property of robustness guarantee (Ben-Tal and Nemirovski 1998; Goodfellow, Shlens, and Szegedy 2014; Madry et al. 2017).

Instead of optimizing the empirical cost on training data, robust learning methods seek to optimize the worst-case cost over an uncertainty set and can be further separated into two main branches named adversarially and distributionally robust learning. In adversarially robust learning, the uncertainty set is constructed point-wisely (Goodfellow, Shlens,

and Szegedy 2014; Papernot et al. 2016; Madry et al. 2017; Ye and Zhu 2018). Adversarial attack is performed independently on each data point within a L_2 or L_∞ norm ball around itself. In distributionally robust learning, on the other hand, the uncertainty set is characterized on a distributional level (Sinha, Namkoong, and Duchi 2018; Esfahani and Kuhn 2018; Duchi and Namkoong 2018). A joint perturbation, typically measured by Wasserstein distance or f -divergence, is applied to the entire distribution entailed by training data. These methods can provide robustness guarantees under distributional shifts when testing distribution is captured in the uncertainty set. However, in real scenarios, to contain the true distribution, the uncertainty set is often overwhelmingly large, which is also referred to as the over-pessimism or the low confidence problem (Frogner et al. 2019; Sagawa et al. 2019). Specifically, with an overwhelmingly large set, the learner optimizes for implausible worst-case scenarios, resulting in meaningless results (e.g. the classifier assigns equal probability to all classes). Such a problem greatly hurts the generalization ability of robust learning methods in practice.

The essential problem of the above methods lies in the construction of the uncertainty set. To address the over-pessimism of the learning algorithm, one should form a more practical uncertainty set which is likely to contain the potential distributional shifts in the future. More specifically, in real applications we observe that different covariates may be perturbed in a non-uniform way, which should be considered in building a practical uncertainty set. Taking the problem of waterbirds and landbirds classification as an example (Wah et al. 2011). There exist two types of covariates where the stable covariates (e.g. representing the bird itself) preserve immutable correlations with the target across different environments, while those unstable ones (e.g. representing the background) are likely to change. Therefore, for the example above, the construction of the uncertainty set should mainly focus on the perturbation of those unstable covariates (e.g. background) to generate more practical and meaningful samples. Following such intuition, there are several works (Bhattad et al. 2019; Vaishnavi et al. 2019) based on the adversarial attack which focus on perturbing the color or background of images to improve the adversarial robustness. However, these methods mainly follow a step by step routine where the segmentation is conducted first to sepa-

rate the background from the foreground and cannot theoretically provide robustness guarantees under unknown distributional shifts, which limits their applications on more general settings.

In this paper, we propose the Stable Adversarial Learning (SAL) algorithm to address this problem in a more principled and unified way, which leverages heterogeneous data sources to construct a more practical uncertainty set. Specifically, we adopt the framework of Wasserstein distributionally robust learning (WDRL) and further characterize the uncertainty set to be anisotropic according to the stability of covariates across the multiple environments, which induces stronger adversarial perturbations on unstable covariates than those stable ones. A synergistic algorithm is designed to jointly optimize the covariates differentiating process as well as the adversarial training process of model’s parameters. Compared with traditional robust learning techniques, the proposed method is able to provide robustness under strong distributional shifts while maintaining enough confidence of the learner. Theoretically, we prove that our method constructs a more compact uncertainty set, which as far as we know is the first analysis of the compactness of adversarial sets in WDRL literature. Empirically, the advantages of our SAL algorithm are demonstrated on both synthetic and real-world datasets in terms of uniformly good performance across distributional shifts.

The SAL Method

We first introduce the Wasserstein Distributionally Robust Learning (WDRL) framework which attempts to learn a model with minimal risk against the worst-case distribution in the uncertainty set characterized by Wasserstein distance:

Definition 1 Let $\mathcal{Z} \subset \mathbb{R}^{m+1}$ and $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, given a transportation cost function $c : \mathcal{Z} \times \mathcal{Z} \rightarrow [0, \infty)$, which is nonnegative, lower semi-continuous and satisfies $c(z, z) = 0$, for probability measures P and Q supported on \mathcal{Z} , the Wasserstein distance between P and Q is :

$$W_c(P, Q) = \inf_{M \in \Pi(P, Q)} \mathbb{E}_{(z, z') \sim M} [c(z, z')] \quad (1)$$

where $\Pi(P, Q)$ denotes the couplings with $M(A, \mathcal{Z}) = P(A)$ and $M(\mathcal{Z}, A) = Q(A)$ for measures M on $\mathcal{Z} \times \mathcal{Z}$.

As mentioned above, the uncertainty set built in WDRL is often overwhelmingly large in wild high-dimensional scenarios. To demonstrate this over pessimism problem of WDRL, we design a toy example in to show the necessity to construct a more practical uncertainty set. Indeed, without any prior knowledge or structural assumptions, it is quite difficult to design a practical set for robustness under distributional shifts. Therefore, we consider a more flexible setting with heterogeneous datasets $D^e = \{X^e, Y^e\}$ from multiple training environments $e \in \mathcal{E}_{tr}$. Specifically, each dataset D^e contains examples identically and independently distributed according to some joint distribution P_{XY}^e on $\mathcal{X} \times \mathcal{Y}$. Then we come up with one basic assumption for our problem. Given the observations that in real scenarios, different covariates have different extents of stability, we propose assumption 1.

