
Environment Inference for Invariant Learning

Elliot Creager^{1,2} Jörn-Henrik Jacobsen^{1,2} Richard Zemel^{1,2}

Abstract

Learning models that gracefully handle distribution shifts is central to research on domain generalization, robust optimization, and fairness. A promising formulation is domain-invariant learning, which identifies the key issue of learning which features are domain-specific versus domain-invariant. An important assumption in this area is that the training examples are partitioned into “domains” or “environments”. Our focus is on the more common setting where such partitions are not provided. We propose EIIL, a general framework for domain-invariant learning that incorporates Environment Inference to directly infer partitions that are maximally informative for downstream Invariant Learning. We show that EIIL outperforms invariant learning methods on the CMNIST benchmark without using environment labels, and significantly outperforms ERM on worst-group performance in the Waterbirds and CivilComments datasets. Finally, we establish connections between EIIL and algorithmic fairness, which enables EIIL to improve accuracy and calibration in a fair prediction problem.

1. Introduction

Machine learning achieves super-human performance on many tasks when the test data is drawn from the same distribution as the training data. However, when the two distributions differ, model performance can severely degrade, even to below-chance predictions (Geirhos et al., 2020). Tiny perturbations can derail classifiers, as shown by adversarial examples (Szegedy et al., 2014) and common image corruptions (Hendrycks & Dietterich, 2019). Even new test sets collected from the same data acquisition pipeline induce distribution shifts that significantly harm performance (Recht et al., 2019; Engstrom et al., 2020). Many approaches have been proposed to overcome the brittleness of supervised

¹University of Toronto ²Vector Institute. Correspondence to: Elliot Creager <creager@cs.toronto.edu>.



(a) **Inferred environment 1** (mostly) landbirds on land, and waterbirds on water
(b) **Inferred environment 2** (mostly) landbirds on water, and waterbirds on land

Figure 1. In the Waterbirds dataset (Sagawa et al., 2020), the two target labels (landbirds and waterbirds) are correlated with their respective typical background habitats (land and water). This spurious correlation causes sub-par performance on the smallest subgroups (e.g. waterbirds on land). Environment Inference for Invariant Learning (EIIL) organizes the training data into two environments that are maximally informative for use by a downstream invariant learner, enabling the use of invariant learning in situations where environment labels are not readily available. By grouping examples where class and background disagree into the same environment, EIIL encourages learning an invariance w.r.t. background features, which improves worst-group test accuracy by 18% relative to standard supervised learning.

learning—e.g. Empirical Risk Minimization (ERM)—in the face of distribution shifts. Robust optimization aims to achieve good performance on any distribution close to the training distribution (Goodfellow et al., 2015; Duchi et al., 2021; Madry et al., 2018). Invariant learning on the other hand tries to go one step further, to generalize to distributions potentially far away from the training distribution.

However, common invariant learning methods typically come at a serious disadvantage: they require datasets to be partitioned into multiple domains or environments.¹ Environment assignments should implicitly define variation the algorithm should become invariant or robust to, but often such environment labels are unavailable at training time, either because they are difficult to obtain or due to privacy

¹We use “domains”, “environments” and “groups”/“subgroups” interchangeably.

limitations. In some cases, relevant side-information or metadata, e.g., human annotations, or device ID used to take a medical image, hospital or department ID, etc., may be abundant, but it remains unclear how best to specify environments based on this information (Srivastava et al., 2020). A similar issue arises in mitigating algorithmic unfairness, where so-called sensitive attributes may be difficult to define in practice (Hanna et al., 2020), or their values may be impossible to collect. We aim to overcome the difficulty of manual environment specification by developing a new method inspired by fairness approaches for unknown group memberships (Kim et al., 2019; Lahoti et al., 2020).

The core idea is to leverage the bias of an ERM-trained reference model to discover useful environment partitions directly from the training data. We derive an environment inference objective that maximizes variability across environments, and is differentiable w.r.t. a distribution over environment assignments. After performing environment inference given a fixed reference classifier, we use the inferred environments to train an invariant learner from scratch.

Our method, **Environment Inference for Invariant Learning (EIIIL)**, discovers environment labels that can then be used to train any off-the-shelf invariant learning algorithm in applications where environment labels are unavailable. This approach can outperform ERM in settings where standard learning tends to focus on spurious features or exhibit performance discrepancies between subgroups of the training data (which need not be specified ahead of time). EIIIL discovers environments capturing spurious correlations hidden in the dataset (see Figure 1), making them readily available for invariant learning. Surprisingly, even when manual specification of environments is available (e.g. the CMNIST benchmark), inferring environments directly from aggregated data may *improve* the quality of invariant learning.

Our main contributions are as follows:

- We propose a general framework for inferring environments from data based on the bias of a reference classifier;
- we provide a theoretical characterization of the dependence on the reference classifier, and when we can expect the method to do well;
- we derive a specific instance of environment inference in this framework using gradients w.r.t. soft environment assignments, which outperforms invariant learning (using environment labels) on the CMNIST benchmark and outperforms ERM on Waterbirds;
- we establish a connection to similar themes in the fairness literature, and show that our method can improve accuracy and calibration in a fair prediction problem.

2. Invariant Learning

This section discusses the problem setting and presents background materials that will be used to formulate our proposed method. Our approach is primarily motivated by recent approaches to learning domain- or environment-invariant representations—which we simply refer to as “invariant learning”—that have been applied to domain adaptation and generalization tasks.

Notation Let \mathcal{X} be the input space, \mathcal{E}^{obs} the set of training environments (a.k.a. “domains”), \mathcal{Y} the target space. Let $x, y, e \sim p^{obs}(x, y, e)$ be observational data, with $x \in \mathcal{X}$, $y \in \mathcal{Y}$, and $e \in \mathcal{E}^{obs}$. \mathcal{H} denotes a representation space, from which a classifier $w \circ \Phi$ (that maps to the logit space of \mathcal{Y} via a linear map w) can be applied. $\Phi : \mathcal{X} \rightarrow \mathcal{H}$ denotes the parameterized mapping or “model” that we optimize. We refer to $\Phi(x) \in \mathcal{H}$ as the “representation” of example x . $\hat{y} \in \mathcal{Y}$ denotes a hard prediction derived from the classifier by stochastic sampling or probability thresholding. $\ell : \mathcal{H} \times \mathcal{Y} \rightarrow \mathbb{R}$ denotes a scalar loss, which guides the learning.

