

Learning Optimal Features via Partial Invariance

Moulik Choraria^{1,*}, Ibtihal Ferwana¹, Ankur Mani², Lav R. Varshney¹

¹University of Illinois at Urbana-Champaign,

²University of Minnesota Twin Cities

{moulikc2, iferwna2, varshney}@illinois.edu, amani@umn.edu

Abstract

Learning models that are robust to distribution shifts is a key concern in the context of their real-life applicability. Invariant Risk Minimization (IRM) is a popular framework that aims to learn robust models from multiple environments. The success of IRM requires an important assumption: the underlying causal mechanisms/features remain invariant across environments. When not satisfied, we show that IRM can over-constrain the predictor and to remedy this, we propose a relaxation via *partial invariance*. In this work, we theoretically highlight the sub-optimality of IRM and then demonstrate how learning from a partition of training domains can help improve invariant models. Several experiments, conducted both in linear settings as well as with deep neural networks on tasks over both language and image data, allow us to verify our conclusions.

1 Introduction

Standard machine learning models trained using classical Empirical Risk Minimization (ERM) can be expected to generalize well to unseen data drawn from the same distribution as the training data (Vapnik 2013). However, distribution shifts during test time (when data is from different sources or under different conditions) can severely degrade model performance (Beery, V. Horn, and Perona 2018; Lake et al. 2017; Marcus 2018). The errors can often be attributed to the model picking up statistically informative but spurious correlations, which in turn limits their real-life applications since in practice, the use-case almost always differs from the training data. Thus, several lines of research explore alternate learning objectives for training robust models.

One particular line of research stems from the Invariant Causal Prediction framework (Peters, Bühlmann, and Meinshausen 2015), where the goal is to learn causal mechanisms that work well under interventions; our work focuses on the similarly inspired Invariant Risk Minimization (IRM) framework, which aims to learn a predictor that relies only on features that are invariant across all training environments. The underlying motivation for invariance is rooted in its strong links with causality (Pearl 2009), with the intuition being that by invariance can help the model distinguish

the causal features from domain-specific spurious features, which it can then discard for better generalization.

A standard assumption in such invariance-based objectives is that of sufficiency (Ahuja et al. 2020b): there exists a predictor, relying solely on the invariant features, which can achieve Bayes optimal risk in all environments. While fairly general, this assumption may not be satisfied for certain classes of distribution shifts. For instance, consider a prediction task with *concept drift*, wherein the relationship of the ‘causal’ features (features that are causally responsible to the label) with the label changes across training environments. Here, a predictor relying solely on invariant features ends up being over-constrained, since it is incentivized to discard non-invariant but informative causal features that are needed for Bayes optimality. Such situations are ubiquitous in practice, for instance in language tasks in which linguistic features can have different connotations within different communities (Gallacher 2021; Mani, Varshney, and Pentland 2021) or in tasks with distribution shifts across time (Luu et al. 2021). Additionally, even when a sufficient representation exists theoretically, it may not be accessible due to shortcomings in the optimization of the learning objective. However, these factors are seldom accounted for when considering the application of IRM (or other invariant learning objectives) for a given use-case (Peyrard et al. 2021; Adragna et al. 2020). Thus, it is important to develop a characterization for the same.

To address this gap in literature, we present a first study to characterize the behaviour of IRM under explicit concept drifts. Then, we take a step further and propose a relaxation for IRM via the *Partial Invariance* (P-IRM) framework. We find that our framework increases the flexibility of invariant models by allowing learning of features that are locally invariant within a partition of the training environments. This flexibility is accompanied with an inherent trade-off; the cost of finding the right partition, in an information-agnostic setting, grows exponentially in the number of environments. However, for certain classes of problems, including the language tasks alluded to previously, readily available meta-information often allows us to easily infer the ‘optimal’ training partition for a given use-case. Notice that in doing so, we move away from the OoD minimax regime, and instead focus on improving generalization conditioned on availability of this meta-information. In this work, we be-

*corresponding author

gin by presenting a theoretical characterization of IRM under concept shifts. Next, we formally quantify the notion of meta-information and assuming access to it, we theoretically and empirically demonstrate how the notion of Partial Invariance can help improve the performance of invariant models. The rest paper is organized as follows: we begin with a literature review in Sec. 2, and motivate P-IRM and present our main results in Sec. 3. We report our empirical evaluations in Sec. 4 and wrap up with some concluding remarks in Sec. 5.

2 Related Work

Many approaches aim to learn deep invariant feature representations: some focus on domain adaptation by finding a representation whose distribution is invariant across source and target distributions (Ben-David et al. 2010; Zhang, Gong, and Schoelkopf 2015), while others focus on conditional domain-invariance (Gong et al. 2016; Li et al. 2018). However, there is evidence that domain adaption approaches are insufficient when the test distribution may lie outside the convex hull of training distributions (Lee and Raginsky 2018; Duchi, Glynn, and Namkoong 2021; Mohri, Sivek, and Suresh 2019). Other approaches include Bayesian Deep Learning (Neal 1996), which tries to account for model uncertainty during test-time, and Robust Optimization (Ben-Tal, El Ghaoui, and Nemirovski 2009), which aims to generalize well to distributions close to training.

Our work focuses particularly on the IRM framework (Arjovsky et al. 2019), which relates to domain generalization wherein access to the test distribution is not assumed. IRM is rooted in the theory of causality (Schölkopf et al. 2012) and proposes invariance for achieving OoD generalization (Peters, Bühlmann, and Meinshausen 2016; Heinze-Deml, Peters, and Meinshausen 2018). In (Ahuja et al. 2020a), the authors reformulate IRM via a game-theoretic approach, wherein the invariant representation corresponds to the Nash equilibrium of a game. While the IRM framework assumes only the invariance of the conditional expectation of the label given the representation, some follow-ups rely on stronger invariance assumptions (Xie et al. 2021; Mahajan, Tople, and Sharma 2021). As mentioned before, this line of work assumes sufficiency of invariant features whereas we specifically focus on distribution shifts when sufficiency is violated.

