Probable Domain Generalization
via Quantile Risk Minimization

Cian Eastwood[*1,2]   **Alexander Robey**[*3]   **Shashank Singh**[1]

Julius von Kügelgen[1,4]   **Hamed Hassani**[3]   **George J. Pappas**[3]   **Bernhard Schölkopf**[1]

[1] Max Planck Institute for Intelligent Systems, Tübingen
[2] University of Edinburgh   [3] University of Pennsylvania   [4] University of Cambridge

Abstract

Domain generalization (DG) seeks predictors which perform well on unseen test distributions by leveraging data drawn from multiple related training distributions or domains. To achieve this, DG is commonly formulated as an average- or worst-case problem over the set of possible domains. However, predictors that perform well on average lack robustness while predictors that perform well in the worst case tend to be overly-conservative. To address this, we propose a new probabilistic framework for DG where the goal is to learn predictors that perform well *with high probability*. Our key idea is that distribution shifts seen during training should inform us of probable shifts at test time, which we realize by explicitly relating training and test domains as draws from the same underlying meta-distribution. To achieve probable DG, we propose a new optimization problem called *Quantile Risk Minimization* (QRM). By minimizing the α-quantile of predictor's risk distribution over domains, QRM seeks predictors that perform well with probability α. To solve QRM in practice, we propose the *Empirical QRM* (EQRM) algorithm, and prove: (i) a generalization bound for EQRM; and (ii) that EQRM recovers the causal predictor as $\alpha \to 1$. In our experiments, we introduce a more holistic quantile-focused evaluation protocol for DG, and demonstrate that EQRM outperforms state-of-the-art baselines on `CMNIST` and several datasets from WILDS and DomainBed.

1 Introduction

Despite remarkable successes in recent years [1–3], machine learning systems often fail calamitously when presented with *out-of-distribution* (OOD) data [4–7]. Evidence of state-of-the-art systems failing in the face of distribution shift is mounting rapidly—be it due to spurious correlations [8–10], changing sub-populations [11–13], changes in location or time [14–16], or other naturally-occurring variations [17–23]. These OOD failures are particularly concerning in safety-critical applications such as medical imaging [24–28] and autonomous driving [29–31], where they represent one of the most significant barriers to the real-world deployment of machine learning systems [32–35].

Domain generalization (DG) seeks to improve a system's OOD performance by leveraging datasets from multiple environments or domains at training time, each collected under different experimental conditions [36–38] (see Fig. 1a). The goal is to build a predictor which exploits invariances across the training domains in the hope that these invariances also hold in related but distinct test domains [38–41]. To realize this goal, DG is commonly formulated as an average- [36, 42, 43] or worst-case [9, 44, 45] optimization problem over the set of possible domains. However, optimizing for average performance provably lacks robustness to OOD data [46], while optimizing for worst-domain performance tends to lead to overly-conservative solutions, with worst-case outcomes unlikely in practice [47, 48].

---

*Equal contribution. Correspondence to `c.eastwood@ed.ac.uk` or `arobey1@seas.upenn.edu`.
Code available at: https://github.com/cianeastwood/qrm

(a)   (b)   (c)

Figure 1: **Overview of Probable Domain Generalization and Quantile Risk Minimization.** (a) In domain generalization, training and test data are drawn from multiple related distributions or domains. For example, in the `iWildCam` dataset [50], which contains camera-trap images of animal species, the domains correspond to the different camera-traps which captured the images. (b) We relate training and test domains as draws from the same underlying (and often unknown) meta-distribution over domains $\mathbb{Q}$. (c) We consider a predictor's estimated risk distribution over training domains, naturally-induced by $\mathbb{Q}$. By minimizing the α-quantile of this distribution, we learn predictors that perform well with high probability ($\approx \alpha$) rather than on average or in the worst case.

In this work, we argue that DG is neither an average-case nor a worst-case problem, but rather a probabilistic one. To this end, we propose a probabilistic framework for DG, which we call *Probable Domain Generalization* (§ 3), wherein the key idea is that distribution shifts seen during training should inform us of *probable* shifts at test time. To realize this, we explicitly relate training and test domains as draws from the same underlying meta-distribution (Fig. 1b), and then propose a new optimization problem called *Quantile Risk Minimization* (QRM). By minimizing the α-quantile of predictor's risk distribution over domains (Fig. 1c), QRM seeks predictors that perform well *with high probability* rather than on average or in the worst case. In particular, QRM leverages the key insight that this α-quantile is an upper bound on the test-domain risk which holds with probability α, meaning that α is an interpretable conservativeness-hyperparameter with $\alpha = 1$ corresponding to the worst-case setting.

To solve QRM in practice, we introduce the *Empirical QRM* (EQRM) algorithm (§ 4). Given a predictor's empirical risks on the training domains, EQRM forms an estimated risk distribution using kernel density estimation (KDE, [49]). Critically, KDE-smoothing ensures that the CDF has a right tail which extends beyond the largest training risk (see Fig. 1c), with this risk "extrapolation" [41] unlocking *invariant prediction* for EQRM (§ 4.1). We then provide theory for EQRM (§ 4.2, § 4.3) and demonstrate empirically that EQRM outperforms state-of-the-art baselines on real and synthetic data (§ 6).

Contributions. To summarize our main contributions:

- *A new probabilistic perspective and objective for DG:* We argue that predictors should be trained and tested based on their ability to perform well *with high probability*. We then propose Quantile Risk Minimization for achieving this *probable* form of domain generalization (§ 3).
- *A new algorithm:* We propose the EQRM algorithm to solve QRM in practice and ultimately learn predictors that generalize with probability α (§ 4). We then provide several analyses of EQRM:
  - *Learning theory:* We prove a uniform convergence bound, meaning the empirical α-quantile risk tends to the population α-quantile risk given sufficiently many domains and samples (Thm. 4.1).
  - *Causality.* We prove EQRM learns predictors with invariant risk as $\alpha \to 1$ (Prop. 4.3), and that this is sufficient to recover the causal predictor under weaker assumptions than prior work (Thm. 4.4).
  - *Experiments:* We demonstrate that EQRM outperforms state-of-the-art baselines on several standard DG benchmarks, including `CMNIST` [9] and datasets from WILDS [12] and DomainBed [38], and highlight the importance of assessing the tail or *quantile performance* of DG algorithms (§ 6).

2   Background: Domain generalization

Setup. In domain generalization (DG), predictors are trained on data drawn from multiple related training distributions or *domains* and then evaluated on related but unseen test domains. For example, in the `iWildCam` dataset [50], the task is to classify animal species in images, and the domains correspond to the different camera-traps which captured the images (see Fig. 1a). More formally, we consider datasets $D^e = \{(x_i^e, y_i^e)\}_{i=1}^{n_e}$ collected from m different training domains or *environments* $\mathcal{E}_{\text{tr}} := \{e_1, \dots, e_m\}$, with each dataset D^e containing data pairs (x_i^e, y_i^e) sampled i.i.d. from

$\mathbb{P}(X^e, Y^e)$. Then, given a suitable function class $\mathcal{F}$ and loss function ℓ, the goal of DG is to learn a predictor $f \in \mathcal{F}$ that generalizes to data drawn from a larger set of all possible domains $\mathcal{E}_{\text{all}} \supset \mathcal{E}_{\text{tr}}$.

Average case. Letting $\mathcal{R}^e(f)$ denote the statistical risk of f in domain e, and $\mathbb{Q}$ a distribution over the domains in $\mathcal{E}_{\text{all}}$, DG was first formulated [36] as the following average-case problem:

$$\min_{f \in \mathcal{F}} \mathbb{E}_{e \sim \mathbb{Q}} \mathcal{R}^e(f) \qquad \text{where} \qquad \mathcal{R}^e(f) := \mathbb{E}_{\mathbb{P}(X^e, Y^e)}[\ell(f(X^e), Y^e)]. \tag{2.1}$$

Worst case. Since predictors that perform well *on average* provably lack robustness [46], i.e. they can perform quite poorly on large subsets of $\mathcal{E}_{\text{all}}$, subsequent works [9, 22, 41, 44, 45, 51] have sought robustness by formulating DG as the following *worst-case* problem:

$$\min_{f \in \mathcal{F}} \max_{e \in \mathcal{E}_{\text{all}}} \mathcal{R}^e(f). \tag{2.2}$$

As we only have access to data from a finite subset of $\mathcal{E}_{\text{all}}$ during training, solving (2.2) is not just challenging but in fact impossible [41, 52, 53] without restrictions on how the domains may differ.

Causality and invariance in DG. Causal works on DG [9, 41, 53–55] describe domain differences using the language of causality and the notion of *interventions* [56, 57]. In particular, they assume all domains share the same underlying *structural causal model* (SCM) [56], with different domains corresponding to different interventions (see Appendix A.1 for formal definitions and a simple example). Assuming the mechanism of Y remains fixed or invariant but all Xs may be intervened upon, recent works have shown that only the causal predictor has invariant: (i) predictive distributions [54], coefficients [9] or risks [41] across domains; and (ii) generalizes to arbitrary interventions on the Xs [9, 54, 55]. These works then leverage some form of invariance across domains to discover causal relationships which, through the invariant mechanism assumption, generalize to new domains.

3 Quantile Risk Minimization

In this section we introduce *Quantile Risk Minimization* (QRM) for achieving *Probable Domain Generalization*. The core idea is to replace the worst-case perspective of (2.2) with a probabilistic one. This approach is founded on a great deal of work in classical fields such as control theory [58, 59] and smoothed analysis [60], wherein approaches that yield high-probability guarantees are used in place of worst-case approaches in an effort to mitigate conservatism and computational limitations. This mitigation is of particular interest in domain generalization since generalizing to arbitrary domains is impossible [41, 52, 53]. Thus, motivated by this classical literature, our goal is to obtain predictors that are robust *with high probability* over domains drawn from $\mathcal{E}_{\text{all}}$, rather than in the worst-case.

A distribution over environments. We start by assuming the existence of a probability distribution $\mathbb{Q}(e)$ over the set of all environments $\mathcal{E}_{\text{all}}$. For instance, in the context of medical imaging, $\mathbb{Q}$ could represent a distribution over potential changes to a hospital's setup or simply a distribution over candidate hospitals. Given that such a distribution $\mathbb{Q}$ exists[2], we can think of the risk $\mathcal{R}^e(f)$ as a *random variable* for each $f \in \mathcal{F}$, where the randomness is engendered by the draw of $e \sim \mathbb{Q}$. This perspective gives rise to the following analogue of the optimization problem in (2.2):

$$\min_{f \in \mathcal{F}} \operatorname*{ess\,sup}_{e \sim \mathbb{Q}} \mathcal{R}^e(f) \quad \text{where} \quad \operatorname*{ess\,sup}_{e \sim \mathbb{Q}} \mathcal{R}^e(f) = \inf \left\{ t \geq 0 : \Pr_{e \sim \mathbb{Q}} \{ \mathcal{R}^e(f) \leq t \} = 1 \right\} \tag{3.1}$$

Here, $\operatorname{ess\,sup}$ denotes the *essential-supremum* operator from measure theory, meaning that for each $f \in \mathcal{F}$, $\operatorname{ess\,sup}_{\mathbb{Q}} \mathcal{R}^e(f)$ is the least upper bound on $\mathcal{R}^e(f)$ that holds for almost every $e \sim \mathbb{Q}$. In this way, the $\operatorname{ess\,sup}$ in (3.1) is the measure-theoretic analogue of the max operator in (2.2), with the subtle but critical difference being that the $\operatorname{ess\,sup}$ in (3.1) can neglect domains of measure zero under $\mathbb{Q}$. For example, for discrete $\mathbb{Q}$, (3.1) ignores domains which are impossible (i.e. have probability zero) while (2.2) does not, laying the foundation for ignoring domains which are *improbable*.

High-probability generalization. Although the minimax problem in (3.1) explicitly incorporates the distribution $\mathbb{Q}$ over environments, this formulation is no less conservative than (2.2). Indeed, in many cases, (3.1) is equivalent to (2.2); see Appendix B for details. Therefore, rather than considering the worst-case problem in (3.1), we propose the following generalization of (3.1) which requires that predictors generalize with probability α rather than in the worst-case:

$$\min_{f \in \mathcal{F}, t \in \mathbb{R}} t \qquad \text{subject to} \qquad \Pr_{e \sim \mathbb{Q}} \{ \mathcal{R}^e(f) \leq t \} \geq \alpha \tag{3.2}$$

---

[2]As $\mathbb{Q}$ is often unknown, our analysis does not rely on using an explicit expression for $\mathbb{Q}$.

The optimization problem in (3.2) formally defines what we mean by *Probable Domain Generalization*. In particular, we say that *a predictor f generalizes with risk t at level α* if f has risk at most t with probability at least α over domains sampled from $\mathbb{Q}$. In this way, the conservativeness parameter α controls the strictness of generalizing to unseen domains.

A distribution over risks. The optimization problem presented in (3.2) offers a principled formulation for generalizing to unseen distributional shifts governed by $\mathbb{Q}$. However, $\mathbb{Q}$ is often unknown in practice and its support $\mathcal{E}_{\text{all}}$ may be high-dimensional or challenging to define [22]. While many previous works have made progress by limiting the scope of possible shift types over domains [19, 22, 45], in practice, such structural assumptions are often difficult to justify and impossible to test. For this reason, we start our exposition of QRM by offering an alternative view of (3.2) which elucidates how a predictor's *risk distribution* plays a central role in achieving probable domain generalization.

To begin, note that for each $f \in \mathcal{F}$, the distribution over domains $\mathbb{Q}$ naturally induces[3] a distribution $\mathbb{T}_f$ over the risks in each domain $\mathcal{R}^e(f)$. In this way, rather than considering the randomness of $\mathbb{Q}$ in the often-unknown and (potentially) high-dimensional space of possible shifts (Fig. 1b), one can consider it in the real-valued space of risks (Fig. 1c). This is analogous to statistical learning theory, where the analysis of convergence of empirical risk minimizers (i.e., of functions) is substituted by that of a weaker form of convergence, namely that of scalar risk functionals—a crucial step for VC theory [61]. From this perspective, the statistics of $\mathbb{T}_f$ can be thought of as capturing the sensitivity of f to different environmental shifts, summarizing the effect of different intervention types, strengths, and frequencies. To this end, (3.2) can be equivalently rewritten in terms of the risk distribution $\mathbb{T}_f$ as follows:

$$\min_{f \in \mathcal{F}} F_{\mathbb{T}_f}^{-1}(\alpha) \quad \text{where} \quad F_{\mathbb{T}_f}^{-1}(\alpha) := \inf \left\{ t \in \mathbb{R} : \Pr_{R \sim \mathbb{T}_f} \{R \leq t\} \geq \alpha \right\}. \qquad \text{(QRM)}$$

Here, $F_{\mathbb{T}_f}^{-1}(\alpha)$ denotes the inverse CDF (or quantile[4]) function of the risk distribution $\mathbb{T}_f$. By means of this reformulation, we elucidate how solving (QRM) amounts to finding a predictor with minimal α-quantile risk. That is, (QRM) requires that a predictor f satisfy the probabilistic constraint for at least an α-fraction of the risks $R \sim \mathbb{T}_f$, or, equivalently, for an α-fraction of the environments $e \sim \mathbb{Q}$. In this way, by varying α, (QRM) recovers several notable objectives for DG as special cases, namely an average-case problem with $\alpha = 0.5$ and a worst-case problem with $\alpha = 1$. Thus, α can be used to interpolate between average and worst-case problems in an interpretable manner.

Proposition 3.1. *For $\alpha = 1$, (QRM) is equivalent to the worst-case problem of (3.1). For $\alpha = 0.5$, it is equivalent to the average-case problem of (2.1) if the mean and median of $\mathbb{T}_f$ coincide $\forall f \in \mathcal{F}$:*

$$\min_{f \in \mathcal{F}} \mathbb{E}_{R \sim \mathbb{T}_f} R = \min_{f \in \mathcal{F}} \mathbb{E}_{e \sim \mathbb{Q}} \mathcal{R}^e(f) \qquad (3.3)$$

Connection to DRO. While fundamentally different in terms of objective and generalization capabilities (see § 4), we draw connections between QRM and distributionally robust optimization (DRO) in Appendix F by considering an alternative problem which optimizes the *superquantile*.

4   Algorithms for Quantile Risk Minimization

In this section we introduce *Empirical Quantile Risk Minimization* (EQRM)—an algorithm for solving (QRM) in practice. We then provide several theoretical analyses of EQRM.

4.1   From QRM to Empirical QRM

In practice, given a predictor f and its empirical risks $\hat{\mathcal{R}}^{e_1}(f), \ldots, \hat{\mathcal{R}}^{e_m}(f)$ on the m training domains, we must form an *estimated* risk distribution $\widehat{\mathbb{T}}_f$. In general, given no prior knowledge about the form of $\mathbb{T}_f$ (e.g. Gaussian), we use *kernel density estimation* (KDE, [49, 63]) with Gaussian kernels and either the Gaussian-optimal rule [64] or Silverman's rule-of-thumb [64] for bandwidth selection. Fig. 1c depicts the PDF and CDF for 10 training risks when using Silverman's rule-of-thumb. Armed

---

[3]$\mathbb{T}_f$ can be formally defined as the push-forward measure of $\mathbb{Q}$ through the risk functional $\mathcal{R}^e(f)$; see App. B.

[4]In financial optimization, when concerned with a distribution over potential losses, the α-quantile value is known as the *value at risk* (VaR) at level α [62].

with a predictor's estimated risk distribution $\widehat{\mathbb{T}}_f$, we can approximately solve (QRM) using the following empirical analogue:

$$\min_{f \in \mathcal{F}} \ F_{\widehat{\mathbb{T}}_f}^{-1}(\alpha) \tag{4.1}$$

Note that (4.1) depends only on known quantities so we can compute and minimize it in practice, as detailed in Alg. 1 of Appendix E.1.

Smoothing permits risk extrapolation. Fig. 2 compares the KDE-smoothed CDF (black) to the unsmoothed empirical CDF (gray). As shown, the latter places zero probability mass on risks greater than our largest training risk, thus implicitly assuming that test risks cannot be larger than training risks. In contrast, the KDE-smoothed CDF permits "risk extrapolation" [41] since its right tail extends beyond our largest training risk, with the estimated α-quantile risk going to infinity as $\alpha \to 1$ (when kernels have full support). Note that different bandwidth-selection methods encode different assumptions about right-tail heaviness and thus about projected OOD risk. In § 4.3, we discuss how, as $\alpha \to 1$, this KDE-smoothing allows EQRM to learn predictors with invariant risk over domains. In Appendix C, we discuss different bandwidth-selection methods for EQRM.

Figure 2: Risk CDFs.

4.2 Theory: Generalization bound

We now give a simplified version of our main generalization bound—Thm. D.1—which states that, given sufficiently many domains and samples, the empirical α-quantile risk is a good estimate of the population α-quantile risk. In contrast to previous results for DG, we bound the *proportion of test data* for which a predictor performs well, rather than the *average error* [36, 42], and make no assumptions about the shift type, e.g. covariate shift [37]. The full version, stated and proved in Appendix D, provides specific finite-sample bounds on ϵ_1 and ϵ_2 below, depending on the hypothesis class $\mathcal{F}$, the empirical estimator $F_{\widehat{\mathbb{T}}_f}^{-1}(\alpha)$, and the assumptions on the possible risk profiles of hypotheses $f \in \mathcal{F}$.

Theorem 4.1 (Simplified form of Thm. D.1, uniform convergence)**.** *Given m domains and n samples in each, then with high probability over the training data,*

$$\sup_{f \in \mathcal{F}} \left| F_{\mathbb{T}_f}^{-1}(\alpha - \epsilon_2) - F_{\widehat{\mathbb{T}}_f}^{-1}(\alpha) \right| \le \epsilon_1, \tag{4.2}$$

where $\epsilon_1 \to 0$ as $n \to \infty$ and $\epsilon_2 \to 0$ as $m \to \infty$.

While many domains are required for this to bound be tight, i.e. for α to *precisely* estimate the true quantile, our empirical results in § 6 demonstrate that EQRM performs well in practice given only a few domains. In such settings, α still controls conservativeness, but with a less precise interpretation.

4.3 Theory: Causal recovery

We now prove that EQRM can recover the causal predictor in two parts. First, we prove that, as $\alpha \to 1$, EQRM learns a predictor with minimal, invariant risk over domains. For Gaussian estimators of the risk distribution $\mathbb{T}_f$, some intuition can be gained from Eq. (A.3) of Appendix A.2.1, noting that $\alpha \to 1$ puts increasing weight on the sample standard deviation of risks over domains $\hat{\sigma}_f$, eventually forcing it to zero. For kernel density estimators, a similar intuition applies so long as the bandwidth has a certain dependence on $\hat{\sigma}_f$, as detailed in Appendix A.2.2. Second, we prove that learning such a *minimal invariant-risk predictor* is sufficient to recover the *causal* predictor under weaker assumptions than Peters et al. [54] and Krueger et al. [41]. Together, these two parts provide the conditions under which EQRM successfully performs "causal recovery", i.e., correctly recovers the true causal coefficients in a linear causal model of the data.

Definition 4.2. A predictor f is said to be an *invariant-risk predictor* if its risk is equal almost surely across domains (i.e., $\text{Var}_{e \sim \mathbb{Q}}[\mathcal{R}^e(f)] = 0$). A predictor is said to be a *minimal invariant-risk predictor* if it achieves the minimal possible risk across all possible invariant-risk predictors.

Proposition 4.3 (EQRM learns a minimal invariant-risk predictor as $\alpha \to 1$, informal version of Props. A.4 and A.5)**.** *Assume: (i) $\mathcal{F}$ contains an invariant-risk predictor with finite training risks;*

and (ii) no arbitrarily-negative training risks. Then, as $\alpha \to 1$, Gaussian and kernel EQRM predictors (the latter with certain bandwidth-selection methods) converge to minimal invariant-risk predictors.

Props. A.4 and A.5 are stated and proved in Appendices A.2.1 and A.2.2 respectively. In addition, for the special case of Gaussian estimators of $\mathbb{T}_f$, Appendix A.2.1 relates our α parameter to the β parameter of VREx [41, Eq. 8]. We next specify the conditions under which learning such a minimal invariant-risk predictor is sufficient to recover the causal predictor.

Theorem 4.4 (The causal predictor is (often) the only minimal invariant-risk predictor). *Assume that: (i) Y is generated from a linear SEM, $Y = \beta^\mathsf{T} X + N$, with coefficients $\beta \in \mathbb{R}^d$ and the risk $\mathbb{E}[(Y - \beta^T X)^2]$ of the causal predictor β invariant across domains; (ii) $\mathcal{F}$ is the class of linear predictors, indexed by $\hat{\beta} \in \mathbb{R}^d$; (iii) all of X is observed; (iv) the loss ℓ is squared-error; and (v) the system of equations*

$$
\begin{aligned}
0 \geq & x^\mathsf{T} \mathrm{Cov}_{X \sim e_1}(X, X) x + 2 x^\mathsf{T} \mathrm{Cov}_{N, X \sim e_1}(X, N) \\
= & \cdots \\
= & x^\mathsf{T} \mathrm{Cov}_{X \sim e_m}(X, X) x + 2 x^\mathsf{T} \mathrm{Cov}_{N, X \sim e_m}(X, N)
\end{aligned}
\tag{4.3}
$$

has the unique solution $x = 0$. If $\hat{\beta}$ is a minimal invariant-risk predictor, then $\hat{\beta} = \beta$.