Assumption 1 There exists a decomposition of all the covariates $X = \{S, V\}$, where S represents the stable covariate set and V represents the unstable one, so that for all environments $e \in \mathcal{E}$, $\mathbb{E}[Y^e | S^e = s, V^e = v] = \mathbb{E}[Y^e | S^e = s] = \mathbb{E}[Y | S = s]$.

Intuitively, assumption 1 indicates that the correlation between stable covariates S and the target Y stays invariant across environments, which is quite similar to the assumption in (Arjovsky et al. 2019; Kuang et al. 2020; Shen et al. 2020). Moreover, assumption 1 also demonstrates that the influence of V on the target Y can be wiped out as long as whole information of S is accessible. Under the assumption 1, the disparity among covariates revealed in the heterogeneous datasets can be leveraged for better construction of the uncertainty set.

Here we propose the Stable Adversarial Learning (SAL) algorithm, which leverages heterogeneous data to build a more practical uncertainty set with covariates differentiated according to their stability. The objective function of our SAL algorithm is:

$$\min_{\theta \in \Theta} \sup_{Q: W_{c_w}(Q, P_0) \leq \rho} \mathbb{E}_{X, Y \sim Q} [\ell(\theta; X, Y)] \quad (2)$$

where $c_w(z_1, z_2) = \|w \odot (z_1 - z_2)\|_2^2$ and (3)

$$w \in \arg \min_{w \in \mathcal{W}} \left\{ \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} \mathcal{L}^e(\theta) + \alpha \max_{e_p, e_q \in \mathcal{E}_{tr}} \mathcal{L}^{e_p} - \mathcal{L}^{e_q} \right\} \quad (4)$$

where P_0 denotes the training distribution, W_{c_w} denotes the Wasserstein distance with transportation cost function c_w defined as equation 3, $\mathcal{W} = \{w : w \in [1, +\infty)^{m+1} \ \&\& \ \min(w^{(1)}, \dots, w^{(m+1)}) = 1\}$ denotes the covariate weight space ($w^{(i)}$ denotes the i th element of w), and \mathcal{L}^e denotes the average loss in environment $e \in \mathcal{E}_{tr}$, α is a hyper-parameter to adjust the tradeoff between average performance and the stability.

Intuitively, w controls the perturbation level of each covariate and formulates an anisotropic uncertainty set compared with the conventional WDRL methods. The objective function of w (equation 4) contains two parts: the average loss in training environments as well as the maximum margin, which aims at learning such w that the resulting uncertainty set leads to a learner with uniformly good performance across environments. Equation 2 is the objective function of model’s parameters via distributionally robust learning with the learnable covariate weight w . During training, the covariate weight w and model’s parameters θ are iteratively optimized. Details of the algorithm are delineated below. We first will introduce the optimization of model’s parameter in section , then the transportation cost function learning procedure in section .

Tractable Optimization

In SAL algorithm, the model’s parameters θ and covariate weight w is optimized iteratively. In each iteration, given current w , the objective function for θ is:

$$\min_{\theta \in \Theta} \sup_{Q: W_{c_w}(Q, P_0) \leq \rho} \mathbb{E}_{X, Y \sim Q} [\ell(\theta; X, Y)] \quad (5)$$

The duality results in lemma 1 show that the infinite-dimensional optimization problem (5) can be reformulated as a finite-dimensional convex optimization problem (Esfahani and Kuhn 2018). Besides, inspired by (Sinha, Namkoong, and Duchi 2018), a Lagrangian relaxation is provided for computation efficiency.

Lemma 1 *Let $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ and $\ell : \Theta \times \mathcal{Z} \rightarrow \mathbb{R}$ be continuous. For any distribution Q and any $\rho \geq 0$, let $s_\lambda(\theta; (x, y)) = \sup_{\xi \in \mathcal{Z}} (\ell(\theta; \xi) - \lambda c_w(\xi, (x, y)))$, $\mathcal{P} = \{Q : W_c(Q, P_0) \leq \rho\}$, we have:*

$$\sup_{Q \in \mathcal{P}} \mathbb{E}_Q[\ell(\theta; x, y)] = \inf_{\lambda \geq 0} \{\lambda \rho + \mathbb{E}_{P_0}[s_\lambda]\} \quad (6)$$

and for any $\lambda \geq 0$, we have:

$$\sup_{Q \in \mathcal{P}} \{\mathbb{E}_Q[\ell(\theta; (x, y))] - \lambda W_{c_w}(Q, P_0)\} = \mathbb{E}_{P_0}[s_\lambda] \quad (7)$$

Notice that there exists only the inner supremum in $\mathbb{E}_{P_0}[s_\lambda(\theta; (x, y))]$, which can be seen as a relaxed Lagrangian penalty function of the original objective function (5). Here we give up the prescribed amount ρ of robustness in equation (5) and focus instead on the relaxed Lagrangian penalty function for efficiency in equation (7). The loss function on empirical distribution \hat{P}_N becomes $\frac{1}{N} \sum_{i=1}^N s_\lambda(\theta; (x_i, y_i))$. We adopt adversarial training procedure proposed in (Sinha, Namkoong, and Duchi 2018) to approximate the supremum for s_λ .

