



UvA-DARE (Digital Academic Repository)

An Information-theoretic Approach to Distribution Shifts

Federici, M.; Forre, P.; Tomioka, R.

DOI

[10.48550/arXiv.2106.03783](https://doi.org/10.48550/arXiv.2106.03783)

Publication date

2022

Document Version

Author accepted manuscript

Published in

35th Conference on Neural Information Processing Systems (NeurIPS 2021)

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Federici, M., Forre, P., & Tomioka, R. (2022). An Information-theoretic Approach to Distribution Shifts. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, & J. Wortman Vaughan (Eds.), *35th Conference on Neural Information Processing Systems (NeurIPS 2021) : online, 6-14 December 2021* (Vol. 21, pp. 17628-17641). (Advances in Neural Information Processing Systems; Vol. 34). Neural Information Processing Systems Foundation. <https://doi.org/10.48550/arXiv.2106.03783>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

An Information-theoretic Approach to Distribution Shifts

Marco Federici
AMLab
University of Amsterdam
m.federici@uva.nl

Ryota Tomioka
Microsoft Research
Cambridge, UK
ryoto@microsoft.com

Patrick Forré
AI4Science Lab, AMLab
University of Amsterdam
p.d.forre@uva.nl

Abstract

Safely deploying machine learning models to the real world is often a challenging process. Models trained with data obtained from a specific geographic location tend to fail when queried with data obtained elsewhere, agents trained in a simulation can struggle to adapt when deployed in the real world or novel environments, and neural networks that are fit to a subset of the population might carry some selection bias into their decision process. In this work, we describe the problem of data shift from a novel information-theoretic perspective by (i) identifying and describing the different sources of error, (ii) comparing some of the most promising objectives explored in the recent domain generalization and fair classification literature. From our theoretical analysis and empirical evaluation, we conclude that the model selection procedure needs to be guided by careful considerations regarding the observed data, the factors used for correction, and the structure of the data-generating process.

1 Introduction

One of the most common assumptions for machine learning models is that the training and test data are independently and identically sampled (IID) from the same distribution. In practice, this assumption does not hold in many practical scenarios (Bengio et al., 2020). A machine learning model trained to recognize land usage from satellite images using pictures from the early 2000s may struggle to recognize the style of modern architectures (Christie et al., 2018), data collected on a limited set of hospitals may not be representative of the variation introduced by the use of different machines or procedures (Zech et al., 2018; Beede et al., 2020). Other kinds of distribution shifts are more subtle and difficult to recognize despite having a noticeable impact on the model’s predictive performance. Examples include under-represented or over-represented population groups (Popejoy & Fullerton, 2016; Buolamwini & Gebu, 2018) or biased annotations collected from crowd-sourcing services (Zhang et al., 2017; Xia et al., 2020).

Different approaches in literature address these issues by using some external source of knowledge such as domain or environment annotations (Wang & Deng, 2018), protected attributes (Mehrabi et al., 2019) or sub-population groups (Santurkar et al., 2020) to reduce bias and minimize the model error outside of the training distribution. Despite the progress in the field of domain generalization literature, Gulrajani & Lopez-Paz (2020) has shown that the effectiveness of some of the most common algorithms heavily relies on the hyper-parameter tuning strategy, revealing limitations of

examined models when compared to a more traditional empirical risk minimization strategy. One of the major issues behind this observed behavior is a lack of clarity and applicability for the underlying assumptions regarding the problem statement and the data-generating process (Zhao et al., 2019; Rosenfeld et al., 2020; Mahajan et al., 2020).

With this work, we aim to present a new perspective on the problem to analyze and clarify the fundamental differences between some of the most common approaches by:

1. Introducing a novel information-theoretical framework to describe the problem of distribution shift and connecting it to the test error and its components (section 2).
2. Analyzing four main families of objectives and describing some of their guarantees and assumptions (section 3).
3. Demonstrating that the effectiveness of different criteria is determined by the structure of the underlying data-generating process (section 4).
4. Showing that the results obtained by popular models designed according to the aforementioned criteria can drastically differ from the theoretically expected performance (section 4).

The analysis of the *bottleneck*, *independence*, *sufficiency*, and *separation* criteria reveals that some of the most popular models can systematically fail to reduce the test error even for simple datasets. No unique objective is simultaneously optimal for every problem, but additional knowledge about the selection procedure and better approximations can help to mitigate the bias.

1.1 Problem Statement

Consider \mathbf{x} and \mathbf{y} as the features and targets respectively with joint density $p(\mathbf{x}, \mathbf{y})$ for the predictive problem of interest. Let t be a binary random variable representing which data is selected for training ($t = 1$) and which is not ($t = 0$). We will refer to $p(\mathbf{x}, \mathbf{y}|t = 1)$ as the joint *Training distribution* and $p(\mathbf{x}, \mathbf{y}|t = 0)$ as the *Test distribution* that are induced by the selection t . Throughout the paper, we will consider selections t that can be expressed as a function of the features \mathbf{x} , targets \mathbf{y} , independent noise ϵ and other variables \mathbf{e} , which represent other factors that are affecting the data collection procedure (e.g. geographic location, time intervals, population groups).

Let $q(\mathbf{y}|\mathbf{x})$ represent a learnable model for the predictive distribution $p(\mathbf{y}|\mathbf{x})$. We use the Kullback-Leibler divergence to express the *Train* and *Test* error respectively¹:

$$\underbrace{D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} || q_{\mathbf{y}|\mathbf{x}})}_{\text{Train error}} := D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t = 1) || q(\mathbf{y}|\mathbf{x})), \quad (1)$$

$$\underbrace{D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0} || q_{\mathbf{y}|\mathbf{x}})}_{\text{Test error}} := D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t = 0) || q(\mathbf{y}|\mathbf{x})). \quad (2)$$

Although only the train distribution can be accessed at training time, we are interested in learning models $q(\mathbf{y}|\mathbf{x})$ that result in small test error.

1.2 Characterizing Distribution Shift

As a first step, we characterize the difference between training and test distribution as a function of the selection. The effect that the selection has in the joint distribution $p(\mathbf{x}, \mathbf{y})$ can be quantified by considering the mutual information between the features-target pair \mathbf{xy} and the selection variable t , which can be also expressed as a Kullback-Leibler divergence: $I(\mathbf{xy}; t) = D_{\text{KL}}(p(\mathbf{x}, \mathbf{y}|t) || p(\mathbf{x}, \mathbf{y}))$. This measure of distribution shift intuitively represents how many bits the selection variable carries about the joint distribution of targets and features or, equivalently, how much the joint density $p(\mathbf{x}, \mathbf{y})$ has changed as a result of the selection.

Using the chain rule of mutual information, one can express distribution shift as the sum of two separate components:

$$\underbrace{I(\mathbf{xy}; t)}_{\text{Distribution shift}} = \underbrace{I(\mathbf{x}; t)}_{\text{Covariate shift}} + \underbrace{I(\mathbf{y}; t|\mathbf{x})}_{\text{Concept shift}}, \quad (3)$$

¹Further details regarding the convention used for conditional KL-divergence can be found in appendix A




Figure 1: (a) Example of distribution shift due to selection bias (chest pain complaints) for the prediction of cholesterol levels (y-axis) given blood pressure (x-axis). A maximum likelihood approach (green dashed line) results in a model that over-estimates y for any given x when compared to the best model in the class (red dashed line). (b) Example of a data generating process in which the features \mathbf{x} are composed of two parts x_1 and x_2 . When both feature components x_1 and x_2 are observed, concept shift can be strictly positive. Discarding information about x_2 will reduce the effect of the selection bias, while removing x_1 might increase it $0 = I(\mathbf{y}; t|x_1) \leq I(\mathbf{y}; t|x_1x_2) \leq I(\mathbf{y}; t|x_2)$.

which can be interpreted as the amount of *covariate shift* (Shimodaira, 2000) and *concept shift* (Widmer & Kubat, 1996; Moreno-Torres et al., 2012), respectively. These two quantities refer to the changes in the predictive distribution $p(\mathbf{y}|\mathbf{x})$ and the marginal features distribution $p(\mathbf{x})$ respectively, which add up to represent the changes in the joint distribution.