Assumptions (i–iv). The assumptions that Y is drawn from a linear SEM, that all causes of Y are observed, and that the loss is squared-error, while restrictive, are needed for all comparable causal recovery results [41, 54]. In fact, these assumptions are weaker than both Peters et al. [54, Thm. 2] (assume a linear *Gaussian* SCM for X *and* Y) and Krueger et al. [41, Thm. 1] (assume a linear SCM for X *and* Y). The assumption that the risk of the causal predictor is invariant across domains, often called *domain homoskedasticity* [41], is necessary for any method inferring causality from the *invariance of risks* across domains. For methods based on the *invariance of functions*, namely the conditional mean $\mathbb{E}[Y|\mathrm{Pa}(Y)]$ [9, 65], this domain-homoskedasticity assumption is not required. See Appendix G.1.2 for a detailed analysis and comparison of these approaches.

Assumption (v). In contrast to both Peters et al. and Krueger et al., we do not require specific types of interventions on the covariates. Instead, we require that a more general condition be satisfied, namely that the system of d-variate quadratic equations in (4.3) has a unique solution. Intuitively, $\mathrm{Cov}(X, X)$ captures how correlated the covariates are and ensures they are sufficiently uncorrelated to distinguish each of their influences on Y, while $\mathrm{Cov}(X, N)$ captures how correlated descendant covariates are with Y (via N). Together, these terms capture the idea that *predicting Y from the causal covariates must result in the minimal invariant-risk*: the first inequality ensures the risk is *minimal* and the subsequent $m - 1$ equalities that it is *invariant*. While this generality comes at the cost of abstraction, Appendix A.2.3 provides several concrete examples with different types of interventions in order to aid understanding and illustrate how this condition generalizes existing causal-recovery results based on invariant risks [41, 54]. Appendix A.2.3 also provides a proof of Thm. 4.4 and further discussion.

5 Related work

Robust optimization in DG. Throughout this paper, we follow an established line of work (see e.g., [9, 41, 51]) which formulates the DG problem through the lens of robust optimization [44]. To this end, various algorithms have been proposed for solving constrained [22] and distributionally robust [45] variants of the worst-case problem in (2.2). Indeed, this robust formulation has a firm foundation in the broader machine learning literature, with notable works in adversarial robustness [66–70] and fair learning [71, 72] employing similar formulations. Unlike these past works, we consider a robust but non-adversarial formulation for DG, where predictors are trained to generalize with high probability rather than in the worst case. Moreover, the majority of this literature—both within and outside of DG—relies on specific structural assumptions (e.g. covariate shift) on the types of possible interventions or perturbations. In contrast, we make the weaker and more flexible assumption of i.i.d.-sampled domains, which ultimately makes use of the observed domain-data to determine the types of shifts that are *probable*. We further discuss this important difference in § 7.

Other approaches to DG. Outside of robust optimization, many algorithms have been proposed for the DG setting which draw on insights from a diverse array of fields, including approaches based on tools from meta-learning [40, 43, 73–75], kernel methods [76, 77], and information theory [51]. Also prominent are works that design regularizers to generalize OOD [78–80] and works that seek

domain-invariant representations [81–83]. Many of these works employ hyperparameters which are difficult to interpret, which has no doubt contributed to the well-established model-selection problem in DG [38]. In contrast, in our framework, α can be easily interpreted in terms of quantiles of the risk distribution. In addition, many of these works do not explicitly relate the training and test domains, meaning they lack theoretical results in the non-linear setting (e.g. [9, 41, 43, 84]). For those which do, they bound either average error over test domains [36, 42, 85] or worst-case error under specific shift types (e.g. covariate [22]). As argued above, the former lacks robustness while the latter can be both overly-conservative and difficult to justify in practice, where shift types are often unknown.

High-probability generalization. As noted in § 3, relaxing worst-case problems in favor of probabilistic ones has a long history in control theory [58, 59, 86–88], operations research [89], and smoothed analysis [60]. Recently, this paradigm has been applied to several areas of machine learning, including perturbation-based robustness [90, 91], fairness [92], active learning [93], and reinforcement learning [94, 95]. However, it has not yet been applied to domain generalization.

Quantile minimization. In financial optimization, the quantile and superquantile functions [62, 96, 97] are central to the literature surrounding portfolio risk management, with numerous applications spanning banking regulations and insurance policies [98, 99]. In statistical learning theory, several recent papers have derived uniform convergence guarantees in terms of alternative risk functionals besides expected risk [93, 100–102]. These results focus on functionals that can be written in terms of expectations over the loss distribution (e.g., the superquantile). In contrast, our uniform convergence guarantee (Theorem D.1) shows uniform convergence of the quantile function, which *cannot* be written as such an expectation; this necessitates stronger conditions to obtain uniform convergence, which ultimately suggest regularizing the estimated risk distribution (e.g. by kernel smoothing).

Invariant prediction and causality. Early work studied the problem of learning from multiple cause-effect datasets that share a functional mechanism but differ in noise distributions [39]. More generally, given (data from) multiple distributions, one can try to identify components which are stable, robust, or *invariant*, and find means to transfer them across problems [103–107]. As discussed in § 2, recent works have leveraged different forms of invariance across domains to discover causal relationships which, under the invariant mechanism assumption [57], generalize to new domains [9, 41, 54, 55, 108–110]. In particular, VREx [41] leveraged *invariant risks* (like EQRM) while IRM [9] leveraged *invariant functions* or coefficients—see Appendix G.1.2 for a detailed comparison of these approaches.

6 Experiments

We now evaluate our EQRM algorithm on synthetic datasets (§ 6.1), real-world datasets of WILDS (§ 6.2), and the few-domain datasets of DomainBed (§ 6.3). Appendices E and G report further results and experimental details respectively.

6.1 Synthetic datasets

Linear regression. We first consider a linear regression dataset based on the following linear SCM:

$$X_1 \leftarrow N_1, \qquad\qquad Y \leftarrow X_1 + N_Y, \qquad\qquad X_2 \leftarrow Y + N_2,$$

with $N_j \sim \mathcal{N}(0, \sigma_j^2)$. Here we have two features: one cause $X_1 = X_{\text{cause}}$ and one effect $X_2 = X_{\text{effect}}$ of Y. By fixing $\sigma_1^2 = 1$ and $\sigma_Y^2 = 2$ across domains but sampling $\sigma_2 \sim \text{LogNormal}(0, 0.5)$, we create a dataset in which X_2 is more predictive of Y than X_1 but less stable. Importantly, as we know the true distribution over domains $\mathbb{Q}(e) = \text{LogNormal}(\sigma_2^e; 0, 0.5)$, we know the true risk quantiles. Fig. 3 depicts results for different α's with $m = 1000$ domains and $n = 100000$ samples in each, using the mean-squared-error (MSE) loss. Here we see that: **A:** for each true quantile (x-axis), the corresponding α has the lowest risk (y-axis), confirming that the empirical α-quantile risk is a good estimate of the population α-quantile risk; **B:** As $\alpha \to 1$, the estimated risk distribution of f_α approaches an invariant (or Dirac delta) distribution centered on the risk of the causal predictor; **C:** the regression coefficients approach those of the causal predictor as $\alpha \to 1$, trading predictive performance for robustness; and **D:** reducing the number of domains m reduces the accuracy of the estimated α-quantile risks. In Appendix G.1, we additionally: (i) depict the risk CDFs corresponding to plot **B** above, and discuss how they depict the predictors' risk-robustness curves (G.1.1); and (ii) discuss the solutions of EQRM on datasets in which σ_1^2, σ_2^2 and/or σ_Y^2 change over domains, compared to existing invariance-seeking algorithms like IRM [9] and VREx [41] (G.1.2).

Figure 3: **EQRM on a toy linear regression dataset (A–D) and on ColoredMNIST (E–F). A:** Test risk at different quantiles or degrees of "OODness". For each quantile (x-axis), the corresponding α has the lowest risk (y-axis). **B:** Estimated risk distributions (corresponding CDFs in Appendix G.1.1). **C:** Regression coefficients approach those of the causal predictor ($\beta_{cause} = 1$, $\beta_{effect} = 0$) as $\alpha \to 1$. **D:** Q-Q plot comparing the "true" risk quantiles (estimated with $m = 1000$) against estimated ones ($m < 1000$), for $\alpha = 0.9$. **E:** Performance of different α's over increasingly OOD test domains, with dots showing training-domain accuracies. **F:** KDE-estimated accuracy-CDFs depicting accuracy-robustness curves. Larger α's make lower accuracies less likely.

ColoredMNIST. We next consider the `ColoredMNIST` or `CMNIST` dataset [9]. Here, the `MNIST` dataset is used to construct a binary classification task (0–4 or 5–9) in which digit color (red or green) is a highly-informative but spurious feature. In particular, the two training domains are constructed such that red digits have an 80% and 90% chance of belonging to class 0, while the single test domain is constructed such that they only have a 10% chance. The goal is to learn an invariant predictor which uses only digit shape—a stable feature having a 75% chance of correctly determining the class in all 3 domains. We compare with IRM [9], GroupDRO [45], SD [111], IGA [112] and VREx [41] using: (i) random initialization (Xavier method [113]); and (ii) random initialization followed by several iterations of ERM. The ERM initialization or pretraining directly corresponds to the delicate penalty "annealing" or warm-up periods used by most penalty-based methods [9, 41, 111, 112]. For all methods, we use a 2-hidden-layer MLP with 390 hidden units, the Adam optimizer, a learning rate of 0.0001, and dropout with $p = 0.2$. We sweep over five penalty weights for the baselines and five α's for EQRM. See Appendix E.2 for more experimental details. Table 1 shows that: (i) all methods struggle without ERM pretraining, explaining the need for penalty-annealing strategies in previous works and corroborating the results of [114, Table 1]; (ii) with ERM pretraining, EQRM matches or outperforms baseline methods, even approaching oracle performance (that of ERM trained on grayscale digits). These results suggest ERM pretraining as an effective strategy for DG methods.

In addition, Fig. 3 depicts the behavior of EQRM with different αs. Here we see that: **E:** increasing α leads to more consistent performance across domains, eventually forcing the model to ignore color and focus on shape for invariant-risk prediction; and **F:** a predictor's (estimated) accuracy-CDF depicts its accuracy-robustness curve, just as its risk-CDF depicts its risk-robustness curve. Note that $\alpha = 0.5$ gives the best worst-case (i.e. worst-domain) risk over the two training domains—the preferred solution of DRO [45]—while $\alpha \to 1$ sacrifices risk for increased invariance or robustness.

6.2 Real-world datasets

We now evaluate our methods on the real-world or *in-the-wild* distribution shifts of WILDS [12]. We focus our evaluation on `iWildCam` [50] and `OGB-MolPCBA` [115, 116]—two large-scale classification datasets which have numerous test domains and thus facilitate a comparison of the test-domain risk distributions and their quantiles. Additional comparisons (e.g. using average accuracy) can be found in Appendix G.3. Our results demonstrate that, across two distinct data types (images and molecular graphs), EQRM offers superior tail or quantile performance.

Table 1: CMNIST test accuracy.

Algorithm	Initialization	
	Rand.	ERM
ERM	27.9 ± 1.5	27.9 ± 1.5
IRM	$\mathbf{52.5 \pm 2.4}$	69.7 ± 0.9
GrpDRO	27.3 ± 0.9	29.0 ± 1.1
SD	49.4 ± 1.5	70.3 ± 0.6
IGA	50.7 ± 1.4	57.7 ± 3.3
V-REx	$\mathbf{55.2 \pm 4.0}$	71.6 ± 0.5
EQRM	53.4 ± 1.7	71.4 ± 0.4
Oracle	72.1 ± 0.7	

Figure 4: Test-domain risk distributions.

Table 2: EQRM test risks on iWildCam.

Alg.	Mean risk	Quantile risk						
		0.0	0.25	0.50	0.75	0.90	0.99	1.0
ERM	1.31	0.015	0.42	0.76	2.25	2.73	4.99	5.25
IRM	1.53	0.098	0.52	1.24	1.86	2.36	6.95	7.46
GroupDRO	1.73	0.091	0.68	1.65	2.18	3.36	5.29	5.54
CORAL	1.27	0.024	0.45	0.73	2.12	2.66	4.50	4.98
$\text{EQRM}_{0.25}$	2.03	0.024	0.46	2.70	3.01	3.48	5.03	5.26
$\text{EQRM}_{0.50}$	1.11	**0.004**	0.24	0.68	1.71	2.15	4.04	4.11
$\text{EQRM}_{0.75}$	1.05	0.009	**0.21**	0.68	1.50	2.35	4.88	5.45
$\text{EQRM}_{0.90}$	**0.98**	0.047	0.28	**0.63**	**1.26**	**1.81**	4.11	4.48
$\text{EQRM}_{0.99}$	0.99	0.12	0.35	0.64	1.30	2.00	**3.44**	**3.55**

Table 3: EQRM test risks on OGB-MolPCBA.

Alg.	Mean risk	Quantile risk						
		0.0	0.25	0.50	0.75	0.90	0.99	1.0
ERM	**0.051**	0.0	0.004	0.017	0.060	0.13	0.49	16.04
IRM	0.073	0.098	0.52	1.24	1.86	2.36	6.95	7.46
GroupDRO	0.21	0.091	0.68	1.65	2.18	3.36	5.29	**5.54**
CORAL	0.055	0.0	0.12	0.32	1.23	2.01	5.76	7.44
$\text{EQRM}_{0.25}$	0.054	0.0	0.003	0.016	0.059	0.13	0.48	15.46
$\text{EQRM}_{0.50}$	0.052	0.0	0.003	0.015	0.059	0.13	0.48	11.33
$\text{EQRM}_{0.75}$	0.052	0.0	0.003	0.015	0.059	0.13	0.47	12.15
$\text{EQRM}_{0.90}$	0.052	0.0	0.003	0.015	0.059	0.12	0.47	10.81
$\text{EQRM}_{0.99}$	0.053	0.0	0.003	**0.014**	**0.055**	**0.11**	**0.46**	7.16

iWildCam. We first consider the iWildCam image-classification dataset, which has 243 training domains and 48 test domains. Here, the label Y is one of 182 different animal species and the domain e is the camera trap which captured the image. In Table 2, we observe that EQRM_α does indeed tend to optimize the α-risk quantile, with larger αs during training resulting in lower test-domain risks at the corresponding quantiles. In the left pane of Fig. 4, we plot the (KDE-smoothed) test-domain risk distribution for ERM and EQRM. Here we see a clear trend: as α increases, the tails of the risk distribution tend to drop below ERM, which corroborates the superior quantile performance reported in Table 2. Note that, in Table 2, EQRM tends to record lower *average* risks than ERM. This has several plausible explanations. First, the number of testing domains (48) is relatively small, which could result in a biased sample with respect to the training domains. Second, the test domains may not represent i.i.d. draws from $\mathbb{Q}$, as WILDS [12] test domains tend to be more challenging.

OGB-MolPCBA. We next consider the OGB-MolPCBA (or OGB) dataset, which is a molecular graph-classification benchmark containing 44,930 training domains and 43,793 test domains with an average of 3.6 samples per domain. Table 3 shows that ERM achieves the lowest *average* test risk on OGB, in contrast to the iWildCam results, while EQRM_α still achieves stronger quantile performance. Of particular note is the fact that our methods significantly outperform ERM with respect to worst-case performance (columns/quantiles labeled 1.0); when QRM_α is run with large values of α, we reduce the worst-case risk by more than a factor of two. In Fig. 4, we again see that the risk distributions of EQRM_α have lighter tails than that of ERM.

A new evaluation protocol for DG. The analysis provided in Tables 2-3 and Fig. 4 diverges from the standard evaluation protocol in DG [12, 38]. Rather than evaluating an algorithm's performance *on average* across test domains, we seek to understand *the distribution of its performance*—particularly in the tails by means of the quantile function. This new evaluation protocol lays bare the importance of multiple test domains in DG benchmarks, allowing predictors' risk distributions to be analyzed and compared. Indeed, as DG is neither an average-case nor a worst-case problem, solely reporting a predictor's average or worst risk over test domains can be misleading when assessing its ability to generalize OOD (as shown in Tables 2-3). This underscores the necessity of incorporating tail- or quantile-risk measures into a more holistic evaluation protocol for DG, ultimately providing a more nuanced and complete picture. In practice, which measure is preferred will depend on the application, e.g. via a human-specified robustness-level or quantile-of-interest.

6.3 DomainBed datasets

Finally, we consider the benchmark datasets of DomainBed [38], in particular VLCS [117], PACS [118], OfficeHome [119], TerraIncognita [5] and DomainNet [120]. As each of these datasets contain just 4 or 6 domains, it is not possible to meaningfully compare tail or quantile performances. Nonetheless, in line with much recent work, and to compare EQRM to a range of standard baselines on *few-domain* datasets, Table 4 reports DomainBed results in terms of the average performance across each choice of test domain. Appendix G.2 reports the full set of results, including per-

Table 4: DomainBed results. Model selection: training-domain validation set.

Algorithm	VLCS	PACS	OfficeHome	TerraIncognita	DomainNet	Avg
ERM	77.5 ± 0.4	85.5 ± 0.2	66.5 ± 0.3	46.1 ± 1.8	40.9 ± 0.1	63.3
IRM	78.5 ± 0.5	83.5 ± 0.8	64.3 ± 2.2	47.6 ± 0.8	33.9 ± 2.8	61.6
GroupDRO	76.7 ± 0.6	84.4 ± 0.8	66.0 ± 0.7	43.2 ± 1.1	33.3 ± 0.2	60.9
Mixup	77.4 ± 0.6	84.6 ± 0.6	68.1 ± 0.3	47.9 ± 0.8	39.2 ± 0.1	63.4
MLDG	77.2 ± 0.4	84.9 ± 1.0	66.8 ± 0.6	47.7 ± 0.9	41.2 ± 0.1	63.6
CORAL	78.8 ± 0.6	86.2 ± 0.3	68.7 ± 0.3	47.6 ± 1.0	41.5 ± 0.1	**64.6**
ARM	77.6 ± 0.3	85.1 ± 0.4	64.8 ± 0.3	45.5 ± 0.3	35.5 ± 0.2	61.7
VREx	78.3 ± 0.2	84.9 ± 0.6	66.4 ± 0.6	46.4 ± 0.6	33.6 ± 2.9	61.9
EQRM	77.8 ± 0.6	86.5 ± 0.2	67.5 ± 0.1	47.8 ± 0.6	41.0 ± 0.3	64.1

dataset performances, performances with test-domain model selection, and comparisons to additional baselines, while Appendix E.3 gives the implementation details. Note that EQRM outperforms most baselines, including ERM. However, as we have argued throughout this paper, DG is not an average-case problem, so comparing algorithms solely in terms of *average performance* can be misleading.

7 Discussion

Interpretable model selection. α approximates the probability with which our predictor will generalize with risk below the associated α-quantile value. Thus, α represents an interpretable parameterization of the risk-robustness trade-off. Such interpretability is critical for model selection in DG, and for practitioners with application-specific requirements on performance and/or robustness.

The assumption of i.i.d. domains. For α to approximate the probability of generalizing, training and test domains must be i.i.d.-sampled. While this is rarely true in practice—e.g. hospitals have shared funders, service providers, etc.—we can better satisfy this assumption by subscribing to a new data collection process in which we collect training-domain data which is representative of how the underlying system tends to change. For example: (i) randomly select 100 US hospitals; (ii) gather and label data from these hospitals; (iii) train our system with the desired α; (iv) deploy our system to all US hospitals, where it will be successful with probability $\approx \alpha$. While this process may seem expensive, time-consuming and vulnerable (e.g. to new hospitals), it offers a promising path to machine learning systems which *generalize with high probability*. Moreover, it is worth noting the alternative: prior works achieve generalization by assuming that only particular types of shifts can occur, e.g. covariate shifts [22, 121, 122], label shifts [122, 123], concept shifts [124], measurement shifts [19], mean shifts [125], shifts which leave the mechanism of Y invariant [9, 39, 41, 54], etc. In real-world settings, where the underlying shift mechanisms are often unknown, such assumptions are both difficult to justify and impossible to test. Future work could look to relax the i.i.d.-domains assumption by leveraging knowledge of domain dependencies (e.g. time).

The wider value of risk distributions. As demonstrated in § 6, a predictor's risk distribution has value beyond quantile-minimization—it estimates the probability associated with each level of risk. Thus, regardless of the algorithm used, risk distributions can be used to analyze trained predictors.

8 Conclusion

We have presented Quantile Risk Minimization for achieving *Probable* Domain Generalization, motivated by the argument that the goal of domain generalization should be to learn predictors which perform well *with high probability* rather than *on-average* or *in the worst case*. By explicitly relating training and test domains as draws from the same underlying meta-distribution, we proposed to learn predictors with minimal α-quantile risk under the training domains. We then introduced the EQRM algorithm, for which we proved a generalization bound and recovery of the causal predictor as $\alpha \to 1$, before demonstrating that EQRM outperforms state-of-the-art baselines on real and synthetic data.

Acknowledgments and Disclosure of Funding

The authors thank Chris Williams and Ian Mason for providing feedback on an earlier draft, as well as Lars Lorch, David Krueger, Francesco Locatello and members of the MPI Tübingen causality group for helpful discussions and comments. This work was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A, 01IS18039B; and by the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

References

[1] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553): 436–444, 2015. 1

[2] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587): 484–489, 2016.

[3] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021. 1

[4] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Proceedings of The IEEE Conference on Computer Vision and Pattern Recognition*, pages 1521–1528, 2011. 1

[5] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision*, pages 456–473, 2018. 9

[6] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.

[7] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2:665–673, 2020. 1

[8] John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl Oermann. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS Medicine*, 15(11), 2018. 1

[9] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv:1907.02893*, 2019. 1, 2, 3, 6, 7, 8, 10, 21, 24, 37, 39, 40, 41

[10] Timothy Niven and Hung Yu Kao. Probing neural network comprehension of natural language arguments. In *Association for Computational Linguistics*, pages 4658–4664, 2020. 1

[11] Shibani Santurkar, Dimitris Tsipras, and Aleksander Madry. Breeds: Benchmarks for subpopulation shift. *arXiv:2008.04859*, 2020. 1

[12] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton A. Earnshaw, Imran S. Haque, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, 2021. 2, 8, 9, 37, 46

[13] Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Nuanced metrics for measuring unintended bias with real data for text classification. In *World Wide Web Conference*, pages 491–500, 2019. 1

[14] Matthew C Hansen, Peter V Potapov, Rebecca Moore, Matt Hancher, Svetlana A Turubanova, Alexandra Tyukavina, David Thau, Stephen V Stehman, Scott J Goetz, Thomas R Loveland, et al. High-resolution global maps of 21st-century forest cover change. *Science*, 342(6160): 850–853, 2013. 1

[15] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *Computer Vision and Pattern Recognition*, pages 6172–6180, 2018.

[16] Vaishaal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. Do image classifiers generalize across time? In *International Conference on Computer Vision*, pages 9661–9669, 2021. 1

[17] Samil Karahan, Merve Kilinc Yildirum, Kadir Kirtac, Ferhat Sukru Rende, Gultekin Butun, and Hazim Kemal Ekenel. How image degradations affect deep cnn-based face recognition? In *International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5, 2016. 1

[18] Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *Journal of Machine Learning Research*, 20:1–25, 2019.