Specifically, given predictor x , we adopt gradient ascent to obtain an approximate maximizer \hat{x} of $\{\ell(\theta; (\hat{x}, y)) - \lambda c_w(\hat{x}, x)\}$ and optimize the model's parameter θ using \hat{x} as: $\hat{\mathcal{L}} = \frac{1}{N} \sum_{i=1}^N \ell(\theta; \hat{x}, y)$. In the following parts, we simply use \hat{X}_A to denote $\{\hat{x}\}_N$, which means the set of maximizers for training data $\{x\}_N$. The convergence guarantee for this optimization can be referred to (Sinha, Namkoong, and Duchi 2018).

Learning for Transportation Cost Function

We introduce the learning for transportation cost function c_w in this section. In supervised scenarios, perturbations are typically only added to predictor X and not target Y . Therefore, we simplify $c_w : \mathcal{Z} \times \mathcal{Z} \rightarrow [0, +\infty)$ ($\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$) to be:

$$c_w(z_1, z_2) = c_w(x_1, x_2) + \infty \times \mathbb{I}(y_1 \neq y_2) \quad (8)$$

$$= \|w \odot (x_1 - x_2)\|_2^2 + \infty \times \mathbb{I}(y_1 \neq y_2) \quad (9)$$

and omit 'y-part' in c_w as well as w , that is $w \in [1, +\infty)^m$ in the following parts. Intuitively, w controls the strength of adversary put on each covariate. The higher the weight is, the weaker perturbation is put on the corresponding covariate. Ideally, we hope the covariate weights on stable covariates are extremely high to protect them from being perturbed and to maintain the stable correlations, while weights on unstable covariates are nearly 1 to encourage perturbations for breaking the harmful spurious correlations. With the goal towards uniformly good performance across environments, we come up with the objective function $R(\theta(w))$ for learning w as:

$$R(\theta(w)) = \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} \mathcal{L}^e(\theta(w)) + \alpha \max_{e_p, e_q \in \mathcal{E}_{tr}} (\mathcal{L}^{e_p} - \mathcal{L}^{e_q}) \quad (10)$$

where α is the hyper-parameter. $R(\theta(w))$ contains two parts: the first is the average loss in multiple training environments; the second reflects the max margin among environments, which reflects the stability of $\theta(w)$, since it is easy to prove that $\max_{e_p, e_q \in \mathcal{E}_{tr}} \mathcal{L}^{e_p}(\theta(w)) - \mathcal{L}^{e_q}(\theta(w)) = 0$ if

and only if the errors among all training environments are same. Here α is used to adjust the tradeoff between average performance and stability.

In order to optimize w , $\partial R(\theta(w))/\partial w$ can be approximated as following.

$$\frac{\partial R(\theta(w))}{\partial w} = \frac{\partial R}{\partial \theta} \frac{\partial \theta}{\partial X_A} \frac{\partial X_A}{\partial w} \quad (11)$$

Note that the first term $\partial R/\partial \theta$ can be calculated easily. The second term can be approximated during the gradient descent process of θ as :

$$\frac{\partial \theta}{\partial X_A} \approx -\epsilon \sum_t \frac{\nabla_{\theta} \hat{\mathcal{L}}(\theta^t; X_A, Y)}{\partial X_A} \quad (12)$$

where $\frac{\nabla_{\theta} \hat{\mathcal{L}}(\theta^t; X_A, Y)}{\partial X_A}$ can be calculated during the training process. The third term $\partial X_A/\partial w$ can be approximated during the adversarial learning process of X_A as:

$$\frac{\partial X_A}{\partial w} \approx -2\epsilon_x \lambda \sum_t \text{Diag}(X_A^t - X) \quad (13)$$

which can be accumulated during the adversarial training process. Then given current θ , we can update w as:

$$w^{t+1} = \text{Proj}_{\mathcal{W}} \left(w^t - \epsilon_w \frac{\partial R(\theta^t)}{\partial w} \right) \quad (14)$$

where $\text{Proj}_{\mathcal{W}}$ means projecting onto the space \mathcal{W} .

Theoretical Analysis

Here we first provide the robustness guarantee for our method, and then we analyze the rationality of our uncertainty set, which also demonstrates the uncertainty set built in our SAL is more practical. First, we provide the robustness guarantee in theorem 1 with the help of lemma 1 and Rademacher complexity (Bartlett and Mendelson 2002).