Whenever the data selection is perfectly IID, the selection variable can be expressed as a function of some independent noise ($t = f(\epsilon)$) and the corresponding distribution shift $I(\mathbf{xy}; t)$ is zero. On the other hand, if the data collection procedure has been influenced by other factors, we do not have such a guarantee, even when the selection does not depend on features and targets directly ($t = f(\mathbf{e}, \epsilon)$).

2 Information-theoretic Framework

Using the quantities described in the previous sections, we can show that the sum train and test error is lower-bounded by the amount concept shift:

Proposition 1. For any model $q(\mathbf{y}|\mathbf{x})$ and $\alpha := \min\{p(t=0), p(t=1)\}$:

$$D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} || q_{\mathbf{y}|\mathbf{x}}) + D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0} || q_{\mathbf{y}|\mathbf{x}}) \geq \frac{1}{1-\alpha} I(\mathbf{y}; t|\mathbf{x}), \quad (4)$$

As a consequence, whenever the selection induces concept shift $I(\mathbf{y}; t|\mathbf{x}) > 0$, any sufficiently flexible model $\hat{q}(\mathbf{y}|\mathbf{x})$ trained using a maximum likelihood approach must incur in strictly positive test error.

Intuitively, whenever the selection induces a change in the predictive distribution ($p(\mathbf{y}|\mathbf{x}) \neq p(\mathbf{y}|\mathbf{x}, t=1)$) fitting the model to the training distribution will incorporate the selection bias into the model prediction, necessarily resulting in errors when evaluated on the test distribution, as shown in the example reported in figure 1a.

The different approaches analyzed in this work are based in the introduction of a *latent representation* \mathbf{z} , which allows for the definition of different regularization strategies. Using the latent representation as an intermediate variable, the model $q(\mathbf{y}|\mathbf{x})$ can be re-parametrized with an *encoder* $q(\mathbf{z}|\mathbf{x})$ and a *latent classifier*² $q(\mathbf{y}|\mathbf{z})$:

$$q(\mathbf{y}|\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x})} [q(\mathbf{y}|\mathbf{z} = \mathbf{z})]. \quad (5)$$

The effect of the introduction of a latent representation can be observed by expressing the training and test error as a function of the encoder and the classifier³.

²The name classifier will be used for convenience to express both classification and regression problems.

³The expressions in equations 6 and 7 hold with equality for deterministic encoders as shown in appendix B.

Proposition 2. For any encoder $q(\mathbf{z}|\mathbf{x})$ and classifier $q(\mathbf{y}|\mathbf{z})$:

$$D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{x}}) \leq \underbrace{I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z})}_{\text{Train information loss}} + \underbrace{D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1}||q_{\mathbf{y}|\mathbf{z}})}_{\text{Latent train error}} \quad (6)$$

$$D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0}||q_{\mathbf{y}|\mathbf{x}}) \leq \underbrace{I_{t=0}(\mathbf{x}; \mathbf{y}|\mathbf{z})}_{\text{Test information loss}} + \underbrace{D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=0}||q_{\mathbf{y}|\mathbf{z}})}_{\text{Latent test error}}. \quad (7)$$

The two terms $I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z})$ and $I_{t=0}(\mathbf{x}; \mathbf{y}|\mathbf{z})$ represent the amount of predictive information that is lost by encoding the features \mathbf{x} into \mathbf{z} on train and test distribution respectively, while $D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1}||q_{\mathbf{y}|\mathbf{z}})$ and $D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=0}||q_{\mathbf{y}|\mathbf{z}})$ refer to the train and test error when using \mathbf{z} instead of the original observations as the predictive features, which will be referred to as *latent training error* and *latent test error* respectively. Test information loss and latent test error capture two intrinsically different kinds of error. The former indicates the increase in the prediction uncertainty as a result of the encoding procedure, while the latter represents the discrepancy between the model $q(\mathbf{y}|\mathbf{z})$ and the latent test predictive distribution $p(\mathbf{y}|\mathbf{z}, t = 0)$.

2.1 Latent test error

Minimizing the latent test error makes the latent predictor $q(\mathbf{y}|\mathbf{z})$ approach the test latent predictive distribution $p(\mathbf{y}|\mathbf{z}, t = 0)$. We can show that the Jensen-Shannon divergence between the two distributions is upper-bounded by a monotonic function of the latent training error and the amount of concept shift in the latent space (*latent concept shift* $I(\mathbf{y}; t|\mathbf{z})$):

Proposition 3. For any $q(\mathbf{z}|\mathbf{x})$, $q(\mathbf{y}|\mathbf{z})$ and any representation \mathbf{z} that satisfies $p(\mathbf{z} = \mathbf{z}|t = 0) > 0$ and $p(\mathbf{z} = \mathbf{z}|t = 1) > 0$:

$$\left(\sqrt{\frac{1}{2\alpha} I(\mathbf{y}; t|\mathbf{z} = \mathbf{z})} + \sqrt{\frac{1}{2} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1}||q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})} \right)^2 \geq D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0}||q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}). \quad (8)$$

Whenever the Jensen-Shannon divergence between the test predictive distribution $p(\mathbf{y}|\mathbf{z}, t = 0)$ and the classifier $q(\mathbf{y}|\mathbf{z})$ is small, the latent test error (measured in terms of KL-divergence) must also be small at least for the regions that have positive probability according to both train $p(\mathbf{x}|t = 1)$ and test $p(\mathbf{x}|t = 0)$ data distributions. Since the train predictive distribution $p(\mathbf{y}|\mathbf{x}, t = 1)$ is not defined for \mathbf{x} that have zero probability on the train distribution, we have no guarantees regarding the model predictions in those regions unless other inductive biases are considered.

The left hand side of the expression in proposition 3 can be minimized by addressing $I(\mathbf{y}; t|\mathbf{z})$ and $D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1}||q_{\mathbf{y}|\mathbf{z}})$ with respect to the encoder and classifier respectively. In other words, we can minimize the latent test error by simultaneously fitting $q(\mathbf{y}|\mathbf{z})$ to the train predictive distribution $p(\mathbf{y}|\mathbf{z}, t = 1)$ and minimizing the latent concept shift induced by the encoder $q(\mathbf{z}|\mathbf{x})$. When the latent concept shift and latent test error approach zero, the latent train predictive distribution $p(\mathbf{y}|\mathbf{z}, t = 1)$ approaches the true (unselected) predictive distribution $p(\mathbf{y}|\mathbf{z})$, and so does the modeled classifier $q(\mathbf{y}|\mathbf{z})$. In this ideal scenario, the only source of test error is the additional uncertainty that is due to the information lost in the encoding procedure.

2.2 Minimizing the information loss

Since losing information generally results in increased test error (equation 7), we will consider objectives that discard the minimal amount of information required to reduce the latent concept shift. This can be done by minimizing the KL-divergence between the training predictive distribution $p(\mathbf{y}|\mathbf{x}; t = 1)$ and the latent classifier model $q(\mathbf{y}|\mathbf{z})$:

$$\begin{aligned} \min_{q(\mathbf{z}|\mathbf{x})} I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z}) &= \min_{q(\mathbf{z}|\mathbf{x}), q(\mathbf{y}|\mathbf{z})} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{z}}) - KL(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{z}}) \\ &\leq \min_{q(\mathbf{z}|\mathbf{x}), q(\mathbf{y}|\mathbf{z})} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{z}}). \end{aligned} \quad (9)$$

In addition to reducing the amount of information lost in the encoding procedure, minimizing the right hand side of the expression in equation 9 makes the model $q(\mathbf{y}|\mathbf{z})$ approach the latent predictive

distribution $p(\mathbf{y}|\mathbf{z}, t = 1)$. Note that only the selected training distribution $p(\mathbf{y}, \mathbf{x}|t = 1)$ is available at training time. In practice, minimizing the train information loss also generally results in decreased test information loss since the same features are usually informative for both train and test distributions.