Several follow-up works attempt to characterize IRM’s performance under different settings and model assumptions. It has been noted that carefully tuned ERM can often outperform state-of-the-art domain generalization approaches, including IRM, across multiple benchmarks (Gulrajani and Lopez-Paz 2020). The failure of IRM may stem from the gap between the proposed framework and its practical “linear” version (IRMv1), which fails to capture natural invariances (Kamath Pritish and Srebro 2021). Indeed, the authors of (Rosenfeld, Ravikumar, and Risteski 2020) demonstrate that a near-optimal solution to the IRMv1 objective, which matches IRM on training environments, does no better than ERM on environments that differ significantly from training. Following these deficiencies, several works propose alternate objectives for achieving invariance (Krueger et al. 2021;

Bellot and van der Schaar 2020; Jin, Barzilay, and Jaakkola 2020; Ahuja et al. 2021; Shui, Wang, and Gagné 2021).

However, unlike previous works that aim to improve the invariance learning objective, we question whether invariance as a constraint can be improved upon for better performance. To that end, our notion of *partial invariance* generalizes not only IRM, but all similar invariance learning objectives. The use of meta-information for invariant learning has been proposed in (Lin, Zhu, and Cui 2022). However, unlike partitioning, the focus therein is to artificially generate environment membership for samples when not available a priori. Finally, a related idea appears in (Yu et al. 2022), which proposes applying different invariance penalty weights for different domains, but with the goal of addressing data quality variance across domains.

3 Theory

In this section, we present the notion of partial invariance.

Notation: We use upper-case boldface U to denote matrix/tensor/vector valued random variables, and lowercase boldface u to denote scalar valued random variables. We use upper-case U to denote matrices/vectors/tensors and lowercase u to denote scalars.

3.1 Invariant Risk Minimization

The IRM setup assumes access to datasets of the form $D_e := \{X_i^e, y_i^e\}_{i=1}^{n_e}$ collected from multiple training environments $e \in \mathcal{E}_{tr}$. The samples in dataset D_e are i.i.d. from the environment’s joint distribution, $P(\mathbf{X}^e, \mathbf{y}^e)$. The task is to estimate a map $f : \mathcal{X} \rightarrow \mathcal{Y}$ or alternatively, the conditional distribution $P(Y|X)$, so that it performs well across unseen environments $\mathcal{E}_{all} \supset \mathcal{E}_{tr}$. Formally, the IRM framework aims to minimize the Out-of-Distribution (OoD) risk: $R^{OoD}(f) = \max_{e \in \mathcal{E}_{all}} R^e(f)$, where $R^e(f) := E_{\mathbf{X}^e, \mathbf{y}^e}[\ell(f(\mathbf{X}^e), \mathbf{y}^e)]$ is the expected risk in environment e . The predictor f is parametrized as $w \circ \Phi$, wherein $\Phi : \mathcal{X} \rightarrow \mathcal{Z}$ represents the learned representation and $w : \mathcal{Z} \rightarrow \mathcal{Y}$ is a linear predictor over said representation. The IRM learning objective is posed as a constrained optimization problem:

$$\begin{aligned} \min_{\Phi, w} \sum_{e \in \mathcal{E}_{obs}} R^e(w \circ \Phi) & \quad (\text{IRM}) \\ \text{s.t. } w \in \arg \min_{\tilde{w}} R^{e'}(\tilde{w} \circ \Phi) \forall e' \in \mathcal{E}_{tr}. & \quad (1) \end{aligned}$$

To avoid the inner optimization, the minimization constraint is replaced by a more tractable gradient penalty:

$$\begin{aligned} \min_{\Phi, w} \sum_{e \in \mathcal{E}_{obs}} R^e(w \circ \Phi) & \quad (\text{IRMgc}) \\ \text{s.t. } w \in \{\tilde{w} : \|\nabla_w R^{e'}(w \circ \Phi)\| = 0 \forall e' \in \mathcal{E}_{tr}\}, & \quad (2) \end{aligned}$$

where IRMgc is shorthand for the gradient constrained IRM. In practice, this constraint is enforced via a regularizer λ :

$$\min_{\Phi} \sum_{e \in \mathcal{E}_{obs}} R^e(\Phi) + \lambda \|\nabla_{w, w=1.0} R^e(\Phi)\|, \quad (\text{IRMv1})$$

where the implicit overparametrization in having a separate classifier and representation map is removed by fixing

a dummy classifier $w = 1.0$. Thus, Φ becomes the entire invariant predictor and the strictness of the gradient norm penalty, which enforces invariance, is via λ . Note that when $\lambda = \infty$, IRMv1 is equivalent to IRMgc, which in turn is the first order approximation for the true IRM objective.

An intrinsic assumption in the IRM learning setup for proving minimax optimality is the ideal scenario of sufficiency i.e. there exists a Φ that is invariant across all $e \in \mathcal{E}_{tr}$ and is sufficient i.e. $\mathbf{y}^e \perp \mathbf{X}^e \mid \Phi(\mathbf{X}^e) \forall e \in \mathcal{E}_{all}$ (Ahuja et al. 2020b). However if sufficiency is violated for an environment, one would expect the IRM model, which relies solely on invariant features, to be sub-optimal for that environment (compared to a model that utilizes non-invariant features along with invariant ones). Such a situation may arise under *concept drift*, wherein the conditional expectation of the label \mathbf{y}_e given “causal” features may change across environments. Thus, in practice, if an invariant Φ that is also sufficient for environments does not exist for the desired use-case, we expect the performance of IRM (or related frameworks) to degrade. We illustrate this with a simple example. **Example 1** We adapt the generative model from (Arjovsky et al. 2019): the goal is to predict target y using $X = [x_1, x_2, x_3]$, in environment e such that $e \in \mathcal{E}_{all}$ can affect the distribution of X as well as the conditional distribution of y given X via a deterministic map $c(e) : \mathcal{E}_{all} \rightarrow \{-1, 1\}$:

$$\begin{aligned} \mathbf{x}_1 &\leftarrow N(0, \sigma(e)^2), \mathbf{x}_2 \leftarrow N(0, \sigma(e)^2), \\ c(e) &\in \{1, -1\}, \epsilon \sim N(0, \sigma(e)^2) \\ \mathbf{y} &\leftarrow \mathbf{x}_1 + c(e)\mathbf{x}_2 + \epsilon, \epsilon \perp \mathbf{x}_1, \epsilon \perp \mathbf{x}_2 \\ \mathbf{x}_3 &\leftarrow \mathbf{y} + N(0, 1), \sigma(e)^2 \in [0, \sigma_{MAX}^2]. \end{aligned}$$

We estimate \mathbf{y} as $\hat{\mathbf{y}} = \alpha_1\mathbf{x}_1 + \alpha_2\mathbf{x}_2 + \alpha_3\mathbf{x}_3$. Within the IRM framework, the only feasible representation Φ (upto scaling) that yields invariant predictors across all e is $\Phi([\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3]) = [\mathbf{x}_1, 0, 0]$, with corresponding regression coefficients $[1, 0, 0]$. Although this minimizes the OoD error for arbitrary e , it does so by discarding the non-invariant but informative \mathbf{x}_2 . However, if our predictor is privy to some knowledge of $c(e)$, we could first partition the set of training environments \mathcal{E}_{tr} into two partitions, such that environments within a partition have the same $c(e)$ value. Then, applying IRM within each partition yields models with better performance that can exploit \mathbf{x}_2 as an invariant feature in the partition. Note that this partial notion of invariance still retains the ability to discard spurious/non-causal \mathbf{x}_3 . Additionally with partitioning, we can improve generalization if information about $c(e_{unseen})$ is available, by choosing the right model/partition for prediction. Next, we study the conditions under which partitioning can improve upon IRM performance and we refer to this method as P-IRM.