Theorem 1 *Let $\Theta = R^m$, $x \in \mathcal{X}$, $y \in \mathcal{Y}$. Assume $|\ell(\theta; z)|$ is bounded by $T_\ell \geq 0$ for all $\theta \in \Theta$, $z = (x, y) \in \mathcal{X} \times \mathcal{Y}$. Let $F : \mathcal{X} \rightarrow \mathcal{Y}$ be a class of prediction functions, then for $\theta \in \Theta$, $\rho \geq 0$, $\lambda \geq 0$, with probability at least $1 - \delta$, for $P \in \{P : W_{c_w}(P, P_0) \leq \rho\}$, we have:*

$$\sup_P \mathbb{E}_P [\ell(\theta; Z)] \leq \lambda \rho + \mathbb{E}_{\hat{P}_n} [s_\lambda(\theta; Z)] + \mathcal{R}_n(\tilde{\ell} \circ F) + k T_\ell \sqrt{\frac{\ln(1/\delta)}{n}} \quad (15)$$

Specially, let $M(\theta; z_0) = \arg \min_{z \in \mathcal{Z}} \{s_\lambda(\theta; z_0)\}$ when

$$\hat{\rho}_n(\theta) = \mathbb{E}_{\hat{P}_n} [c_w(M(\theta; Z), Z)], \text{ for } P \in \{P : W_{c_w}(P, P_0) \leq \hat{\rho}_n(\theta)\},$$

$$\sup_P \mathbb{E}_P [\ell(\theta; Z)] = \sup_P \mathbb{E}_P [\ell(\theta; Z)] + \mathcal{R}_n(\tilde{\ell} \circ F) + k T_\ell \sqrt{\frac{\ln(1/\delta)}{n}} \quad (16)$$

with probability at least $1 - \delta$, where $\tilde{\ell} \circ F = \{(x, y) \mapsto \ell(f(x), y) - \ell(0, y) : f \in F\}$ and \mathcal{R}_n denotes the Rademacher complexity (Bartlett and Mendelson 2002) and k is a numerical constant no less than 0.

Theorem 1 is the standard result on Rademacher complexity as in previous distributionally robust optimization literature. It proves our empirical loss given by our optimization method can control the original worst-case cost of the uncertainty set in SAL.

Then we analyze the rationality of our method in theorem 2, where our major theoretical contribution lies on. As far as we know, it is the first analysis of the compactness of adversary sets in WDRL literature.

Assumption 2 Given $\rho > 0$, $\exists Q_0 \in \mathcal{P}_0$ that satisfies:

- (1) $\forall \epsilon > 0$, $\left| \inf_{M \in \Pi(P_0, Q_0)} \mathbb{E}_{(z_1, z_2 \sim M)} [c(z_1, z_2)] \right| \leq \epsilon$, we refer to the couple minimizing the expectation as M_0 .
- (2) $\mathbb{E}_{M \in \Pi(P_0, Q_0) - M_0} [c(z_1, z_2)] \geq \rho$, where $\Pi(P_0, Q_0) - M_0$ means excluding M_0 from $\Pi(P_0, Q_0)$.
- (3) $Q_{0 \# S} \neq P_{0 \# S}$, where $S = \{i : w^{(i)} > 1\}$ and $w^{(i)}$ denotes the i th element of w and $P_{\# S}$ denotes the marginal distribution on dimensions S .

Assumption 2 describes the boundary property of the original uncertainty set $\mathcal{P}_0 = \{Q : W_c(Q, P_0) \leq \rho\}$, which assumes that there exists at least one distribution on the boundary whose marginal distribution on S is not the same as the center distribution P_0 's and is easily satisfied. Based on this assumption, we come up with the following theorem.

Theorem 2 Under assumption 2, assume the transportation cost function in Wasserstein distance takes form of $c(x_1, x_2) = \|x_1 - x_2\|_1$ or $c(x_1, x_2) = \|x_1 - x_2\|_2^2$. Then, given observed distribution P_0 supported on \mathcal{Z} and $\rho \geq 0$, for the adversary set $\mathcal{P} = \{Q : W_{c_w}(Q, P_0) \leq \rho\}$ and the original $\mathcal{P}_0 = \{Q : W_c(Q, P_0) \leq \rho\}$, given c_w where $\min(w^{(1)}, \dots, w^{(m)}) = 1$ and $\max(w^{(1)}, \dots, w^{(m)}) > 1$, we have $\mathcal{P} \subset \mathcal{P}_0$. Furthermore, for the set $U = \{i | w^{(i)} = 1\}$, $\exists Q_0 \in \mathcal{P}$ that satisfies $W_{c_w}(P_{0 \# U}, Q_{0 \# U}) = \rho$.

Theorem 2 proves that the constructed uncertainty set of our method is smaller than the original. Intuitively, in adversarial learning paradigm, if stable covariates are perturbed, the target should also change correspondingly to maintain the underlying relationship. However, we have no access to the target value corresponding to the perturbed stable covariates in practice, so optimizing under an isotropic uncertainty set (e.g. P_0) which contains perturbations on both stable and unstable covariates would generally lower the confidence of the learner and produce meaningless results. Therefore, from this point of view, by adding high weights on stable covariates in the cost function, we may construct a more reasonable and practical uncertainty set in which the ineffective perturbations are avoided.