The empirical risk minimization (ERM) solution is found by minimizing the global risk, expressed as the expected loss over the observational distribution:

$$C^{ERM}(\Phi) = \mathbb{E}_{p^{obs}(x,y,e)}[\ell(\Phi(x), y)]. \quad (1)$$

Representation Learning with Environment Labels

Domain generalization is concerned with achieving low error rates on unseen test distributions $p(x, y|e_{test})$ for $e_{test} \notin \mathcal{E}^{obs}$. Domain adaption is a related problem where model parameters can be adapted at test time using unlabeled data. Recently, *Invariant Learning* approaches such as Invariant Risk Minimization (IRM) (Arjovsky et al., 2019) and Risk Extrapolation (REx) (Krueger et al., 2021) were proposed to overcome the limitations of adversarial domain-invariant representation learning (Zhao et al., 2019) by discovering invariant relationships between inputs and targets across domains. Invariance serves as a proxy for causality, as features representing “causes” of target labels rather than effects will generalize well under intervention. In IRM, a representation $\Phi(x)$ is learned that performs optimally within each environment—and is thus invariant to the choice of environment $e \in \mathcal{E}^{obs}$ —with the ultimate goal of generalizing to an unknown test dataset $p(x, y|e_{test})$. Because optimal classifiers under standard loss functions can be realized via a conditional label distribution ($f^*(x) = \mathbb{E}[y|x]$), an invariant representation $\Phi(x)$ must satisfy the following *Environment Invariance Constraint*:

$$\begin{aligned} \mathbb{E}[y|\Phi(x) = h, e_1] &= \mathbb{E}[y|\Phi(x) = h, e_2] \\ \forall h \in \mathcal{H} \quad \forall e_1, e_2 \in \mathcal{E}^{obs}. \end{aligned} \quad (\text{EIC})$$

Intuitively, the representation $\Phi(x)$ encodes features of the input x that induce the same conditional distribution over

labels across each environment. This is closely related to the notion of “group sufficiency” studied in the fairness literature (Liu et al., 2019) (see Appendix C).

Because trivial representations such as mapping all x onto the same value may satisfy environment invariance, other objectives must be introduced to encourage the predictive utility of Φ . Arjovsky et al. (2019) propose IRM as a way to satisfy (EIC) while achieving a good overall risk. As a practical instantiation, the authors introduce IRMv1, a regularized objective enforcing simultaneous optimality of the same classifier $w \circ \Phi$ in all environments;² here w.l.o.g. $w = \bar{w}$ is a constant scalar multiplier of 1.0 for each output dimension. Denoting by $R^e = \mathbb{E}_{p^{obs}(x,y|e)}[\ell]$ the per-environment risk, the objective to be minimized is

$$C^{IRM}(\Phi) = \sum_{e \in \mathcal{E}^{obs}} R^e(\Phi) + \lambda \|\nabla_{\bar{w}} R^e(\bar{w} \circ \Phi)\|. \quad (\text{IRMv1})$$

Robust Optimization Another approach at generalizing beyond the training distribution is robust optimization (Ben-Tal et al., 2009), where one aims to minimize the worst-case loss for every subset of the training set, or other well-defined perturbation sets around the data (Duchi et al., 2021; Madry et al., 2018). Rather than optimizing a notion of invariance, Distributionally Robust Optimization (DRO) (Duchi et al., 2021) seeks good performance for all nearby distributions by minimizing the worst-case loss: $\max_q \mathbb{E}_q[\ell]$ s.t. $D(q||p) < \epsilon$, where D denotes similarity between two distributions (e.g., χ^2 divergence) and ϵ is a hyperparameter. The objective can be computed as an expectation over p via per-example importance weights $\gamma_i = \frac{q(x_i, y_i)}{p(x_i, y_i)}$. GroupDRO operationalizes this principle by sharing importance weights across training examples, using environment labels to define relevant groups for this parameter sharing. This can be expressed as an expected risk under a worst-case distribution over group proportions:

$$C^{GroupDRO}(\Phi) = \max_g \mathbb{E}_{g(e)}[R^e(\Phi)]$$

This is a promising approach towards tackling distribution shift with deep nets (Sagawa et al., 2020), and we show in our experiments how environment inference enables application of GroupDRO to improve over standard learning without requiring group labels.

Limitations of Invariant Learning While the use of invariant learning to tackle domain generalization is still relatively nascent, several known limitations merit discussion. IRM can provably find an invariant predictor that generalizes OOD, but only under restrictive assumptions, such

² $w \circ \Phi$ yields a classification decision via linear weighting on the representation features.

as linearity of the data generative process and access to many environments (Arjovsky et al., 2019). However, most benchmark datasets are in the non-linear regime; Rosenfeld et al. (2021) demonstrated that for some non-linear datasets, the IRMv1 penalty term induces multiple optima, not all of which yield invariant predictors. Nevertheless, IRM has found empirical success in some high dimensional non-linear classification tasks (e.g. CMNIST) using just a few environments (Arjovsky et al., 2019; Koh et al., 2021). On the other hand, it was recently shown that, using careful and fair model selection strategies across a suite of image classification tasks, neither IRM nor other invariant learners consistently beat ERM in OOD generalization (Gulrajani & Lopez-Paz, 2021). This last study underscores the importance of *model selection* in any domain generalization approach, which we discuss further below.

3. Invariance Without Environment Labels

In this section we propose a novel invariant learning framework that does not require a priori domain/environment knowledge. This framework is useful in algorithmic fairness scenarios when demographic makeup is not directly observed; it is also applicable in standard machine learning settings when relevant environment information is either unavailable or not clearly identified. In both cases, a method that sorts training examples \mathcal{D} into environments that maximally separate the spurious features—i.e. inferring populations $\mathcal{D}_1 \cup \mathcal{D}_2 = \mathcal{D}$ —can facilitate effective invariant learning.