3.2 Model

For our analysis, we consider a simple regression task to succinctly capture our intuition about the conditions under which partitioning is feasible. To begin with, we assume access to the underlying causal features and instead, focus on understanding the nature of the IRM solution set under distribution shifts. In the next part, we extend this analysis to

study learning under partial invariance.

We consider the following generative model: we observe samples (X_i^e, y_i^e) in environment e and the goal is to predict y_i^e from X_i^e . X_i^e 's are samples corresponding to the random variable $\mathbf{X}^e \sim P(\mathbf{X}^e)$, as described below:

$$\mathbf{X}^e = [\mathbf{x}_1^e, \mathbf{x}_2^e, \dots, \mathbf{x}_c^e]^\top \sim P(\mathbf{X}^e),$$

where each \mathbf{x}_i^e denotes an individual feature. To simplify our initial analysis, we assume that the individual features are independent of each other and are normalized i.e. $E[\mathbf{X}^e] = \mathbf{0}$ and $E[\mathbf{X}^e \mathbf{X}^{e\top}] = \mathbf{I} \forall e$. The target \mathbf{y}^e for given \mathbf{X}^e can be characterized as:

$$\begin{aligned} \mathbf{y}^e &= \langle \mathbf{W}^e, \mathbf{X}^e \rangle + \epsilon_y, \\ \mathbf{W}^e &= [w_1^e, w_2^e, \dots, w_c^e] \in \mathbb{R}^c, \epsilon_y \sim N(0, \sigma_y^2(e)). \end{aligned} \quad (3)$$

where weights \mathbf{W}^e encode the conditional distribution of observing label \mathbf{y}^e given \mathbf{X}^e in environment e and are fixed for that environment, and $\langle \cdot, \cdot \rangle$ denotes the standard inner product in \mathbb{R}^c . For a given feature \mathbf{x}_i^e in environment e , the corresponding feature weight w_i^e is independently and uniformly sampled from set A_i for each environment e . Once sampled however, these weights remain fixed for that environment. Additionally $|A_1| = 1$, so that feature weight w_1^e is fixed and thus \mathbf{x}_1^e is invariant for all e :

$$\begin{aligned} \mathbf{W}^e &= [w_1^e, w_2^e, \dots, w_c^e], \\ \text{where } w_i^e &\sim \text{Unif}(\{A_i\}) \forall i \in \{1, 2, \dots, c\}, \\ |A_1| &= 1, |A_i| > 1 \forall i > 1. \end{aligned} \quad (4)$$

We make note of some important aspects. As per our model, \mathbf{x}_1^e is the only *truly invariant feature* since $E[\mathbf{y}^e | \mathbf{x}_1^e] = w_{inv} \cdot \mathbf{x}_1^e$, is fixed for all e , where $A_1 = \{w_{inv}\}$ is a singleton, and w_{inv} denotes the invariant feature weight. Additionally, the cardinality of set $A_i, |A_i|$ defines an implicit notion of the variance of feature x_i , with a higher cardinality indicating that the feature weight is more likely to change across environments and is thus, *less invariant*.

With our generative model in place, we next consider the task of predicting \mathbf{y}^e given \mathbf{X}^e , under the mean squared loss. Recall that the IRM framework considers predictors of the form $w \circ \Phi$, where the transformation Φ extracts a suitable representation and w is the linear predictor acting on that representation. Due to the implicit overparametrization, we fix $w = 1.0$ to a scalar value as proposed in (Arjovsky et al. 2019) and analyze the corresponding IRM solutions with $\Phi \in \mathbb{R}^c$. For simplicity, we ignore finite sample effects and consider the objective in (IRMv1) when $\lambda = \infty$, or equivalently, the gradient penalty constraint in (2) which ideally approximates the true IRM objective. Additionally, we assume the following for training environments \mathcal{E}_{tr} .

Assumption 1 (Sufficiency for IRM). Assume \exists an environment $e \in \mathcal{E}_{tr}$ for which the truly invariant predictor is sufficient, i.e. the corresponding feature weights satisfy $w_1^e = w_{inv}$ and $w_i^e = 0 \forall i \in \{2, \dots, c\}$.

In other words, we assume existence of a training environment in which the invariant predictor that only recovers the invariant feature \mathbf{x}_1^e achieves optimal MSE risk, which is a standard assumption in related literature (Ahuja et al. 2020b).

Lemma 1. *As per above parametrization (with $w = 1.0$), under Assumption 1, the values for Φ that satisfies the IRM solution constraints in (2) is a singleton and the value of the corresponding predictor equates to $\Phi = [w_{inv}, 0, 0 \dots, 0]$, the predictor only recovers the invariant feature.*

The proof of the lemma is included in the Appendix and relies on showing that any predictor which assigns non-zero weights to any of non-invariant features would violate the gradient penalty constraints. More importantly, the previous lemma roughly says that any non-invariant feature will be discarded by the IRM predictor. Note that while this is a desirable property for minimax optimality, we ask whether we can do better given additional contextual information.

We formalize the notion of contextual information explicitly by defining an oracle $\omega(e) = \mathbf{1}[\|W^{e_{ref}} - W^e\|_0 \leq \delta]$, that provides us a notion of distance between environments, from a fixed reference environment e_{ref} . Alternatively, it identifies whether environment e is close to e^{ref} .

Remark 1: The choice of the ℓ_0 metric for the oracle suits our combinatorial setting, since we do make any assumptions on the individual elements in the feature weight sets (i.e. A_i 's).