[19] Cian Eastwood, Ian Mason, Christopher K. I. Williams, and Bernhard Schölkopf. Source-free adaptation to measurement shift via bottom-up feature restoration. In *International Conference on Learning Representations*, 2021. 4, 10

[20] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021.

[21] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021.

[22] Alexander Robey, George J. Pappas, and Hamed Hassani. Model-based domain generalization. In *Advances in Neural Information Processing Systems*, 2021. 3, 4, 6, 7, 10, 29

[23] Allan Zhou, Fahim Tajwar, Alexander Robey, Tom Knowles, George J Pappas, Hamed Hassani, and Chelsea Finn. Do deep networks transfer invariances across classes? *arXiv preprint arXiv:2203.09739*, 2022. 1

[24] Jorge Jovicich, Silvester Czanner, Xiao Han, David Salat, Andre van der Kouwe, Brian Quinn, Jenni Pacheco, Marilyn Albert, Ronald Killiany, Deborah Blacker, et al. MRI-derived measurements of human subcortical, ventricular and intracranial brain volumes: reliability effects of scan sessions, acquisition sequences, data analyses, scanner upgrade, scanner vendors and field strengths. *Neuroimage*, 46(1):177–192, 2009. 1

[25] Ehab A AlBadawy, Ashirbani Saha, and Maciej A Mazurowski. Deep learning for segmentation of brain tumors: Impact of cross-institutional training and testing. *Medical Physics*, 45(3): 1150–1158, 2018.

[26] David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen van der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical Image Analysis*, 58:101–544, 2019.

[27] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M. Vardoulakis. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, page 1–12. Association for Computing Machinery, 2020.

[28] Christian Wachinger, Anna Rieckmann, Sebastian Pölsterl, Alzheimer's Disease Neuroimaging Initiative, et al. Detect and correct bias in multi-site neuroimaging datasets. *Medical Image Analysis*, 67:101879, 2021. 1

[29] Dengxin Dai and Luc Van Gool. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *International Conference on Intelligent Transportation Systems*, pages 3819–3824, 2018. 1

[30] Georg Volk, Stefan Müller, Alexander von Bernuth, Dennis Hospach, and Oliver Bringmann. Towards robust CNN-based object detection through augmentation with synthetic rain variations. In *International Conference on Intelligent Transportation Systems*, pages 285–292, 2019.

[31] C. Michaelis, B. Mitzkus, R. Geirhos, E. Rusak, O. Bringmann, A. S. Ecker, M. Bethge, and W. Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. In *Machine Learning for Autonomous Driving Workshop, NeurIPS 2019*, 2019. 1

[32] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144, 2016. 1

[33] Battista Biggio and Fabio Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84:317–331, 2018.

[34] Gustav Mårtensson, Daniel Ferreira, Tobias Granberg, Lena Cavallin, Ketil Oppedal, Alessandro Padovani, Irena Rektorova, Laura Bonanni, Matteo Pardini, Milica G Kramberger, et al. The reliability of a deep learning model in clinical out-of-distribution MRI data: a multicohort study. *Medical Image Analysis*, 66:101714, 2020.

[35] Daniel C Castro, Ian Walker, and Ben Glocker. Causality matters in medical imaging. *Nature Communications*, 11(1):1–10, 2020. 1

[36] Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. In *Advances in Neural Information Processing Systems*, volume 24, 2011. 1, 3, 5, 7

[37] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18, 2013. 5

[38] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2020. 1, 2, 7, 9, 30, 37

[39] Bernhard Schölkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris M Mooij. On causal and anticausal learning. In *ICML*, 2012. 7, 10

[40] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, pages 3490–3497, 2018. 6

[41] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (REx). In *International Conference on Machine Learning*, volume 139, pages 5815–5826, 2021. 1, 2, 3, 5, 6, 7, 8, 10, 21, 22, 24, 25, 37, 39, 40, 41

[42] Gilles Blanchard, Aniket Anand Deshmukh, Ürun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *The Journal of Machine Learning Research*, 22(1):46–100, 2021. 1, 5, 7

[43] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. *Advances in Neural Information Processing Systems*, 34:23664–23678, 2021. 1, 6, 7

[44] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. Robust optimization. In *Robust optimization*. Princeton University Press, 2009. 1, 3, 6

[45] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2019. 1, 3, 4, 6, 8, 38, 39

[46] Vaishnavh Nagarajan, Anders Andreassen, and Behnam Neyshabur. Understanding the failure modes of out-of-distribution generalization. In *International Conference on Learning Representations*, 2021. 1, 3

[47] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2019. 1

[48] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. Adversarial training can hurt generalization. In *ICML 2019 Workshop on Identifying and Understanding Deep Learning Phenomena*, 2019. 1

[49] Emanuel Parzen. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 1962. 2, 4

[50] Sara Beery, Arushi Agarwal, Elijah Cole, and Vighnesh Birodkar. The iWildCam 2021 competition dataset. *arXiv preprint arXiv:2105.03494*, 2021. 2, 8

[51] Kartik Ahuja, Ethan Caballero, Dinghuai Zhang, Jean-Christophe Gagnon-Audet, Yoshua Bengio, Ioannis Mitliagkas, and Irina Rish. Invariance principle meets information bottleneck for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34, 2021. 3, 6

[52] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79 (1):151–175, 2010. 3

[53] Rune Christiansen, Niklas Pfister, Martin Emil Jakobsen, Nicola Gnecco, and Jonas Peters. A causal framework for distribution generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 3, 21

[54] Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 947–1012, 2016. 3, 5, 6, 7, 10, 21, 24, 25, 26, 40

[55] Mateo Rojas-Carulla, Bernhard Schölkopf, Richard Turner, and Jonas Peters. Invariant models for causal transfer learning. *The Journal of Machine Learning Research*, 19(1):1309–1342, 2018. 3, 7, 21

[56] Judea Pearl. *Causality*. Cambridge University Press, 2009. 3, 21

[57] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017. 3, 7

[58] Marco C Campi and Simone Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008. 3, 7

[59] Federico Alessandro Ramponi. Consistency of the scenario approach. *SIAM Journal on Optimization*, 28(1):135–162, 2018. 3, 7

[60] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004. 3, 7

[61] Vladimir Vapnik. *The nature of statistical learning theory*. Springer Science & Business Media, 1999. 4

[62] Darrell Duffie and Jun Pan. An overview of value at risk. *Journal of derivatives*, 4(3):7–49, 1997. 4, 7

[63] Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, pages 832–837, 1956. 4

[64] Bernard W Silverman. *Density Estimation for Statistics and Data Analysis*, volume 26. CRC Press, 1986. 4, 23, 30, 37

[65] Mingzhang Yin, Yixin Wang, and David M Blei. Optimization-based causal estimation from heterogenous environments. *arXiv preprint arXiv:2109.11990*, 2021. 6, 24, 25, 26

[66] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 6

[67] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

[68] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pages 7472–7482. PMLR, 2019.

[69] Alexander Robey, Luiz Chamon, George J Pappas, Hamed Hassani, and Alejandro Ribeiro. Adversarial robustness with semi-infinite constrained learning. *Advances in Neural Information Processing Systems*, 34:6198–6215, 2021.

[70] Jia-Jie Zhu, Christina Kouridi, Yassine Nemmour, and Bernhard Schölkopf. Adversarially robust kernel smoothing. *arXiv preprint arXiv:2102.08474*, 2021. 6

[71] Natalia L Martinez, Martin A Bertran, Afroditi Papadaki, Miguel Rodrigues, and Guillermo Sapiro. Blind pareto fairness and subgroup robustness. In *International Conference on Machine Learning*, pages 7492–7501. PMLR, 2021. 6

[72] Emily Diana, Wesley Gill, Michael Kearns, Krishnaram Kenthapadi, and Aaron Roth. Minimax group fairness: Algorithms and experiments. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 66–76, 2021. 6

[73] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *Advances in neural information processing systems*, 31, 2018. 6

[74] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32, 2019.

[75] Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9624–9633, 2021. 6

[76] Abhimanyu Dubey, Vignesh Ramanathan, Alex Pentland, and Dhruv Mahajan. Adaptive methods for real-world domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14340–14349, 2021. 6

[77] Aniket Anand Deshmukh, Yunwen Lei, Srinagesh Sharma, Urun Dogan, James W Cutler, and Clayton Scott. A generalization error bound for multi-class domain generalization. *arXiv preprint arXiv:1905.10392*, 2019. 6

[78] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Advances in Neural Information Processing Systems*, 33:16096–16107, 2020. 6

[79] Haoliang Li, YuFei Wang, Renjie Wan, Shiqi Wang, Tie-Qiang Li, and Alex Kot. Domain generalization for medical imaging classification with linear-dependency regularization. *Advances in Neural Information Processing Systems*, 33:3118–3129, 2020.

[80] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9619–9628, 2021. 6

[81] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. 7

[82] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 624–639, 2018.

[83] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *European Conference on Computer Vision*, pages 124–140. Springer, 2020. 7

[84] Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 23(5):828–841, 2019. 7

[85] Vikas Garg, Adam Tauman Kalai, Katrina Ligett, and Steven Wu. Learn to expect the unexpected: Probably approximately correct domain generalization. In *International Conference on Artificial Intelligence and Statistics*, pages 3574–3582. PMLR, 2021. 7

[86] Roberto Tempo, Giuseppe Calafiore, and Fabrizio Dabbene. *Randomized algorithms for analysis and control of uncertain systems: with applications*. Springer, 2013. 7

[87] Lars Lindemann, Nikolai Matni, and George J Pappas. Stl robustness risk over discrete-time stochastic processes. *arXiv preprint arXiv:2104.01503*, 2021.

[88] Lars Lindemann, Alena Rodionova, and George J. Pappas. Temporal robustness of stochastic signals. In *25th ACM International Conference on Hybrid Systems: Computation and Control*, pages 1–11, 2022. 7

[89] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczynski. *Lectures on stochastic programming: modeling and theory*. SIAM, 2021. 7, 38

[90] Alexander Robey, Luiz FO Chamon, George J Pappas, and Hamed Hassani. Probabilistically robust learning: Balancing average- and worst-case performance. *arXiv preprint arXiv:2202.01136*, 2022. 7

[91] Leslie Rice, Anna Bair, Huan Zhang, and J Zico Kolter. Robustness between the worst and average case. *Advances in Neural Information Processing Systems*, 34, 2021. 7

[92] Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. Tilted empirical risk minimization. *arXiv preprint arXiv:2007.01162*, 2020. 7

[93] Sebastian Curi, Kfir Y Levy, Stefanie Jegelka, and Andreas Krause. Adaptive sampling for stochastic risk-averse learning. *Advances in Neural Information Processing Systems*, 33: 1036–1047, 2020. 7, 38

[94] Santiago Paternain, Miguel Calvo-Fullana, Luiz FO Chamon, and Alejandro Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *IEEE Transactions on Automatic Control*, 2022. 7

[95] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120, 2017. 7

[96] R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000. 7, 38

[97] Pavlo Krokhmal, Jonas Palmquist, and Stanislav Uryasev. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of risk*, 4:43–68, 2002. 7

[98] David Wozabal. Value-at-risk optimization using the difference of convex algorithm. *OR spectrum*, 34(4):861–883, 2012. 7

[99] Philippe Jorion. *Value at risk: the new benchmark for controlling market risk*. Irwin Professional Pub., 1997. 7

[100] Jaeho Lee, Sejun Park, and Jinwoo Shin. Learning bounds for risk-sensitive learning. *Advances in Neural Information Processing Systems*, 33:13867–13879, 2020. 7

[101] Justin Khim, Liu Leqi, Adarsh Prasad, and Pradeep Ravikumar. Uniform convergence of rank-weighted learning. In *International Conference on Machine Learning*, pages 5254–5263. PMLR, 2020.

[102] John C Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, 2021. 7

[103] Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pages 819–827. PMLR, 2013. 7

[104] E. Bareinboim and J. Pearl. Transportability from multiple environments with limited experiments: Completeness results. In *Advances in Neural Information Processing Systems 27*, pages 280–288, 2014.

[105] Kun Zhang, Mingming Gong, and Bernhard Schölkopf. Multi-source domain adaptation: A causal view. In *Twenty-ninth AAAI Conference on Artificial Intelligence*, 2015.

[106] Mingming Gong, Kun Zhang, Tongliang Liu, Dacheng Tao, Clark Glymour, and Bernhard Schölkopf. Domain adaptation with conditional transferable components. In *International Conference on Machine Learning*, pages 2839–2848. PMLR, 2016.

[107] B. Huang, K. Zhang, J. Zhang, R. Sanchez-Romero, C. Glymour, and B. Schölkopf. Behind distribution shift: Mining driving forces of changes and causal arrows. In *IEEE 17th International Conference on Data Mining (ICDM 2017)*, pages 913–918, 2017. 7

[108] Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2), 2018. 7

[109] Niklas Pfister, Peter Bühlmann, and Jonas Peters. Invariant causal prediction for sequential data. *Journal of the American Statistical Association*, 114(527):1264–1276, 2019.

[110] Juan L Gamella and Christina Heinze-Deml. Active invariant causal prediction: Experiment selection through stability. *Advances in Neural Information Processing Systems*, 33:15464–15475, 2020. 7

[111] Mohammad Pezeshki, Oumar Kaba, Yoshua Bengio, Aaron C Courville, Doina Precup, and Guillaume Lajoie. Gradient starvation: A learning proclivity in neural networks. *Advances in Neural Information Processing Systems*, 34:1256–1272, 2021. 8

[112] Masanori Koyama and Shoichiro Yamaguchi. Out-of-distribution generalization with maximal invariant predictor. https://openreview.net/forum?id=FzGiUKN4aBp, 2020. 8

[113] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Atatistics*, pages 249–256. PMLR, 2010. 8

[114] Jianyu Zhang, David Lopez-Paz, and Léon Bottou. Rich feature construction for the optimization-generalization dilemma. *arXiv preprint arXiv:2203.15516*, 2022. 8, 37

[115] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in Neural Information Processing Systems*, 33:22118–22133, 2020. 8

[116] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, 2018. 8

[117] Chen Fang, Ye Xu, and Daniel N. Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2013. 9

[118] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 9

[119] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 9

[120] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. 9

[121] Joaquin Quiñonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. MIT Press, 2008. 10, 29

[122] Amos J Storkey. When training and test sets are different: characterising learning transfer. In *Dataset Shift in Machine Learning*, pages 3–28. MIT Press, 2009. 10

[123] Zachary Lipton, Yu-Xiang Wang, and Alexander Smola. Detecting and correcting for label shift with black box predictors. In *International Conference on Machine Learning*, pages 3122–3130, 2018. 10

[124] Jose G Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V Chawla, and Francisco Herrera. A unifying view on dataset shift in classification. *Pattern Recognition*, 45:521–530, 2012. 10

[125] Dominik Rothenhäusler, Nicolai Meinshausen, Peter Bühlmann, and Jonas Peters. Anchor regression: Heterogeneous data meet causality. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(2):215–246, 2021. 10

[126] Maximilian Ilse, Jakub M Tomczak, Christos Louizos, and Max Welling. Diva: Domain invariant variational autoencoders. In *Medical Imaging with Deep Learning*, pages 322–348. PMLR, 2020. 30

[127] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel principal component analysis. In *International Conference on Artificial Neural Networks*, pages 583–588, 1997. 30

[128] Ichiro Takeuchi, Quoc V. Le, Timothy D. Sears, and Alexander J. Smola. Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7(45):1231–1264, 2006.

[129] Bharath K Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Gert Lanckriet, and Bernhard Schölkopf. Kernel choice and classifiability for RKHS embeddings of probability distributions. *Advances in Neural Information Processing Systems*, 22, 2009. 30

[130] Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to statistical learning theory. In *Summer School on Machine Learning*, pages 169–207. Springer, 2003. 31

[131] Pascal Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pages 1269–1283, 1990. 31

[132] Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer, 2004. 33

[133] Ronald A DeVore and George G Lorentz. *Constructive approximation*, volume 303. Springer Science & Business Media, 1993. 34

[134] Keith Ball et al. An elementary introduction to modern convex geometry. *Flavors of geometry*, 31(1-58):26, 1997. 34

[135] JM Blair, CA Edwards, and J Howard Johnson. Rational Chebyshev approximations for the inverse of the error function. *Mathematics of Computation*, 30(136):827–830, 1976. 37

[136] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 37

[137] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 37

[138] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018. 37

[139] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272, 2017. 37

[140] Núria Armengol Urpí, Sebastian Curi, and Andreas Krause. Risk-averse offline reinforcement learning. *arXiv preprint arXiv:2102.05371*, 2021. 38

[141] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004. 39

Appendices

Table of Contents

A Causality

A.1 Definitions and example

As in previous causal works on DG [9, 41, 53–55], our causality results assume all domains share the same underlying *structural causal model* (SCM) [56], with different domains corresponding to different interventions. For example, the different camera-trap deployments depicted in Fig. 1a may induce changes in (or interventions on) equipment, lighting, and animal-species prevalence rates.

Definition A.1. An *SCM*[5] $\mathcal{M} = (\mathcal{S}, \mathbb{P}_N)$ consists of a collection of d *structural assignments*

$$\mathcal{S} = \{X_j \leftarrow g_j(\mathrm{Pa}(X_j), N_j)\}_{j=1}^d, \tag{A.1}$$

where $\mathrm{Pa}(X_j) \subseteq \{X_1, \dots, X_d\} \setminus \{X_j\}$ are the *parents* or *direct causes* of X_j, and $\mathbb{P}_N = \prod_{j=1}^d \mathbb{P}_{N_j}$, a joint distribution over the (jointly) independent noise variables $N_1, \dots, N_d$. An SCM $\mathcal{M}$ induces a ("causal") graph $\mathcal{G}$ which is obtained by creating a node for each X_j and then drawing a directed edge from each parent in $\mathrm{Pa}(X_j)$ to X_j. We assume this graph to be acyclic.

We can draw samples from the *observational distribution* $\mathbb{P}_{\mathcal{M}}(X)$ by first sampling a noise vector $n \sim \mathbb{P}_N$, and then using the structural assignments to generate a data point $x \sim \mathbb{P}_{\mathcal{M}}(X)$, recursively computing the value of every node X_j whose parents' values are known. We can also manipulate or *intervene* upon the structural assignments of $\mathcal{M}$ to obtain a related SCM $\mathcal{M}^e$.

Definition A.2. An *intervention e* is a modification to one or more of the structural assignments of $\mathcal{M}$, resulting in a new SCM $\mathcal{M}^e = (\mathcal{S}^e, \mathbb{P}_N^e)$ and (potentially) new graph $\mathcal{G}^e$, with structural assignments

$$\mathcal{S}^e = \{X_j^e \leftarrow g_j^e(\mathrm{Pa}^e(X_j^e), N_j^e)\}_{j=1}^d. \tag{A.2}$$

We can draw samples from the *intervention distribution* $\mathbb{P}_{\mathcal{M}^e}(X^e)$ in a similar manner to before, now using the modified structural assignments. We can connect these ideas to DG by noting that each intervention e creates a new domain or *environment e* with interventional distribution $\mathbb{P}(X^e, Y^e)$.

Example A.3. Consider the following linear SCM, with $N_j \sim \mathcal{N}(0, \sigma_j^2)$:

$$X_1 \leftarrow N_1, \qquad\qquad Y \leftarrow X_1 + N_Y, \qquad\qquad X_2 \leftarrow Y + N_2.$$

Here, interventions could replace the structural assignment of X_1 with $X_1^e \leftarrow 10$ and change the noise variance of X_2, resulting in a set of training environments $\mathcal{E}_{\mathrm{tr}} = \{\text{fix } X_1 \text{ to } 10, \text{ replace } \sigma_2 \text{ with } 10\}$.

A.2 EQRM recovers the causal predictor

Overview. We now prove that EQRM recovers the causal predictor in two stages. First, we prove the formal versions of Prop. 4.3, i.e. that EQRM learns a minimal invariant-risk predictor as $\alpha \to 1$ when using the following estimators of $\mathbb{T}_f$: (i) a Gaussian estimator (Prop. A.4 of Appendix A.2.1); and (ii) kernel-density estimators with certain bandwidth-selection methods (Prop. A.5 of Appendix A.2.2). Second, we prove Thm. 4.4, i.e. that learning a minimal invariant-risk predictor is sufficient to recover the causal predictor under weaker assumptions than those of Peters et al. [54, Thm 2] and Krueger et al. [41, Thm 1] (Appendix A.2.3). Throughout this section, we consider the "population" setting within each domain (i.e., $n \to \infty$); in general, with only finitely-many observations from each domain, only approximate versions of these results are possible.

Notation. Given m training risks $\{\mathcal{R}^{e_1}(f), \dots, \mathcal{R}^{e_m}(f)\}$ corresponding to the risks of a fixed predictor f on m training domains, let

$$\hat{\mu}_f = \frac{1}{m} \sum_{i=1}^m \mathcal{R}^{e_i}(f)$$

denote the sample mean and

$$\hat{\sigma}_f^2 = \frac{1}{m-1} \sum_{i=1}^m (\mathcal{R}^{e_i}(f) - \hat{\mu}_f)^2$$

the sample variance of the risks of f.

---

[5]A Non-parametric Structural Equation Model with Independent Errors (NP-SEM-IE) to be precise.

A.2.1 Gaussian estimator

When using a Gaussian estimator for $\widehat{\mathbb{T}}_f$, we can rewrite the EQRM objective of (4.1) in terms of the standard-Normal inverse CDF Φ^{-1} as

$$\hat{f}_\alpha := \arg\min_{f \in \mathcal{F}} \; \hat{\mu}_f + \Phi^{-1}(\alpha) \cdot \hat{\sigma}_f. \tag{A.3}$$

Informally, we see that $\alpha \to 1 \implies \Phi^{-1}(\alpha) \to \infty \implies \hat{\sigma}_f \to 0$. We formally, we now show that, as $\alpha \to 1$, minimizing (A.3) leads to a predictor with minimal invariant-risk:

Proposition A.4 (Gaussian QRM learns a minimal invariant-risk predictor as $\alpha \to 1$). *Assume*

1. *$\mathcal{F}$ contains an invariant-risk predictor $f_0 \in \mathcal{F}$ with finite mean risk (i.e., $\hat{\sigma}_{f_0} = 0$ and $\hat{\mu}_{f_0} < \infty$), and*
2. *there are no arbitrarily negative mean risks (i.e., $\mu_* := \inf_{f \in \mathcal{F}} \mu_f > -\infty$).*

Then, for the Gaussian QRM predictor $\hat{f}_\alpha$ given in Eq. (A.3),

$$\lim_{\alpha \to 1} \hat{\sigma}_{\hat{f}_\alpha} = 0 \quad \text{and} \quad \limsup_{\alpha \to 1} \hat{\mu}_{\hat{f}_\alpha} \leq \hat{\mu}_{f_0}.$$

Prop. A.4 essentially states that, if an invariant-risk predictor exists, then Gaussian EQRM equalizes risks across the m domains, to a value at most the risk of the invariant-risk predictor. As we discuss in Appendix A.2.3, an invariant-risk predictor f_0 (Assumption 1. of Prop. A.4 above) exists under the assumption that the mechanism generating the labels Y does not change between domains and is contained in the hypothesis class $\mathcal{F}$, together with a homoscedasticity assumption (see Appendix G.1.2). Meanwhile, Assumption 2. of Prop. A.4 above is quite mild and holds automatically for most loss functions used in supervised learning (e.g., squared loss, cross-entropy, hinge loss, etc.). We now prove Prop. A.4.