3.1. Environment Inference for Invariant Learning

Our aim is to find environments that maximally violate the invariant learning principle. We can then evaluate the quality of these inferred environments by utilizing them in an invariant learning method. Our overall algorithm EIIL is a two-stage process: (1) Environment Inference (EI): infer the environment assignments; and (2) Invariant Learning (IL): run invariant learning given these assignments.

The primary goal of invariant-learning is to find features that are domain-invariant, i.e. that reliably predict the true class regardless of the domain. The EI phase aims to identify domains that help uncover these features. This phase depends on a reference classifier $\tilde{\Phi}$; which maps inputs X to outputs Y , and defines a putative set of invariant features. This model could be found using ERM on $p^{obs}(x, y)$, for example. Environments are then derived that partition the mapping of the reference model which maximally violate the invariance principle, i.e., where for the reference classifier the same feature vector is associated with examples of different classes. While any of the aforementioned invariant learning objectives can be incorporated into the EI phase, the invariance principle or group-sufficiency—as expressed

in (EIC)—is a natural fit, since it explicitly depends on learned feature representations Φ .

To realize an EI phase focused on the invariance principle, we utilize the IRM objective (IRMv1). We begin by noting that the per-environment risk R^e depends implicitly on the manual environment labels from the dataset. For a given environment e' , we denote $\mathbb{1}(e_i = e')$ as an indicator that example i is assigned to that environment, and re-express the per-environment risk as:

$$R^e(\Phi) = \frac{1}{\sum_{i'} \mathbb{1}(e_{i'} = e)} \sum_i \mathbb{1}(e_i = e) \ell(\Phi(x_i), y_i) \quad (2)$$

Now we relax this risk measure to search over the space of environment assignments. We replace the manual assignment indicator $\mathbb{1}(e_i = e')$, with a probability distribution $\mathbf{q}_i(e') := q(e'|x_i, y_i)$, representing a soft assignment of the i -th example to the e' -th environment. To *infer* environments, we optimize $q(e|x_i, y_i)$ so that it captures the worst-case environments for a fixed classifier Φ . This corresponds to maximizing w.r.t. \mathbf{q} the following soft relaxation of the regularizer³ from C^{IRM} :

$$C^{EI}(\Phi, \mathbf{q}) = \|\nabla_{\bar{w}} \tilde{R}^e(\bar{w} \circ \Phi, \mathbf{q})\|, \quad (3)$$

$$\tilde{R}^e(\Phi, \mathbf{q}) = \frac{1}{\sum_{i'} \mathbf{q}_{i'}(e)} \sum_i \mathbf{q}_i(e) \ell(\Phi(x_i), y_i) \quad (4)$$

where \tilde{R}^e represents a soft per-environment risk that can pass gradients to the environment assignments \mathbf{q} . See Algorithm 1 in Appendix A for pseudocode.

To summarize, EIIL involves the following sequential⁴ approach:

1. Input *reference model* $\tilde{\Phi}$;
2. Fix $\Phi \leftarrow \tilde{\Phi}$ and optimize the EI objective to infer environments: $\mathbf{q}^* = \arg \max_{\mathbf{q}} C^{EI}(\tilde{\Phi}, \mathbf{q})$;
3. Fix $\tilde{\mathbf{q}} \leftarrow \mathbf{q}^*$ and optimize the IL objective to yield the new model: $\Phi^* = \arg \min_{\Phi} C^{IL}(\Phi, \tilde{\mathbf{q}})$

In our experiments we consider binary environments and parameterize the \mathbf{q} as a vector of probabilities for each example in the training data.⁵ EIIL is applicable more broadly to

³ We omit the average risk term as we are focused on maximally violating (EIC) regardless of the risk.

⁴ We also tried jointly training Φ and \mathbf{q} using alternating updates, as in GAN training, but did not find empirical benefits. This formulation introduces optimization and conceptual difficulties, e.g. ensuring that invariances apply to all environments discovered throughout learning.

⁵ Note that under this parameterization, when optimizing the inner loop with fixed Φ the number of parameters equals the number of data points (which is small relative to standard neural net training). We leave amortization of q to future work.

any environment-based invariant learning objective through the choice of C^{IL} in Step 3. We present experiments using $C^{IL} \in \{C^{IRM}, C^{GroupDRO}\}$, and leave a more complete exploration to future work.

3.2. Analyzing the Inferred Environments

To characterize the ability of EIIL to generalize to unseen test data, we now examine the inductive bias for generalization provided by the reference model $\tilde{\Phi}$. We state the main result here and defer the proofs to Appendix B. Consider a dataset with some feature(s) z which are spurious, and other(s) v which are valuable/invariant/causal w.r.t. the label y . Our proof considers binary features/labels and two environments, but the same argument extends to other cases. Our goal is to find a model Φ whose representation $\Phi(v, z)$ is invariant w.r.t. z and focuses solely on v .

Proposition 1 *Consider environments that differ in the degree to which the label y agrees with the spurious features z : $\mathbb{P}(\mathbb{1}(y = z)|e_1) \neq \mathbb{P}(\mathbb{1}(y = z)|e_2)$: then a reference model $\tilde{\Phi} = \Phi_{Spurious}$ that is invariant to valuable features v and solely focuses on spurious features z maximally violates the invariance principle (EIC). Likewise, consider the case with fixed representation Φ that focuses on the spurious features: then a choice of environments that maximally violates (EIC) is $e_1 = \{v, z, y|\mathbb{1}(y = z)\}$ and $e_2 = \{v, z, y|\mathbb{1}(y \neq z)\}$.*

If environments are split according to agreement of y and z , then the constraint from (EIC) is satisfied by a representation that ignores z : $\Phi(x) \perp z$. Unfortunately this requires a priori knowledge of either the spurious feature z or a reference model $\tilde{\Phi} = \Phi_{Spurious}$ that extracts it. When the suboptimal solution $\Phi_{Spurious}$ is not a priori known, it will sometimes be recovered directly from the training data; for example in CMNIST we find that Φ_{ERM} approximates Φ_{Color} . This allows EIIL to find environment partitions providing the starkest possible contrast for invariant learning.