Next we characterize our objective to utilize this information. Suppose we know that our test environment shares the feature weight with the reference environment for a given feature x_i^e . Then we can define the goal of minimizing the risk w.r.t. to the predictor f , conditioned on this information:

$$R^{cond}(f) = \mathbf{E}_e \text{ s.t. } w_i^{e_{ref}} = w_i^e R^e(f),$$

where the expectation is over the draw of environments as per the uniform sampling. We note that a predictor that accounts for the prior condition (reference feature) will improve performance (i.e. with a lower MSE risk R^{cond}), as compared to the truly invariant predictor in the previous lemma. However, to obtain the required feature as a feasible solution via IRM constraints, we need to first isolate a subset of training environments $\mathcal{E}_{partition} \subseteq \mathcal{E}_{tr}$ such that within this set, w_i^e is invariant and secondly, that we avoid learning the rest of the non-invariant features to avoid feature weight mismatches in unseen environments. It turns out that with access to the oracle and under certain mild conditions, we can ensure exactly that in our uniform distribution shift model. Before stating the result, we require a similar sufficiency assumption for the *partially invariant predictor*.

Assumption 2 (Sufficiency for P-IRM). *Assume \exists an environment $e \in \mathcal{E}_{tr}$ for which the partially invariant predictor is sufficient, i.e. the corresponding feature weights satisfy $w_1^e = w_{inv}$, $w_i^e = w_i^{e_{ref}}$ and $w_j^e = 0 \forall j \in \{2, \dots, c\} \setminus \{i\}$.*

Theorem 1. *Under the model (4), under Assumption 2, with access to oracle $\omega(e) = \mathbf{1}[\|W^{e_{ref}} - W^e\|_0 \leq \delta]$ and $\delta < (c - 2)/2$, isolate $\mathcal{E}_{partition} := \{e \in \mathcal{E}_{tr} | \omega(e) = 1\} \cup \{e_{ref}\} \subseteq \mathcal{E}_{tr}$. Next, let $|A_i| = k$, where A_i is the set corresponding to the feature weight $w_i^{e_{ref}}$ of interest. Then, if the sets $\{A_j\} \forall j \in \{2, \dots, c\} \setminus \{i\}$ satisfy $|A_j| > \alpha k$ for some $\alpha > 1$, we have with probability greater or equal to $(\frac{p}{p+1})^{|\mathcal{E}_{partition}|}$, where $p \geq \frac{(c-1-\delta)\alpha}{\delta}$, the IRM solution over set $\mathcal{E}_{partition}$ will recover the feature of interest $w_i^{e_{ref}}$.*

The proof, available in the Appendix, relies on showing that within the partition that satisfies the oracle condition, the probability of successfully isolating the required feature is high. Then the result follows as a consequence of Lemma 1. In words, the theorem says that if we can identify a partition in which the environments are not too different, then with high probability, the IRM solution will recover features which do not vary too much (i.e. non-invariant but still close to invariant). Note that in case of erroneous partitioning, the solution set allowed by the non-convex penalty becomes harder to characterize due to the presence of other feature weights besides the reference. Nevertheless, if the conditions are such that probability of that happening is sufficiently low, we can safely assume that partitioning will achieve a better expected risk. Additionally, it suggests that P-IRM becomes feasible as the oracle becomes more precise and the feature of interest is close to invariant.

Remark 2: While P-IRM does improve upon the IRM solution, both variants are likely to be outperformed by ERM in this setting. However, we point out that this is a simplified setting wherein access to causal features is assumed. In more general settings when the causal features need to be inferred from complex data, ERM may be susceptible invariance to confounders/anti-causal variables and thus, we require invariance as a means to make the solution robust.

3.3 Partitioning and Partial Invariance

Next, we study P-IRM in a general setup, using previous results to characterize the required number of training environments as in IRM. As before, we assume access to the oracle, ω to identify the partition, i.e. $\mathcal{E}_{partition} \subseteq \mathcal{E}_{tr}$.

Learning Setup: We consider the same causal mechanism for regression task (x^e, y^e) from before. The goal is to find a partition using the oracle such that a feature of interest corresponding to the reference environment, $w_i^{e_{ref}}$ is retained. Note that since we want to retain only the invariant features denoted as $\mathbf{X}_{inv}^e = [x_1^e, x_i^e]$, and discard the non-invariant (or non-partially invariant) features, we encapsulate them into the noise term as $\tilde{\epsilon}_y = \epsilon_y + (\mathbf{X}_{\{1 \dots c\} \setminus \{1, i\}}^e)^\top W_{\{1 \dots c\} \setminus \{1, i\}}^e$. Then, notice that we still have $\tilde{\epsilon}_y \perp \mathbf{X}_{inv}^e$ and that $\mathbf{E}[\tilde{\epsilon}_y] = 0$, due to feature independence and centering assumptions. Next, we consider a realistic learning setup where we observe a scrambled version $\tilde{\mathbf{X}}^e$ of the true causal features \mathbf{X}^e :

$$\begin{aligned} \mathbf{y}^e &= (\mathbf{X}_{inv}^e)^\top W_{inv}^e + \tilde{\epsilon}_y, \quad \tilde{\epsilon}_y \perp \mathbf{X}_{inv}^e, \quad \mathbf{E}[\tilde{\epsilon}_y] = 0 \\ \tilde{\mathbf{X}}^e &= S(\mathbf{X}^e, \mathbf{X}'^e). \end{aligned} \tag{5}$$

Here, $\mathbf{X}^e = [\mathbf{X}_{inv}^e, \mathbf{X}_{\{1 \dots c\} \setminus \{1, i\}}^e] \in \mathbb{R}^c$ denote the causal features with respect to the label, $\mathbf{X}'^e \in \mathbb{R}^q$, $\tilde{\mathbf{X}}^e = S(x^e, d^e) \in \mathbb{R}^d$ with $S \in \mathbb{R}^{d \times (c+q)}$. The variable \mathbf{X}'^e may be arbitrarily correlated with \mathbf{X}_{inv}^e , $\tilde{\epsilon}_y$ or the label \mathbf{y}^e and is intended to represent the spurious correlations in data. However, we require S to be such that $\exists \tilde{S}$ s.t. $\tilde{S}(S(\mathbf{X}^e, \mathbf{X}'^e)) = \mathbf{X}_{inv}^e$ i.e. an inverse map such that the recovery of the desired features is feasible.