2.3 A General Loss function

To summarize, the overall objective consists in finding the maximally informative representation that minimizes latent concept shift so that the same learned predictor $q(\mathbf{y}|\mathbf{z})$ can perform similarly on both train and test settings. This is achieved by (i) minimizing the amount of latent concept shift induced by the representation, (ii) maximizing the amount of predictive information in the representation, and (iii) matching $q(\mathbf{y}|\mathbf{z})$ and the training latent predictive distribution $p(\mathbf{y}|\mathbf{z}, t = 1)$. The three requirements can be enforced by considering a loss function with the following form:

$$\mathcal{L}(q_{\mathbf{z}|\mathbf{x}}, q_{\mathbf{y}|\mathbf{z}}; \lambda) = D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} || q_{\mathbf{y}|\mathbf{z}}) + \lambda \mathcal{R}(q_{\mathbf{z}|\mathbf{x}}), \quad (10)$$

in which, the first term addresses (ii) and (iii), while the second term represent a regularization term that acts on the encoder $q(\mathbf{z}|\mathbf{x})$ to minimize latent concept shift $I(\mathbf{y}; t|\mathbf{z})$, following (i). For a sufficiently flexible family of latent predictors, requirement (iii) depends only $q(\mathbf{y}|\mathbf{z})$, while the hyper-parameter λ defines the trade-off between latent concept shift (i) and predictive information loss (ii).

3 Regularization Criteria

Since only selected data ($t = 1$) is accessible at training time, the latent concept shift can not be computed or minimized directly. Most of the approaches considered in this analysis make use of a regularization $\mathcal{R}(q_{\mathbf{z}|\mathbf{x}})$ that is based on the observation of an additional variable \mathbf{e} which relates to the selection criteria. This variable is usually referred to as *domain* or *environment*⁴ in the domain adaptation and generalization literature, while the name *protected attribute* is used in the context of fair classification. We will refer to this variable \mathbf{e} as *environmental factors* in the following sections.

We analyze four families of criteria proposed in the representation learning (Tishby & Zaslavsky, 2015), domain generalization (Koyama & Yamaguchi, 2020) and fair classification (Barocas et al., 2018) literature focusing on their underlying assumptions and theoretical guarantees. The different regularization strategies and models can be seen as specific instance of the loss function in equation 10. An empirical comparison between instances of the different criteria can be found in section 4, proofs are reported in appendix B, while the relation between the reported criteria is further discussed in appendix B.9.

3.1 Information Bottleneck Criterion

Combining the results from propositions 1 and 2 we can infer that reducing the latent test error necessarily requires the representation to discard some train predictive information ($I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z}) > 0$). This is because a lossless encoder $\hat{q}(\mathbf{z}|\mathbf{x})$ together with an optimal latent classifier $\hat{q}(\mathbf{y}|\mathbf{z})$ would result in an overall model $\hat{q}(\mathbf{y}|\mathbf{x})$ that matches the train distribution $p(\mathbf{y}|\mathbf{x}, t = 1)$, and, therefore, results in positive latent test error (proposition 1).

The *Information Bottleneck* criterion (Tishby & Zaslavsky, 2015) introduces a regularization term $\mathcal{R}(q_{\mathbf{z}|\mathbf{x}}) = I_{t=1}(\mathbf{x}; \mathbf{z})$ to control the amount of information in the representation, defining a lossy compression scheme (Alemi et al., 2017) that can be regulated using the hyper-parameter λ . Note that this criterion does not use any environmental information and blindly discards data-features depending on the regularization strength λ . As shown in the example reported in figure 1b, discarding information without any additional constraint can increase the amount of latent concept shift depending on the structure of the underlying data generating process. This is because discarding information is a necessary but not sufficient condition to reduce the latent concept shift.

⁴In contrast with the domain adaptation and generalization literature, we will consider the more general case in which \mathbf{e} is represented by a vector.

3.2 Independence Criterion

Whenever the data selection t depends only on some observed variable \mathbf{e} ($t = f(\mathbf{e}; \epsilon)$), the most intuitive approach to reduce the latent test error is to make the representation \mathbf{z} independent of the environmental factors \mathbf{e} . This can be done by minimizing mutual information between \mathbf{e} and \mathbf{z} : $\mathcal{R}(q_{\mathbf{z}|\mathbf{x}}) := I_{t=1}(\mathbf{e}; \mathbf{z})$. This criterion, known as *independence* or *statistical parity* in the fair classification literature (Dwork et al., 2011; Corbett-Davies et al., 2017), aims to remove any environmental information from \mathbf{z} , resulting in a representation that satisfies $p(\mathbf{z}|t = 1) = p(\mathbf{z}|\mathbf{e}, t = 1)$.

Despite the usefulness of this criterion in the fairness and differential privacy literature, enforcing independence does not necessarily reduce the test error (Zhao et al., 2019; Johansson et al., 2019). This is because a consistent marginal across different environments does not imply a consistent predictive distribution ($p(\mathbf{y}|\mathbf{z}, \mathbf{e}, t = 1) \neq p(\mathbf{y}|\mathbf{z}, t = 1)$), and enforcing independence may even increase the latent concept shift and the test error (as shown in figure 3).

3.3 Sufficiency Criterion

Instead of enforcing a property on the marginal feature distribution, one can consider stable properties of the joint distribution of features \mathbf{z} and labels \mathbf{y} . The requirement of creating a representation that yields a stable classifier for different values of the environmental factor \mathbf{e} can be captured by the following regularization: $\mathcal{R}(q_{\mathbf{z}|\mathbf{x}}) := I_{t=1}(\mathbf{y}; \mathbf{e}|\mathbf{z})$. Intuitively minimizing $I_{t=1}(\mathbf{y}; \mathbf{e}|\mathbf{z})$ corresponds to minimizing the distance between the predictive distribution $p(\mathbf{y}|\mathbf{e}, \mathbf{z}, t = 1)$ for each one of the observed environmental conditions. We can show the following:

Proposition 4. *Whenever the selection t can be expressed as a function of \mathbf{e} and some independent noise ϵ , the latent concept shift induced by a representation \mathbf{z} of \mathbf{x} is upper-bounded by $I(\mathbf{y}; \mathbf{e}|\mathbf{z})$:*

$$\exists f : t = f(\mathbf{e}, \epsilon) \implies I(\mathbf{y}; t|\mathbf{z}) \leq I(\mathbf{y}; \mathbf{e}|\mathbf{z}). \quad (11)$$

In other words, for a given selection t , it is possible to remove the effect of concept shift by enforcing the *sufficiency* constraint using the variable (or variables) \mathbf{e} that are responsible for that selection.

The result from proposition 4 is applicable only if the two following conditions are met: (i) a sufficient representation must exist (Koyama & Yamaguchi, 2020); (ii) enforcing sufficiency on the selected train distribution results in sufficiency for the overall joint distribution ($I_{t=1}(\mathbf{y}; \mathbf{e}|\mathbf{z}) = 0 \implies I(\mathbf{y}; \mathbf{e}|\mathbf{z}) = 0$) (Rosenfeld et al., 2020). Even if assumptions (i) and (ii) are often acceptable in practice, lack of environment variety at training time or direct dependencies between environmental factors and targets (such as in the y-CMNIST dataset in section 4.2) can compromise the effectiveness of the sufficiency criterion. An in depth discussion with a simple example is reported in appendix G.