Even if environment partitions are available, it may be possible to improve performance by inferring new partitions from scratch. It can be shown (see Appendix B.3) that the environments provided in the CMNIST dataset (Arjovsky et al., 2019) do not maximally violate (EIC) for a reference model $\tilde{\Phi} = \Phi_{Color}$, and are thus not maximally informative for learning to ignore color. Accordingly, EIIL improves test accuracy for IRM compared with the hand-crafted environments (Table 2).

If $\tilde{\Phi} = \Phi_{ERM}$ focuses on a mix of z and v , EIIL may still find environment partitions that enable effective invariant learning, as we find in the Waterbirds dataset, but they are not guaranteed to maximally violate (EIC).

3.3. Binned Environment Invariance

We can derive a heuristic algorithm for EI that maximizes violations of the invariance principle by stratifying examples into discrete bins (i.e. confidence bins for 1-D representations), then sorting them into environments within each bin. This algorithm provides insight into both the EI task and the relationship between the IRMv1 regularizer and the invariance principle. We define bins in the space of the learned representation $\Phi(x)$, indexed by b ; s_{ib} indicates whether example i is in bin b . The intuition behind the algorithm is that a simple approach can separate the examples in a bin to achieve the maximal value of the (EIC).

The degree to which the environment assignments violate (EIC) can be expressed as follows, which can then be approximated in terms of the bins:

$$\begin{aligned} \Delta\text{EIC} &= (E[y|\Phi(x), e_1] - E[y|\Phi(x), e_2])^2 \\ &\approx \sum_b \left(\sum_i s_{ib} y_i \mathbf{q}_i(e=e_1) - \sum_i s_{ib} y_i \mathbf{q}_i(e=e_2) \right)^2 \end{aligned}$$

Inspection of this objective leads to a simple algorithm: assign all the $y = 1$ examples to one environment, and $y = -1$ examples to the other. This results in the expected values of y equal to ± 1 , which achieves the maximum possible value of ΔEIC per bin.⁶

This binning leads to an important insight into the relationship between the IRMv1 regularizer and (EIC). Despite the analysis in Arjovsky et al. (2019), this link is not completely clear (Kamath et al., 2021; Rosenfeld et al., 2021). However, in the situation considered here, with binary classes, we can use this binning approach to show a tight link between the two objectives: finding an environment assignment that maximizes the violation of our softened IRMv1 regularizer (Equation 3) also maximizes the violation of the softened Environment Invariance Constraint (ΔEIC); see Appendix B.2 for the proof. This binning approach highlights the dependence on the reference model, as the bins are defined in its learned Φ space; the reference model also played a key role in the analysis above. We analyze it empirically in Section 5.3.

4. Related Work

Domain adaptation and generalization Beyond the methods discussed above, a variety of recent works have approached the domain generalization problem from the lens of learning invariances in the training data. Adversarial training is a popular approach for learning representations invariant (Zhang et al., 2017; Hoffman et al., 2018; Ganin et al., 2016) or conditionally invariant (Li et al., 2018) to

⁶Multiple global optima exist, this heuristic is not the only possible solution. For very confident reference models, where few confidence bins are populated, this relates to partitioning based on the error cases.

the environment. However, this approach has limitations in settings where distribution shift affects the marginal distribution over labels (Zhao et al., 2019).

Arjovsky et al. (2019) proposed IRM to mitigate the effect of test-time label shift, which was inspired by applications of causal inference to select invariant features (Peters et al., 2016). Krueger et al. (2021) proposed the related Risk Extrapolation (REx) principle, which dictates a stronger preference to exactly equalize $R^e \forall e$ (e.g. by penalizing variance across e as in their practical algorithm V-REx), which is shown to improve generalization in several settings.⁷

Recently, Ahmed et al. (2021) proposed a new invariance regularizer based on matching class-conditioned average predictive distributions across environments, which we note is closely related to the equalized odds criterion commonly used in fair classification (Hardt et al., 2016). Moreover, they deploy this training on top of environments inferred by our EI method, showing that the overall EIIL approach can effectively handle “systematic” generalization (Bahdanau et al., 2019) on a semi-synthetic foreground/background task similar to the Waterbirds dataset that we study.

Several large-scale benchmarks have recently been proposed to highlight difficult open problems in this field, including the use of real-world data (Koh et al., 2021), handling sub-population shift (Santurkar et al., 2021), and model selection (Gulrajani & Lopez-Paz, 2021).

Leveraging a reference classifier A number of methods have recently been proposed that improve performance by exploiting the mistakes of a pre-trained auxiliary model, as we do when inferring environments for the invariant learner using $\tilde{\Phi}$. Nam et al. (2020) jointly train a “biased” model f_B and a “debiased” model f_D , where the relative cross-entropy losses of f_B and f_D on each training example determine their importance weights in the overall training objective for f_D . Sohoni et al. (2020) infer a different set of “hidden subclasses” for each class label $y \in \mathcal{Y}$, Subclasses computed in this way are then used as group labels for training a GroupDRO model, so the overall two-step process corresponds to certain choices of EI and IL objectives.

Liu et al. (2021) and Dagaev et al. (2021) concurrently proposed to compute importance weights for the primary model using an ERM reference, which can be seen as a form of distributionally robust optimization where the worst-case distribution only updates once. Dagaev et al. (2021) use the confidence of the reference model to assign importance weights to each training example. Liu et al. (2021) split the training examples into two disjoint groups based on the errors of ERM, akin to our EI step, with the per-group

⁷Analogous to V-REx, Williamson & Menon (2019) adapt Conditional Variance at Risk (CVaR) (Rockafellar & Uryasev, 2002) to equalize risk across demographic groups.