Next, we define $\gamma = \frac{1}{k\sqrt{2n}} \exp(-nD(\delta/n \|1/\alpha k\|))$, where as before, δ is the oracle distance parameter, k is the car-

dinality of the set A_i , $|A_i|$, α is as defined in Theorem 1, $n = c - 2$ and $D(m||n)$ denotes KL divergence between $\text{Bern}(m)$ and $\text{Bern}(n)$. Intuitively, γ estimates the lower bound on the probability of sampling an environment under the generative model that satisfies the oracle condition of close distance to the reference environment. Then we have the following sample complexity on the number of required environments.

Theorem 2 (Informal). *Assume we observe $(\tilde{\mathbf{X}}^e, \mathbf{y}^e)$ as per (5), with environments $e \in \mathcal{E}_{tr}$ sampled as per (4) and let $\mathcal{E}_{partition} := \{e \in \mathcal{E}_{tr} | \omega(e) = 1\} \cup e_{ref} \subseteq \mathcal{E}_{tr}$. Let $\Phi \in \mathbb{R}^{d \times d}$ have rank $r > 0$. Then sampling $|\mathcal{E}_{tr}| > \frac{1}{\gamma}(d - r + d/r) \log(1/\epsilon)$ ensures partition cardinality $|\mathcal{E}_{partition}| > d - r + d/r$ with probability $> 1 - \epsilon$. Furthermore, if $e \in \mathcal{E}_{partition}$ lie in linear general position of degree r (Assumption 3 in Appendix), then with probability greater than or equal to $(\frac{p}{p+1})^{|\mathcal{E}_{partition}|}$, where $p \geq \frac{(c-1-\delta)\alpha}{\delta}$, the oracle identifies $\mathcal{E}_{partition}$ such that the predictor $w \circ \Phi$ learnt via IRM within that partition recovers the desired features/weights and corresponding prediction $(\mathbf{X}_{inv}^e)^\top W_{inv}^e$, $\forall e \in \mathcal{E}_{all}$ which satisfy $w_i^e = w_i^{e_{ref}}$.*

The proof along with the formal statement is included in the Appendix and follows from our previous results by applying concentration bounds on the draw of environments, and subsequently using prior generalization results for IRM. In words, Theorem 2 states that if the obtained partition is accurate, is of sufficient cardinality and is sufficiently diverse, then Φ recovers the partially invariant features. However, notice that the required number of environments grows inversely with γ , meaning that we need stronger priors (i.e. sample environments close to the reference) to obtain feasible sample complexities in the number of required environments.

3.4 Partial Invariance in Practice

Next, we state the P-IRM objective more formally. We first assume a distance metric d between environments (known directly or via contextual information). Then, our goal is to identify a subset of training environments $\mathcal{E}_{partition} \subseteq \mathcal{E}_{tr}$ such that its average distance w.r.t. a reference environment e^{ref} roughly satisfies:

$$\frac{1}{|\mathcal{E}_{partition}|} \sum_{e \in \mathcal{E}_{partition}} d(e, e^{ref}) < \frac{1}{|\mathcal{E}_{tr}|} \sum_{e \in \mathcal{E}_{tr}} d(e, e^{ref}).$$

Thus, the predictor is trained on a subset of observed environments. However, discarding environments is not data-efficient and can lead to lower fidelity and worse generalization, especially in high-complexity models. To avoid this, we introduce the notion of *conditional invariance* as an alternative. Formally, consider the set of observed training environments \mathcal{E}_{tr} and a subset corresponding to the partition $\mathcal{E}_{partition}$ (chosen suitably via d), satisfying $\mathcal{E}_{partition} \subseteq \mathcal{E}_{tr}$. We propose the following two variants of P-IRM:

$$\min_{\Phi, w} \sum_{e \in \mathcal{E}_1} R_e(w \circ \Phi) \text{ s.t. } w \in \arg \min_{\tilde{w}} R_{e'}(\tilde{w} \circ \Phi) \forall e' \in \mathcal{E}_2,$$

if $\mathcal{E}_1 = \mathcal{E}_2 = \mathcal{E}_{partition}$, **(P-IRM (Partitioning))**

if $\mathcal{E}_1 = \mathcal{E}_{tr}$ & $\mathcal{E}_2 = \mathcal{E}_{partition}$ **(P-IRM (Conditioning))**

where the empirical risk minimization objective is over environments in \mathcal{E}_1 and the IRM invariance constraint is applied on environments in \mathcal{E}_2 . For P-IRM (Conditioning), note that while the model uses data from all environments, the invariance penalty is applied only to environments within the chosen partition, which mitigates the issue of having fewer data samples. Intuitively, it serves as a relaxation of the IRM objective to allow for partially invariant features. Next, we qualitatively discuss some potential issues in the application of P-IRM. Firstly, fulfilling the requirements as per Theorem 2, for the required worst case number of environments is infeasible. Fortunately, in practice, IRM can pick up the required invariances from just two environments and we expect P-IRM to overcome that issue as well.

Next, we revisit the idea of the distance oracle. While a precise characterization of the distance between causal features of different domains is essentially unobtainable in practice, certain situations allow for inferring the nature of the distribution shift via available contextual information which, while often discarded by practitioners, can serve as an effective pseudo-metric for the same. For instance, authors of (Luu et al. 2021) pointed out that temporal misalignments of distributions in language tasks leads to performance degradation, noting that degradation increases with an increase in the time duration between test and train environments. Thus, learning from only the recent past could yield a larger and more relevant set of invariant features for a use-case on future data.

4 Experiments

We start with a basic sanity check via a synthetic experiment, as an extension of the example presented earlier to visualize how IRM can end up suppressing non-invariant causal features, leading to performance degradation. We then evaluate the efficacy of the P-IRM framework (both partitioning or conditioning) on four tasks: a regression task for housing price prediction, an image classification task on the MetaShift dataset (Liang and Zou 2022), an entity recognition task for scientific texts on the SciERC dataset (Luan et al. 2018) dataset, and a text classification task for prediction of venues of scientific papers. Within image classification, we consider two sub-tasks: Domain Generalization and Sub-population shifts. We defer the synthetic experiment on IRM, along with the text classification and Sub-population shift tasks to the Appendix.

For baselines besides IRM, we evaluate the results for standard ERM as well Information Bottleneck IRM (IB-IRM) (Ahuja et al. 2021). In addition, we include experiments in the image and language tasks to empirically characterize the effect of partitioning on ERM and IB-IRM, which we dub as P-ERM and P-IB-IRM respectively.