Experiments

In this section, we validate the effectiveness of our method on simulation data and real-world data.

Baselines We compare our proposed SAL with the following methods.

- Empirical Risk Minimization(ERM): $\min_{\theta} \mathbb{E}_{P_0} [\ell(\theta; X, Y)]$

- Wasserstein Distributionally Robust Learning(WDRL): $\min_{\theta} \sup_{Q \in \mathcal{W}(Q, P_0) \leq \rho} \mathbb{E}_Q [\ell(\theta; X, Y)]$
- Invariant Risk Minimization(IRM(Arjovsky et al. 2019)): $\min_{\theta} \sum_{e \in \mathcal{E}} \mathcal{L}^e + \lambda \|\nabla_{w|_{w=1.0}} \mathcal{L}^e(w \cdot \theta)\|^2$

For ERM and WDRL, we simply pool the multiple environments data for training. For fairness, we search the hyper-parameter λ in $\{0.01, 0.1, \dots, 1e0, 1e1, \dots, 1e4\}$ for IRM and the hyper-parameter ρ in $\{1, 5, 10, 20, 50, 80, 100\}$ for WDRL, and select the best hyper-parameter according to the validation performance.

Evaluation Metrics To evaluate the prediction performance, we use Mean_Error defined as $\text{Mean_Error} = \frac{1}{|\mathcal{E}_{te}|} \sum_{e \in \mathcal{E}_{te}} \mathcal{L}^e$ and Std_Error defined as $\text{Std_Error} = \sqrt{\frac{1}{|\mathcal{E}_{te}|-1} \sum_{e \in \mathcal{E}_{te}} (\mathcal{L}^e - \text{Mean_Error})^2}$, which are the mean and standard deviation error across testing environments $e \in \mathcal{E}_{te}$.

Imbalanced Mixture In our experiments, we perform a non-uniform sampling among different environments in training set which follows the natural phenomena that empirical data follow a power-law distribution. It is widely accepted that only a few environments/subgroups are common and the rest majority are rare(Shen et al. 2018; Sagawa et al. 2019, 2020).


Simulation Data

Firstly, we design one toy example to demonstrate the over pessimism problem of conventional WDRL. Then, we design two mechanisms to simulate the varying correlations of unstable covariates across environments, named by selection bias and anti-causal effect.

Toy Example In this setting, we have $Y = 5 * S + S^2 + \epsilon$, $V = \alpha Y + \epsilon$, where the effect of S on Y stays invariant, but the correlation between V and Y , i.e. the parameter α , varies across environments. In training, we generate 180 data points with $\alpha = 1$ for environment 1 and 20 data points with $\alpha = -0.1$ for environment 2. We compared methods for linear regression across testing environments with $\alpha \in \{-2.0, -1.5, \dots, 1.5, 2.0\}$.

We first set the radius for WDRL and SAL to be 20.0, and the results are shown in Figure 1(a). We find the ERM induces high estimation error as it puts high regression coefficient on V . Therefore, it performs poor in terms of prediction error when there are distribution shifts. While WDRL achieves more robust performance than ERM across environments, the prediction error is much higher than the others. Our method SAL achieves not only the smallest prediction error, but also the most robust performance across environments.

Furthermore, we train SAL and WDRL for linear regression with a varying radius $\rho \in \{0.0, 0.01, \dots, 20.0\}$. From the results shown in Figure 1(b), we can see that, with the radius growing larger, the robustness of WDRL becomes better, but meanwhile, its performance maintains poor in terms of high Mean_Error and much worse than ERM ($\rho = 0$). This further verifies the limitation of WDRL with respect to



(a) Testing performance for each environment. (b) Testing performance with respect to radius (c) The learned coefficient value of S and V with respect to radius

Figure 1: Results of the toy example. The left figure shows the testing performance in different environments under fixed radius, where RMSE is root mean square error for the prediction. The middle and right denotes the prediction error and the learned coefficients of WDRL and SAL with respect to radius respectively.

the overwhelmingly-large adversary distribution set. In contrast, SAL achieves not only better prediction performance but also better robustness across environments. The plausible reason for the performance difference between WDRL and SAL can be explained by Figure 1(c). As the radius ρ grows larger, WDRL tends to conservatively estimate small coefficients for both S and V so that the model can produce robust prediction performances over the overwhelmingly-large uncertainty set. Comparatively, as our SAL provides a mechanism to differentiate covariates and focus on the robustness optimization over unstable ones, the learned coefficient of unstable covariate V is gradually decreased to improve robustness, while the coefficient of stable covariate S does not change much to guarantee high prediction accuracy.