Environment Inference for Invariant Learning

Statistic to match/optimize	e known?	Dom-Gen method	Fairness method
match $\mathbb{E}[\ell e] \forall e$	yes	REx (Krueger et al., 2021),	CVaR Fairness (Williamson & Menon, 2019)
min $\max_e \mathbb{E}[\ell e]$	yes	Group DRO (Sagawa et al., 2020)	
min $\max_q \mathbb{E}_q[\ell]$	no	DRO (Duchi et al., 2021)	Fairness without Demographics (Hashimoto et al., 2018; Lahoti et al., 2020)
match $\mathbb{E}[y \Phi(x), e] \forall e$	yes	IRM (Arjovsky et al., 2019)	Group Sufficiency (Chouldechova, 2017; Liu et al., 2019)
match $\mathbb{E}[y \Phi(x), e] \forall e$	no	EIIL (ours)	EIIL (ours)
match $\mathbb{E}[\hat{y} \Phi(x), e, y = y'] \forall e$	yes	C-DANN (Li et al., 2018) PGI (Ahmed et al., 2021)	Equalized Odds (Hardt et al., 2016)

Table 1. Domain Generalization (Dom-Gen) and Fairness methods can be understood as matching or optimizing some statistic across conditioning variable e , representing “environment” or “domains” in Dom-Gen and “sensitive” group membership in the Fairness.

importance weights treated as a hyperparameter for model selection (which requires a subgroup-labeled validation set). We note that the implementation of EIIL using binning heuristic, discussed in 3.3, can also realize an error splitting behavior when the reference classifier is very confident. In this case, both methods use the same disjoint groups of training examples towards slightly different ends: we train an invariant learner, whereas Liu et al. (2021) train a cross-entropy classifier with fixed per-group importance weights.

Algorithmic fairness Our work draws inspiration from a rich body of recent work on learning fair classifiers in the absence of demographic labels (Hébert-Johnson et al., 2018; Kearns et al., 2018; Hashimoto et al., 2018; Kim et al., 2019; Lahoti et al., 2020). Generally speaking, these works seek a model that performs well for group assignments that are the worst case according to some fairness criterion. Table 1 enumerates several of these criteria, and draw analogies to domain generalization methods that match or optimize similar statistics.⁸ Environment inference serves a similar purpose for our method, but with a slightly different motivation: rather than learn a fair model in an online way that provides favorable in-distribution predictions, we learn discrete data partitions as an intermediary step, which enables use of invariant learning methods to tackle distribution shift.

Adversarially Reweighted Learning (ARL) (Lahoti et al., 2020) is most closely related to ours, since they emphasize subpopulation shift as a key motivation. Whereas ARL uses a DRO objective that prioritizes stability in the loss space, we explore environment inference to encourage invariance in the learned representation space. We see these as complementary approaches that are each suited to different types of distribution shift, as we discuss in the experiments.

⁸ We refer the interested reader to Appendix C for a more in-depth discussion of the relationships between domain generalization and fairness methods.

5. Experiments

For lack of space we defer a proof-of-concept synthetic regression experiment to Appendix F.1. We proceed by describing the remaining datasets under study in Section 5.1. We then present the main results measuring the ability of EIIL to handle distribution shift in Section 5.2, and offer a more detailed analysis of the EIIL solution and its dependence on the reference model in Section 5.3. See <https://github.com/ecreager/eiil> for code.

Model selection Tuning hyperparameters when train and test distributions differ is a difficult open problem (Krueger et al., 2021; Gulrajani & Lopez-Paz, 2021). Where possible, we reuse effective hyperparameters for IRM and GroupDRO found by previous authors. Because these works allowed limited validation samples for hyperparameter tuning (all baseline methods benefit fairly from this strategy), these results represent an optimistic view on the ability for invariant learning. As discussed above, the choice of reference classifier is of crucial importance when deploying EIIL; if worst-group performance can be measured on a validation set, this could be used to tune the hyperparameters of the reference model (i.e. model selection subsumes reference model selection). See Appendix E for further discussion.

5.1. Datasets

CMNIST CMNIST is a noisy digit recognition task⁹ where color is a spurious feature that correlates with the label at train time but anti-correlates at test time, with the correlation strength at train time varying across environments (Arjovsky et al., 2019). In particular, the two training environments have $\text{Corr}(\text{color}, \text{label}) \in \{0.8, 0.9\}$ while the test environment has $\text{Corr}(\text{color}, \text{label}) = 0.1$. Cru-

⁹MNIST digits are grouped into $\{0, 1, 2, 3, 4\}$ and $\{5, 6, 7, 8, 9\}$ so the CMNIST target label y is binary.

cially, label noise is applied by flipping y with probability $\theta_y = 0.25$. This implies that shape (the invariant feature) is marginally less reliable than color in the training data, so ERM ignores shape to focus on color and suffers from below-chance test performance.

Waterbirds To evaluate whether EIIL can infer useful environments in a more challenging setting with high-dimensional images, we turn to the Waterbirds dataset (Sagawa et al., 2020). Waterbirds is a composite dataset that combines 4,795 bird images from the CUB dataset (Welinder et al., 2010) with background images from the Places dataset (Zhou et al., 2017). It examines the proposition (which frequently motivates invariant learning approaches) that modern networks often learn spurious background features (e.g. green grass in pictures of cows) that are predictive of the label at train time but fail to generalize in new contexts (Beery et al., 2018; Geirhos et al., 2020). The target labels are two classes of birds—“landbirds” and “waterbirds” respectively coming from dry or wet habitats—superimposed on land and water backgrounds. At training time, landbirds and waterbirds are most frequently paired with land and water backgrounds, respectively, but at test time the 4 subgroup combinations are uniformly sampled. To mitigate failure under distribution shift, a robust representation should learn primarily features of the bird itself, since these are invariant, rather than features of the background. Beyond the increase in dimensionality, this task differs from CMNIST in that the ERM solution does not fail catastrophically at test time, and in fact can achieve 97.3% average accuracy. However, because ERM optimizes average loss, it suffers in performance on the worst-case subgroup (waterbirds on land, which has only 56 training examples).

Adult-Confounded To assess the ability of EIIL to address worst-case group performance without group labels, we construct a variant of the UCI Adult dataset,¹⁰ which comprises 48,842 census records collected from the USA in 1994. The task commonly used as an algorithmic fairness benchmark is to predict a binarized income indicator (thresholded at \$50,000) as the target label, possibly considering sensitive attributes such as age, sex, and race.