An underlying thread for our experiments is the availability to meta-information that allows us to estimate a notion of distance or similarity between environments, which P-IRM can then exploit to construct the required partitions. Specifically, in both housing price prediction and entity recognition task, our environments are partitioned across time and due to distribution shifts, we expect environments closer in time to have higher similarity. Similarly in MetaShift, meta-labels

for each image is made available within the data-set, that allows an explicit notion of the distance between training and testing environments. In all our experiments we employ the train-domain validation strategy (Gulrajani and Lopez-Paz 2020) for hyper-parameter tuning. The code is available at <https://github.com/IbtihalFerwana/pirm> and other implementation details are deferred to Appendix.

4.1 Linear Regression

We consider a regression task to predict house prices based on house features¹, built across years [1910-2010]. Each data point consists of 79 predictive features (for instance, number of bedrooms or house area) and a corresponding target, which is the house price. As pre-processing, we drop all non-numerical features, samples with missing values and normalize each feature and price labels to zero mean, unit variance with the samples, $\{X_i, y_i\}_i \in (\mathbb{R}^{32} \times \mathbb{R})$.

Experiment Setup To adapt this task to OoD prediction, following (Lin, Zhu, and Cui 2022), we manually split the training data-set into 10-year segments and use the house year built as a meta-data for partitioning, with the intuition being that factors affecting house prices change over time.

For prediction, we consider a linear regression model for the task. Since the IRM framework $w \circ \Phi$ is inherently over-parametrized, we fix $w = 1.0 \in \mathbb{R}$ and we consider $\Phi \in \mathbb{R}^{32}$ (prediction $(\Phi^\top X)$) with the Adam optimizer (Kingma and Ba 2015). We consider 6 training environments corresponding to years [1910-1970], while the test samples draw from 4 OoD environments [1970-2010]. We expect partitions closer to the test set to yield better predictors.

Results We report the test MSE error (both average and worst group) over the set of testing OoD environments, averaged over 5 random seeds in Table 1. We find that P-IRM significantly improves the average and worst group OoD error over IRM. Partitioning also benefits ERM, showing more evidence of a distribution shift over time, evidence presented in the Appendix. Finally, note that for the two variants for P-IRM, partitioning performs much better where we have more samples than parameters.

Model	Training	Avg. MSE	Worst Group MSE
ERM	1910-1970	0.475 (0.000)	1.037 (0.000)
ERM	1930-1970	0.431 (0.000)	0.963 (0.000)
IRM	1910-1970	0.522 (0.015)	1.129 (0.038)
P-IRM (partitioned)	1930-1970	0.427 (0.009)	0.873 (0.024)
P-IRM (conditioned)	1930-1970	0.490 (0.014)	1.035 (0.034)

Table 1: House Prices Shifts, partitioning demonstrates improvement for both ERM and IRM, test set is 4 OoD environments consisting of houses built between 1970-2010.

4.2 Image Classification

We evaluate P-IRM on a binary image classification task on the MetaShift dataset (Liang and Zou 2022).

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

Dataset In MetaShift dataset, each image is associated with a set of tags that describe the image context (e.g., cat on a rug, cat beside a chair). Thus, for each given tag (e.g. rug, chair), there is an associated set of images and these sets can overlap if an image has multiple tags. This structure naturally induces a graph, with each image context C_i denotes a node (or *community*) in the graph. This graph is weighted and the weights between nodes is determined by the number of images that are shared between the communities. The weights between each pair of communities, C_i and C_j , estimate the similarity between two communities and are calculated using the Szymkiewicz-Simpson coefficient, which yields the corresponding adjacency matrix G :

$$G(i, j) = \frac{|C_i \cap C_j|}{\min(|C_i|, |C_j|)} \quad (6)$$

Having access to such an undirected weighted graph over sets of images thus allows us to derive an implicit notion of distance between the corresponding communities.

Notion of Distance To introduce partitioning, we develop a notion of distance, which then allows us to quantify the relatedness between training and testing environments. These environments are assumed to be sets of communities. To estimate the distance d between any two given nodes/communities, given that our data is structured as a weighted graph, we can make use of the *spectral embeddings* (Belkin and Niyogi 2001). Spectral embeddings are based on graph Laplacian connectivity (Ng, Jordan, and Weiss 2001). The graph Laplacian L is calculated by $L = D_{diag} - G$, where D_{diag} is a diagonal degree matrix of the graph G . The corresponding eigenvectors of L , u_1, \dots, u_k , computed and normalized to form the matrix U , are the corresponding embeddings for the graph. Once we calculate the spectral embeddings, we measure d between communities as the euclidean distance between the corresponding spectral embeddings of each community node. With our notion of distance, we can partition the graph based on distances between sets of communities and identify subsets of training communities which are closer to the test environment.

Experiment Setup For all our experiments, we consider the same set of training communities as in (Liang and Zou 2022), which are split into two environments in the IRM setting. To introduce partitioning, we assume distances d between the training environments and the test communities is known/can be estimated via the meta-labels. For learning the P-IRM model, we consider the training environment for IRM which is closer to the test set on average, and split it into two sub-environments. Note that under this split, P-IRM has access to roughly only half the training samples compared to IRM. To remedy this, we consider additional data splits wherein we add samples from communities in the other IRM training environment, that are close to the test set. These additional samples amount to a percentage p of samples in that environment, allowing P-IRM access to a slightly larger portion of the training set. Following (Liang and Zou 2022), we fix the test community to be *dog(shelf)* and vary distance d between dog train vs test communities. The cat training set remains unchanged.

Results For all experiments, we report the binary classification accuracy averaged over 3 seeds, with the random-

	$d = 0.17$	$d = 0.54$	$d = 0.81$	$d = 0.92$	Avg. Performance
ERM	0.777(0.078)	0.560(0.179)	0.493(0.119)	0.667(0.114)	0.62425
P-ERM ($p = 0$)	0.823(0.045)	0.790(0.086)	0.387(0.074)	0.663(0.192)	0.66575
P-ERM ($p = 10$)	0.820(0.098)	0.770(0.057)	0.493(0.141)	0.663(0.128)	0.6865
IRM	0.757(0.231)	0.477(0.172)	0.757(0.110)	0.687(0.309)	0.6695
P-IRM ($p = 0$)	0.960(0.050)	0.817(0.045)	0.487(0.083)	0.650(0.142)	0.7285
P-IRM ($p = 10$)	0.710(0.107)	0.813(0.147)	0.727(0.087)	0.690(0.184)	0.735
IB_IRM	0.647(0.197)	0.740(0.171)	0.750(0.155)	0.303(0.241)	0.61
P-IB_IRM ($p = 0$)	0.663(0.242)	0.643(0.137)	0.437(0.289)	0.617(0.059)	0.59
P-IB_IRM ($p = 10$)	0.690(0.340)	0.790(0.070)	0.377(0.214)	0.837(0.160)	0.6735

Table 2: Domaing Generalization in Metashift. Training environments are d away from testing community $dog(shelf)$, with additional samples up to percentage $p \in \{0, 10\}$ for partitioned models. Results for $p = 25$ are in Table 5 (in Appendix)

ness solely arising from the learning algorithm. We compare the performance of P-IRM against IRM, as well other benchmarks and their corresponding partitioned versions in table 2. In most experiments, especially those with higher deviation between the training and testing data, models with partitioning tend to perform better.