Proof. By definitions of $\hat{f}_\alpha$ and f_0,

$$\hat{\mu}_{\hat{f}_\alpha} + \Phi^{-1}(\alpha) \cdot \hat{\sigma}_{\hat{f}_\alpha} \leq \hat{\mu}_{f_0} + \Phi^{-1}(\alpha) \cdot \hat{\sigma}_{f_0} = \hat{\mu}_{f_0}. \tag{A.4}$$

Since for $\alpha \geq 0.5$ we have that $\Phi^{-1}(\alpha)\hat{\sigma}_{\hat{f}_\alpha} \geq 0$, it follows that $\hat{\mu}_{\hat{f}_\alpha} \leq \hat{\mu}_{f_0}$. Moreover, rearranging and using the definition of μ_*, we obtain

$$\hat{\sigma}_{\hat{f}_\alpha} \leq \frac{\hat{\mu}_{f_0} - \hat{\mu}_{\hat{f}_\alpha}}{\Phi^{-1}(\alpha)} \leq \frac{\hat{\mu}_{f_0} - \mu_*}{\Phi^{-1}(\alpha)} \to 0 \quad \text{as} \quad \alpha \to 1.$$

$\square$

Connection to VREx. For the special case of using a Gaussian estimator for $\widehat{\mathbb{T}}_f$, we can equate the EQRM objective of (A.3) with the $\mathcal{R}_{\text{VREx}}$ objective of [41, Eq. 8]. To do so, we rewrite $\mathcal{R}_{\text{VREx}}$ in terms of the sample mean and variance:

$$\arg\min_{f \in \mathcal{F}} \; \mathcal{R}_{\text{VREx}}(f) = \arg\min_{f \in \mathcal{F}} \; m \cdot \hat{\mu}_f + \beta \cdot \hat{\sigma}_f^2. \tag{A.5}$$

Note that as $\beta \to \infty$, $\mathcal{R}_{\text{VREx}}$ learns a minimal invariant-risk predictor under the same assumptions, and by the same argument, as Prop. A.4. Dividing this objective by the positive constant $m > 0$, we can rewrite it in a form that allows a direct comparison of our α parameter and this β parameter:

$$\arg\min_{f \in \mathcal{F}} \; \hat{\mu}_f + \left(\frac{\beta \cdot \hat{\sigma}_f}{m}\right) \cdot \hat{\sigma}_f. \tag{A.6}$$

Comparing (A.6) and (A.3), we note the relation $\beta = m \cdot \Phi^{-1}(\alpha)/\hat{\sigma}_f$ for a fixed f. For different fs, a particular setting of our parameter α corresponds to different settings of Krueger et al.'s β parameter, depending on the sample standard deviation over training risks $\hat{\sigma}_f$.

A.2.2 Kernel density estimator

We now consider the case of using a kernel density estimate, in particular,

$$\hat{F}_{\text{KDE},f}(x) = \frac{1}{m} \sum_{i=1}^{m} \Phi\left(\frac{x - R^{e_i}(f)}{h_f}\right) \tag{A.7}$$

to estimate the cumulative risk distribution.

Proposition A.5 (Kernel EQRM learns a minimal risk-invariant predictor as $\alpha \to 1$). *Let*

$$\hat{f}_\alpha := \arg\min_{f \in \mathcal{F}} \hat{F}^{-1}_{KDE,f}(\alpha),$$

be the kernel EQRM predictor, where $\hat{F}^{-1}_{KDE,f}$ denotes the quantile function computed from the kernel density estimate over (empirical) risks of f with a standard Gaussian kernel. Suppose we use a data-dependent bandwidth h_f such that $h_f \to 0$ implies $\hat{\sigma}_f \to 0$ (e.g., the "Gaussian-optimal" rule $h_f = (4/3m)^{0.2} \cdot \hat{\sigma}_f$ [64]). As in Proposition A.4, suppose also that

1. *$\mathcal{F}$ contains an invariant-risk predictor $f_0 \in \mathcal{F}$ with finite training risks (i.e., $\hat{\sigma}_{f_0} = 0$ and each $R^{e_i}(f_0) < \infty$), and*
2. *there are no arbitrarily negative training risks (i.e., $R_* := \inf_{f \in \mathcal{F}, i \in [m]} R^{e_i}(f) > -\infty$).*

For any $f \in \mathcal{F}$, let $R^*_f := \min_{i \in [m]} R^{e_i}(f)$ denote the smallest of the (empirical) risks of f across domains. Then,

$$\lim_{\alpha \to 1} \hat{\sigma}_{\hat{f}_\alpha} = 0 \quad and \quad \limsup_{\alpha \to 1} R^*_{\hat{f}_\alpha} \le R^*_{f_0}.$$

As in Prop. A.4, Assumption 1 depends on invariance of the label-generating mechanism across domains (as discussed further in Appendix A.2.3 below), while Assumption 2 automatically holds for most loss functions used in supervised learning. We now prove Prop. A.5.

Proof. By our assumption on the choice of bandwidth, it suffices to show that, as $\alpha \to 1$, $h_{\hat{f}_\alpha} \to 0$.

Let Φ denote the standard Gaussian CDF. Since Φ is non-decreasing, for all $x \in \mathbb{R}$,

$$\hat{F}_{\text{KDE},\hat{f}_\alpha}(x) = \frac{1}{m} \sum_{i=1}^{m} \Phi\left(\frac{x - R^{e_i}(\hat{f}_\alpha)}{h_{\hat{f}_\alpha}}\right) \le \Phi\left(\frac{x - R^*_{\hat{f}_\alpha}}{h_{\hat{f}_\alpha}}\right).$$

In particular, for $x = \hat{F}^{-1}_{\text{KDE},\hat{f}_\alpha}(\alpha)$, we have

$$\alpha = \hat{F}_{\text{KDE},\hat{f}_\alpha}(\hat{F}^{-1}_{\text{KDE},\hat{f}_\alpha}(\alpha)) \le \Phi\left(\frac{\hat{F}^{-1}_{\text{KDE},\hat{f}_\alpha}(\alpha) - R^*_f}{h_{\hat{f}_\alpha}}\right).$$

Inverting Φ and rearranging gives

$$R^*_f + h_{\hat{f}_\alpha} \cdot \Phi^{-1}(\alpha) \le \hat{F}^{-1}_{\text{KDE},\hat{f}_\alpha}(\alpha).$$

Hence, by definitions of $\hat{f}_\alpha$ and f_0,

$$R^*_f + h_{\hat{f}_\alpha} \cdot \Phi^{-1}(\alpha) \le \hat{F}^{-1}_{\text{KDE},\hat{f}_\alpha}(\alpha) \le \hat{F}^{-1}_{\text{KDE},f_0}(\alpha) = R^*_{f_0}. \tag{A.8}$$

Since, for $\alpha \ge 0.5$ we have that $h_{\hat{f}_\alpha} \cdot \Phi^{-1}(\alpha) \ge 0$, it follows that $R^*_{\hat{f}_\alpha} \le R^*_{f_0}$. Moreover, rearranging Inequality (A.8) and using the definition of R_*, we obtain

$$h_{\hat{f}_\alpha} \le \frac{R^*_{f_0} - R^*_{\hat{f}_\alpha}}{\Phi^{-1}(\alpha)} \le \frac{R^*_{f_0} - R_*}{\Phi^{-1}(\alpha)} \to 0$$

as $\alpha \to 1$. $\qquad\square$

A.2.3 Causal recovery

We now discuss and prove our main result, Thm. 4.4, regarding the conditions under which the causal predictor is the only minimal invariant-risk predictor. Together with Props. A.4 and A.5, this provides the conditions under which EQRM successfully performs "causal recovery", i.e., correctly recovers the true causal coefficients in a linear causal model of the data. As discussed in Appendix G.1.2, EQRM recovers the causal predictor by seeking *invariant risks* across domains, which differs from seeking *invariant functions* or coefficients (as in IRM [9]), with each approach having its pros and cons depending on the end goal. As we discuss below, Thm. 4.4 generalizes related results in the literature regarding causal recovery based on *invariant risks* [41, 54].

Assumptions (i–iv). The assumptions that Y is drawn from a linear SEM, that all causes of Y are observed, and that the loss is squared-error, while restrictive, are needed for all comparable causal recovery or identifiability results (of which we are aware) in the existing literature, namely Theorem 1 of Krueger et al. [41] and Theorem 2 of Peters et al. [54]. In fact, these assumptions are weaker than both Krueger et al. [41] (who assume a linear SEM for X *and* Y) and Peters et al. [54] (who assume a linear *Gaussian* SEM for X *and* Y). The assumption that the risk of the causal predictor is invariant across domains, also referred to as *domain homoskedasticity*, is necessary for any method inferring causality from the *invariance of risks* across domains. This is closely related to, but slightly weaker than, assuming that all interventions on the system that affect Y are captured by the distribution of X (i.e., that we observe all relevant causes of Y). For methods based on the *invariance of functions*, in particular the invariance of the conditional mean $\mathbb{E}[Y|\mathrm{Pa}(Y)]$ across domains [9, 65], this domain-homoskedasticity assumption is not required. Appendix G.1.2 compares methods based on invariant risks and to those based on invariant functions.

Assumption (v). In contrast to both Peters et al. [54] and Krueger et al. [41], we do not require specific types of interventions on the covariates. In particular, our main assumption on the distributions of the covariates across domains, namely that the system of d-variate quadratic equations in (4.3) has a unique solution, is more general than these comparable results. For example, whereas both Peters et al. [54] and Krueger et al. [41] require one or more separate interventions for *every* covariate X_j, Example 4 below shows that we only require interventions on the subset of covariates that are effects of Y, while weaker conditions suffice for other covariates. Although this generality comes at the cost of abstraction, we now provide some concrete examples with different types of interventions to aid understanding. Note that, to simply calculations and provide a more intuitive form, (4.3) of Thm. 4.4 assumes, without loss of generality, that all covariates are standardized to have mean 0 and variance 1, except where interventions change these. We can, however, rewrite (4.3) of Thm. 4.4 in a slightly more general form which does not require this assumption of standardized covariates:

$$
\begin{aligned}
0 \geq & x^\mathsf{T} \mathbb{E}_{X \sim e_1}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N, X \sim e_1}[XN] \\
= & \cdots \\
= & x^\mathsf{T}\mathbb{E}_{X \sim e_m}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N, X \sim e_m}[XN].
\end{aligned}
\tag{A.9}
$$

We now present a number of concrete examples or special cases in which Assumption (v) of Thm. 4.4 would be satisfied, using this slightly more general form. In each example, we assume that variables are generated according to an SCM with an acyclic causal graph, as described in Appendix A.1.

1. *No effects of Y.* In the case that there are no effects of Y (i.e., no X_j is a causal descendant of Y, and hence each X_j is uncorrelated with N), it suffices for there to exists at least one environment e_i in which the covariance $\mathrm{Cov}_{X \sim e}[X]$ has full rank. These are standard conditions for identifiability in linear regression. More generally, it suffices for $\sum_{i=1}^{m} \mathrm{Cov}_{X \sim e_i}[X]$ to have full rank; this is the same condition one would require if simply performing linear regression on the pooled data from all m environments. Intuitively, this full-rank condition guarantees that the observed covariate values are sufficiently uncorrelated to distinguish the effect of each covariate on the response Y. However, it does not necessitate interventions on the covariates, which are necessary to identify the *direction of causation* in a linear model; hence, this full-rank condition fails to imply causal recovery in the presence of effects of Y. See Appendix G.1.2 for a concrete example of this failure.

2. Hard *interventions.* For each covariate X_j, compared to some baseline environment e_0, there is some environment e_{X_j} arising from a hard single-node intervention $do(X_j = z)$, with $z \neq 0$. If X_j is any leaf node in the causal DAG, then in e_{X_j}, X_j is uncorrelated with N and with each X_k

$(k \neq j)$, so the inequality in (A.9) gives

$$0 \geq x^\mathsf{T}\mathbb{E}_{X\sim e_{X_j}}[XX^\mathsf{T}]x = x_j^2 z^2 + x_{-j}^\mathsf{T}\mathbb{E}_{X\sim e_0}[XX^\mathsf{T}]x_{-j}.$$

Since the matrix $\mathbb{E}_{X\sim e}[XX^\mathsf{T}]$ is positive semidefinite (and $z \neq 0$ implies $z^2 > 0$), it follows that $x_j = 0$. The terms in (A.9) containing x_j thus vanish, and iterating this argument for parents of leaf nodes in the causal DAG, and so on, gives $x = 0$. This condition is equivalent to that in Theorem 2(a) of Peters et al. [54] and is a strict improvement over Corollary 2 of Yin et al. [65] and Theorem 1 of Krueger et al. [41], which respectively require two and three distinct hard interventions on each variable.

3. *Shift interventions.* For each covariate X_j, compared to some baseline environment e_0, there is some environment e_{X_j} consisting of the shift intervention $X_j \leftarrow g_j(\mathrm{Pa}(X_j), N_j) + z$, for some $z \neq 0$. Recalling that we assumed each covariate was centered (i.e., $\mathbb{E}_{X\sim e_0}[X_k] = 0$) in e_0, if X_j is any leaf node in the causal DAG, then every other covariate remains centered in e_{X_j} (i.e., $\mathbb{E}_{X\sim e_{X_j}}[X_k] = 0$ for each $k \neq j$). Hence, the excess risk is

$$x^\mathsf{T}\mathbb{E}_{X\sim e_{X_j}}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_{X_j}}[XN] = x_j^2 z^2 + x^\mathsf{T}\mathbb{E}_{X\sim e_0}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_0}[XN].$$

Since, by (A.9),

$$x^\mathsf{T}\mathbb{E}_{X\sim e_0}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_0}[XN] = x^\mathsf{T}\mathbb{E}_{X\sim e_{X_j}}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_{X_j}}[XN],$$

it follows that $x_j^2 z^2 = 0$, and so, since $z \neq 0$, $x_j = 0$. As above, the terms in (A.9) containing x_j thus vanish, and iterating this argument for parents of leaf nodes in the causal DAG, and so on, gives $x = 0$. This condition is equivalent to the additive setting of Theorem 2(b) of Peters et al. [54].

4. *Noise interventions.* Suppose that each covariate is related to its causal parents through an additive noise model; i.e.,

$$X_j = g_j(\mathrm{Pa}(X_j)) + N_j,$$

where $\mathbb{E}[N_j] = 0$ and $0 < \mathbb{E}[N^2] < \infty$. Theorem 2(b) of Peters et al. [54] considers "noise" interventions, of the form

$$X_j \leftarrow g_j(\mathrm{Pa}(X_j)) + \sigma N_j,$$

where $\sigma^2 \neq 1$. Suppose that, for each covariate X_j, compared to some baseline environment e_0, there exists an environment e_{X_j} consisting of the above noise intervention. If X_j is any leaf node in the causal DAG, then, since we assumed $\mathbb{E}_{X\sim e_0}[X_j^2] = 1$,

$$\begin{aligned}
&x^\mathsf{T}\mathbb{E}_{X\sim e_{X_j}}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_{X_j}}[XN] \\
&= (\sigma^2 - 1)x_j^2\mathbb{E}[N_j^2] + x^\mathsf{T}\mathbb{E}_{X\sim e_0}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_0}[XN].
\end{aligned}$$

Hence, the system (A.9) implies $0 = (\sigma^2 - 1)x_j^2\mathbb{E}[N_j^2]$. Since $\sigma^2 \neq 1$ and $\mathbb{E}[N_j^2] > 0$, it follows that $x_j = 0$.

5. *Scale interventions.* For each covariate X_j, compared to some baseline environment e_0, there exist two environments $e_{X_j,i}$ ($i \in \{1,2\}$) consisting of scale interventions $X_j \leftarrow \sigma_i g_j(\mathrm{Pa}(X_j), N_j)$, for some $\sigma_i \neq \pm 1$, with $\sigma_1 \neq \sigma_2$. If X_j is any leaf node in the causal DAG, then, since we assumed $\mathbb{E}_{X\sim e_0}[X_j^2] = 1$,

$$\begin{aligned}
&x^\mathsf{T}\mathbb{E}_{X\sim e_{X_j}}[XX^\mathsf{T}]x + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_{X_j}}[XN] \\
&= (\sigma_i^2 - 1)x_j^2 + 2(\sigma_i - 1)x_j\mathbb{E}_{X\sim e_0}[X_j X_{-j}^\mathsf{T}]x_{-j} + x^\mathsf{T}\mathbb{E}_{X\sim e_0}[XX^\mathsf{T}]x \\
&\quad + 2(\sigma_i - 1)x_j\mathbb{E}_{N,X\sim e_0}[X_j N] + 2x^\mathsf{T}\mathbb{E}_{N,X\sim e_0}[XN].
\end{aligned}$$

Hence, the system (A.9) implies

$$0 = (\sigma_i^2 - 1)x_j^2 + 2(\sigma_i - 1)x_j\left(\mathbb{E}_{X\sim e_0}[X_j X_{-j}^\mathsf{T}]x_{-j}^\mathsf{T} + \mathbb{E}_{N,X\sim e_0}[X_j N]\right).$$

Figure 5: Example causal DAG with various types of covariates. X_1 and X_3 are the parents of Y, and so the true causal coefficient β has only two non-zero coordinates β_1 and β_3. X_1, X_2, and X_3 are ancestors of Y. X_5, X_6, and X_7 are effects, or descendants, of Y and are the only covariates for which $\mathbb{E}[X_j N]$ can be nonzero; hence, X_5, X_6, and X_7 are the only covariates on which interventions are generally necessary.

Since $\sigma_i^2 \neq 1$, if $x_j \neq 0$, then solving for x_j gives

$$x_j = -2\frac{\mathbb{E}_{X\sim e_0}[X_j X_{-j}^{\mathsf{T}}]x_{-j}^{\mathsf{T}} + \mathbb{E}_{N,X\sim e_0}[X_j N]}{\sigma_i + 1}.$$

Since $\sigma_1 \neq \sigma_2$, this is possible only if $x_j = 0$. This provides an example where a single intervention per covariate would be insufficient to guarantee causal recovery, but two distinct interventions per covariate suffice.

6. *Sufficiently uncorrelated causes and intervened-upon effects.* Suppose that, within the true causal DAG, $\mathrm{De}(Y) \subseteq [d]$ indexes the *descendants*, or *effects* of Y (e.g., in Figure 5, $\mathrm{De}(Y) = \{5,6,7\}$). Suppose that for every $j \in \mathrm{De}(Y)$, compared to a single baseline environment e_0, there is a environment e_{X_j} consisting of either a $do(X_j = z)$ intervention or a shift intervention $X_j \leftarrow g_j(\mathrm{Pa}(X_j), N_j) + z$, with $z \neq 0$ and that the matrix

$$\sum_{i=1}^{m} \mathrm{Cov}_{X\sim e_i}\left[X_{[d]\setminus\mathrm{De}(Y)}\right] \tag{A.10}$$

has full rank. Then, as argued in the previous two cases, for each $j \in \mathrm{De}(Y)$, $x_j = 0$. Moreover, for any $j \in [d]\setminus\mathrm{De}(Y)$, $\mathbb{E}[X_j N] = 0$, and so the system of equations (A.9) reduces to

$$0 \geq x_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}} \mathbb{E}_{X\sim e_1}\left[X_{[d]\setminus\mathrm{De}(Y)} X_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}}\right] x_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}}$$
$$= \cdots$$
$$= x_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}} \mathbb{E}_{X\sim e_m}\left[X_{[d]\setminus\mathrm{De}(Y)} X_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}}\right] x_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}}.$$

Since each $\mathbb{E}_{X\sim e_m}\left[X_{[d]\setminus\mathrm{De}(Y)} X_{[d]\setminus\mathrm{De}(Y)}^{\mathsf{T}}\right]$ is positive semidefinite, the solution $x = 0$ to this reduced system of equations is unique if (and only if) the matrix (A.10) has full rank. This example demonstrates that interventions are only needed for effect covariates, while a weaker full-rank condition suffices for the remaining ones. In many practical settings, it may be possible to determine *a priori* that a particular covariate X_j is not a descendant of Y; in this case, the practitioner need not intervene on X_j, as long as sufficiently diverse observational data on X_j is available. To the best of our knowledge, this does not follow from any existing results in the literature, such as Theorem 2 of Peters et al. [54] or Corollary 2 of [65].

We conclude this section with the proof of Thm. 4.4:

Proof. Under the linear SEM setting with squared-error loss, for any estimator $\hat{\beta}$,

$$\mathcal{R}^e(\hat{\beta}) = \mathbb{E}_{N,X\sim e}\left[\left((\beta - \hat{\beta})^{\mathsf{T}}X + N\right)^2\right]$$
$$= \mathbb{E}_{X\sim e}\left[\left((\beta - \hat{\beta})^{\mathsf{T}}X\right)^2\right] + 2\mathbb{E}_{N,X\sim e}\left[(\beta - \hat{\beta})^{\mathsf{T}}XN\right] + \mathbb{E}_N\left[N^2\right].$$

Since the second moment of the noise term $\mathbb{E}_N[N^2]$ is equal to the risk $\mathbb{E}_{(X,Y)\sim e}[(Y-\beta^T X)^2]$ of the causal predictor β, by the definition of $Y = \beta^T X + N$, we have that $\mathbb{E}_N[N^2]$ is invariant across environments. Thus, minimizing the squared error risk $\mathcal{R}^e(\hat{\beta})$ is equivalent to minimizing the excess risk

$$\mathbb{E}_{X\sim e}\left[\left((\beta-\hat{\beta})^\intercal X\right)^2\right] + 2\mathbb{E}_{N,X\sim e}\left[(\beta-\hat{\beta})^\intercal XN\right]$$
$$= (\beta-\hat{\beta})^\intercal \mathbb{E}_{X\sim e}[XX^\intercal](\beta-\hat{\beta}) + 2(\beta-\hat{\beta})^\intercal\mathbb{E}_{N,X\sim e}[XN]$$

over estimators $\hat{\beta}$. Since the true coefficient β is an invariant-risk predictor with 0 excess risk, if $\hat{\beta}$ is a minimal invariant-risk predictor, it has at most 0 invariant excess risk; i.e.,

$$0 \geq (\beta-\hat{\beta})^\intercal\mathbb{E}_{X\sim e_1}[XX^\intercal](\beta-\hat{\beta}) + 2(\beta-\hat{\beta})^\intercal\mathbb{E}_{N,X\sim e_1}[XN]$$
$$= \cdots$$
$$= (\beta-\hat{\beta})^\intercal\mathbb{E}_{X\sim e_m}[XX^\intercal](\beta-\hat{\beta}) + 2(\beta-\hat{\beta})^\intercal\mathbb{E}_{N,X\sim e_m}[XN]. \tag{A.11}$$

By Assumption (v), the unique solution to this is $\beta - \hat{\beta} = 0$; i.e., $\hat{\beta} = \beta$. $\qquad\square$

B   On the equivalence of different DG formulations

In Section 3, we claimed that under mild conditions, the minimax domain generalization problem in (2.2) is equivalent to the essential supremum problem in (3.1). In this subsection, we formally describe the conditions under which these two problems are equivalent. We also highlight several examples in which the assumptions needed to prove this equivalence hold.