Selection Bias In this setting, the correlations between unstable covariates and the target are perturbed through selection bias mechanism. According to assumption 1, we assume $X = [S, V]^T$ and $Y = f(S) + \epsilon$ and $P(Y|S)$ remains invariant across environments while $P(Y|V)$ can arbitrarily change. For simplicity, we select data points according to a certain unstable covariate v_0 .

$$\hat{P}(x) = |r|^{-5*|f(s)-\text{sign}(r)*v_0|} \quad (17)$$

where $|r| > 1$ and $\hat{P}(x)$ denotes the probability of point x to be selected. Intuitively, r eventually controls the strengths and direction of the spurious correlation between v_0 and Y (i.e. if $r > 0$, a data point whose v_0 is close to its y is more probably to be selected.). The larger value of $|r|$ means the stronger spurious correlation between v_0 and Y , and $r \geq 0$ means positive correlation and vice versa. Therefore, here we use r to define different environments.

In training, we generate n data points, where κn points from environment e_1 with a predefined r and $(1-\kappa)n$ points from e_2 with $r = -1.1$. In testing, we generate data points for 10 environments with $r \in [-3, -2, -1.7, \dots, 1.7, 2, 3]$. β is set to 1.0.

We compare our SAL with ERM, IRM and WDRL for Linear Regression. We conduct extensive experiments with different settings on r , n , and κ . In each setting, we carry out the procedure 15 times and report the average results. The results are shown in Table ??.

From the results, we have the following observations and analysis: **ERM** suffers from the distributional shifts in testing and yields poor performance in most of the settings. Compared with ERM, the other three robust learning methods achieve better average performance due to the consideration of robustness during the training process. When the distributional shift becomes serious as r grows, **WDRL** suffers from the overwhelmingly-large distribution set and performs poorly in terms of prediction error, which is consistent with our analysis. **IRM** has stable performances across testing environments, while its average error is higher than SAL, which reveals that IRM may harm the average performance for stability. Compared with other robust learning baselines, our **SAL** achieves nearly perfect performance with respect to average performance and stability, especially the variance of losses across environments close to 0, which reflects the effectiveness of assigning different weights to covariates for constructing the uncertainty set.

Anti-causal Effect Inspired by (Arjovsky et al. 2019), in this setting, we introduce the spurious correlation by using anti-causal relationship from the target Y to the unstable covariates V . In this experiment, we assume $X = [S, V]^T$, and firstly sample S from mixture Gaussian distribution characterized as $\sum_{i=1}^k z_k \mathcal{N}(\mu_i, I)$ and the target $Y = \theta_s^T S + \beta S_1 S_2 S_3 + \mathcal{N}(0, 0.3)$. Then the unstable covariates V are generated by anti-causal effect from Y as

$$V = \theta_v Y + \mathcal{N}(0, \sigma(\mu_i)^2) \quad (18)$$

where $\sigma(\mu_i)$ means the Gaussian noise added to V depends on which component the stable covariates S belong to. Intuitively, in different Gaussian components, the corresponding correlations between V and Y are varying due to the different value of $\sigma(\mu_i)$. The larger the $\sigma(\mu_i)$ is, the weaker correlation between V and Y . We use the mixture weight $Z = [z_1, \dots, z_k]^T$ to define different environments, where different mixture weights represent different overall strength of the effect Y on V .

In this experiment, we set $\beta = 0.1$ and build 10 environments with varying σ and the dimension of S, V , the first three for training and the last seven for testing. The average prediction errors are shown in Table ??, where the first

Scenario 1: varying selection bias rate r ($n = 2000, p = 10, \kappa = 0.95$)						
r	$r = 1.5$		$r = 1.7$		$r = 2.0$	
Methods	Mean_Error	Std_Error	Mean_Error	Std_Error	Mean_Error	Std_Error
ERM	0.484	0.058	0.561	0.124	0.572	0.140
WDRL	0.482	0.044	0.550	0.114	0.532	0.112
IRM	0.475	0.014	0.464	0.015	0.477	0.015
SAL	0.450	0.019	0.449	0.015	0.452	0.017
Scenario 2: varying ratio κ and sample size n ($p = 10, r = 1.7$)						
κ, n	$\kappa = 0.90, n = 500$		$\kappa = 0.90, n = 1000$		$\kappa = 0.975, n = 4000$	
Methods	Mean_Error	Std_Error	Mean_Error	Std_Error	Mean_Error	Std_Error
ERM	0.580	0.103	0.562	0.113	0.555	0.110
WDRL	0.563	0.101	0.527	0.083	0.536	0.108
IRM	0.460	0.014	0.464	0.015	0.459	0.014
SAL	0.454	0.015	0.451	0.015	0.448	0.014

Table 1: Results in selection bias simulation experiments of different methods with varying selection bias r , ratio κ and sample size n of training data, and each result is averaged over ten times runs.

three environments are used for training and the last seven are not captured in training with weaker correlation between V and Y . **ERM** and **IRM** achieve the best training performance with respect to their prediction errors on training environments e_1, e_2, e_3 , while their performances in testing are poor. **WDRL** performs worst due to its over pessimism problem. **SAL** achieves nearly uniformly good performance in training environments as well as the testing ones, which validates the effectiveness of our method and proves the excellent generalization ability of SAL.

Real Data

Regression In this experiment, we use a real-world regression dataset (Kaggle) of house sales prices from King County, USA, which includes the houses sold between May 2014 and May 2015¹. The target variable is the transaction price of the house and each sample contains 17 predictive variables such as the built year of the house, number of bedrooms, and square footage of home, etc. We normalize all the predictive covariates to get rid of the influence by their original scales.