4.3 Named Entity Recognition (NER)

Distributional shifts are common in language tasks, given that societal changes are known to influence language usage over time. These changes are also reflected in word embeddings (words vectors to represent language) (Garg et al. 2018). Within this context, we explore effects of partitioning (Lazaridou et al. 2021; Luu et al. 2021).

Experiment Setup We consider the SciERC (Luan et al. 2018) dataset, which consists of CS publications from 1980 to 2016. The specific task is *Named Entity Recognition*, a multi-class classification task, that labels each scientific mention in a sentence into six possible categories (*Task, Method, Evaluation Metric, Material, Other-Scientific-Term, or Generic*). The training set comprises of years from 1980-2009 and we test the model on data obtained between 2010-2016, with an intention to study distribution shift over time. For creating the training environments, we split training years into smaller intervals, 1990-2009, 2000-2009 and 2005-2009, such that each interval has roughly the same number of samples. For partitioning, we consider contiguous partitions of time intervals, based on the intuition that vocabularies in text have higher overlap when closer in time (Gururangan et al. 2020). For building the model, we train a classifier over the BERT pretrained language model (Devlin et al. 2019). Due to high sample complexity, we also consider the conditioned P-IRM method that makes use of all training environments.

Results We report the classification accuracy, averaged over 3 seeds in table 3. We find that both variants of P-IRM indeed improve performance over IRM. Additionally, we find that leveraging more training data using conditioned P-IRM leads to marginally better predictors, when compared against standard partitioning. Comparisons against IB_IRM as well as ERM demonstrate that partitioning can improve efficacy of other learning algorithms as well.

Model	# envs	Training	Accuracy (2010-2016)
ERM	4	1980-2009	0.800 (0.012)
P-ERM	3	1990-2009	0.804 (0.020)
P-ERM	2	2000-2009	0.804 (0.016)
IRM	4	1980-2009	0.795 (0.005)
P-IRM (partitioned)	3	1990-2009	0.795 (0.017)
P-IRM (partitioned)	2	2000-2009	0.807 (0.005)
P-IRM (conditioned)	3	1990-2009	0.812 (0.008)
P-IRM (conditioned)	2	2000-2009	0.807 (0.015)
IB_IRM	4	1980-2009	0.800 (0.010)
P-IB_IRM (partitioned)	3	1990-2009	0.800 (0.015)
P-IB_IRM (partitioned)	2	2000-2009	0.794 (0.015)
P-IB_IRM (conditioned)	3	1990-2009	0.807 (0.008)
P-IB_IRM (conditioned)	2	2000-2009	0.805 (0.020)

Table 3: Language Shifts in SciERC dataset. Partitioning improves performance, with (1990-2009) consistently optimal across all learning algorithms.

5 Discussion

In this work, we propose partial invariance, as a relaxation of IRM objective, which allows us to explore a subtle trade-off in invariant models, namely accessing more domains at the cost of a smaller permissible invariant feature set. We then verify, with experiments across multiple domains, that when feasible, partitioning can indeed improve upon IRM as well as other learning frameworks.

We note that the proposed framework is naturally limited by the available information about training/deployment domains. While distribution shifts across time allows for partitions to be contiguous time intervals, finding appropriate partitions is non-trivial under complex shift topologies. In that sense, our work is the first step towards understanding the need for training domain selection in invariant learning. Thus, developing general heuristics for identifying the right partition is an important direction of future work. Second, we note that the conditional variant of P-IRM provides tangible gains in low data regimes, and it is of interest to study the nature of the accessible feature set as well as the associated sample complexities.