Lahoti et al. (2020) demonstrate the benefit of per-example loss reweighting on Adult using their method ARL to improve predictive performance for undersampled subgroups. Following Lahoti et al. (2020), we consider the effect of four sensitive *subgroups*—defined by composing binarized race and sex labels—on model performance, assuming the model does not know a priori which features are sensitive. However, we focus on a distinct generalization problem where a pernicious dataset bias confounds the training data,

¹⁰<https://archive.ics.uci.edu/ml/datasets/adult>

making subgroup membership predictive of the label on the training data. At test time these correlations are reversed, so a predictor that infers subgroup membership to make predictions will perform poorly at test time (see Appendix D for details). This large test-time distribution shift can be understood as a controlled *audit* to determine if the classifier uses subgroup information to predict the target label. We call our dataset variant Adult-Confounded.

CivilComments-WILDS We apply EIIL to a large and challenging comment toxicity prediction task with important fairness implications (Borkan et al., 2019), where ERM performs poorly on comments associated with certain identity groups. We follow the procedure and data splits of Koh et al. (2021) to finetune DistilBERT embeddings (Sanh et al., 2019). EIIL uses an ERM reference classifier and its inferred environments are fed to a GroupDRO invariant learner. Because the large training set ($N_{train} = 269,038$) increases the convergence time for gradient-based EI, we deploy the binning heuristic discussed in Section 3.3, which in this instance finds environments that correspond to the error and non-error cases of the reference classifier. While ERM and EIIL do not have access to the sensitive group labels, we note that worst-group validation accuracy is used to tune hyperparameters for all methods. See Appendix E for details. We also compare against a GroupDRO (oracle) learner that has access to group labels.

5.2. Results

Method	Handcrafted Environments	Train	Test
ERM	✗	86.3 ± 0.1	13.8 ± 0.6
IRM	✓	71.1 ± 0.8	65.5 ± 2.3
EIIL	✗	73.7 ± 0.5	68.4 ± 2.7

Table 2. Accuracy (%) on CMNIST, a digit classification task where color is a spurious feature correlated with the label during training but anti-correlated at test time. EIIL exceeds test-time performance of IRM *without* knowledge of pre-specified environment labels, by instead finding worst-case environments using aggregated data and a reference classifier.

CMNIST IRM was previously shown to learn an invariant representation on this dataset, allowing it to generalize relatively well to the held-out test set whereas ERM fails dramatically (Arjovsky et al., 2019). It is worth noting that label noise makes the problem challenging, so even an oracle classifier can achieve at most 75% test accuracy on this binary classification task. To realize EIIL in our experiments, we discard the environment labels, and run the procedure described in Section 3.1 with ERM as the reference model and IRM as the invariant learner used in the final stage. We find that EIIL’s environment labels are

very effective for invariant learning, ultimately *outperforming* standard IRM using the environment labels provided in the dataset (Table 2). This suggests that in this case the EIIL solution approaches the *maximally informative* set of environments discussed in Proposition 1.

Waterbirds Sagawa et al. (2020) demonstrated that ERM suffers from poor worst-group performance on this dataset, and that GroupDRO can mitigate this performance gap if group labels are available. In this dataset, group labels should be considered as oracle information, meaning that the relevant baseline for EIIL is standard ERM. The main contribution of Sagawa et al. (2020) was to show *how* deep nets can be optimized for the GroupDRO objective using their online algorithm that adaptively adjusts per-group importance weights. In our experiment, we combine this insight with our EIIL framework to show that distributionally robust neural nets can be realized without access to oracle information. We follow the same basic procedure as above,¹¹ in this case using GroupDRO as the downstream invariant learner for which EIIL’s inferred labels will be used.

Method	Train (avg)	Test (avg)	Test (worst group)
ERM	100.0	97.3	60.3
EIIL	99.6	96.9	78.7
GroupDRO (oracle)	99.1	96.6	84.6

Table 3. Accuracy (%) on the Waterbirds dataset. EIIL strongly outperforms ERM on worst-group performance, approaching the performance of the GroupDRO algorithm proposed by Sagawa et al. (2020), which requires oracle access to group labels. In this experiment we feed environments inferred by EIIL into a GroupDRO learner.

EIIL is significantly more robust than the ERM baseline (Table 3), raising worst-group test accuracy by 18% with only a 1% drop in average accuracy. In Figure 2 we plot the distribution of subgroups for each inferred environment, showing that the minority subgroups (landbirds on water and waterbirds on land) are mostly organized in the same inferred environment. This suggests the possibility of leveraging environment inference for interpretability to automatically discover a model’s performance discrepancies on subgroups, which we leave for future work.

Adult-Confounded Using EIIL—to first infer worst-case environments then ensure invariance across them—performs favorably on the audit test set, compared with ARL and an ERM baseline (Table 4). We also find that, without access

¹¹For this dataset, environment inference worked better with reference models that were not fully trained. We suspect this is because ERM focuses on the easy-to-compute features like background color early in training, precisely the type of bias EIIL can exploit to learn informative environments.

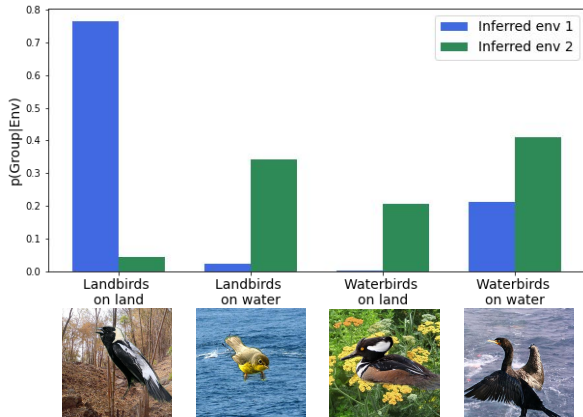


Figure 2. After using EIIL to directly infer two environments from the Waterbirds dataset, we examine the proportion of each subgroup (available in the original dataset but not used by EIIL) present in the inferred environment.

to sensitive group information, the solution found by EIIL achieves significantly better calibration on the test distribution (Figure 3). Because the train and test distributions differ based on the correlation pattern of small subgroups to the target label, this suggests that EIIL has achieved favorable group sufficiency (Liu et al., 2019) in this setting. See Appendix F.3 for a discussion of this point, as well as an ablation showing that all components of the EIIL approach are needed to achieve the best performance.