Specifically, this appendix is organized as follows. First, in § B.1 we offer a more formal analysis of the equivalence between the probable domain general problems in (3.2) and (QRM). Next, in § B.2, we connect the domain generalization problem in (2.2) to the essential supremum problem in (3.1).

B.1   Connecting formulations for QRM via a push-forward measure

To begin, we consider the abstract measure space $(\mathcal{E}_{\text{all}}, \mathcal{A}, \mathbb{Q})$, where $\mathcal{A}$ is a σ-algebra defined on the subsets of $\mathcal{E}_{\text{all}}$. Recall that in our setting, the domains $e \in \mathcal{E}_{\text{all}}$ are assumed to be drawn from the distribution $\mathbb{Q}$. Given this setting, in § 3 we introduced the probable domain generalization problem in (3.2), which we rewrite below for convenience:

$$\min_{f\in\mathcal{F}, t\in\mathbb{R}} t \qquad \text{subject to} \qquad \Pr_{e\sim\mathbb{Q}}\{\mathcal{R}^e(f) \leq t\} \geq \alpha. \tag{B.1}$$

Our objective is to formally show that this problem is equivalent to (QRM). To do so, for each $f \in \mathcal{F}$, let consider a second measurable space $(\mathbb{R}_+, \mathcal{B})$, where $\mathbb{R}_+$ denotes the set of non-negative real numbers and $\mathcal{B}$ denotes the Borel σ-algebra over this space. For each $f \in \mathcal{F}$, we can now define the $(\mathbb{R}_+, \mathcal{B})$-valued random variable[6] $G_f : \mathcal{E}_{\text{all}} \to \mathbb{R}_+$ via

$$G_f : e \mapsto \mathcal{R}^e(f) = \mathbb{E}_{\mathbb{P}(X^e, Y^e)}[\ell(f(X^e), Y^e)]. \tag{B.2}$$

Concretely, G_f maps an domain e to the corresponding risk $\mathcal{R}^e(f)$ of f in that domain. In this way, G_f effectively summarizes e by its effect on our predictor's risk, thus projecting from the often-unknown and potentially high-dimensional space of possible distribution shifts or interventions to the one-dimensional space of observed, real-valued risks. However, note that G_f is not necessarily injective, meaning that two domains e_1 and e_2 may be mapped to the same risk value under G_f.

The utility of defining G_f is that it allows us to formally connect (3.2) with (QRM) via a push-forward measure through G_f. That is, given any $f \in \mathcal{F}$, we can define the measure[7]

$$\mathbb{T}_f =^d G_f \# \mathbb{Q} \tag{B.3}$$

---

[6]For brevity, we will assume that G_f is always measurable with respect to the underlying σ-algebra $\mathcal{A}$.
[7]Here $\mathbb{T}_f$ is defined over the induced measurable space $(\mathbb{R}_+, \mathcal{B})$.

where # denotes the *push-forward* operation and $=^d$ denotes equality in distribution. Observe that the relationship in (B.3) allows us to explicitly connect $\mathbb{Q}$—the often unknown distribution over (potentially high-dimensional and/or non-Euclidean) domain shifts in Fig. 1b—to $\mathbb{T}_f$—the distribution over real-valued risks in Fig. 1c, from which we can directly observe samples. In this way, we find that for each $f \in \mathcal{F}$,

$$\Pr_{e \sim \mathbb{Q}} \{ \mathcal{R}^e(f) \leq t \} = \Pr_{R \sim \mathbb{T}_f} \{ R \leq t \}. \tag{B.4}$$

This relationship lays bare the connection between (3.2) and (QRM), in that the domain or environment distribution $\mathbb{Q}$ can be replaced by a distribution over risks $\mathbb{T}_f$.

B.2 Connecting (2.2) to the essential supremum problem (3.1)

We now study the relationship between (2.2) and (3.1). In particular, in § B.2.1 and § B.2.2, we consider the distinct settings wherein $\mathcal{E}_{\text{all}}$ comprises continuous and discrete spaces respectively.

B.2.1 Continuous domain sets $\mathcal{E}_{\text{all}}$

When $\mathcal{E}_{\text{all}}$ is a continuous space, it can be shown that (2.2) and (3.1) are *equivalent* whenever: (a) the map G_f defined in Section B.1 is continuous; and (b) the measure $\mathbb{Q}$ satisfies very mild regularity conditions.

The case when $\mathbb{Q}$ is the Lebesgue measure. Our first result concerns the setting in which $\mathcal{E}_{\text{all}}$ is a subset of Euclidean space and where $\mathbb{Q}$ is chosen to be the Lebesgue measure on $\mathcal{E}_{\text{all}}$. We summarize this result in the following proposition.

Proposition B.1. *Let us assume that the map G_f is continuous for each $f \in \mathcal{F}$. Further, let $\mathbb{Q}$ denote the Lebesgue measure over $\mathcal{E}_{\text{all}}$; that is, we assume that domains are drawn uniformly at random from $\mathcal{E}_{\text{all}}$. Then (2.2) and (3.1) are equivalent.*

Proof. To prove this claim, it suffices to show that under the assumptions in the statement of the proposition, it holds for any $f \in \mathcal{F}$ that

$$\sup_{e \in \mathcal{E}_{\text{all}}} R^e(f) = \operatorname*{ess\,sup}_{e \sim \mathbb{Q}} R^e(f). \tag{B.5}$$

To do so, let us fix an arbitrary $f \in \mathcal{F}$ and write

$$A := \sup_{e \in \mathcal{E}_{\text{all}}} R^e(f) \quad \text{and} \quad B := \operatorname*{ess\,sup}_{e \sim \mathbb{Q}} R^e(f). \tag{B.6}$$

At a high-level, our approach is to show that $B \leq A$, and then that $A \leq B$, which together will imply the result in (B.5). To prove the first inequality, observe that by the definition of the supremum, it holds that $R^e(f) \leq A \ \forall e \in \mathcal{E}_{\text{all}}$. Therefore, $\mathbb{Q}\{e \in \mathcal{E}_{\text{all}} : R^e(f) > A\} = 0$, which directly implies that $B \leq A$. Now for the second inequality, let $\epsilon > 0$ be arbitrarily chosen. Consider that due to the continuity of G_f, there exists an $e_0 \in \mathcal{E}_{\text{all}}$ such that

$$R^{e_0}(f) + \epsilon > A. \tag{B.7}$$

Now again due to the continuity of G_f, we can choose a ball $\mathcal{B}_\epsilon \subset \mathcal{E}_{\text{all}}$ centered at e_0 such that $|R^e(f) - R^{e_0}(f)| \leq \epsilon \ \forall e \in \mathcal{B}_\epsilon$. Given such a ball, observe that $\forall e \in \mathcal{B}_\epsilon$, it holds that

$$R^e(f) \geq R^{e_0}(f) - \epsilon > A - 2\epsilon \tag{B.8}$$

where the first inequality follows from the reverse triangle inequality and the second inequality follows from (B.7). Because $\mathbb{Q}\{e \in \mathcal{B}_\epsilon : R^e(f) > A - 2\epsilon\} > 0$, it directly follows that $A - 2\epsilon \leq B$. As $\epsilon > 0$ was chosen arbitrarily, this inequality holds for any $\epsilon > 0$, and thus we can conclude that $A \leq B$, completing the proof. $\qquad\square$

Generalizing Prop. B.1 to other measure $\mathbb{Q}$. We note that this proof can be generalized to measures $\mathbb{Q}$ other than the Lebesgue measure. Indeed, the result holds for any measure $\mathbb{Q}$ taking support on $\mathcal{E}_{\text{all}}$ for which it holds that $\mathbb{Q}$ places non-zero probability mass on any closed subset of $\mathcal{E}_{\text{all}}$. This would be the case, for instance, if $\mathbb{Q}$ was a truncated Gaussian distribution with support on $\mathcal{E}_{\text{all}}$. Furthermore, if we let $\mathbb{L}$ denote the Lebesgue measure on $\mathcal{E}_{\text{all}}$, then another more general instance of this property occurs whenever $\mathbb{L}$ is absolutely continuous with respect to $\mathbb{Q}$, i.e., whenever $\mathbb{L} \ll \mathbb{Q}$.

Corollary B.2. *Let us assume that $\mathbb{Q}$ places nonzero mass on every open ball with radius strictly larger than zero. Then under the continuity assumptions of Prop. B.1, it holds that (2.2) and (3.1) are equivalent.*

Proof. The proof of this fact follows along the same lines as that of Prop. B.1. In particular, the same argument shows that $B \leq A$. Similarly, to show that $A \leq B$, we can use the same argument, noting that $\mathbb{Q}\{e \in \mathcal{B}_\epsilon : R^e(f) > A - 2\epsilon\}$ continues to hold, due to our assumption that $\mathbb{Q}$ places nonzero mass on $\mathcal{B}_\epsilon$. $\qquad\square$

Examples. We close this subsection by considering several real-world examples in which the conditions of Prop. B.1 hold. In particular, we focus on examples in the spirit of "Model-Based Domain Generalization" [22]. In this setting, it is assumed that the variation from domain to domain is parameterized by a *domain transformation model* $x^e \mapsto D(x^e, e') =: x^{e'}$, which maps the covariates x^e from a given domain $e \in \mathcal{E}_{\mathrm{all}}$ to another domain $e' \in \mathcal{E}_{\mathrm{all}}$. As discussed in [22], domain transformation models cover settings in which inter-domain variation is due to *domain shift* [121, §1.8]. Indeed, under this model (formally captured by Assumptions 4.1 and 4.2 in [22]), the domain generalization problem in (2.2) can be equivalently rewritten as

$$\min_{f \in \mathcal{F}} \max_{e \in \mathcal{E}_{\mathrm{all}}} \mathbb{E}_{(X,Y)}[\ell(f(D(X,e)),Y)]. \tag{B.9}$$

For details, see Prop. 4.3 in [22]. In this problem, (X,Y) denote an underlying pair of random variables such that

$$\mathbb{P}(X^e) =^d D \#(\mathbb{P}(X), \delta(e)) \quad \text{and} \quad \mathbb{P}(Y^e) =^d \mathbb{P}(Y) \tag{B.10}$$

for each $e \in \mathcal{E}_{\mathrm{all}}$ where $\delta(e)$ is a Dirac measure placed at $e \in \mathcal{E}_{\mathrm{all}}$. Now, turning our attention back to Prop. B.1, we can show the following result for (B.9).

Remark B.3. Let us assume that the map $e \mapsto D(\cdot, e)$ is continuous with respect to a metric $d_{\mathcal{E}_{\mathrm{all}}}(e, e')$ on $\mathcal{E}_{\mathrm{all}}$ and that $x \mapsto \ell(x, \cdot)$ is continuous with respect to the absolute value. Further, assume that each predictor $f \in \mathcal{F}$ is continuous in the standard Euclidean metric on $\mathbb{R}^d$. Then (2.2) and (3.1) are equivalent.

Proof. By Prop. B.1, it suffices to show that the map

$$G_f : e \mapsto \mathbb{E}_{(X,Y)}[\ell(f(D(X,e)),Y)] \tag{B.11}$$

is a continuous function. To do so, recall that the composition of continuous functions is continuous, and therefore we have, by the assumptions listed in the above remark, that the map $e \mapsto \ell(f(D(x,e)),y)$ is continuous for each $(x,y) \sim (X,Y)$. To this end, let us define the function $h_f(x,y,e) := \ell(f(D(x,e)),y)$ and let $\epsilon > 0$. By the continuity of h_f in e, there exists a $\delta = \delta(\epsilon) > 0$ such that $|h_f(x,y,e) - h_f(x,y,e')| < \epsilon$ whenever $d_{\mathcal{E}_{\mathrm{all}}}(e,e') < \delta$. Now observe that

$$\left| \mathbb{E}_{(X,Y)}[h_f(X,Y,e)] - \mathbb{E}_{(X,Y)}[h_f(X,Y,e')] \right| \tag{B.12}$$

$$= \left| \int_{\mathcal{E}_{\mathrm{all}}} h_f(X,Y,e) \mathrm{d}\mathbb{P}(X,Y) - \int_{\mathcal{E}_{\mathrm{all}}} h_f(X,Y,e') \mathrm{d}\mathbb{P}(X,Y) \right| \tag{B.13}$$

$$= \left| \int_{\mathcal{E}_{\mathrm{all}}} (h_f(X,Y,e) - h_f(X,Y,e')) \mathrm{d}\mathbb{P}(X,Y) \right| \tag{B.14}$$

$$\leq \int_{\mathcal{E}_{\mathrm{all}}} \left| h_f(X,Y,e) - h_f(X,Y,e') \right| \mathrm{d}\mathbb{P}(X,Y). \tag{B.15}$$

Therefore, whenever $d_{\mathcal{E}_{\mathrm{all}}}(e,e') < \delta$ it holds that

$$\left| \mathbb{E}_{(X,Y)}[h_f(X,Y,e)] - \mathbb{E}_{(X,Y)}[h_f(X,Y,e')] \right| \leq \int_{\mathcal{E}_{\mathrm{all}}} \epsilon \mathrm{d}\mathbb{P}(X,Y) = \epsilon \tag{B.16}$$

by the monotonicity of expectation. This completes the proof that G_f is continuous. $\qquad\square$

In this way, provided that the risks in each domain vary in a continuous way through e, (2.2) and (3.1) are equivalent. As a concrete example, consider an image classification setting in which the variation from domain to domain corresponds to different rotations of the images. This is the case, for instance, in the `RotatedMNIST` dataset [38, 126], wherein the training domains correspond to different rotations of the `MNIST` digits. Here, a domain transformation model D can be defined by

$$D(x,e) = R(e)x \quad \text{where} \quad e \in \mathcal{E}_{\text{all}} \subseteq [0, 2\pi), \tag{B.17}$$

and where $R(e)$ is a rotation matrix. In this case, it is clear that D is a continuous function of e (in fact, the map is *linear*), and therefore the result in (B.3) holds.

B.2.2 Discrete domain sets $\mathcal{E}_{\text{all}}$

When $\mathcal{E}_{\text{all}}$ is a discrete set, the conditions we require for (2.2) and (3.1) to be equivalent are even milder. In particular, the only restriction we place on the problems is that $\mathbb{Q}$ must place non-zero mass on each element of $\mathcal{E}_{\text{all}}$; that is, $\mathbb{Q}(e) > 0 \ \forall e \in \mathcal{E}_{\text{all}}$. We state this more formally below.

Proposition B.4. *Let us assume that $\mathcal{E}_{all}$ is discrete, and that $\mathbb{Q}$ is such that $\forall e \in \mathcal{E}_{all}$, it holds that $\mathbb{Q}(e) > 0$. Then it holds that (2.2) and (3.1) are equivalent.*

C Notes on KDE bandwidth selection

In our setting, we are interested in bandwidth-selection methods which: (i) work well for 1D distributions and small sample sizes m; and (ii) guarantee recovery of the causal predictor as $\alpha \to 1$ by satisfying $h_f \to 0 \implies \hat{\sigma}_f \to 0$, where h_f is the data-dependent bandwidth and $\hat{\sigma}_f$ is the sample standard deviation (see Appendices A.2.2 and A.2.3). We thus investigated three popular bandwidth-selection methods: (1) the Gaussian-optimal rule [64], $h_f = (4/3m)^{0.2} \cdot \hat{\sigma}_f$; (2) Silverman's rule-of-thumb [64], $h_f = m^{-0.2} \cdot \min(\hat{\sigma}_f, \frac{\text{IQR}}{1.34})$, with IQR the interquartile range; and (3) the median-heuristic [127–129], which sets the bandwidth to be the median pairwise-distance between data points. Note that many sensible methods exist, as do more complete studies on bandwidth selection—see e.g. [64].

For (i), we found Silverman's rule-of-thumb [64] to perform very well, the Gaussian-optimal rule [64] to perform well, and the median-heuristic [127–129] to perform poorly. For (ii), only the Gaussian-optimal rule satisfies $h_f \to 0 \implies \hat{\sigma}_f \to 0$. Thus, in practice, we use either the Gaussian-optimal rule (particularly when causal predictor's are sought as $\alpha \to 1$), or Silverman's rule-of-thumb.

D Generalization bounds

This appendix states and proves our main generalization bound, Theorem D.1. Theorem D.1 applies for many possible estimates $\widehat{\mathbb{T}}_f$, and we further show how to apply Theorem D.1 to the specific case of using a kernel density estimate.

D.1 Main generalization bound and proof

Suppose that, from each of N IID environments $e_1, ..., e_N \sim \mathbb{P}(e)$, we observe n IID labeled samples $(X_{i,1}, Y_{i,1}), ..., (X_{n,1}, Y_{n,1}) \sim \mathbb{P}(X^e, Y^e)$. Fix a hypothesis class $\mathcal{F}$ and confidence level $\alpha \in [0, 1]$. For any hypothesis $f : \mathcal{X} \to \mathcal{Y}$, define the *empirical risk on environment e_i* by

$$\widehat{\mathcal{R}}^{e_i}(f) := \frac{1}{n} \sum_{j=1}^{n} \ell\left(Y_{i,j}, f(X_{i,j})\right), \quad \text{for each} \quad i \in [N].$$

Throughout this section, we will abbreviate the distribution $F_{\mathbb{T}_f}(t) = \Pr_e[\mathcal{R}^e(f) \le t]$ of f's risk by $F_f(t)$ and its estimate $F_{\widehat{\mathbb{T}}_f}$, computed from the observed empirical risks $\widehat{\mathcal{R}}^{e_1}(f), ..., \widehat{\mathcal{R}}^{e_N}(f)$, by $\widehat{F}_f$.

We propose to select a hypothesis by minimizing this over our hypothesis class:

$$\widehat{f} := \underset{f \in \mathcal{F}}{\arg\min} \, F_{\widehat{\mathbb{T}}_f}^{-1}(\alpha). \tag{D.1}$$

In this section, we prove a uniform generalization bound, which in particular, provides conditions under which the estimator (D.1) generalizes uniformly over $\mathcal{F}$. Because the novel aspect of the present paper is the notion of generalizing *across* environments, we will take for granted that the hypothesis class $\mathcal{F}$ generalizes uniformly *within* each environments (i.e., that each $\sup_{f \in \mathcal{F}} \mathcal{R}^{e_i}(f) - \widehat{\mathcal{R}}^{e_i}(f)$ can be bounded with high probability); myriad generalization bounds from learning theory can be used to show this.

Theorem D.1. *Let $\mathcal{G} := \{\widehat{F}(\mathcal{R}^{e_1}(f), \mathcal{R}^{e_2}(f), ..., \mathcal{R}^{e_N}(f)) : f \in \mathcal{F}, e_1, ..., e_n \in \mathcal{E}_{all}\}$ denote the class of possible estimated risk distributions over N environments, and, for any $\epsilon > 0$, let $\mathcal{N}_\epsilon(\mathcal{G})$ denote the ϵ-covering number of $\mathcal{G}$ under $\mathcal{L}_\infty(\mathbb{R})$. Suppose the class $\mathcal{F}$ generalizes uniformly within environments; i.e., for any $\delta > 0$, there exists $t_{n,\delta,\mathcal{F}}$ such that*

$$\operatorname*{ess\,sup}_{e} \Pr_{\{(X_j, Y_j)\}_{j=1}^n \sim \mathbb{P}(X^e, Y^e)} \left[\sup_{f \in \mathcal{F}} R^e(f) - \widehat{\mathcal{R}}^e(f) > t_{n,\delta,\mathcal{F}} \right] \leq \delta.$$

Let

$$\operatorname{Bias}(\mathcal{F}, \widehat{F}) := \sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f(t) - \mathbb{E}_{e_1, ..., e_N}[\widehat{F}_f(t)]$$

denote the worst-case bias of the estimator $\widehat{F}$ over the class f. Noting that $\widehat{F}_f$ is a function of the empirical risk CDF

$$\widehat{Q}_f(t) := \frac{1}{N} \sum_{i=1}^N 1\{\mathcal{R}^{e_i}(f) \leq t\},$$

suppose that the function $\widehat{Q}_f \mapsto \widehat{F}_f$ is L-Lipschitz under $\mathcal{L}_\infty(\mathbb{R})$. Then, for any $\epsilon, \delta > 0$,

$$\Pr_{\substack{e_1, ..., e_N \\ \{(X_j, Y_j)\}_{j=1}^n \sim \mathbb{P}(X^{e_i}, Y^{e_i})}} \left[\sup_{f \in \mathcal{F}} F_f^{-1}\left(\alpha - B(\mathcal{F}, \widehat{F}) - \epsilon\right) - \widehat{F}_f^{-1}(\alpha) > t_{n, \frac{\delta}{N}, \mathcal{F}} \right] \leq \delta + 8\mathcal{N}_{\epsilon/16}(\mathcal{G}) e^{-\frac{N\epsilon^2}{64L}}.$$

$$(D.2)$$

The key technical observation of Theorem D.1 is that we can pull the supremum over $\mathcal{F}$ outside the probability by incurring a $\mathcal{N}_{\epsilon/16}(\mathcal{G})$ factor increase in the probability of failure. To ensure $\mathcal{N}_{\epsilon/16}(\mathcal{G}) < \infty$, we need to limit the space of possible empirical risk profiles $\mathcal{G}$ (e.g., by kernel smoothing), incurring an additional bias term $B(\mathcal{F}, \widehat{F})$. As we demonstrate later, for common distribution estimators, such as kernel density estimators, one can bound the covering number $\mathcal{N}_{\epsilon/16}(\mathcal{G})$ in Inequality (D.2) by standard methods, and the Lipschitz constant L is typically 1. Under mild (e.g., smoothness) assumptions on the family of possible true risk profiles, one can additionally bound the Bias Term, again by standard arguments.

Before proving Theorem D.1, we state two standard lemmas used in the proof:

Lemma D.2 (Symmetrization; Lemma 2 of [130]). *Let X and X' be independent realizations of a random variable with respect to which $\mathcal{F}$ is a family of integrable functions. Then, for any $\epsilon > 0$,*

$$\Pr\left[\sup_{f \in \mathcal{F}} f(X) - \mathbb{E}\, f(X) > \epsilon \right] \leq 2 \Pr\left[\sup_{f \in \mathcal{F}} f(X) - f(X') > \frac{\epsilon}{2} \right].$$

Lemma D.3 (Dvoretzky–Kiefer–Wolfowitz (DKW) Inequality; Corollary 1 of [131]). *Let $X_1, ..., X_n$ be IID $\mathbb{R}$-valued random variables with CDF P. Then, for any $\epsilon > 0$,*

$$\Pr\left[\sup_{t \in \mathbb{R}} \left| F_f(t) - \frac{1}{n} \sum_{i=1}^n 1\{X_i \leq t\} \right| > \epsilon \right] \leq 2e^{-2n\epsilon^2}.$$

We now prove our main result, Theorem D.1.