3.4 Separation Criterion

The last family of objectives and design principles includes approaches that aim to capture the stability in the latent feature distribution when the target is observed across different environmental conditions $p(\mathbf{z}|\mathbf{y}) = p(\mathbf{z}|\mathbf{y}, \mathbf{e})$ (Chouldechova, 2017; Li et al., 2018b). This requirement can be enforced by minimizing the dependency between environmental factors and the representation when the label is observed: $\mathcal{R}(q_{\mathbf{z}|\mathbf{x}}) := I_{t=1}(\mathbf{e}; \mathbf{z}|\mathbf{y})$. The resulting *separation* criterion (since \mathbf{y} separates \mathbf{z} and \mathbf{e}) can be used to identify stable properties of the joint distributions even when the selection t depends on both targets \mathbf{y} and environmental factors \mathbf{e} :

Proposition 5. *If the selection t can be expressed as a function of \mathbf{e} , \mathbf{y} and some independent noise ϵ , the latent concept shift of a representation \mathbf{z} of \mathbf{x} is upper-bounded by the sum of prior-shift $I(\mathbf{y}; t)$ and $I(\mathbf{e}; \mathbf{z}|\mathbf{y})$:*

$$\exists f : t = f(\mathbf{e}, \mathbf{y}, \epsilon) \implies I(\mathbf{y}; t|\mathbf{z}) \leq I(\mathbf{y}; t) + I(\mathbf{e}; \mathbf{z}|\mathbf{y}). \quad (12)$$

When selection t and targets \mathbf{y} are marginally independent, proposition 5 guarantees that the latent concept shift of a representation that enforces separation is zero. Furthermore, whenever the marginal distribution $p(\mathbf{y}|t = 0)$ is known, it is possible to adjust the prediction of the latent classifier on the test distribution⁵.

⁵Further details on the re-weighting procedure can be found in appendix B.8

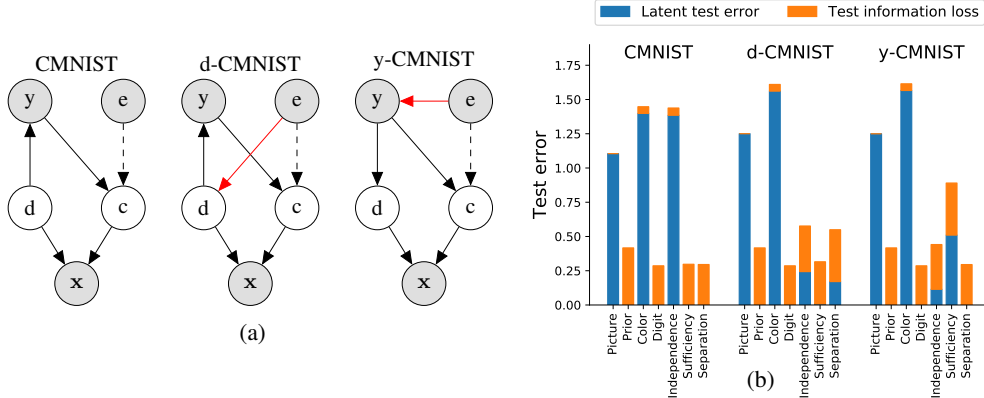


Figure 2: Graphical models (a) and error components (b) for the CMNIST, d-CMNIST and y-CMNIST data distributions. (a) Dashed lines are used to underline marginal independence between color c and environment e , while red arrows denote dependencies added to the original CMNIST distribution. (b) Models trained with strong regularization ($\lambda \approx 10^7$) for the different criteria are compared against the classifiers trained using only color, digit, picture, or prior information. The colors show the proportion of the test error (in nats) due to the predictive information loss ($I_{t=0}(\mathbf{x}; \mathbf{y}|\mathbf{z})$, in orange) and the latent test error ($D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=0} || q_{\mathbf{y}|\mathbf{z}})$, in blue) according to the decomposition in equation 7.

Similarly to the other criteria, one needs to assume that enforcing separation on train ($I_{t=1}(\mathbf{e}; \mathbf{z}|\mathbf{y}) = 0$) suffices to guarantee $I(\mathbf{e}; \mathbf{z}|\mathbf{y}) = 0$. Although a representation that enforces separation always exists, the effectiveness of the separation criteria depends on the data-generating process, since the requirement could be exclusively satisfied by a constant representation.

4 Experiments

We evaluate the effectiveness of the criteria presented in section 3 and some of their most popular implementations on multiple versions of the CMNIST dataset (Arjovsky et al., 2019) produced by altering the data-generating process (figure 2a) to underline the shortcomings of the different methods.

The main advantage of using a synthetic dataset such as CMNIST lies in the possibility to directly optimize for the different criteria using differentiable discrete mutual information measurements. This is possible since the joint occurrence of color and digit $\hat{\mathbf{x}} := [c, d]$ is a low-dimensional discrete sufficient statistic of the pictures \mathbf{x} , and, consequently, all the mutual information quantities of interest involving $\mathbf{x} \in [0, 1]^{28 \times 28 \times 2}$ can be computed using $\hat{\mathbf{x}} \in \{0, 1\} \times \{0, \dots, 9\}$ instead. Further details regarding the direct optimization of the criteria are reported in appendix D.

In figure 3, the theoretical performance for each criterion is compared against the results obtained by training different models on the pictures \mathbf{x} using neural network architectures to parametrize the encoder $q(\mathbf{z}|\mathbf{x})$ and classifier $q(\mathbf{y}|\mathbf{z})$ for different regularization strength λ . For the comparison, we consider diverse popular models designed according to the criteria defined in section 3:

- Variational Information Bottleneck (VIB) (Alemi et al., 2017): a variational tractable approximation of the **Information Bottleneck** criterion;
- Domain Adversarial Neural Network (DANN) (Ganin et al., 2016): an adversarial model based on a min-max game with discriminator $d(\mathbf{e}|\mathbf{z})$ that is optimized to predict environment information from the representation \mathbf{z} to enforce the **Independence** criterion;
- Invariant Risk Minimization (IRM) (Arjovsky et al., 2019): A model designed following the **Sufficiency** criterion that aims to create a representation from which the same latent classifier is simultaneously optimal on all environments.
- Conditional Domain Adversarial Neural Network (CDANN) (Li et al., 2018b): an adversarial approximation of the **Separation** criterion in which, analogously to DANN, a discriminator tries to predict the environment when the representation \mathbf{z} and the true label \mathbf{y} are given.




Figure 3: Comparison of training cross-entropy (x-axis) and test cross-entropy (y-axis) on the CMNIST, d-CMNIST and y-CMNIST datasets for representations trained using different criteria directly (top row) and corresponding approximations/model (bottom row). Each trajectory describes the trade-off between the two errors for a wide range of regularization strength values λ . The white symbols in each plot represent the error of models that consider only color ($p(y|c, t = 1)$), digit ($p(y|d, t = 1)$), picture ($p(y|x, t = 1)$) or prior ($p(y|t = 1)$) information. The shaded area represents the standard deviation observed across three runs with different seeds.

- Variance-based Risk Extrapolation (VREx) (Krueger et al., 2020): a model designed following the **Sufficiency** and **Independence** criteria, based on the minimization of the training error variance across the different environments;

For a better comparison, all the models use the same encoder and classifier neural network architectures. Each model has been trained by slowly increasing the regularization strength after an initial pre-training with small λ . Further details regarding the neural network architectures, objectives, optimization and specific hyper-parameters can be found in appendix E.⁶

4.1 Evaluation metric

The measure of test error, information loss and latent test error reported in figure 2b can be computed only for discrete variables. For neural network models the training and test error defined in section 2 can be estimated up to a constant entropy by considering the expected negative log-likelihood (empirical cross-entropy):

$$D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} || q_{\mathbf{y}|\mathbf{x}}) \approx -\frac{1}{N} \sum_{i=1}^N \log q(\mathbf{y} = \mathbf{y}_i | \mathbf{x} = \mathbf{x}_i) - \underbrace{H_{t=1}(\mathbf{y}|\mathbf{x})}_{\text{constant}} \quad (13)$$

The samples $\mathbf{x}_i, \mathbf{y}_i$ are obtained from $p(\mathbf{x}, \mathbf{y} | t = 1)$ and $p(\mathbf{x}, \mathbf{y} | t = 0)$ for train and test cross-entropy respectively. By computing the training and test cross-entropy values for different regularization strength λ , each model defines a trajectory from a maximum likelihood solution ($\lambda = 0$) to the results obtained when the corresponding independence constraint is enforced (large λ), as shown in figure 3. Contrarily to accuracy measurements, the use of cross-entropy allows us to detect under-confident or over-confident predictive distributions.