To test the stability of different algorithms, we simulate different environments according to the built year of the house. It is fairly reasonable to assume the correlations between parts of the covariates and the target may vary along time, due to the changing popular style of architecture. Specifically, the houses in this dataset were built between 1900 \sim 2015 and we split the dataset into 6 periods, where each period approximately covers a time span of two decades. In training, we train all methods on the first and second decade where *built year* \in [1900, 1910) and [1910, 1920) respectively and validate on 100 data points sampled from the second period.

From the results shown in figure 2(a), we can find that **SAL** achieves not only the smallest Mean_Error but also the lowest Std_Error compared with baselines. From figure 2(b), we can find that from period 4 and so on, where large distribution shifts occurs, **ERM** performs poorly and has

larger prediction errors. **IRM** performs stably across the first 4 environments but it also fails on the last two, whose distributional shifts are stronger. **WDRL** maintains stable across environments while the mean error is high, which is consistent with our analysis in that WDRL equally perturbs all covariates and sacrifices accuracy for robustness. From figure 2(b), we can find that from period 3 and so on, **SAL** performs better than ERM, IRM and WDRL, especially when distributional shifts are large. In periods 1-2 with slight distributional shift, the SAL method incurs a performance drop compared with IRM and WDRL, while SAL performs much better when larger distributional shifts occur, which is consistent with our intuition that our method sacrifice a little performance in nearly I.I.D. setting for its superior robustness under unknown distribution shifts.

Classification Finally, we validate the effectiveness of our SAL on classification tasks, including an income prediction task and colored MNIST classification task.

Income Prediction In this task we use the Adult dataset (Dua and Graff 2017) which involves predicting personal income levels as above or below \$50,000 per year based on personal details. We split the dataset into 10 environments according to demographic attributes, among which distributional shifts might exist. In training phase, we train all methods on 693 data points from environment 1 and 200 points from the second respectively and validate on 100 points sampled from both. We normalize all the predictive covariates to get rid of the influence by their original scales. In testing phase, we test all methods on the 10 environments and report the mis-classification rate on all environments in figure 3. From the results shown in figure 3, we can find that the **SAL** outperforms baselines on almost all environments except a slight drop on the first. However, our SAL outperforms the others in the rest 8 environments where agnostic distributional shifts occur.


Colored MNIST In this task we build a synthetic binary classification task derived from MNIST. The goal is to predict a binary label assigned to each image based on the digit. We color each image either red or green which spuriously correlates with the label similar to (Arjovsky et al. 2019).

¹<https://www.kaggle.com/c/house-prices-advanced-regression-techniques/data>


Scenario 1: $S \in R^5, V \in R^5$										
e	Training environments			Testing environments						
Methods	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8	e_9	e_{10}
ERM	0.281	0.305	0.341	0.461	0.555	0.636	0.703	0.733	0.765	0.824
IRM	0.287	0.293	0.329	0.345	0.382	0.420	0.444	0.461	0.478	0.504
WDRL	0.282	0.331	0.399	0.599	0.750	0.875	0.983	1.030	1.072	1.165
SAL	0.324	0.329	0.331	0.357	0.380	0.403	0.425	0.435	0.446	0.458

Scenario 2: $S \in R^9, V \in R^1$										
e	Training environments			Testing environments						
Methods	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8	e_9	e_{10}
ERM	0.272	0.278	0.298	0.362	0.411	0.460	0.504	0.526	0.534	0.580
IRM	0.306	0.312	0.325	0.328	0.343	0.358	0.365	0.374	0.377	0.394
WDRL	0.300	0.314	0.332	0.396	0.441	0.483	0.529	0.545	0.555	0.596
SAL	0.290	0.284	0.288	0.287	0.288	0.287	0.290	0.284	0.293	0.294

Table 2: Results of the anti-causal effect experiment. The average prediction errors of 15 runs are reported.



(a) Mean Error and Std. Error.



(b) Prediction error with respect to build year.




Figure 3: Results of the Adult dataset.

Figure 2: Results of the real regression dataset. RMSE refers to the Root Mean Square Error.

The direction of the correlation is reversed in the testing environment, which ruins the method relying on such spurious correlation to predict. Specifically, we generate the color id by flipping the label with probability μ , where $\mu = 1.0$ in the first environment, $\mu = 0.3$ in the second and $\mu = -1.0$ in testing. Furthermore, we induce noisy labels by randomly flipping the label with probability 0.2.

In this experiment, we consider the imbalanced mixture which is a more challenging and practical problem. Specifically, we sample 20000 images from environment 1 and 500 from 2 as training data and 10000 images from environment 3 for testing. For our SAL and WDRL, we conduct a two-stage optimization which firstly uses a three-layer CNN to extract the representation of 128 dimensions as the input covariates. For ERM and IRM, we use the same architecture and do the end-to-end optimization. We select the hyperparameters according to the performance on the validation set sampled from training environments. From the results in Table 3, ERM performs terribly because of the spurious correlations and IRM and WDRL are closed to random guess. Our SAL outperforms all baselines, which shows that our method can handle more complicated data such as vision and lingual data with a feature extractor(e.g. deep neural network).