CivilComments-WILDS Without knowledge of which comments are associated with which groups, EIIL improves worst-group accuracy over ERM with only a modest cost in average accuracy, approaching the oracle GroupDRO solution (which requires group labels).

5.3. Influence of the reference model

As discussed in Section 3.2, the ability of EIIL to find useful environments—partitions yielding an invariant representation when used by an invariant learner—depends on its ability to exploit variation in the predictive distribution of a reference classifier. Here we study the influence of the reference classifier on the final EIIL solution. We return to

Method	Train accs	Test accs
ERM	92.7 ± 0.5	31.1 ± 4.4
ARL (Lahoti et al., 2020)	72.1 ± 3.6	61.3 ± 1.7
EIIL	69.7 ± 1.6	78.8 ± 1.4

Table 4. Accuracy on Adult-Confounded, a variant of the UCI Adult dataset where some sensitive subgroups correlate to the label at train time, and this correlation pattern is reversed at test time.

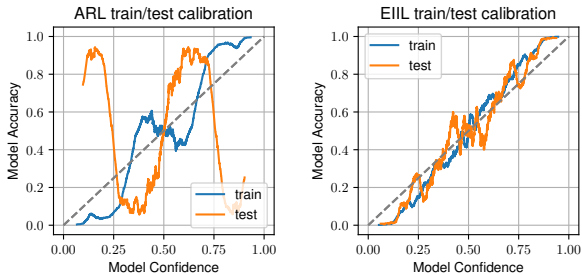


Figure 3. By inferring environments that maximally violate the invariance principle, and then applying invariant learning across the inferred environments, EIIL finds a solution that is well calibrated on the test set (right), compared with ARL (left).

the CMNIST dataset, which provides a controlled sampling setup where particular ERM solutions can be induced to serve as reference for EIIL. In Appendix F.1, we discuss a similar experiment in a synthetic regression setting.

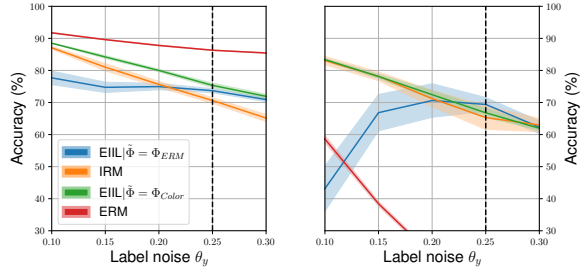
EIIL was shown to outperform IRM *without access to environment labels* in the standard CMNIST dataset (Table 2), which has label noise of $\theta_y = 0.25$. Because $\text{Corr}(\text{color}, \text{label})$ is 0.85 (on average) for the train set, this amount of label noise implies that color is the most predictive feature on aggregated training set (although its predictive power varies across environments). ERM, even with access to infinite data, will focus on color given this amount of label noise to achieve an average train accuracy of 85%. However we can implicitly control the ERM solution Φ_{ERM} by tuning θ_y , an insight that we use to study the dependence of EIIL on the reference model $\tilde{\Phi} = \Phi_{ERM}$.

Figure 4 shows the results of our study. We find that EIIL generalizes better than IRM with sufficiently high label noise $\theta_y > .2$, but generalizes poorly under low label noise. This is precisely because ERM learns the color feature under high label noise, and the shape feature under low label noise. We verify this conclusion by evaluating EIIL when $\tilde{\Phi} = \Phi_{Color}$, i.e. a hand-coded color-based predictor as reference, which does well across all settings of θ_y .

We saw in the Waterbirds experiment that it is not a strict requirement that ERM fail completely in order for EIIL to

Method	Train (avg)	Test (avg)	Test (worst group)
ERM	96.0 \pm 1.5	92.0 \pm 0.4	61.6 \pm 1.3
EIIL	97.0 \pm 0.8	90.5 \pm 0.2	67.0 \pm 2.4
GroupDRO (oracle)	93.6 \pm 1.3	89.0 \pm 0.3	69.8 \pm 2.4

Table 5. EIIL improves worst-group accuracy in the CivilComments-WILDS toxicity prediction task, without access to group labels.



(a) Train accuracy.

(b) Test accuracy

Figure 4. CMNIST with varying label noise θ_y . Under high label noise ($\theta_y > .2$), where the spurious feature color correlates to label *more* than shape on the train data, EIIL matches or exceeds the test performance of IRM *without relying on hand-crafted environments*. Under medium label noise ($.1 < \theta_y < .2$), EIIL is worse than IRM but better than ERM, the logical approach if environments are not available. Under low label noise ($\theta_y < .1$), where color is *less* predictive than shape at train time, ERM performs well and EIIL fails. The vertical dashed black line indicates the default setting of $\theta_y = 0.25$, which we report in Table 2.

succeed. However, this controlled study highlights the importance of the reference model in the ability of EIIL to find environments that emphasize the right invariances, which leaves open the question of how to effectively choose a reference model for EIIL in general. One possible way forward is by using validation data that captures the *kind* of distribution shift we expect at test time, without exactly producing the test distribution, e.g. as in the WILDS benchmark (Koh et al., 2021). In this case we could choose to run EIIL with a reference model that exhibits a large generalization gap between the training and validation distributions.

6. Conclusion

We introduced EIIL, a new method that infers environment partitions of aggregated training data for invariant learning. Without access to environment labels, EIIL can outperform or approach invariant learning methods that require environment labels. EIIL has implications for domain generalization and fairness alike, because in both cases it can be hard to specify meaningful environments or sensitive subgroups.