Proof of Theorem D.1. For convenience, let $F_f(t) := \mathbb{P}_{e \sim \mathbb{P}(e)}[R^e(f) \leq t]$. In preparation for Symmetrization, for any $f \in \mathcal{F}$, let $\widehat{F}'_f$ denote $\widehat{F}_f$ computed on an independent "ghost" sample

$e'_1, ..., e'_N \sim \mathbb{P}(e)$. Then,

$$\Pr_{e_1,...,e_N} \left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} \mathbb{E}_{e_1,...,e_N} \left[\widehat{F}_f(t) \right] - \widehat{F}_f(t) > \epsilon \right] \tag{D.3}$$

$$\leq 2 \Pr_{\substack{e_1,...,e_N \\ e'_1,...,e'_N}} \left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} \widehat{F}'_f(t) - \widehat{F}_f(t) > \epsilon/2 \right] \tag{D.4}$$

$$\leq 2 \Pr_{\substack{e_1,...,e_N \\ e'_1,...,e'_N}} \left[\sup_{f \in \mathcal{F}} \left\| \widehat{F}'_f - \widehat{F}_f \right\|_\infty > \epsilon/2 \right] \tag{D.5}$$

$$\leq 2 \Pr_{\substack{e_1,...,e_N \\ e'_1,...,e'_N}} \left[\sup_{f \in \mathcal{F}} \epsilon/8 + \left\| D\widehat{F}'_f - D\widehat{F}_f \right\|_\infty > \epsilon/2 \right] \tag{D.6}$$

$$\leq 2\mathcal{N}_{\epsilon/16} \sup_{f \in \mathcal{F}} \Pr_{\substack{e_1,...,e_N \\ e'_1,...,e'_N}} \left[\epsilon/8 + \left\| D\widehat{F}'_f - D\widehat{F}_f \right\|_\infty > \epsilon/2 \right] \tag{D.7}$$

$$\leq 2\mathcal{N}_{\epsilon/16} \sup_{f \in \mathcal{F}} \Pr_{\substack{e_1,...,e_N \\ e'_1,...,e'_N}} \left[\epsilon/4 + \left\| \widehat{F}'_f - \widehat{F}_f \right\|_\infty > \epsilon/2 \right] \tag{D.8}$$

$$= 2\mathcal{N}_{\epsilon/16} \sup_{f \in \mathcal{F}} \Pr_{\substack{e_1,...,e_N \\ e'_1,...,e'_N}} \left[\left\| \widehat{F}'_f - \widehat{F}_f \right\|_\infty > \epsilon/4 \right] \tag{D.9}$$

$$\leq 4\mathcal{N}_{\epsilon/16} \sup_{f \in \mathcal{F}} \Pr_{e_1,...,e_N} \left[\left\| \mathbb{E}_{e_1,...,e_N} \left[\widehat{F}_f \right] - \widehat{F}_f \right\|_\infty > \epsilon/8 \right] \tag{D.10}$$

$$\leq 4\mathcal{N}_{\epsilon/16} \sup_{f \in \mathcal{F}} \Pr_{e_1,...,e_N} \left[\sup_{t \in \mathbb{R}} \left| F_f(t) - \frac{1}{N} \sum_{i=1}^N 1\{\mathcal{R}^e(f) \leq t\} \right| > \frac{\epsilon}{8L} \right] \tag{D.11}$$

$$\leq 8\mathcal{N}_{\epsilon/16} \exp\left(-\frac{N\epsilon^2}{64L} \right). \tag{D.12}$$

Here, line (D.4) follows from the Symmetrization Lemma (Lemma D.2), lines (D.6) and (D.8) follow from the definition of D, line (D.7) is a union bound over $\widehat{\mathcal{P}}_{\epsilon/16}$, line (D.10) follows from the triangle inequality, line (D.11) follows from the Lipschitz assumption, and line (D.12) follows from the DKW Inequality (Lemma D.3).

Since $\sup_x f(x) - \sup_x g(x) \leq \sup_x f(x) - g(x)$,

$$\Pr_{e_1,...,e_N} \left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f(t) - \widehat{F}_f(t) > \epsilon + \text{Bias}(\mathcal{F}, \widehat{F}) \right]$$

$$= \Pr_{e_1,...,e_N} \left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f(t) - \widehat{F}_f(t) > \epsilon + \sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f(t) - \mathbb{E}_{e_1,...,e_N} \left[\widehat{F}_f(t) \right] \right]$$

$$\leq \Pr_{e_1,...,e_N} \left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} \mathbb{E}_{e_1,...,e_N} \left[\widehat{F}_f(t) \right] - \widehat{F}_f(t) > \epsilon \right]$$

$$\leq 8\mathcal{N}_{\epsilon/16} \exp\left(-\frac{N\epsilon^2}{64L} \right), \tag{D.13}$$

by (D.12). Meanwhile, applying the presumed uniform bound on within-environment generalization error together with a union bound over the N environments, gives us a high-probability bound on the maximum generalization error of f within any of the N environments:

$$\Pr_{\substack{\{e_i\}_{i=1}^N \sim \mathbb{P}(e) \\ \{(X_{i,j}, Y_{i,j})\}_{j=1}^n \sim \mathbb{P}(X^{e_i}, Y^{e_i})}} \left[\max_{i \in [N]} \sup_{f \in \mathcal{F}} \mathcal{R}^{e_i}(f) - \widehat{R}^{e_i}(f) \leq t_{n, \frac{\delta}{2N}, \mathcal{F}} \right] \leq \delta/2,$$

It follows that, with probability at least $1 - \delta/2$, for all $f \in \mathcal{F}$ and $t \in \mathbb{R}$,

$$\widehat{F}_f\left(t + t_{n,\frac{\delta}{2N},\mathcal{F}}\right) \leq \widehat{F}_{\widehat{\mathcal{R}}^{e_1}(f),\dots,\widehat{\mathcal{R}}^{e_1}(f)}(t),$$

where $\widehat{F}_{\widehat{\mathcal{R}}^{e_1}(f),\dots,\widehat{\mathcal{R}}^{e_1}(f)}(t)$ is the actually empirical estimate $\widehat{F}_f(t)$ of computed using the N empirical risks $\widehat{\mathcal{R}}^{e_1}(f),\dots,\widehat{\mathcal{R}}^{e_N}(f)$. Plugging this into the left-hand side of Inequality (D.13),

$$\Pr_{e_1,\dots,e_N}\left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f\left(t + t_{n,\frac{\delta}{2N},\mathcal{F}}\right) - \widehat{F}_{\widehat{\mathcal{R}}^{e_1}(f),\dots,\widehat{\mathcal{R}}^{e_1}(f)}(t) > \epsilon + \text{Bias}(\mathcal{F},\widehat{F})\right] \leq 8\mathcal{N}_{\epsilon/16}\exp\left(-\frac{N\epsilon}{64L}\right).$$

Setting $t = \widehat{F}_{\widehat{\mathcal{R}}^{e_1}(f),\dots,\widehat{\mathcal{R}}^{e_1}(f)}^{-1}(\alpha)$ and applying the non-decreasing function F_f^{-1} gives the desired result:

$$\Pr_{e_1,\dots,e_N}\left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f^{-1}\left(\alpha - \epsilon - \text{Bias}(\mathcal{F},\widehat{F})\right) - \widehat{F}_{\widehat{\mathcal{R}}^{e_1}(f),\dots,\widehat{\mathcal{R}}^{e_1}(f)}^{-1}(\alpha) + \geq t_{n,\frac{\delta}{2N},\mathcal{F}}\right] \leq 8\mathcal{N}_{\epsilon/16}\exp\left(-\frac{N\epsilon}{64L}\right).$$

$\square$

D.2 Kernel density estimator

In this section, we apply our generalization bound Theorem (D.1) to the kernel density estimator (KDE)

$$\widehat{F}_h(t) = \int_{-\infty}^{t} \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{\tau - X_i}{h}\right) d\tau$$

of the cumulative risk distribution under the assumptions that:

1. the loss ℓ takes values in a bounded interval $[a, b] \subseteq \mathbb{R}$, and
2. for all $f \in \mathcal{F}$, the true risk profile F_f is β-Hölder continuous with constant L, for any $\beta > 0$.

We also make standard integrability and symmetry assumptions on the kernel $K : \mathbb{R} \to \mathbb{R}$ (see Section 1.2.2 [132] for discussion of these assumptions):

$$\int_{\mathbb{R}} |K(u)|\, du < \infty, \quad \int_{\mathbb{R}} K(u)\, du = 1, \quad \int_{\mathbb{R}} |u|^\beta |K(u)\, du < \infty,$$

and, for each positive integer $j < \beta$,

$$\int_{\mathbb{R}} u^j K(u)\, du = 0. \tag{D.14}$$

We will use Theorem D.1 to show that, for an appropriately chosen bandwidth h,

$$\sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f(t) - \widehat{F}_f(t) \in O_P\left(\left(\frac{\log N}{N}\right)^{\frac{\beta}{2\beta+1}}\right).$$

We start by bounding the bias term $B(\mathcal{F},\widehat{F})$. Since

$$\mathbb{E}_{X_1,\dots,X_n}\left[\int_{-\infty}^{t}\left|\frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{\tau - X_i}{h}\right)\right|\right] d\tau \leq \frac{1}{h}\mathbb{E}_X\left[\int_{-\infty}^{\infty}\left|K\left(\frac{\tau - X_i}{h}\right)\right|\right] d\tau$$
$$\leq \|K\|_1 < \infty,$$

applying Fubini's theorem, linearity of expectation, the change of variables $x \mapsto \tau + xh$, Fubini's theorem again, and the fact that $\int_{\mathbb{R}} K(u)\,dx = 1$,

$$
\begin{aligned}
F_f(t) - \mathbb{E}_{X_1,\ldots,X_n}\left[\widehat{F}_h(t)\right] &= F_f(t) - \mathbb{E}_{e_1,\ldots,e_N}\left[\int_{-\infty}^{t} \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{\tau - X_i}{h}\right)\right] \\
&= F_f(t) - \int_{-\infty}^{t} \mathbb{E}_{X_1,\ldots,X_n}\left[\frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{\tau - X_i}{h}\right)\right] \\
&= F_f(t) - \int_{-\infty}^{t}\int_{\mathbb{R}} \frac{1}{h}K\left(\frac{\tau - x}{h}\right) p(x)\,dx\,d\tau \\
&= F_f(t) - \int_{-\infty}^{t}\int_{\mathbb{R}} K(x) p(\tau + xh)\,dx\,d\tau \\
&= F_f(t) - \int_{\mathbb{R}} K(x)\int_{-\infty}^{t} p(\tau + xh)\,d\tau\,dx \\
&= \int_{\mathbb{R}} K(x)\left(F_f(t) - F(t + xh)\right)\,dx.
\end{aligned}
$$

By Taylor's theorem for some $\pi \in [0,1]$,

$$
F(t + xh) = \sum_{j=0}^{\lfloor\beta\rfloor - 1} \frac{(xh)^j}{j!}\frac{d^j}{dt^j}F_f(t) + \frac{(xh)^{\lfloor\beta\rfloor}}{\lfloor\beta\rfloor!}\frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F(t + \pi xh).
$$

Hence, by the assumption (D.14),

$$
\begin{aligned}
F_f(t) - \mathbb{E}_{X_1,\ldots,X_n}\left[\widehat{F}_h(t)\right] &= \int_{\mathbb{R}} K(x)\left(F_f(t) - \sum_{j=0}^{\lfloor\beta\rfloor - 1} \frac{(xh)^j}{j!}\frac{d^j}{dt^j}F_f(t) + \frac{(xh)^{\lfloor\beta\rfloor}}{\lfloor\beta\rfloor!}\frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F(t + \pi xh)\right)\,dx \\
&= \int_{\mathbb{R}} K(x)\left(\frac{(xh)^{\lfloor\beta\rfloor}}{\lfloor\beta\rfloor!}\frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F(t + \pi xh)\right)\,dx \\
&= \int_{\mathbb{R}} K(x)\frac{(xh)^{\lfloor\beta\rfloor}}{\lfloor\beta\rfloor!}\left(\frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F(t + \pi xh) - \frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F_f(t)\right)\,dx.
\end{aligned}
$$

Thus, by the Hölder continuity assumption,

$$
\begin{aligned}
\left|F_f(t) - \mathbb{E}_{X_1,\ldots,X_n}\left[\widehat{F}_h(t)\right]\right| &\leq \int_{\mathbb{R}} K(x)\frac{(xh)^{\lfloor\beta\rfloor}}{\lfloor\beta\rfloor!}\left|\frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F(t + \pi xh) - \frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}F_f(t)\right|\,dx \\
&\leq \int_{\mathbb{R}} K(x)\frac{(xh)^{\lfloor\beta\rfloor}}{\lfloor\beta\rfloor!}L(\pi xh)^{\beta - \lfloor\beta\rfloor}\,dx \leq Ch^\beta, \qquad\text{(D.15)}
\end{aligned}
$$

where $C := \frac{L}{\lfloor\beta\rfloor!}\int_{\mathbb{R}}|x|^\beta|K(x)|\,dx$ is a constant.

Next, since, by the Fundamental Theorem of Calculus,

$$
\frac{d^{\lfloor\beta+1\rfloor}}{dt^{\lfloor\beta+1\rfloor}}\widehat{F}_f(t) = \frac{d^{\lfloor\beta+1\rfloor}}{dt^{\lfloor\beta+1\rfloor}}\int_{-\infty}^{t}\frac{1}{nh}\sum_{i=1}^{N} K\left(\frac{\tau - X_i}{h}\right)\,d\tau = \frac{1}{nh}\sum_{i=1}^{N}\frac{d^{\lfloor\beta\rfloor}}{dt^{\lfloor\beta\rfloor}}K\left(\frac{t - X_i}{h}\right),
$$

$\|F_f\|_{\mathcal{C}^{\beta+1}} \leq \|K_h\|_{\mathcal{C}^\beta} = h^{-(\beta+1)}\|K\|_{\mathcal{C}^\beta}$. Hence, by standard bounds on the covering number of Hölder continuous functions [133], there exists a constant $c > 0$ depending only on β such that

$$
\mathcal{N}_{\epsilon/16}(\mathcal{N}) \leq \exp\left(c(b-a)\left(\frac{\|K\|_{\mathcal{C}^\beta}}{h^{\beta+1}\epsilon}\right)^{\frac{1}{\beta+1}}\right) = \exp\left(c\frac{(b-a)}{h}\left(\frac{\|K\|_{\mathcal{C}^\beta}}{\epsilon}\right)^{\frac{1}{\beta+1}}\right). \qquad\text{(D.16)}
$$

Finally, since $\widehat{F}_h = \widehat{Q} * K_h$ (where $*$ denotes convolution), by linearity of the convolution and Young's convolution inequality [134, p.34],

$$
\left\|\widehat{F}_h - \widehat{F}_h'\right\|_\infty \leq \left\|\widehat{Q} - \widehat{Q}'\right\|_\infty \|K_h\|_1.
$$

Since, by a change of variables, $\|K_h\|_1 = \|K\|_1 = 1$, the KDE is a 1-Lipschitz function of the empirical CDF, under $\mathcal{L}_\infty(\mathbb{R})$.

Thus, plugging Inequality (D.15), Inequality (D.16), and $L = 1$ into Theorem D.1 and taking $n \to \infty$ gives, for any $\epsilon > 0$,

$$\Pr_{e_1,\dots,e_N} \left[\sup_{f \in \mathcal{F}} F_f^{-1} \left(\alpha - Ch^\beta - \epsilon \right) - \widehat{F}_f^{-1}(\alpha) > 0 \right] \le 8 \exp \left(c \frac{b-a}{h} \left(\frac{\|K\|_{\mathcal{C}^\beta}}{\epsilon} \right)^{\frac{1}{\beta+1}} \right) e^{-\frac{N\epsilon^2}{64}}.$$

Plugging in $\epsilon = \sqrt{\frac{\log \frac{1}{\delta} + c\frac{b-a}{h}}{N}}$ gives

$$\Pr_{e_1,\dots,e_N} \left[\sup_{f \in \mathcal{F}} F_f^{-1} \left(\alpha - Ch^\beta - \sqrt{\frac{\log \frac{1}{\delta} + c\frac{b-a}{h}}{N}} \right) - \widehat{F}_f^{-1}(\alpha) > 0 \right] \le \delta.$$

This bound is optimized by $h \asymp \left((b-a) \frac{\log N}{N} \right)^{\frac{1}{2\beta+1}}$, giving an overall bound of

$$\Pr_{e_1,\dots,e_N} \left[\sup_{f \in \mathcal{F}, t \in \mathbb{R}} F_f(t) - \widehat{F}_f(t) > ch^{\frac{\beta}{2\beta+1}} \right] \le \delta$$

$$\Pr_{e_1,\dots,e_N} \left[\sup_{f \in \mathcal{F}} F_f^{-1} \left(\alpha - ch^{\frac{\beta}{2\beta+1}} + \sqrt{\frac{\log \frac{1}{\delta}}{N}} \right) - \widehat{F}_f^{-1}(\alpha) > 0 \right] \le \delta.$$

for some $c > 0$. In particular, as $N, n \to \infty$, the EQRM estimate $\widehat{f}$ satisfies

$$F_{\widehat{f}}^{-1}(\alpha) \to \inf_{f \in \mathcal{F}} F_f^{-1}(\alpha).$$

E Further implementation details

E.1 Algorithm

Below we detail the EQRM algorithm. Note that: (i) any distribution estimator may be used in place of DIST so long as the functions DIST.ESTIMATE_PARAMS and DIST.ICDF are differentiable; (ii) other bandwidth-selection methods may be used on line 14, with the Gaussian-optimal rule serving as the default; and (iii) the bisection method BISECT on line 20 requires an additional parameter, the maximum number of steps, which we always set to 32.

Algorithm 1: Empirical Quantile Risk Minimization (EQRM).

Input: Predictor f_θ, loss function ℓ, desired probability of generalization α, learning rate η, distribution estimator DIST, M datasets with $D^m = \{(x_i^m, y_i^m)\}_{i=1}^{n_m}$.

1 Initialize f_θ;

2 **while** *not converged* **do**

    /* Get per-domain risks (i.e. average losses)   */

3    $L^m \leftarrow \frac{1}{n_m} \sum_{i=1}^{n_m} \ell(f_\theta(x_i^m), y_i^m)$, for $m = 1, \ldots, M$;

    /* Estimate the parameters of $\widehat{\mathbb{T}}_f$ */

4   DIST.ESTIMATE_PARAMS($\mathbf{L}$) ;

    /* Compute the α-quantile of $\widehat{\mathbb{T}}_f$ */

5    $q \leftarrow$ DIST.ICDF(α) ;

    /* Update f_θ */

6    $\theta \leftarrow \theta - \eta \cdot \nabla_\theta q$;

Output: f_θ

---

7 **Procedure** GAUSS.ESTIMATE_PARAMS*($\mathbf{L}$)*

    /* Compute the sample mean and variance   */

8    $\hat{\mu} \leftarrow \frac{1}{M} \sum_{m=1}^{M} L^m$;

9    $\hat{\sigma}^2 \leftarrow \frac{1}{M-1} \sum_{m=1}^{M} (L^m - \hat{\mu})^2$;

10 **Procedure** GAUSS.ICDF*(α)*

11   **return** $\hat{\mu} + \hat{\sigma} \cdot \Phi^{-1}(\alpha)$;

12 **Procedure** KDE.ESTIMATE_PARAMS*($\mathbf{L}$)*

    /* Set bandwidth h (Gaussian-optimal rule used as default)   */

13    $\hat{\sigma}^2 \leftarrow \frac{1}{M-1} \sum_{m=1}^{M} (L^m - \frac{1}{M} \sum_{j=1}^{M} L^j)^2$;

14    $h \leftarrow (\frac{4}{3M})^{0.2} \cdot \hat{\sigma}$

15 **Procedure** KDE.ICDF*(α)*

    /* Define the CDF when using M Gaussian kernels   */

16    $F_m(x') \leftarrow L^m + h \cdot \Phi(x')$;

17    $F(x') \leftarrow \frac{1}{M} \sum_{m=1}^{M} F_m(x')$;

    /* Invert the CDF via bisection   */

18   mn $\leftarrow \min_m F_m^{-1}(\alpha)$;

19   mx $\leftarrow \max_m F_m^{-1}(\alpha)$;

20   **return** BISECT(F, α, mn, mx) ;

E.2 ColoredMNIST

For the CMNIST results of § 6.1, we used full batches (size 25000), 400 steps for ERM pretraining, 600 total steps for IRM, VREx, EQRM, and 1000 total steps for GroupDRO, SD, and IGA. We used the original MNIST training set to create training and validation sets for each domain, and the original MNIST test set for the test sets of each domain. We also decayed the learning rate using cosine annealing/scheduling. We swept over penalty weights in $\{50, 100, 500, 1000, 5000\}$ for IRM,

VREx and IGA, penalty weights in $\{0.001, 0.01, 0.1, 1\}$ for SD, η's in $\{0.001, 0.01, 0.1, 0.5, 1.0\}$ for GroupDRO, and α's in $1 - \{e^{-100}, e^{-250}, e^{-500}, e^{-750}, e^{-1000}\}$ for EQRM. To allow these values of α, which are *very* close to 1, we used an asymptotic expression for the Normal inverse CDF, namely $\Phi^{-1}(\alpha) \approx \sqrt{-2\ln(1-\alpha)}$ as $\alpha \to 1$ [135]. This allowed us to parameterize $\alpha = 1 - e^{-1000}$ as $\ln(1-\alpha) = \ln(e^{-1000}) = -1000$, avoiding issues with floating-point precision. As is the standard for CMNIST, we used a test-domain validation set to select the best settings (after the total number of steps), then reported the mean and standard deviation over 10 random seeds on a test-domain test set. As in previous works, the hyperparameter ranges of all methods were selected by peeking at test-domain performance. While not ideal, this is quite difficult to avoid with CMNIST and highlights the problem of model selection more generally in DG—as discussed by many previous works [9, 38, 41, 114]. Finally, we note several observations from our CMNIST, WILDS and DomainBed experiments which, despite not being thoroughly investigated with their own set of experiments (yet), may prove useful for future work: (i) ERM pretraining seems an effective strategy for DG methods, and can likely replace the more delicate penalty-annealing strategies (as also observed in [114]); (ii) lowering the learning rate after ERM pretraining seems to stabilize DG methods; and (iii) EQRM often requires a lower learning rate than other DG methods after ERM pretraining, with its loss and gradients tending to be significantly larger.

E.3 DomainBed

For EQRM, we used the default algorithm setup: a kernel-density estimator of the risk distribution with the "Gaussian-optimal" rule [64] for bandwidth selection. We used the standard hyperparameter-sampling procedure of Domainbed, running over 3 trials for 20 randomly-sampled hyperparameters per trial. For EQRM, this involved:

Hparam	Default	Sampling
α	0.75	$U(0.5, 0.99)$
Burn-in/anneal iters	2500	10^k, with $k \sim U(2.5, 3.5)$
EQRM learning rate (post burn-in)	10^{-6}	10^k, with $k \sim U(-7, -5)$

All other all hyperparameters remained as their DomainBed-defaults, while the baseline results were taken directly from the most up-to-date DomainBed tables[8]. See our code for further details.

E.4 WILDS

We considered two WILDS datasets: iWildCam and OGB-MolPCBA (henceforth OGB). For both of these datasets, we used the architectures use in the original WILDS paper [12]; that is, for iWildCam we used a ResNet-50 architecture [136] pretrained on ImageNet [137], and for OGB, we used a Graph Isomorphism Network [138] combined with virtual nodes [139]. To perform model-selection, we followed the guidelines provided in the original WILDS paper [12]. In particular, for each of the baselines we consider, we performed grid searches over the hyperparameter ranges listed in [12] with respect to the given validation sets; see [12, Appendices E.1.2 and E.4.2] for a full list of these hyperparameter ranges.