Algorithm	ERM	WDRL	IRM	SAL	Random
Test Acc	0.085	0.48	0.51	0.57	0.50

Table 3: Results of the colored MNIST experiment. We report the average results of 10 runs.

Conclusion

In this paper, we address a practical problem of overwhelmingly-large uncertainty set in robust learning, which often results in unsatisfactory performance under distributional shifts in real situations. We propose the Stable Adversarial Learning (SAL) algorithm that anisotropically considers each covariate to achieve more realistic robustness. We theoretically show that our method constructs a better uncertainty set. Empirical studies validate the effectiveness of our methods in terms of uniformly good performance across different distributed data. We temporarily focus our method at raw feature level for solid theoretical guarantees, and we leave the extension of combining representation learning into our framework as the future work.

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, 61772304, 61531006, U1611461), Beijing Academy of Artificial Intelligence (BAAI), and a grant from the Institute for Guo Qiang, Tsinghua University. Kun Kuang's research was supported in part by National Natural Science Foundation of China (No. 62006207), National Key Research and Development Program of China (No. 2018AAA0101900), the Fundamental Research Funds for the Central Universities. Bo Li's research was supported by the Tsinghua University Initiative Scientific Research Grant, No. 2019THZWJC11; National Natural Science Foundation of China, No. 71490723 and No. 71432004; Science Foundation of Ministry of Education of China, No. 16JJD630006.

References

- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.
- Bartlett, P. L.; and Mendelson, S. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3(Nov): 463–482.
- Ben-Tal, A.; and Nemirovski, A. 1998. Robust convex optimization. *Mathematics of operations research* 23(4): 769–805.
- Berk, R.; Heidari, H.; Jabbari, S.; Kearns, M.; and Roth, A. 2018. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research* 0049124118782533.
- Bhattad, A.; Chong, M. J.; Liang, K.; Li, B.; and Forsyth, D. A. 2019. Big but Imperceptible Adversarial Perturbations via Semantic Manipulation. *CoRR* abs/1904.06347. URL <http://arxiv.org/abs/1904.06347>.
- Daume, H.; and Marcu, D. 2006. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research* 26(1): 101–126.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. URL <http://archive.ics.uci.edu/ml>.
- Duchi, J.; and Namkoong, H. 2018. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*.
- Esfahani, P. M.; and Kuhn, D. 2018. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* 171(1-2): 115–166.
- Frogner, C.; Claiici, S.; Chien, E.; and Solomon, J. 2019. Incorporating Unlabeled Data into Distributionally Robust Learning. *arXiv preprint arXiv:1912.07729*.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Huval, B.; Wang, T.; Tandon, S.; Kiske, J.; Song, W.; Pazhayampallil, J.; Andriluka, M.; Rajpurkar, P.; Migimatsu, T.; Cheng-Yue, R.; et al. 2015. An empirical evaluation of deep learning on highway driving. *arXiv preprint arXiv:1504.01716*.
- 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.
- Kukar, M. 2003. Transductive reliability estimation for medical diagnosis. *Artificial Intelligence in Medicine* 29(1-2): 81–106.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Papernot, N.; McDaniel, P.; Jha, S.; Fredrikson, M.; Celik, Z. B.; and Swami, A. 2016. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*, 372–387. IEEE.
- Rudin, C.; and Ustun, B. 2018. Optimized scoring systems: Toward trust in machine learning for healthcare and criminal justice. *Interfaces* 48(5): 449–466.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2019. Distributionally Robust Neural Networks for Group Shifts: On the Importance of Regularization for Worst-Case Generalization. *arXiv preprint arXiv:1911.08731*.
- Sagawa, S.; Raghunathan, A.; Koh, P. W.; and Liang, P. 2020. An Investigation of Why Overparameterization Exacerbates Spurious Correlations.
- Shen, Z.; Cui, P.; Kuang, K.; Li, B.; and Chen, P. 2018. Causally Regularized Learning with Agnostic Data Selection Bias. In *2018 ACM Multimedia Conference*.
- Shen, Z.; Cui, P.; Liu, J.; Zhang, T.; Li, B.; and Chen, Z. 2020. 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 Kuang, K. 2019. Stable Learning via Sample Reweighting. *arXiv: Learning*.
- Sinha, A.; Namkoong, H.; and Duchi, J. 2018. Certifying Some Distributional Robustness with Principled Adversarial Training. *International Conference on Learning Representations*.
- Torralba, A.; and Efros, A. A. 2011. Unbiased look at dataset bias 1521–1528.
- Vaishnavi, P.; Cong, T.; Eykholt, K.; Prakash, A.; and Rahmati, A. 2019. Can Attention Masks Improve Adversarial Robustness? *arXiv preprint arXiv:1911.11946*.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.

Ye, N.; and Zhu, Z. 2018. Bayesian adversarial learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 6892–6901. Curran Associates Inc.