⁶The implementation of the models in this work is available at <https://github.com/mfederici/dsit>

4.2 Datasets

The two variants of the CMNIST dataset considered in this work have been designed by minimally changing the original distribution to underline the strengths and weaknesses of the different criteria. Across different experiments, each MNIST picture x of a digit d is associated with a color c that depends on a binary target y and an environment e . The d-CMNIST and y-CMNIST datasets are created by adding a dependency from the environment to digit d (d-CMNIST) and label y (y-CMNIST), which results in a correlation between targets and environment ($I(e; y) > 0$). Further details regarding the conditional distributions used to produce the different datasets can be found in appendix C.

The strong correlation between color and label across the different CMNIST versions can be seen as an artifact introduced by the selection $t = f(e)$, and models that consider only digit information (white triangles in figure 3) outperform the ones that capture color information (white crosses) or both (white circle) in terms of test cross-entropy (y-axis).

CMNIST The CMNIST dataset was originally designed to underline the weaknesses of maximum-likelihood and Empirical Risk Minimization strategies (Arjovsky et al., 2019). The results displayed in the first column of figure 3 confirm that both the sufficiency and separation criteria manage to effectively reduce the test cross entropy error for sufficiently strong regularization λ , minimizing the latent test error while retaining more predictive information than a constant representation (first column in figure 2b). On the CMNIST dataset, most of the models in analysis closely follow the trajectory estimated by optimizing the corresponding criterion directly. The trajectory defined by the VIB model differs from the one described by the Information Bottleneck criterion since, in absence of additional constraints, the nature of the information that is discarded (either color, digit or style) depends on the inductive bias introduced by the specific architecture.

d-CMNIST The d-CMNIST dataset adds a dependency between environment e and digit d , increasing the frequency of specific digits for some environments. Digit information makes environment and label conditionally independent (y and e are d -separated in figure 2a). As claimed in Arjovsky et al. (2019), models based on the sufficiency criterion manage to reduce the latent test error (second column of figure 2b), while the separation criterion fails because of the direct dependency between environment and digits. The independence criterion does not improve the test performance since both color and digit correlate with the environment. Both DANN and CDANN architectures results in trajectories that are similar to the ones obtained optimizing for the corresponding criteria. Note that enforcing separation or independence does not improve upon the result that can be obtained by blindly discarding information using VIB. Despite the effectiveness of the sufficiency criterion, the model trained with the IRM objective lies far from the optimal solution. We believe that this is due to the relaxations and approximation introduced by the optimized objective as discussed in appendix F.

y-CMNIST Adding a dependency between environment e and label y simulates the scenario in which some labels are more frequent in some environments. The arrow between d and y flips when compared to the d-CMNIST dataset to represent stable $p(d|y)$ across different environments. The corresponding y-CMNIST dataset includes a path from e to y that can not be blocked since no representation z can achieve sufficiency ($I(e; y|z) = 0$). Despite the optimality of the separation criterion (red line, third column of figure 3), the model trained with the CDANN objective struggles to minimize the test error due to instabilities of the adversarial training procedure, as discussed in appendix F. Once again, the trajectory described by the VREx objective is more favorable when compared to VIB, while the model trained using the IRM objective is far from optimality. Figure 2b confirms that, on the y-CMNIST dataset, the separation criterion is the only one that minimizes the latent test error without discarding the entirety of the picture information. This underlines that the effectiveness of sufficiency and separation criteria strongly depends on the structure of the underlying graphical model.

5 Related work

The problem of distribution (or dataset) shift (Quionero-Candela et al., 2009; Koh et al., 2020) has been explored in different areas in the machine learning literature ranging domain adaptation and generalization (Wang & Deng, 2018) to sub-population shift and fair classification (Mehrabi et al.,

2019; Santurkar et al., 2020). Although goals, availability of test features and environments at training time, and data selection criteria (Koh et al., 2020) may differ, most recent approaches focus on extracting features with some desired properties through three different families of objectives.

The independence criterion has been used to find the transformation which minimizes the distance between the distribution of the encoded features across different environments by using adversarial training (Ganin et al., 2016; Xie et al., 2017; Li et al., 2018a), features adjustment (Lum & Johndrow, 2016), kernel methods (Muandet et al., 2013), or variational approaches (Louizos et al., 2016; Moyer et al., 2018; Ilse et al., 2020). Despite the success of the independence criterion, Zhao et al. (2019); Johansson et al. (2019) have shown that enforcing stability of the marginal feature distribution is not sufficient to guarantee generalization.

Li et al. (2018b) extends the independence criterion using adversarial training to create representations that are independent of the environment when the label is observed. The corresponding separation criterion has been explored in the context of fair classification as *equalized odds* (Hardt et al., 2016) and some tractable relaxations (Darlington, 1971; Woodworth et al., 2017; Chouldechova, 2017).

The idea of considering features that lead to a consistent predictive distribution has been explored in the causality literature (Peters et al., 2016; Magliacane et al., 2018) for linear models as a feature selection criteria. Other work approached the problem in the non-linear case by considering gradients of the classifier (Arjovsky et al., 2019; Koyama & Yamaguchi, 2020; Parascandolo et al., 2020) or penalizing the variance of error (Krueger et al., 2020; Xie et al., 2020) across different environments. Other relaxations of the sufficiency criterion such as *predictive parity* (Chouldechova, 2017) have been considered in the fairness literature.

Despite the promising directions of research, Gulrajani & Lopez-Paz (2020) have shown that the performance of the most popular approaches in the literature strongly depends on the hyper-parameter tuning strategy, underlying the problem of the lack of (i) standardized benchmark procedures and (ii) clarity regarding guarantees and hidden assumptions for the different algorithms. Although the problem of distribution shift has been widely studied in the classic domain generalization literature (Mansour et al., 2009; Schölkopf et al., 2012; Gong et al., 2016), this work aims to characterize the problem from a representation learning and information theoretical perspective, providing different insights and identifying potential issues and benefits of several popular models.

6 Discussion and Conclusion

In this work, we characterize an information-theoretic framework to analyze the distribution shift problem by relating it to the train and out of distribution test error. We identify two sources of errors for models based on latent representations, and we show that different criteria explored in literature can be seen as different strategies to minimize concept shift in the latent space using extra information about the data selection procedure. We demonstrate both theoretically and empirically that their effectiveness depends on the structure of the underlying graphical model with respect to the observed variables, training data, and data selection criteria other than the chosen approximation and relaxations for optimization.

Although test information loss and latent test error are challenging to estimate for real-world datasets, we argue that the presented analysis can be useful to better understand and mitigate the effect of selection bias and distribution shift in machine learning models. An accurate estimation of the error decomposition reported in this work could be further used to guide processes of causal discovery.

Acknowledgments and Disclosure of Funding

We thank Sindy Löwe and David Ruhe for their insightful comments and feedback. This work was supported by the Microsoft Research PhD Scholarship Programme.

References

Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.

- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *CoRR*, abs/1907.02893, 2019.
- Barocas, S., Hardt, M., and Narayanan, A. *Fairness and Machine Learning*. fairmlbook.org, 2018. <http://www.fairmlbook.org>.
- Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., and Vardoulakis, L. M. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In Bernhaupt, R., Mueller, F. F., Verweij, D., Andres, J., McGrenere, J., Cockburn, A., Avellino, I., Goguy, A., Bjøn, P., Zhao, S., Samson, B. P., and Kocielnik, R. (eds.), *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*, pp. 1–12. ACM, 2020.
- Bengio, Y., Deleu, T., Rahaman, N., Ke, N. R., Lachapelle, S., Bilaniuk, O., Goyal, A., and Pal, C. J. A meta-transfer objective for learning to disentangle causal mechanisms. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- Binette, O. A note on reverse pinsker inequalities. *IEEE Trans. Inf. Theory*, 2019.
- Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In Friedler, S. A. and Wilson, C. (eds.), *Conference on Fairness, Accountability and Transparency, FAT 2018, 23-24 February 2018, New York, NY, USA*, volume 81 of *Proceedings of Machine Learning Research*, pp. 77–91. PMLR, 2018.
- Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2):153–163, 2017.
- Christie, G., Fendley, N., Wilson, J., and Mukherjee, R. Functional map of the world. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 6172–6180. IEEE Computer Society, 2018.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*, pp. 797–806. ACM, 2017.
- Darlington, R. B. Another look at "cultural fairness". *Journal of Educational Measurement*, 8:71–82, 1971. ISSN 00220655, 17453984.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. S. Fairness through awareness. *CoRR*, abs/1104.3913, 2011.
- Endres, D. M. and Schindelin, J. E. A new metric for probability distributions. *IEEE Trans. Inf. Theory*, 49(7):1858–1860, 2003.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. S. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.*, 17: 59:1–59:35, 2016.
- Gilardoni, G. L. On pinsker’s and vajda’s type inequalities for csiszar’s f-divergences. *IEEE Trans. Inf. Theory*, 2010.
- Gong, M., Zhang, K., Liu, T., Tao, D., Glymour, C., and Schölkopf, B. Domain adaptation with conditional transferable components. In Balcan, M. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 2839–2848. JMLR.org, 2016.
- Gulrajani, I. and Lopez-Paz, D. In search of lost domain generalization. *CoRR*, abs/2007.01434, 2020.

- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3315–3323, 2016.
- Huszár, F. Invariant risk minimization: An information theoretic view, 2019. URL <https://www.inference.vc/invariant-risk-minimization/>.
- Ilse, M., Tomczak, J. M., Louizos, C., and Welling, M. DIVA: domain invariant variational autoencoders. In Arbel, T., Ayed, I. B., de Bruijne, M., Descoteaux, M., Lombaert, H., and Pal, C. (eds.), *International Conference on Medical Imaging with Deep Learning, MIDL 2020, 6-8 July 2020, Montréal, QC, Canada*, volume 121 of *Proceedings of Machine Learning Research*, pp. 322–348. PMLR, 2020.
- Johansson, F. D., Sontag, D. A., and Ranganath, R. Support and invertibility in domain-invariant representations. In Chaudhuri, K. and Sugiyama, M. (eds.), *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 527–536. PMLR, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Beery, S., Leskovec, J., Kundaje, A., Pierson, E., Levine, S., Finn, C., and Liang, P. WILDS: A benchmark of in-the-wild distribution shifts. *CoRR*, 2020.
- Koyama, M. and Yamaguchi, S. Out-of-distribution generalization with maximal invariant predictor. *CoRR*, 2020.
- Krueger, D., Caballero, E., Jacobsen, J., Zhang, A., Binas, J., Priol, R. L., and Courville, A. C. Out-of-distribution generalization via risk extrapolation (rex). *CoRR*, 2020.
- Li, H., Pan, S. J., Wang, S., and Kot, A. C. Domain generalization with adversarial feature learning. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 5400–5409. IEEE Computer Society, 2018a.
- Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., and Tao, D. Deep domain generalization via conditional invariant adversarial networks. In Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y. (eds.), *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV*, volume 11219 of *Lecture Notes in Computer Science*, pp. 647–663. Springer, 2018b.
- Louizos, C., Swersky, K., Li, Y., Welling, M., and Zemel, R. S. The variational fair autoencoder. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- Lum, K. and Johndrow, J. E. A statistical framework for fair predictive algorithms. *CoRR*, abs/1610.08077, 2016.
- Magliacane, S., van Ommen, T., Claassen, T., Bongers, S., Versteeg, P., and Mooij, J. M. Domain adaptation by using causal inference to predict invariant conditional distributions. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 10869–10879, 2018.
- Mahajan, D., Tople, S., and Sharma, A. Domain generalization using causal matching. *CoRR*, abs/2006.07500, 2020.
- Mansour, Y., Mohri, M., and Rostamizadeh, A. Domain adaptation: Learning bounds and algorithms. In *COLT 2009 - The 22nd Conference on Learning Theory, Montreal, Quebec, Canada, June 18-21, 2009*, 2009.

- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. A survey on bias and fairness in machine learning. *CoRR*, 2019.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- Moreno-Torres, J. G., Raeder, T., Alaíz-Rodríguez, R., Chawla, N. V., and Herrera, F. A unifying view on dataset shift in classification. *Pattern Recognit.*, 45(1):521–530, 2012.
- Moyer, D., Gao, S., Brekelmans, R., Galstyan, A., and Steeg, G. V. Invariant representations without adversarial training. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 9102–9111, 2018.
- Muandet, K., Balduzzi, D., and Schölkopf, B. Domain generalization via invariant feature representation. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, pp. 10–18. JMLR.org, 2013.
- Parascandolo, G., Neitz, A., Orvieto, A., Gresele, L., and Schölkopf, B. Learning explanations that are hard to vary. *CoRR*, abs/2009.00329, 2020.
- Peters, J., Bühlmann, P., and Meinshausen, N. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pp. 947–1012, 2016.
- Popejoy, A. and Fullerton, S. Genomics is failing on diversity. *Nature*, 538:161–164, 2016.
- Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. D. *Dataset Shift in Machine Learning*. The MIT Press, 2009. ISBN 0262170051.
- Rosenfeld, E., Ravikumar, P., and Risteski, A. The risks of invariant risk minimization. *CoRR*, abs/2010.05761, 2020.
- Santurkar, S., Tsipras, D., and Madry, A. BREEDS: benchmarks for subpopulation shift. *CoRR*, abs/2008.04859, 2020.
- Sason, I. and Verdú, S. f-divergence inequalities. *IEEE Trans. Inf. Theory*, 62(11):5973–6006, 2016.
- Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., and Mooij, J. M. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012.
- Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 2000.
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop, ITW 2015, Jerusalem, Israel, April 26 - May 1, 2015*, pp. 1–5. IEEE, 2015.
- Wang, M. and Deng, W. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- Widmer, G. and Kubat, M. Learning in the presence of concept drift and hidden contexts. *Mach. Learn.*, 23(1):69–101, 1996.
- Woodworth, B. E., Gunasekar, S., Ohannessian, M. I., and Srebro, N. Learning non-discriminatory predictors. In Kale, S. and Shamir, O. (eds.), *Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7-10 July 2017*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1920–1953. PMLR, 2017.

- Xia, M., Field, A., and Tsvetkov, Y. Demoting racial bias in hate speech detection. In Ku, L. and Li, C. (eds.), *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media, SocialNLP@ACL 2020, Online, July 10, 2020*, pp. 7–14. Association for Computational Linguistics, 2020.
- Xie, C., Chen, F., Liu, Y., and Li, Z. Risk variance penalization: From distributional robustness to causality. *CoRR*, abs/2006.07544, 2020.
- Xie, Q., Dai, Z., Du, Y., Hovy, E. H., and Neubig, G. Controllable invariance through adversarial feature learning. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 585–596, 2017.
- Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., and Oermann, E. K. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLOS Medicine*, 2018.
- Zhang, J., Sheng, V. S., Li, Q., Wu, J., and Wu, X. Consensus algorithms for biased labeling in crowdsourcing. *Inf. Sci.*, 382-383:254–273, 2017.
- Zhao, H., des Combes, R. T., Zhang, K., and Gordon, G. J. On learning invariant representations for domain adaptation. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7523–7532. PMLR, 2019.