EQRM. For both datasets, we ran EQRM with KDE using the Gaussian-optimal bandwidth-selection method. All EQRM models were initialized with the same ERM checkpoint, which is obtained by training ERM using the code provided by [12]. Following [12], for iWildCam, we trained ERM for 12 epochs, and for OGB, we trained ERM for 100 epochs. We again followed [12] by using a batch size of 32 for iWildCam and 8 groups per batch. For OGB, we performed grid searches over the batch size in the range $B \in \{32, 64, 128, 256, 512, 1024, 2048\}$, and we used $0.25B$ groups per batch. We selected the learning rate for EQRM from $\eta \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}\}$.

Computational resources. All experiments on the WILDS datasets were run across two four-GPU workstations, comprising a total of eight Quadro RTX 5000 GPUs.

---

[8] https://github.com/facebookresearch/DomainBed/tree/main/domainbed/results/2020_
10_06_7df6f06

F Connections between QRM and DRO

In this appendix we draw connections between quantile risk minimization (QRM) and distributionally robust optimization (DRO) by considering an alternative optimization problem which we call *superquantile risk minimization* [9]:

$$\min_{f \in \mathcal{F}} \mathrm{SQ}_\alpha(R; \mathbb{T}_f) \qquad \text{where} \qquad \mathrm{SQ}_\alpha(R; \mathbb{T}_f) := \mathbb{E}_{R \sim \mathbb{T}_f} \left[R \mid R \geq F_{\mathbb{T}_f}^{-1}(\alpha) \right]. \qquad \text{(F.1)}$$

Here, SQ_α represents the *superquantile*—also known as the *conditional value-at-risk* (CVaR) or *expected tail loss*—at level α, which can be seen as the conditional expectation of a random variable R subject to R being larger than the α-quantile $F^{-1}(\alpha)$. In our case, where R represents the statistical risk on a randomly-sampled environment, SQ_α can be seen as the expected risk in the worst $100 \cdot (1 - \alpha)\%$ of cases/domains. Below, we exploit the well-known duality properties of CVaR to formally connect (QRM) and GroupDRO [45]; see Prop. F.1 for details.

F.1 Notation for this appendix

Throughout this appendix, for each $f \in \mathcal{F}$, we will let the risk random variable R be a defined on the probability space $(\mathbb{R}_+, \mathcal{B}, \mathbb{T}_f)$, where $\mathbb{R}_+$ denotes the nonnegative real numbers and $\mathcal{B}$ denotes the Borel σ-algebra on $\mathbb{R}_+$. We will also consider the Lebesgue spaces $L^p := L^p(\mathbb{R}_+, \mathcal{B}, \mathbb{T}_f)$ of functions h for which $\mathbb{E}_{r \sim \mathbb{T}_f}[|h(r)|^p]$ is finite. For conciseness, we will use the notation

$$\langle g(r), h(r) \rangle := \int_{r \geq 0} g(r) h(r) \mathrm{d}r \qquad \text{(F.2)}$$

to denote the standard inner product on $\mathbb{R}_+$. Furthermore, we will use the notation $\mathbb{U} \ll \mathbb{V}$ to signify that $\mathbb{U}$ is *absolutely continuous* with respect to $\mathbb{V}$, meaning that if $\mathbb{U}(A) = 0$ for every set A for which $\mathbb{V}(A) = 0$. We also use the abbreviation "a.e." to mean "almost everywhere." Finally, the notation $\Pi_{[a,b]}(c)$ denotes the projection of a number c into the real interval $[a, b]$.

F.2 (Strong) Duality of the superquantile

We begin by proving that strong duality holds for the superquantile function SQ_α. We note that this duality result is well-known in the literature (see, e.g., [89]), and has been exploited in the context of adaptive sampling [93] and offline reinforcement learning [140]. We state this result and proof for the sake of exposition.

Proposition F.1 (Dual representation of SQ_α). *If $R \in L^P$ for some $p \in (1, \infty)$, then*

$$\mathrm{SQ}_\alpha(R; \mathbb{T}_f) = \max_{\mathbb{U} \in \mathcal{U}_f(\alpha)} \mathbb{E}_{\mathbb{U}}[R] \qquad \text{(F.3)}$$

where the uncertainty set $\mathcal{U}_f(\alpha)$ is defined as

$$\mathcal{U}_f(\alpha) := \left\{ \mathbb{U} \in L^q : \mathbb{U} \ll \mathbb{T}_f, \mathbb{U} \in [0, 1/1-\alpha] \text{ a.e. }, ||U||_{L^1} = 1 \right\}. \qquad \text{(F.4)}$$

Proof. Note that the primal objective can be equivalently written as

$$\mathrm{SQ}_\alpha(R; \mathbb{T}_f) = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1 - \alpha} \langle (R - t)_+, \mathbb{T}_f \rangle \right\} \qquad \text{(F.5)}$$

where $(z)_+ = \max\{0, z\}$ [96], which in turn has the following epigraph form:

$$\min_{t \in \mathbb{R}, \, s \in L_+^p} \quad t + \frac{1}{1 - \alpha} \langle s, \mathbb{T}_f \rangle \qquad \text{(F.6)}$$

$$\text{subject to} \quad R(r) - t \leq s(r) \text{ a.e. } r \in \mathbb{R}_+. \qquad \text{(F.7)}$$

---

[9]This definition assumes that $\mathbb{T}_f$ is continuous; for a more general treatment, see [96].

When written in Lagrangian form, we can express this problem as

$$\min_{t\in\mathbb{R},\, s\in L^p_+}\max_{\lambda\in L^q_+}\left\{t(1-\langle 1,\lambda\rangle)+\left\langle s,\frac{1}{1-\alpha}\mathbb{T}_f-\lambda\right\rangle+\langle R,\lambda\rangle\right\}. \tag{F.8}$$

Note that this objective is *linear* in t, s, and λ, and therefore due to the strong duality of linear programs, we can optimize over s, t, and λ in any order [141]. Minimizing over t reveals that the problem is unbounded unless $\int_{r\geq 0}\lambda(r)\mathrm{d}r=1$, meaning that λ is a probability distribution since $\lambda(r)\geq 0$ almost everywhere. Thus, the problem can be written as

$$\min_{s\in L^p_+}\max_{\lambda\in\mathcal{P}(\mathbb{R}_+)}\left\{\left\langle s,\frac{1}{1-\alpha}\mathbb{T}_f-\lambda\right\rangle+\langle R,\lambda\rangle\right\} \tag{F.9}$$

where $\mathcal{P}^q(\mathbb{R}_+)$ denotes the subspace of L^q of probability distributions on $\mathbb{R}_+$.

Now consider the maximization over s. Note that if there is a set $A\subset\mathcal{E}_{\text{all}}$ of nonzero Lebesgue measure on which $\lambda(A)\geq(1/1-\alpha)\mathbb{T}_f(A)$, then the problem is unbounded below because $s(A)$ can be made arbitrarily large. Therefore, it must be the case that $\lambda\leq(1/1-\alpha)\mathbb{T}_f$ almost everywhere. On the other hand, if $\lambda(A)\leq(1/1-\alpha)\mathbb{T}_f(A)$, then $s(A)=0$ minimizes the first term in the objective. Therefore, s can be eliminated provided that $\lambda\leq(1/1-\alpha)\mathbb{T}_f$ almost everywhere. Thus, we can write the problem as

$$\max_{\lambda\in\mathcal{P}^q(\mathbb{R}_+)}\quad\langle R,\lambda\rangle=\mathbb{E}_\lambda[R] \tag{F.10}$$

$$\text{subject to}\quad\lambda(r)\leq\frac{1}{1-\alpha}\mathbb{T}_f(r)\text{ a.e. }r\geq 0. \tag{F.11}$$

Now observe that the constraint in the above problem is equivalent to $\lambda\ll\mathbb{Q}$. Thus, by defining $\mathbb{U}=\mathrm{d}\lambda/\mathrm{d}\mathbb{T}_f$ to be the Radon-Nikodym derivative of λ with respect to $\mathbb{Q}$, we can write the problem in the form of (F.3), completing the proof. $\quad\square$

Succinctly, this proposition shows that provided that R is sufficiently smooth (i.e., an element of L^p), it holds that minimizing the superquantile function is equivalent to solving

$$\min_{f\in\mathcal{F}}\max_{\mathbb{U}\in\mathcal{U}_f(\alpha)}\mathbb{E}_\mathbb{U}[R] \tag{F.12}$$

which is a distributionally robust optimization (DRO) problem with uncertainty set $\mathcal{U}_f(\alpha)$ as defined in (F.4). In plain terms, for any $\alpha\in(0,1)$, this uncertainty set contains probability distributions on $\mathbb{R}_+$ which can place no larger than $1/1-\alpha$ on any risk value.

At an intuitive level, this shows that by varying α in Eq. (F.1), one can interpolate between a range DRO problems. In particular, at level $\alpha=1$, we recover the problem in (3.1), which can be viewed as a DRO problem which selects a Dirac distribution which places solely on the essential supremum of $R\sim\mathbb{T}_f$. On the other hand, at level $\alpha=0$, we recover a problem which selects a distribution that equally weights each of the risks in different domains equally. A special case of this is the GroupDRO formulation in [45], wherein under the assumption that the data is partitioned into m groups, the inner maximum in (F.12) is taken over the $(m-1)$-dimensional simplex Δ_m (see, e.g., equation (7) in [45]).

G   Additional analyses and experiments

G.1   Linear regression

In this section we extend § 6.1 to provide further analyses and discussion of EQRM using linear regression datasets based on Ex. A.3. In particular, we: (i) extend Fig. 3 to include plots of the predictors' risk CDFs (G.1.1); and (ii) discuss the ability of EQRM to recover the causal predictor when σ_1^2, σ_2^2 and/or σ_Y^2 change over environments, compared to IRM [9] and VREx [41] (G.1.2).

Table 5: Recovering the causal predictor for linear regression tasks based on Ex. A.3. A tick means that it is *possible* to recover the causal predictor, under further assumptions.

Changing	Domain Scedasticity	Invariant Risk	Invariant Function (β_{cause})	IRM	VREx	EQRM
σ_1	*Homo*scedastic	✓	✓	✓	✓	✓
σ_2	*Homo*scedastic	✓	✓	✓	✓	✓
σ_Y	*Hetero*scedastic	✗	✓	✓	✗	✗

G.1.1 Risk CDFs as risk-robustness curves

As an extension of Fig. 3, in particular the PDFs in Fig. 3 **B**, Fig. 6 depicts the risk CDFs for different predictors. Here we see that a predictor's risk CDF depicts its risk-robustness curve, and also that each α results in a predictor f_α with minimal α-quantile risk. That is, for each desired level of robustness (i.e. probability of the upper-bound on risk holding, y-axis), the corresponding α has minimal risk (x-axis).

Figure 6: **Extension of Fig. 3 showing the risk CDFs (i.e. risk-robustness curves) for different predictors.** For each risk upper-bound (x), we see the corresponding probability of it holding under the training domains (y). Note that, for each level of robustness (y, i.e. probability that the risk upper-bound holds), the corresponding α has the lowest upper-bound on risk (x). Also note that these CDFs correspond to the PDFs of Fig. 3 (**B**).

G.1.2 Invariant risks vs. invariant functions

We now compare seeking invariant *risks* to seeking invariant *functions* by analyzing linear regression datasets, based on Ex. A.3, in which σ_1^2, σ_2^2 and/or σ_Y^2 change over domains. This is turn allows us to compare EQRM (invariant risks), VREx [41] (invariant risks), and IRM [9] (invariant functions).

Domain-skedasticity. For recovering the causal predictor, the key difference between using invariant *risks* and invariant *functions* lies in the assumption about *domain-skedasticity*, i.e. the "predicatability" of Y across domains. In particular, the causal predictor only has invariant risks in *domain-homoskedastic* cases and not in *domain-heteroskedastic* cases, the latter describing scenarios in which the predictability of Y (i.e. the amount of irreducible error or intrinsic noise) varies across domains, meaning that the risk of the causal predictor will be smaller on some domains than others. Thus, methods seeking the causal predictor through invariant risks must assume domain homoskedasticity [41, 54]. In contrast, methods seeking the causal predictor through invariant *functions* need not make such a domain-homoskedasticity assumption, but instead the slightly weaker assumption of the conditional mean $\mathbb{E}[Y|\text{Pa}(Y)]$ being invariant across domains. As explained in the next paragraph and summarized in Table 5, this translates into the coefficient β_{cause} being invariant across domains for the linear SEM of Ex. A.3.

Mathematical analysis. We now analyze the risk-invariant solutions of Ex. A.3. We start by expanding the structural equations of Ex. A.3 as:

$$X_1 = N_1,$$
$$Y = N_1 + N_Y,$$
$$X_2 = N_1 + N_Y + N_2.$$

We then note that the goal is to learn a model $\widehat{Y} = \beta_1 \cdot X_1 + \beta_2 \cdot X_2$, which has residual error

$$\widehat{Y} - Y = \beta_1 \cdot N_1 + \beta_2 \cdot (N_1 + N_Y + N_2) - N_1 - N_Y$$
$$= (\beta_1 + \beta_2 - 1) \cdot N_1 + (\beta_2 - 1) \cdot N_Y + \beta_2 \cdot N_2.$$

Then, since all variables have zero mean and the noise terms are independent, the risk (i.e. the MSE loss) is simply the variance of the residuals, which can be written as

$$\mathbb{E}[(\widehat{Y} - Y)^2] = (\beta_1 + \beta_2 - 1)^2 \cdot \sigma_1^2 + (\beta_2 - 1)^2 \cdot \sigma_Y^2 + \beta_2^2 \cdot \sigma_2^2.$$

Here, we have that, when:

- **Only σ_1 changes:** the only way to keep the risk invariant across domains is to set $\beta_1 + \beta_2 = 1$. The minimal invariant-risk solution then depends on σ_y and σ_2:
  - if $\sigma_y < \sigma_2$, the minimal invariant-risk solution sets $\beta_1 = 1$ and $\beta_2 = 0$ (causal predictor);
  - if $\sigma_y > \sigma_2$, the minimal invariant-risk solution sets $\beta_1 = 0$ and $\beta_2 = 1$ (anti-causal predictor);
  - if $\sigma_y = \sigma_2$, then any solution $(\beta_1, \beta_2) = (c, 1-c)$ with $c \in [0,1]$ is a minimal invariant-risk solution, including the causal predictor $c = 1$, anti-causal predictor $c = 0$, and everything in-between.
- **Only σ_2 changes:** the invariant-risk solutions set $\beta_2 = 0$, with the minimal invariant-risk solution also setting $\beta_1 = 1$ (causal predictor).
- **σ_1 and σ_2 change:** *the* invariant-risk solution sets $\beta_1 = 1, \beta_2 = 0$ (causal predictor).
- **Only σ_Y changes:** the invariant-risk solutions set $\beta_2 = 1$, with the minimal invariant-risk solution also setting $\beta_1 = 0$ (anti-causal predictor).
- **σ_1 and σ_Y change:** *the* invariant-risk solution sets $\beta_1 = 0, \beta_2 = 1$ (anti-causal predictor).
- **σ_2 and σ_Y change:** there is no invariant-risk solution.
- **σ_1, σ_2 and σ_Y change:** there is no invariant-risk solution.

Empirical analysis. To see this empirically, we refer the reader to Table 5 of Krueger et al. [41, App. G.2], which compares the invariant-risk solution of VREx to the invariant-function solution of IRM on the synthetic linear-SEM tasks of Arjovsky et al. [9, Sec. 5.1], which calculate the MSE between the estimated coefficients $(\hat{\beta}_1, \hat{\beta}_2)$ and those of the causal predictor $(1, 0)$.

Different goals, solutions, and advantages. We end by emphasizing the fact that the invariant-risk and invariant-function solutions have different pros and cons depending both on the goal and the assumptions made. If the goal is the recover the causal predictor or causes of Y, then the invariant-function solution has the advantage due to weaker assumptions on domain skedasticity. However, if the goal is learn predictors with stable or invariant performance, such that they perform well on new domains with high probability, then the invariant-risk solution has the advantage. For example, in the domain-heteroskedastic cases above where σ_Y changes or σ_Y and σ_1 change, the invariant-function solution recovers the causal predictor $\beta_1 = 1, \beta_2 = 0$ and thus has arbitrarily-large risk as $\sigma_Y \to \infty$ (i.e. in the worst-case). In contrast, the invariant-risk solution recovers the anti-causal predictor $\beta_1 = 0, \beta_2 = 1$ and thus has fixed risk σ_2^2 in all domains.

G.2 DomainBed

In this section, we include the full per-dataset DomainBed results. We consider the two most common model-selection methods of the DomainBed package—training-domain validation set and test-domain validation set (oracle)—and compare EQRM to a range of baselines. Implementation details for these experiments are provided in § E.3 and our open-source code.

G.2.1 Model selection: training-domain validation set

VLCS

Algorithm	C	L	S	V	Avg
ERM	97.7 ± 0.4	64.3 ± 0.9	73.4 ± 0.5	74.6 ± 1.3	77.5
IRM	98.6 ± 0.1	64.9 ± 0.9	73.4 ± 0.6	77.3 ± 0.9	78.5
GroupDRO	97.3 ± 0.3	63.4 ± 0.9	69.5 ± 0.8	76.7 ± 0.7	76.7
Mixup	98.3 ± 0.6	64.8 ± 1.0	72.1 ± 0.5	74.3 ± 0.8	77.4
MLDG	97.4 ± 0.2	65.2 ± 0.7	71.0 ± 1.4	75.3 ± 1.0	77.2
CORAL	98.3 ± 0.1	66.1 ± 1.2	73.4 ± 0.3	77.5 ± 1.2	78.8
MMD	97.7 ± 0.1	64.0 ± 1.1	72.8 ± 0.2	75.3 ± 3.3	77.5
DANN	99.0 ± 0.3	65.1 ± 1.4	73.1 ± 0.3	77.2 ± 0.6	78.6
CDANN	97.1 ± 0.3	65.1 ± 1.2	70.7 ± 0.8	77.1 ± 1.5	77.5
MTL	97.8 ± 0.4	64.3 ± 0.3	71.5 ± 0.7	75.3 ± 1.7	77.2
SagNet	97.9 ± 0.4	64.5 ± 0.5	71.4 ± 1.3	77.5 ± 0.5	77.8
ARM	98.7 ± 0.2	63.6 ± 0.7	71.3 ± 1.2	76.7 ± 0.6	77.6
VREx	98.4 ± 0.3	64.4 ± 1.4	74.1 ± 0.4	76.2 ± 1.3	78.3
RSC	97.9 ± 0.1	62.5 ± 0.7	72.3 ± 1.2	75.6 ± 0.8	77.1
EQRM	98.3 ± 0.0	63.7 ± 0.8	72.6 ± 1.0	76.7 ± 1.1	77.8

PACS

Algorithm	A	C	P	S	Avg
ERM	84.7 ± 0.4	80.8 ± 0.6	97.2 ± 0.3	79.3 ± 1.0	85.5
IRM	84.8 ± 1.3	76.4 ± 1.1	96.7 ± 0.6	76.1 ± 1.0	83.5
GroupDRO	83.5 ± 0.9	79.1 ± 0.6	96.7 ± 0.3	78.3 ± 2.0	84.4
Mixup	86.1 ± 0.5	78.9 ± 0.8	97.6 ± 0.1	75.8 ± 1.8	84.6
MLDG	85.5 ± 1.4	80.1 ± 1.7	97.4 ± 0.3	76.6 ± 1.1	84.9
CORAL	88.3 ± 0.2	80.0 ± 0.5	97.5 ± 0.3	78.8 ± 1.3	86.2
MMD	86.1 ± 1.4	79.4 ± 0.9	96.6 ± 0.2	76.5 ± 0.5	84.6
DANN	86.4 ± 0.8	77.4 ± 0.8	97.3 ± 0.4	73.5 ± 2.3	83.6
CDANN	84.6 ± 1.8	75.5 ± 0.9	96.8 ± 0.3	73.5 ± 0.6	82.6
MTL	87.5 ± 0.8	77.1 ± 0.5	96.4 ± 0.8	77.3 ± 1.8	84.6
SagNet	87.4 ± 1.0	80.7 ± 0.6	97.1 ± 0.1	80.0 ± 0.4	86.3
ARM	86.8 ± 0.6	76.8 ± 0.5	97.4 ± 0.3	79.3 ± 1.2	85.1
VREx	86.0 ± 1.6	79.1 ± 0.6	96.9 ± 0.5	77.7 ± 1.7	84.9
RSC	85.4 ± 0.8	79.7 ± 1.8	97.6 ± 0.3	78.2 ± 1.2	85.2
EQRM	86.5 ± 0.4	82.1 ± 0.7	96.6 ± 0.2	80.8 ± 0.2	86.5

OfficeHome

Algorithm	A	C	P	R	Avg
ERM	61.3 ± 0.7	52.4 ± 0.3	75.8 ± 0.1	76.6 ± 0.3	66.5
IRM	58.9 ± 2.3	52.2 ± 1.6	72.1 ± 2.9	74.0 ± 2.5	64.3
GroupDRO	60.4 ± 0.7	52.7 ± 1.0	75.0 ± 0.7	76.0 ± 0.7	66.0
Mixup	62.4 ± 0.8	54.8 ± 0.6	76.9 ± 0.3	78.3 ± 0.2	68.1
MLDG	61.5 ± 0.9	53.2 ± 0.6	75.0 ± 1.2	77.5 ± 0.4	66.8
CORAL	65.3 ± 0.4	54.4 ± 0.5	76.5 ± 0.1	78.4 ± 0.5	68.7
MMD	60.4 ± 0.2	53.3 ± 0.3	74.3 ± 0.1	77.4 ± 0.6	66.3
DANN	59.9 ± 1.3	53.0 ± 0.3	73.6 ± 0.7	76.9 ± 0.5	65.9
CDANN	61.5 ± 1.4	50.4 ± 2.4	74.4 ± 0.9	76.6 ± 0.8	65.8
MTL	61.5 ± 0.7	52.4 ± 0.6	74.9 ± 0.4	76.8 ± 0.4	66.4
SagNet	63.4 ± 0.2	54.8 ± 0.4	75.8 ± 0.4	78.3 ± 0.3	68.1
ARM	58.9 ± 0.8	51.0 ± 0.5	74.1 ± 0.1	75.2 ± 0.3	64.8
VREx	60.7 ± 0.9	53.0 ± 0.9	75.3 ± 0.1	76.6 ± 0.5	66.4
RSC	60.7 ± 1.4	51.4 ± 0.3	74.8 ± 1.1	75.1 ± 1.3	65.5
EQRM	60.5 ± 0.1	56.0 ± 0.2	76.1 ± 0.4	77.4 ± 0.3	67.5