Acknowledgements

We are grateful to David Madras, Robert Adragna, Silviu Pitis, Will Grathwohl, Jesse Bettencourt, and Eleni Triantafillou for their feedback on this manuscript. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute (www.vectorinstitute.ai/partners).

References

- Ahmed, F., Bengio, Y., van Seijen, H., and Courville, A. Systematic generalisation with group invariant predictions. In *International Conference on Learning Representations*, 2021.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Bahdanau, D., Murty, S., Noukhovitch, M., Nguyen, T. H., de Vries, H., and Courville, A. Systematic generalization: what is required and can it be learned? In *International Conference of Machine Learning*, 2019.
- Beery, S., Van Horn, G., and Perona, P. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision*, pp. 456–473, 2018.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010.
- Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. *Robust Optimization*, volume 28. Princeton University Press, 2009.
- Borkan, D., Dixon, L., Sorensen, J., Thain, N., and Vasserman, L. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion proceedings of the 2019 world wide web conference*, pp. 491–500, 2019.
- Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2):153–163, 2017.
- Chouldechova, A. and Roth, A. The frontiers of fairness in machine learning. *Communications of the ACM*, 63(5): 82–89, 2020.
- Dagaev, N., Roads, B. D., Luo, X., Barry, D. N., Patil, K. R., and Love, B. C. A too-good-to-be-true prior to reduce shortcut reliance. *arXiv preprint arXiv:2102.06406*, 2021.
- Duchi, J. C., Glynn, P. W., and Namkoong, H. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pp. 214–226, 2012.
- Edwards, H. and Storkey, A. Censoring representations with an adversary. In *International Conference for Machine Learning*, 2016.
- Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Steinhardt, J., and Madry, A. Identifying statistical bias in dataset replication. In *International Conference on Machine Learning*, 2020.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., and Wichmann, F. A. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- Gulrajani, I. and Lopez-Paz, D. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021.
- Hanna, A., Denton, E., Smart, A., and Smith-Loud, J. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 conference on Fairness, Accountability, and Transparency*, pp. 501–512, 2020.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pp. 3315–3323, 2016.
- Hashimoto, T. B., Srivastava, M., Namkoong, H., and Liang, P. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, 2018.
- Hébert-Johnson, U., Kim, M. P., Reingold, O., and Rothblum, G. N. Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, 2018.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A., and Darrell, T. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, pp. 1989–1998, 2018.
- Kamath, P., Tangella, A., Sutherland, D. J., and Srebro, N. Does invariant risk minimization capture invariance? In *International Conference on Artificial Intelligence and Statistics*, 2021.

- Kearns, M., Neel, S., Roth, A., and Wu, Z. S. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*, pp. 2564–2572, 2018.
- Kim, M. P., Ghorbani, A., and Zou, J. Multiaccuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 247–254, 2019.
- Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Beery, S., et al. WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, 2021.
- Krueger, D., Caballero, E., Jacobsen, J.-H., Zhang, A., Binas, J., Priol, R. L., and Courville, A. Out-of-distribution generalization via risk extrapolation (REx). In *International Conference on Machine Learning*, 2021.
- Lahoti, P., Beutel, A., Chen, J., Lee, K., Prost, F., Thain, N., Wang, X., and Chi, E. H. Fairness without demographics through adversarially reweighted learning. In *Neural Information Processing Systems*, 2020.
- Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., and Tao, D. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European Conference on Computer Vision*, pp. 624–639, 2018.
- Liu, E., Haghighi, B., Chen, A., Raghunathan, A., Koh, P. W., Sagawa, S., Liang, P., and Finn, C. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, 2021.
- Liu, L. T., Simchowitz, M., and Hardt, M. The implicit fairness criterion of unconstrained learning. In *International Conference on Machine Learning*, 2019.
- Long, M., Cao, Z., Wang, J., and Jordan, M. I. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pp. 1640–1650, 2018.
- Louizos, C., Swersky, K., Li, Y., Welling, M., and Zemel, R. The variational fair autoencoder. In *International Conference on Learning Representations*, 2016.
- Madras, D., Creager, E., Pitassi, T., and Zemel, R. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*, 2018.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- Nam, J., Cha, H., Ahn, S., Lee, J., and Shin, J. Learning from failure: Training debiased classifier from biased classifier. In *Neural Information Processing Systems 2020*, 2020.
- Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.
- Peters, J., Bühlmann, P., and Meinshausen, N. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):947–1012, 2016.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pp. 5389–5400. PMLR, 2019.
- Rockafellar, R. T. and Uryasev, S. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.
- Rosenfeld, E., Ravikumar, P., and Risteski, A. The risks of invariant risk minimization. In *International Conference on Learning Representations*, 2021.
- Sagawa, S., Koh, P. W., Hashimoto, T. B., and Liang, P. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *International Conference on Learning Representations*, 2020.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *The 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing*, 2019.
- Santurkar, S., Tsipras, D., and Madry, A. Breeds: Benchmarks for subpopulation shift. In *International Conference on Learning Representations*, 2021.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., and Vertesi, J. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 59–68, 2019.
- Sohoni, N. S., Dunnmon, J. A., Angus, G., Gu, A., and Ré, C. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. In *Neural Information Processing Systems 2020*, 2020.
- Srivastava, M., Hashimoto, T., and Liang, P. Robustness to spurious correlations via human annotations. In *International Conference on Machine Learning*, pp. 9109–9119. PMLR, 2020.

- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 7167–7176, 2017.
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., and Perona, P. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- Williamson, R. C. and Menon, A. K. Fairness risk measures. In *International Conference on Machine Learning*, 2019.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In *International Conference on Machine Learning*, pp. 325–333, 2013.
- Zhang, B. H., Lemoine, B., and Mitchell, M. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340, 2018.
- Zhang, Y., Barzilay, R., and Jaakkola, T. Aspect-augmented adversarial networks for domain adaptation. *Transactions of the Association for Computational Linguistics*, 5:515–528, 2017.
- Zhao, H., Des Combes, R. T., Zhang, K., and Gordon, G. On learning invariant representations for domain adaptation. In *International Conference on Machine Learning*, pp. 7523–7532. PMLR, 2019.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017.