A Notation

For a given joint distribution $p(\mathbf{x}, \mathbf{y}, \mathbf{e})$ on \mathbf{x} , \mathbf{y} and \mathbf{e} , a binary selection $t = f(\mathbf{x}, \mathbf{y}, \mathbf{e}, \epsilon)$, and a conditional distribution $q(\mathbf{y}|\mathbf{x})$ we use the following notation

A.1 Conditional Kullback-Leibler divergence

$$D_{\text{KL}}(p(\mathbf{y}|\mathbf{x})||q(\mathbf{y}|\mathbf{x})) := \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p(\mathbf{x}, \mathbf{y})} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})}{q(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})} \right]. \quad (14)$$

Note that the expectation is always considered with respect to the joint distribution for the first argument of the KL-divergence. The notation $D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}||q_{\mathbf{y}|\mathbf{x}})$ will be used to abbreviate the same quantity.

A.2 Jensen-Shannon divergence

$$D_{\text{JSD}}(p(\mathbf{x})||q(\mathbf{x})) := \frac{1}{2} D_{\text{KL}}(p(\mathbf{x})||m(\mathbf{x})) + \frac{1}{2} D_{\text{KL}}(q(\mathbf{x})||m(\mathbf{x})), \quad (15)$$

with $m(\mathbf{x}) = 1/2 p(\mathbf{x}) + 1/2 q(\mathbf{x})$.

A.3 Mutual Information

$$I(\mathbf{x}; \mathbf{y}) := D_{\text{KL}}(p(\mathbf{x}, \mathbf{y})||p(\mathbf{x})p(\mathbf{y})) \quad (16)$$

$$= D_{\text{KL}}(p(\mathbf{y}|\mathbf{x})||p(\mathbf{y})) \quad (17)$$

$$= D_{\text{KL}}(p(\mathbf{x}|\mathbf{y})||p(\mathbf{x})). \quad (18)$$

The subscript $t = 1$ is used to indicate that both joint and marginal distribution are conditioned on $t = 1$:

$$I_{t=1}(\mathbf{x}; \mathbf{y}) := D_{\text{KL}}(p(\mathbf{x}, \mathbf{y}|t = 1)||p(\mathbf{x}|t = 1)p(\mathbf{y}|t = 1)) \quad (19)$$

Conditional mutual information is defined as:

$$I(\mathbf{x}; \mathbf{y}|\mathbf{e}) := D_{\text{KL}}(p(\mathbf{x}, \mathbf{y}|\mathbf{e})||p(\mathbf{x}|\mathbf{e})p(\mathbf{y}|\mathbf{e})) \quad (20)$$

$$:= D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, \mathbf{e})||p(\mathbf{y}|\mathbf{e})) \quad (21)$$

$$:= D_{\text{KL}}(p(\mathbf{x}|\mathbf{y}, \mathbf{e})||p(\mathbf{x}|\mathbf{e})). \quad (22)$$

Note that all the mutual information terms in this work are implicitly expressed in terms of the distribution p .

B Proofs and Additional details

The proofs for the statement in the main text are reported in this section.

B.1 Proof for Proposition 1

Proof. Consider $D_{\text{KL}}(p(\mathbf{y}|\mathbf{x})||q(\mathbf{y}|\mathbf{x}))$. By multiplying and dividing by $p(\mathbf{y}|\mathbf{x}, t)$, we obtain:

$$\begin{aligned} D_{\text{KL}}(p(\mathbf{y}|\mathbf{x})||q(\mathbf{y}|\mathbf{x})) &= \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p(\mathbf{x}, \mathbf{y})} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})}{q(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})} \right] \\ &= \mathbb{E}_{\mathbf{x}, \mathbf{y}, t \sim p(\mathbf{x}, \mathbf{y}, t)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t = t)}{q(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})} \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})}{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t = t)} \right] \\ &= D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t)||q(\mathbf{y}|\mathbf{x})) - I(\mathbf{y}; t|\mathbf{x}). \end{aligned} \quad (23)$$

Splitting the expectation on t , we express the first term as:

$$\begin{aligned}
D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t)||q(\mathbf{y}|\mathbf{x})) &= p(t=1)\mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p(\mathbf{x}, \mathbf{y}|t=1)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t=1)}{q(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})} \right] \\
&\quad + p(t=0)\mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p(\mathbf{x}, \mathbf{y}|t=0)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t=0)}{q(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x})} \right] \\
&= p(t=1)D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t=1)||q(\mathbf{y}|\mathbf{x})) + p(t=0)D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t=0)||q(\mathbf{y}|\mathbf{x})) \\
&= p(t=1)D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{x}}) + p(t=0)D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0}||q_{\mathbf{y}|\mathbf{x}}). \tag{24}
\end{aligned}$$

Since the KL-divergence between two distributions is always positive, using the result from equations 23 and 24, we have:

$$\begin{aligned}
I(\mathbf{y}; t|\mathbf{x}) &\leq D_{\text{KL}}(p(\mathbf{y}|\mathbf{x}, t)||q(\mathbf{y}|\mathbf{x})) \\
&= p(t=1)D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{x}}) + p(t=0)D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0}||q_{\mathbf{y}|\mathbf{x}}) \\
&\leq (1 - \alpha) \left(D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{x}}) + D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0}||q_{\mathbf{y}|\mathbf{x}}) \right), \tag{25}
\end{aligned}$$

with $\alpha := \min\{p(t=0), p(t=1)\}$.

Whenever $\alpha = 1$, the inequality in equation 25 holds since $I(\mathbf{y}; t|\mathbf{x}) = 0$ and KL-divergence is always positive. For $\alpha < 1$ we can divide both sides by $1 - \alpha$, obtaining:

$$D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{x}}) + D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0}||q_{\mathbf{y}|\mathbf{x}}) \geq \frac{1}{1 - \alpha} I(\mathbf{y}; t|\mathbf{x})$$

□

B.2 Proof for Proposition 2

Proof. Using Jensen's inequality we have:

$$\begin{aligned}
D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{x}}) &= \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim (\mathbf{x}, \mathbf{y}|t=1)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t=1)}{\mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x}=\mathbf{x})}[q(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})]} \right] \\
&\leq^* \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim (\mathbf{x}, \mathbf{y}|t=1)} \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x}=\mathbf{x})} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t=1)}{q(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \right] \\
&= D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{z}}), \tag{26}
\end{aligned}$$

in which * holds with equality when $q(\mathbf{z}|\mathbf{x})$ is a delta distribution (deterministic encoder).

Secondly, by multiplying and dividing by $p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, t=1)$:

$$\begin{aligned}
D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1}||q_{\mathbf{y}|\mathbf{z}}) &= \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim (\mathbf{x}, \mathbf{y}|t=1)} \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x}=\mathbf{x})} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{x} = \mathbf{x}, t=1)}{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, t=1)} \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, t=1)}{q(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \right] \\
&= I(\mathbf{y}; \mathbf{x}|\mathbf{z}) + D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1}||q_{\mathbf{y}|\mathbf{z}}). \tag{27}
\end{aligned}$$

The bound for $D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=0}||q(\mathbf{y}|\mathbf{x}))$ is obtained analogously.