TerraIncognita

Algorithm	L100	L38	L43	L46	Avg
ERM	49.8 ± 4.4	42.1 ± 1.4	56.9 ± 1.8	35.7 ± 3.9	46.1
IRM	54.6 ± 1.3	39.8 ± 1.9	56.2 ± 1.8	39.6 ± 0.8	47.6
GroupDRO	41.2 ± 0.7	38.6 ± 2.1	56.7 ± 0.9	36.4 ± 2.1	43.2
Mixup	59.6 ± 2.0	42.2 ± 1.4	55.9 ± 0.8	33.9 ± 1.4	47.9
MLDG	54.2 ± 3.0	44.3 ± 1.1	55.6 ± 0.3	36.9 ± 2.2	47.7
CORAL	51.6 ± 2.4	42.2 ± 1.0	57.0 ± 1.0	39.8 ± 2.9	47.6
MMD	41.9 ± 3.0	34.8 ± 1.0	57.0 ± 1.9	35.2 ± 1.8	42.2
DANN	51.1 ± 3.5	40.6 ± 0.6	57.4 ± 0.5	37.7 ± 1.8	46.7
CDANN	47.0 ± 1.9	41.3 ± 4.8	54.9 ± 1.7	39.8 ± 2.3	45.8
MTL	49.3 ± 1.2	39.6 ± 6.3	55.6 ± 1.1	37.8 ± 0.8	45.6
SagNet	53.0 ± 2.9	43.0 ± 2.5	57.9 ± 0.6	40.4 ± 1.3	48.6
ARM	49.3 ± 0.7	38.3 ± 2.4	55.8 ± 0.8	38.7 ± 1.3	45.5
VREx	48.2 ± 4.3	41.7 ± 1.3	56.8 ± 0.8	38.7 ± 3.1	46.4
RSC	50.2 ± 2.2	39.2 ± 1.4	56.3 ± 1.4	40.8 ± 0.6	46.6
EQRM	47.9 ± 1.9	45.2 ± 0.3	59.1 ± 0.3	38.8 ± 0.6	47.8

DomainNet

Algorithm	clip	info	paint	quick	real	sketch	Avg
ERM	58.1 ± 0.3	18.8 ± 0.3	46.7 ± 0.3	12.2 ± 0.4	59.6 ± 0.1	49.8 ± 0.4	40.9
IRM	48.5 ± 2.8	15.0 ± 1.5	38.3 ± 4.3	10.9 ± 0.5	48.2 ± 5.2	42.3 ± 3.1	33.9
GroupDRO	47.2 ± 0.5	17.5 ± 0.4	33.8 ± 0.5	9.3 ± 0.3	51.6 ± 0.4	40.1 ± 0.6	33.3
Mixup	55.7 ± 0.3	18.5 ± 0.5	44.3 ± 0.5	12.5 ± 0.4	55.8 ± 0.3	48.2 ± 0.5	39.2
MLDG	59.1 ± 0.2	19.1 ± 0.3	45.8 ± 0.7	13.4 ± 0.3	59.6 ± 0.2	50.2 ± 0.4	41.2
CORAL	59.2 ± 0.1	19.7 ± 0.2	46.6 ± 0.3	13.4 ± 0.4	59.8 ± 0.2	50.1 ± 0.6	41.5
MMD	32.1 ± 13.3	11.0 ± 4.6	26.8 ± 11.3	8.7 ± 2.1	32.7 ± 13.8	28.9 ± 11.9	23.4
DANN	53.1 ± 0.2	18.3 ± 0.1	44.2 ± 0.7	11.8 ± 0.1	55.5 ± 0.4	46.8 ± 0.6	38.3
CDANN	54.6 ± 0.4	17.3 ± 0.1	43.7 ± 0.9	12.1 ± 0.7	56.2 ± 0.4	45.9 ± 0.5	38.3
MTL	57.9 ± 0.5	18.5 ± 0.4	46.0 ± 0.1	12.5 ± 0.1	59.5 ± 0.3	49.2 ± 0.1	40.6
SagNet	57.7 ± 0.3	19.0 ± 0.2	45.3 ± 0.3	12.7 ± 0.5	58.1 ± 0.5	48.8 ± 0.2	40.3
ARM	49.7 ± 0.3	16.3 ± 0.5	40.9 ± 1.1	9.4 ± 0.1	53.4 ± 0.4	43.5 ± 0.4	35.5
VREx	47.3 ± 3.5	16.0 ± 1.5	35.8 ± 4.6	10.9 ± 0.3	49.6 ± 4.9	42.0 ± 3.0	33.6
RSC	55.0 ± 1.2	18.3 ± 0.5	44.4 ± 0.6	12.2 ± 0.2	55.7 ± 0.7	47.8 ± 0.9	38.9
EQRM	56.1 ± 1.3	19.6 ± 0.1	46.3 ± 1.5	12.9 ± 0.3	61.1 ± 0.0	50.3 ± 0.1	41.0

Averages

Algorithm	VLCS	PACS	OfficeHome	TerraIncognita	DomainNet	Avg
ERM	77.5 ± 0.4	85.5 ± 0.2	66.5 ± 0.3	46.1 ± 1.8	40.9 ± 0.1	63.3
IRM	78.5 ± 0.5	83.5 ± 0.8	64.3 ± 2.2	47.6 ± 0.8	33.9 ± 2.8	61.6
GroupDRO	76.7 ± 0.6	84.4 ± 0.8	66.0 ± 0.7	43.2 ± 1.1	33.3 ± 0.2	60.9
Mixup	77.4 ± 0.6	84.6 ± 0.6	68.1 ± 0.3	47.9 ± 0.8	39.2 ± 0.1	63.4
MLDG	77.2 ± 0.4	84.9 ± 1.0	66.8 ± 0.6	47.7 ± 0.9	41.2 ± 0.1	63.6
CORAL	78.8 ± 0.6	86.2 ± 0.3	68.7 ± 0.3	47.6 ± 1.0	41.5 ± 0.1	64.6
MMD	77.5 ± 0.9	84.6 ± 0.5	66.3 ± 0.1	42.2 ± 1.6	23.4 ± 9.5	63.3
DANN	78.6 ± 0.4	83.6 ± 0.4	65.9 ± 0.6	46.7 ± 0.5	38.3 ± 0.1	62.6
CDANN	77.5 ± 0.1	82.6 ± 0.9	65.8 ± 1.3	45.8 ± 1.6	38.3 ± 0.3	62.0
MTL	77.2 ± 0.4	84.6 ± 0.5	66.4 ± 0.5	45.6 ± 1.2	40.6 ± 0.1	62.9
SagNet	77.8 ± 0.5	86.3 ± 0.2	68.1 ± 0.1	48.6 ± 1.0	40.3 ± 0.1	64.2
ARM	77.6 ± 0.3	85.1 ± 0.4	64.8 ± 0.3	45.5 ± 0.3	35.5 ± 0.2	61.7
VREx	78.3 ± 0.2	84.9 ± 0.6	66.4 ± 0.6	46.4 ± 0.6	33.6 ± 2.9	61.9
EQRM	77.8 ± 0.6	86.5 ± 0.2	67.5 ± 0.1	47.8 ± 0.6	41.0 ± 0.3	64.1

G.2.2   Model selection: test-domain validation set (oracle)

VLCS

Algorithm	C	L	S	V	Avg
ERM	97.6 ± 0.3	67.9 ± 0.7	70.9 ± 0.2	74.0 ± 0.6	77.6
IRM	97.3 ± 0.2	66.7 ± 0.1	71.0 ± 2.3	72.8 ± 0.4	76.9
GroupDRO	97.7 ± 0.2	65.9 ± 0.2	72.8 ± 0.8	73.4 ± 1.3	77.4
Mixup	97.8 ± 0.4	67.2 ± 0.4	71.5 ± 0.2	75.7 ± 0.6	78.1
MLDG	97.1 ± 0.5	66.6 ± 0.5	71.5 ± 0.1	75.0 ± 0.9	77.5
CORAL	97.3 ± 0.2	67.5 ± 0.6	71.6 ± 0.6	74.5 ± 0.0	77.7
MMD	98.8 ± 0.0	66.4 ± 0.4	70.8 ± 0.5	75.6 ± 0.4	77.9
DANN	99.0 ± 0.2	66.3 ± 1.2	73.4 ± 1.4	80.1 ± 0.5	79.7
CDANN	98.2 ± 0.1	68.8 ± 0.5	74.3 ± 0.6	78.1 ± 0.5	79.9
MTL	97.9 ± 0.7	66.1 ± 0.7	72.0 ± 0.4	74.9 ± 1.1	77.7
SagNet	97.4 ± 0.3	66.4 ± 0.4	71.6 ± 0.1	75.0 ± 0.8	77.6
ARM	97.6 ± 0.6	66.5 ± 0.3	72.7 ± 0.6	74.4 ± 0.7	77.8
VREx	98.4 ± 0.2	66.4 ± 0.7	72.8 ± 0.1	75.0 ± 1.4	78.1
RSC	98.0 ± 0.4	67.2 ± 0.3	70.3 ± 1.3	75.6 ± 0.4	77.8
EQRM	98.2 ± 0.2	66.8 ± 0.8	71.7 ± 1.0	74.6 ± 0.3	77.8

PACS

Algorithm	A	C	P	S	Avg
ERM	86.5 ± 1.0	81.3 ± 0.6	96.2 ± 0.3	82.7 ± 1.1	86.7
IRM	84.2 ± 0.9	79.7 ± 1.5	95.9 ± 0.4	78.3 ± 2.1	84.5
GroupDRO	87.5 ± 0.5	82.9 ± 0.6	97.1 ± 0.3	81.1 ± 1.2	87.1
Mixup	87.5 ± 0.4	81.6 ± 0.7	97.4 ± 0.2	80.8 ± 0.9	86.8
MLDG	87.0 ± 1.2	82.5 ± 0.9	96.7 ± 0.3	81.2 ± 0.6	86.8
CORAL	86.6 ± 0.8	81.8 ± 0.9	97.1 ± 0.5	82.7 ± 0.6	87.1
MMD	88.1 ± 0.8	82.6 ± 0.7	97.1 ± 0.5	81.2 ± 1.2	87.2
DANN	87.0 ± 0.4	80.3 ± 0.6	96.8 ± 0.3	76.9 ± 1.1	85.2
CDANN	87.7 ± 0.6	80.7 ± 1.2	97.3 ± 0.4	77.6 ± 1.5	85.8
MTL	87.0 ± 0.2	82.7 ± 0.8	96.5 ± 0.7	80.5 ± 0.8	86.7
SagNet	87.4 ± 0.5	81.2 ± 1.2	96.3 ± 0.8	80.7 ± 1.1	86.4
ARM	85.0 ± 1.2	81.4 ± 0.2	95.9 ± 0.3	80.9 ± 0.5	85.8
VREx	87.8 ± 1.2	81.8 ± 0.7	97.4 ± 0.2	82.1 ± 0.7	87.2
RSC	86.0 ± 0.7	81.8 ± 0.9	96.8 ± 0.7	80.4 ± 0.5	86.2
EQRM	88.3 ± 0.6	82.1 ± 0.5	97.2 ± 0.4	81.6 ± 0.5	87.3

OfficeHome

Algorithm	A	C	P	R	Avg
ERM	61.7 ± 0.7	53.4 ± 0.3	74.1 ± 0.4	76.2 ± 0.6	66.4
IRM	56.4 ± 3.2	51.2 ± 2.3	71.7 ± 2.7	72.7 ± 2.7	63.0
GroupDRO	60.5 ± 1.6	53.1 ± 0.3	75.5 ± 0.3	75.9 ± 0.7	66.2
Mixup	63.5 ± 0.2	54.6 ± 0.4	76.0 ± 0.3	78.0 ± 0.7	68.0
MLDG	60.5 ± 0.7	54.2 ± 0.5	75.0 ± 0.2	76.7 ± 0.5	66.6
CORAL	64.8 ± 0.8	54.1 ± 0.9	76.5 ± 0.4	78.2 ± 0.4	68.4
MMD	60.4 ± 1.0	53.4 ± 0.5	74.9 ± 0.1	76.1 ± 0.7	66.2
DANN	60.6 ± 1.4	51.8 ± 0.7	73.4 ± 0.5	75.5 ± 0.9	65.3
CDANN	57.9 ± 0.2	52.1 ± 1.2	74.9 ± 0.7	76.2 ± 0.2	65.3
MTL	60.7 ± 0.8	53.5 ± 1.3	75.2 ± 0.6	76.6 ± 0.6	66.5
SagNet	62.7 ± 0.5	53.6 ± 0.5	76.0 ± 0.3	77.8 ± 0.1	67.5
ARM	58.8 ± 0.5	51.8 ± 0.7	74.0 ± 0.1	74.4 ± 0.2	64.8
VREx	59.6 ± 1.0	53.3 ± 0.3	73.2 ± 0.5	76.6 ± 0.4	65.7
RSC	61.7 ± 0.8	53.0 ± 0.9	74.8 ± 0.8	76.3 ± 0.5	66.5
EQRM	60.0 ± 0.8	54.4 ± 0.7	76.5 ± 0.4	77.2 ± 0.5	67.0

TerraIncognita

Algorithm	L100	L38	L43	L46	Avg
ERM	59.4 ± 0.9	49.3 ± 0.6	60.1 ± 1.1	43.2 ± 0.5	53.0
IRM	56.5 ± 2.5	49.8 ± 1.5	57.1 ± 2.2	38.6 ± 1.0	50.5
GroupDRO	60.4 ± 1.5	48.3 ± 0.4	58.6 ± 0.8	42.2 ± 0.8	52.4
Mixup	67.6 ± 1.8	51.0 ± 1.3	59.0 ± 0.0	40.0 ± 1.1	54.4
MLDG	59.2 ± 0.1	49.0 ± 0.9	58.4 ± 0.9	41.4 ± 1.0	52.0
CORAL	60.4 ± 0.9	47.2 ± 0.5	59.3 ± 0.4	44.4 ± 0.4	52.8
MMD	60.6 ± 1.1	45.9 ± 0.3	57.8 ± 0.5	43.8 ± 1.2	52.0
DANN	55.2 ± 1.9	47.0 ± 0.7	57.2 ± 0.9	42.9 ± 0.9	50.6
CDANN	56.3 ± 2.0	47.1 ± 0.9	57.2 ± 1.1	42.4 ± 0.8	50.8
MTL	58.4 ± 2.1	48.4 ± 0.8	58.9 ± 0.6	43.0 ± 1.3	52.2
SagNet	56.4 ± 1.9	50.5 ± 2.3	59.1 ± 0.5	44.1 ± 0.6	52.5
ARM	60.1 ± 1.5	48.3 ± 1.6	55.3 ± 0.6	40.9 ± 1.1	51.2
VREx	56.8 ± 1.7	46.5 ± 0.5	58.4 ± 0.3	43.8 ± 0.3	51.4
RSC	59.9 ± 1.4	46.7 ± 0.4	57.8 ± 0.5	44.3 ± 0.6	52.1
EQRM	57.0 ± 1.5	49.5 ± 1.2	59.0 ± 0.3	43.4 ± 0.6	52.2

DomainNet

Algorithm	clip	info	paint	quick	real	sketch	Avg
ERM	58.6 ± 0.3	19.2 ± 0.2	47.0 ± 0.3	13.2 ± 0.2	59.9 ± 0.3	49.8 ± 0.4	41.3
IRM	40.4 ± 6.6	12.1 ± 2.7	31.4 ± 5.7	9.8 ± 1.2	37.7 ± 9.0	36.7 ± 5.3	28.0
GroupDRO	47.2 ± 0.5	17.5 ± 0.4	34.2 ± 0.3	9.2 ± 0.4	51.9 ± 0.5	40.1 ± 0.6	33.4
Mixup	55.6 ± 0.1	18.7 ± 0.4	45.1 ± 0.5	12.8 ± 0.3	57.6 ± 0.5	48.2 ± 0.4	39.6
MLDG	59.3 ± 0.1	19.6 ± 0.2	46.8 ± 0.2	13.4 ± 0.2	60.1 ± 0.4	50.4 ± 0.3	41.6
CORAL	59.2 ± 0.1	19.9 ± 0.2	47.4 ± 0.2	14.0 ± 0.4	59.8 ± 0.2	50.4 ± 0.4	41.8
MMD	32.2 ± 13.3	11.2 ± 4.5	26.8 ± 11.3	8.8 ± 2.2	32.7 ± 13.8	29.0 ± 11.8	23.5
DANN	53.1 ± 0.2	18.3 ± 0.1	44.2 ± 0.7	11.9 ± 0.1	55.5 ± 0.4	46.8 ± 0.6	38.3
CDANN	54.6 ± 0.4	17.3 ± 0.1	44.2 ± 0.7	12.8 ± 0.2	56.2 ± 0.4	45.9 ± 0.5	38.5
MTL	58.0 ± 0.4	19.2 ± 0.2	46.2 ± 0.1	12.7 ± 0.2	59.9 ± 0.1	49.0 ± 0.0	40.8
SagNet	57.7 ± 0.3	19.1 ± 0.1	46.3 ± 0.5	13.5 ± 0.4	58.9 ± 0.4	49.5 ± 0.2	40.8
ARM	49.6 ± 0.4	16.5 ± 0.3	41.5 ± 0.8	10.8 ± 0.1	53.5 ± 0.3	43.9 ± 0.4	36.0
VREx	43.3 ± 4.5	14.1 ± 1.8	32.5 ± 5.0	9.8 ± 1.1	43.5 ± 5.6	37.7 ± 4.5	30.1
RSC	55.0 ± 1.2	18.3 ± 0.5	44.4 ± 0.6	12.5 ± 0.1	55.7 ± 0.7	47.8 ± 0.9	38.9
EQRM	55.5 ± 1.8	19.6 ± 0.1	45.9 ± 1.9	12.9 ± 0.3	61.1 ± 0.0	50.3 ± 0.1	40.9

Averages

Algorithm	VLCS	PACS	OfficeHome	TerraIncognita	DomainNet	Avg
ERM	77.6 ± 0.3	86.7 ± 0.3	66.4 ± 0.5	53.0 ± 0.3	41.3 ± 0.1	65.0
IRM	76.9 ± 0.6	84.5 ± 1.1	63.0 ± 2.7	50.5 ± 0.7	28.0 ± 5.1	60.6
GroupDRO	77.4 ± 0.5	87.1 ± 0.1	66.2 ± 0.6	52.4 ± 0.1	33.4 ± 0.3	63.3
Mixup	78.1 ± 0.3	86.8 ± 0.3	68.0 ± 0.2	54.4 ± 0.3	39.6 ± 0.1	65.4
MLDG	77.5 ± 0.1	86.8 ± 0.4	66.6 ± 0.3	52.0 ± 0.1	41.6 ± 0.1	64.9
CORAL	77.7 ± 0.2	87.1 ± 0.5	68.4 ± 0.2	52.8 ± 0.2	41.8 ± 0.1	65.6
MMD	77.9 ± 0.1	87.2 ± 0.1	66.2 ± 0.3	52.0 ± 0.4	23.5 ± 9.4	61.4
DANN	79.7 ± 0.5	85.2 ± 0.2	65.3 ± 0.8	50.6 ± 0.4	38.3 ± 0.1	63.8
CDANN	79.9 ± 0.2	85.8 ± 0.8	65.3 ± 0.5	50.8 ± 0.6	38.5 ± 0.2	64.1
MTL	77.7 ± 0.5	86.7 ± 0.2	66.5 ± 0.4	52.2 ± 0.4	40.8 ± 0.1	64.8
SagNet	77.6 ± 0.1	86.4 ± 0.4	67.5 ± 0.2	52.5 ± 0.4	40.8 ± 0.2	65.0
ARM	77.8 ± 0.3	85.8 ± 0.2	64.8 ± 0.4	51.2 ± 0.5	36.0 ± 0.2	63.1
VREx	78.1 ± 0.2	87.2 ± 0.6	65.7 ± 0.3	51.4 ± 0.5	30.1 ± 3.7	62.5
RSC	77.8 ± 0.6	86.2 ± 0.5	66.5 ± 0.6	52.1 ± 0.2	38.9 ± 0.6	64.3
EQRM	77.8 ± 0.2	87.3 ± 0.2	67.0 ± 0.4	52.2 ± 0.7	40.9 ± 0.3	65.1

G.3 WILDS

In Figure 7, we visualize the test-time risk distributions of IRM and GroupDRO relative to ERM, as well as $EQRM_\alpha$ for select values[10] of α. In each of these figures, we see that IRM and GroupDRO tend to have heavier tails than any of the other algorithms.

Figure 7: **Baseline test risk distributions on iWildCam and OGB-MolPCBA.** We supplement Figure 4 by providing comparisons to two baseline algorithms: IRM and GroupDRO. In each case, $EQRM_\alpha$ tends to display superior tail performance relative to ERM, IRM, and GroupDRO.

Other performance metrics. In the main text, we studied the tails of the *risk* distributions of predictors trained on `iWildCam` and `OGB`. However, in the broader DG literature, there are a number of other metrics that are used to assess performance or OOD-generalization. In particular, for `iWildCam`, past work has used the macro F_1 score as well as the average accuracy across domains to assess OOD generalization; for `OGB`, the standard metric is a predictor's average precision over test domains [12]. In Tables 6 and 7, we report these metrics and compare the performance of our algorithms to ERM, IRM, and GroupDRO. Below, we discuss the results in each of these tables.

To begin, consider Table 6. Observe that ERM achieves the best *in-distribution* (ID) scores relative to any of the other algorithms. However, when we consider the *out-of-distribution* columns, we see that EQRM offers better performance with respect to both the macro F_1 score and the mean accuracy. Thus, although our algorithms are not explicitly trained to optimize these metrics, their strong performance on the tails of the risk distribution appears to be correlated with strong OOD performance with these alternative metrics. We also observe that relative to ERM, EQRM suffers smaller accuracy drops between ID and OOD mean accuracy. Specifically, ERM dropped 5.50 points, whereas EQRM dropped by an average of 2.38 points.

Next, consider Table 7. Observe again that ERM is the strongest-performing *baseline* (first band of the table). Also observe that EQRM performs similarly to ERM, with validation and test precision tending to cluster around 28 and 27 respectively. However, we stress that these metrics are *averaged* over their respective domains, whereas in Tables 2 and 3, we showed that EQRM performed well on the more difficult domains, i.e. when using *tail* metrics.

Table 6: WILDS metrics on `iWildCam`.

Algorithm	Macro F_1 ($\uparrow$)		Mean accuracy ($\uparrow$)	
	ID	OOD	ID	OOD
ERM	**49.8**	30.6	**77.0**	71.5
IRM	23.4	15.2	59.6	64.1
GroupDRO	34.3	22.1	66.7	67.7
$QRM_{0.25}$	18.3	11.4	54.3	58.3
$QRM_{0.50}$	48.1	33.8	76.2	73.5
$QRM_{0.75}$	49.5	31.8	76.1	72.0
$QRM_{0.90}$	48.6	32.9	77.1	73.3
$QRM_{0.99}$	45.9	30.8	76.6	71.3

Table 7: WILDS metrics on `OGB-MolPCBA`.

Algorithm	Mean precision ($\uparrow$)	
	Validation	Test
ERM	28.1	27.3
IRM	15.4	15.5
GroupDRO	23.5	22.3
$QRM_{0.25}$	28.1	27.3
$QRM_{0.50}$	**28.3**	**27.4**
$QRM_{0.75}$	28.1	27.1
$QRM_{0.90}$	27.9	27.2
$QRM_{0.99}$	28.1	27.4

---

[10]We display results for fewer values of α in Figure 7 to keep the plots uncluttered.

H Limitations of our work

As discussed in the first paragraph of § 7, the main limitation of our work is that, for α to *precisely* approximate the probability of generalizing with risk below the associated α-quantile value, we must have a large number of i.i.d.-sampled domains. Currently, this is rarely satisfied in practice, although § 7 describes how new data-collection procedures could help to better-satisfy this assumption. We believe that our work, and its promise of machine learning systems that generalize with high probability, provides sufficient motivation for collecting real-world datasets with a large number of i.i.d.-sampled domains. In addition, we hope that future work can explore ways to relax this assumption, e.g., by leveraging knowledge of domain dependencies like time.