References

- Adragna, R.; Creager, E.; Madras, D.; and Zemel, R. S. 2020. Fairness and Robustness in Invariant Learning: A Case Study in Toxicity Classification. *CoRR*, abs/2011.06485.
- Ahuja, K.; Caballero, E.; Zhang, D.; Gagnon-Audet, J.-C.; Bengio, Y.; Mitliagkas, I.; and Rish, I. 2021. Invariance Principle Meets Information Bottleneck for Out-of-Distribution Generalization. In *Advances in Neural Information Processing Systems*.
- Ahuja, K.; Shanmugam, K.; Varshney, K. R.; and Dhurandhar, A. 2020a. Invariant Risk Minimization Games. In *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*, volume 119, 145–155. PMLR.
- Ahuja, K.; Wang, J.; Dhurandhar, A.; Shanmugam, K.; and Varshney, K. R. 2020b. Empirical or Invariant Risk Minimization? A Sample Complexity Perspective. In *Proceeding of the 8th International Conference on Learning Representations (ICLR'20)*.
- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant Risk Minimization. arXiv:1907.02893 [stat.ML].
- Beery, S.; V. Horn, G.; and Perona, P. 2018. Recognition in Terra Incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 456–473.
- Belkin, M.; and Niyogi, P. 2001. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in neural information processing systems*, 14.
- Bellot, A.; and van der Schaar, M. 2020. Accounting for Unobserved Confounding in Domain Generalization. arXiv:2007.10653 [stat.ML].
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. 2010. A theory of learning from different domains. *Machine Learning*, 79: 151–175.
- Ben-Tal, A.; El Ghaoui, L.; and Nemirovski, A. 2009. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Duchi, J.; Glynn, P.; and Namkoong, H. 2021. Statistics of Robust Optimization: A Generalized Empirical Likelihood Approach. *Mathematics of Operations Research*, 46(3).
- Gallacher, J. D. 2021. Leveraging cross-platform data to improve automated hate speech detection. arXiv:2102.04895 [CS.CL].
- Garg, N.; Schiebinger, L.; Jurafsky, D.; and Zou, J. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16): E3635–E3644.
- Gong, M.; Zhang, K.; Liu, T.; Tao, D.; Glymour, C.; and Schölkopf, B. 2016. Domain Adaptation with Conditional Transferable Components. In *Proceedings of the 33rd International Conference on Machine Learning (ICML'16)*, volume 48, 2839–2848.
- Gulrajani, I.; and Lopez-Paz, D. 2020. In Search of Lost Domain Generalization. In *Proceeding of the 8th International Conference on Learning Representations (ICLR'20)*.
- Gulrajani, I.; and Lopez-Paz, D. 2020. In Search of Lost Domain Generalization. *CoRR*, abs/2007.01434.
- Gururangan, S.; Marasović, A.; Swayamdipta, S.; Lo, K.; Beltagy, I.; Downey, D.; and Smith, N. A. 2020. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8342–8360.
- Heinze-Deml, C.; Peters, J.; and Meinshausen, N. 2018. Invariant Causal Prediction for Nonlinear Models. *Journal of Causal Inference*, 6(2).
- Jin, W.; Barzilay, R.; and Jaakkola, T. S. 2020. Domain Extrapolation via Regret Minimization. *CoRR*, abs/2006.03908.
- Kamath Pritish, D. S., Akilesh Tangella; and Srebro, N. 2021. Does Invariant Risk Minimization Capture Invariance? In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 4069–4077. PMLR.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Krueger, D.; Caballero, E.; Jacobsen, J.-H.; Zhang, A.; Binas, J.; Zhang, D.; Le Priol, R.; and Courville, A. 2021. Out-of-Distribution Generalization via Risk Extrapolation. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*, 5815–5826. PMLR.
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40: e253.
- Lazaridou, A.; Kuncoro, A.; Gribovskaya, E.; Agrawal, D.; Liska, A.; Terzi, T.; Gimenez, M.; de Masson d'Autume, C.; Ruder, S.; Yogatama, D.; Cao, K.; Kociský, T.; Young, S.; and Blunsom, P. 2021. Pitfalls of Static Language Modelling. arXiv:2102.01951 [cs.CL].
- Lee, J.; and Raginsky, M. 2018. Minimax Statistical Learning with Wasserstein Distances. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS'18)*, 2692–2701.
- Li, Y.; Gong, M.; Tian, X.; Liu, T.; and Tao, D. 2018. Domain Generalization via Conditional Invariant Representation. *Proceedings of the 32nd Association for the Advancement of Artificial Intelligence (AAAI'18)*, 31(1).
- Liang, W.; and Zou, J. 2022. Metashift: A dataset of datasets for evaluating contextual distribution shifts and training conflicts. In *International Conference on Learning Representations, ICLR 2022*.
- Lin, Y.; Zhu, S.; and Cui, P. 2022. ZIN: When and How to Learn Invariance by Environment Inference? *arXiv preprint arXiv:2203.05818*.

- Luan, Y.; He, L.; Ostendorf, M.; and Hajishirzi, H. 2018. Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction. In *Proc. Conf. Empirical Methods Natural Language Process. (EMNLP)*.
- Luu, K.; Khashabi, D.; Gururangan, S.; Mandyam, K.; and Smith, N. A. 2021. Time Waits for No One! Analysis and Challenges of Temporal Misalignment. ArXiv preprint arXiv:2111.07408.
- Mahajan, D.; Tople, S.; and Sharma, A. 2021. Domain Generalization using Causal Matching. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*, volume 139, 7313–7324. PMLR.
- Mani, A.; Varshney, L. R.; and Pentland, A. 2021. Quantization Games on Social Networks and Language Evolution. *IEEE Transactions on Signal Processing*, 69: 3922–3934.
- Marcus, G. 2018. Deep Learning: A Critical Appraisal. arXiv:1801.00631 [CS.AI].
- Mohri, M.; Sivek, G.; and Suresh, A. T. 2019. Agnostic Federated Learning. In *Proceedings of the 36th International Conference on Machine Learning (ICML'19)*, volume 97, 4615–4625. PMLR.
- Neal, R. M. 1996. *Bayesian Learning for Neural Networks*. Berlin, Heidelberg: Springer-Verlag. ISBN 0387947248.
- Ng, A.; Jordan, M.; and Weiss, Y. 2001. On spectral clustering: Analysis and an Algorithm. *Advances in Neural Information Processing Systems*, 14.
- Pearl, J. 2009. Causal inference in statistics: An overview. *Statistics Surveys*, 3(none): 96 – 146.
- Peters, J.; Bühlmann, P.; and Meinshausen, N. 2016. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 78(5): 947–1012.
- Peters, J.; Bühlmann, P.; and Meinshausen, N. 2015. Causal inference using invariant prediction: identification and confidence intervals. Preprint.
- Peyrard, M.; Ghotra, S. S.; Josifoski, M.; Agarwal, V.; Patra, B.; Carignan, D.; Kiciman, E.; and West, R. 2021. Invariant Language Modeling. *CoRR*, abs/2110.08413.
- Rosenfeld, E.; Ravikumar, P. K.; and Risteski, A. 2020. The Risks of Invariant Risk Minimization. In *Proceeding of the 8th International Conference on Learning Representations (ICLR'20)*.
- Schölkopf, B.; Janzing, D.; Peters, J.; Sgouritsa, E.; Zhang, K.; and Mooij, J. 2012. On causal and Anticausal Learning. In *Proceedings of the 29th International Conference on Machine Learning (ICML'12)*, 1255–1262.
- Shui, C.; Wang, B.; and Gagné, C. 2021. On the benefits of representation regularization in invariance based domain generalization. *CoRR*, abs/2105.14529.
- Vapnik, V. 2013. *The Nature of Statistical Learning Theory*. Springer Science and Business Media.
- Xie, C.; Ye, H.; Chen, F.; Liu, Y.; Sun, R.; and Li, Z. 2021. Risk Variance Penalization. arXiv:2006.07544 [cs.LG].
- Yu, R.; Zhu, H.; Li, K.; Hong, L.; Zhang, R.; Ye, N.; Huang, S.-L.; and He, X. 2022. Regularization Penalty Optimization for Addressing Data Quality Variance in OoD Algorithms. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8): 8945–8953.
- Zhang, K.; Gong, M.; and Schoelkopf, B. 2015. Multi-Source Domain Adaptation: A Causal View. In *Proceedings of the 29th Association for the Advancement of Artificial Intelligence Conference (AAAI'15)*.