□

B.3 Proof for Proposition 3

Proof. We first define the average latent predictive distribution $\bar{p}(\mathbf{y}|\mathbf{z}) := \frac{1}{2}p(\mathbf{y}|\mathbf{z}, t = 0) + \frac{1}{2}p(\mathbf{y}|\mathbf{z}, t = 1)$, then

$$\begin{aligned}
I(\mathbf{y}; \mathbf{t}|\mathbf{z}) &= D_{\text{KL}}(p(\mathbf{y}|\mathbf{z}, t) \| p(\mathbf{y}|\mathbf{z})) \\
&= p(t = 0)D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}}) + p(t = 1)D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1} \| p_{\mathbf{y}|\mathbf{z}}) \\
&\geq 2\alpha \left(\frac{1}{2}D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}}) + \frac{1}{2}D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1} \| p_{\mathbf{y}|\mathbf{z}}) \right) \\
&= 2\alpha \mathbb{E}_{t \sim \text{Ber}(1/2)} \mathbb{E}_{\mathbf{z}, \mathbf{y} \sim p(\mathbf{z}, \mathbf{y}|t=t)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, t = t)}{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \right] \\
&= 2\alpha \mathbb{E}_{t \sim \text{Ber}(1/2)} \mathbb{E}_{\mathbf{z}, \mathbf{y} \sim p(\mathbf{z}, \mathbf{y}|t=t)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, t = t) \bar{p}(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})}{\bar{p}(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}) p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \right] \\
&= 2\alpha D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}}^{t=1}) + D_{\text{KL}}(\bar{p}(\mathbf{y}|\mathbf{z}) \| p(\mathbf{y}|\mathbf{z})) \\
&\geq 2\alpha D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}}^{t=1}), \tag{28}
\end{aligned}$$

where $\text{Ber}(1/2)$ refers to a Bernoulli distribution with parameter $p = 1/2$.

Secondly, since the square root of the Jensen-Shannon divergence is a metric (Endres & Schindelin, 2003), using triangle inequality we have:

$$\sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})} + \sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1})} \geq \sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})}. \tag{29}$$

Note that the conditional distributions $p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0}$ and $p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1}$ are defined only for latent vectors \mathbf{z} that have positive probability according to $p(\mathbf{z}|\mathbf{t} = 0)$ and $p(\mathbf{z}|\mathbf{t} = 1)$ respectively. Therefore inequality is defined only for \mathbf{z} that have simultaneously strictly positive probability on both train and test distributions.

Lastly, The Jensen-Shannon divergence can be upper-bounded by a function of the corresponding Kullback-Leibler divergence:

$$D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}) = \frac{1}{2}D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}) - D_{\text{KL}}(m_{\mathbf{y}|\mathbf{z}} \| q_{\mathbf{y}|\mathbf{z}}) \leq \frac{1}{2}D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}), \tag{30}$$

with $m_{\mathbf{y}|\mathbf{z}} := 1/2p(\mathbf{y}|\mathbf{z}, t = 1) + 1/2q(\mathbf{y}|\mathbf{z})$.

Using the results from 28, 29 and 30 we have

$$\begin{aligned}
D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) &\stackrel{(29)}{\leq} \left(\sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})} + \sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1})} \right)^2 \\
&\stackrel{(28)}{\leq} \left(\sqrt{\frac{1}{2\alpha} I(\mathbf{y}; \mathbf{t}|\mathbf{z} = \mathbf{z})} + \sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1})} \right)^2 \\
&\stackrel{(30)}{\leq} \left(\sqrt{\frac{1}{2\alpha} I(\mathbf{y}; \mathbf{t}|\mathbf{z} = \mathbf{z})} + \sqrt{\frac{1}{2} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})} \right)^2 \tag{31}
\end{aligned}$$

□

B.4 Upper-bounding the test error

We show two inequalities, that relate the Jensen-Shannon divergence in equation 31 to the latent test error (defined as a Kullback-Leibler divergence).

B.4.1 Ratio of f -divergences

Since both Jensen-Shannon and Kullback Leibler divergence are instances of f -divergences, we can use the result from theorem 1 in Sason & Verdú (2016) to write the following inequality:

$$D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) \geq K(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0}, q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}), \quad (32)$$

with

$$K(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0}, q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) = \sup_{\mathbf{y}} k \left(\frac{p(\mathbf{y} = \mathbf{y} | \mathbf{z} = \mathbf{z}, t = 0)}{q(\mathbf{y} = \mathbf{y} | \mathbf{z} = \mathbf{z})} \right),$$

in which

$$k(r) = \begin{cases} \frac{f(r)}{g(r)} & \text{for } r \in (0, 1) \cup (1, +\infty) \\ 1 & \text{for } r = 1 \end{cases},$$

and

$$\begin{aligned} f(r) &= r \log r \\ g(r) &= \frac{1}{2} \left[r \log r - (r + 1) \log \frac{r + 1}{2} \right] \end{aligned}$$

B.4.2 Pinsker-type bound

First, since the Jensen-Shannon divergence is a f -divergence, we can use the result from theorem 3 in Gilardoni (2010) to express a lower-bound as a function of the total variation D_{TV} :

$$\begin{aligned} \sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})} &\geq \sqrt{\frac{f''(1)}{2}} D_{TV}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) \\ &= \sqrt{2} D_{TV}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}), \end{aligned} \quad (33)$$

with $f''(1)$ as the second derivative of $f(t) = \frac{1}{2} \left[(t + 1) \log \frac{2}{t+1} + t \log t \right]$ evaluated in $t = 1$.

Secondly, since the Kullback-Leibler divergence is also a f -divergence, we use the result of theorem 1 from Binette (2019) to upper-bound the latent test error as a function of the total variation:

$$\begin{aligned} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) &\leq \left(\frac{f(m(\mathbf{z}))}{1 - m(\mathbf{z})} + \frac{f(M(\mathbf{z}))}{M(\mathbf{z}) - 1} \right) D_{TV}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) \\ &= \left(\frac{m(\mathbf{z}) \log m(\mathbf{z})}{1 - m(\mathbf{z})} + \frac{M(\mathbf{z}) \log M(\mathbf{z})}{M(\mathbf{z}) - 1} \right) D_{TV}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}), \end{aligned} \quad (34)$$

with $m(\mathbf{z}) := \min_{\mathbf{y}} \frac{p(\mathbf{y}=\mathbf{y}|\mathbf{z}=\mathbf{z}, t=0)}{q(\mathbf{y}=\mathbf{y}|\mathbf{z}=\mathbf{z})}$, $M(\mathbf{z}) := \max_{\mathbf{y}} \frac{p(\mathbf{y}=\mathbf{y}|\mathbf{z}=\mathbf{z}, t=0)}{q(\mathbf{y}=\mathbf{y}|\mathbf{z}=\mathbf{z})}$, and $f(t) = t \log t$.

Combining the two bounds we obtain:

$$D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) \leq \frac{1}{\sqrt{2}} \left(\frac{m(\mathbf{z}) \log m(\mathbf{z})}{1 - m(\mathbf{z})} + \frac{M(\mathbf{z}) \log M(\mathbf{z})}{M(\mathbf{z}) - 1} \right) \sqrt{D_{\text{JSD}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{t=0} \| q_{\mathbf{y}|\mathbf{z}=\mathbf{z}})}. \quad (35)$$

Note that the right-hand side of this last inequality is finite whenever the support of the model $q(\mathbf{y}|\mathbf{z} = \mathbf{z})$ contains the support of $p(\mathbf{y}|\mathbf{z} = \mathbf{z}, t = 0)$ (and therefore $M(\mathbf{z})$ is finite). When this does not happen, both left-hand and right hand side of the inequality are unbounded.

B.5 Bound in expression 9

Proof. Consider the amount of predictive information lost by encoding \mathbf{x} into \mathbf{z} :

$$\begin{aligned} I_{t=1}(\mathbf{x}; \mathbf{y} | \mathbf{z}) &= I_{t=1}(\mathbf{z}\mathbf{x}; \mathbf{y}) - I_{t=1}(\mathbf{z}; \mathbf{y}) \\ &= I_{t=1}(\mathbf{x}; \mathbf{y}) + I_{t=1}(\mathbf{z}; \mathbf{y} | \mathbf{x}) - I_{t=1}(\mathbf{z}; \mathbf{y}) \\ &= I_{t=1}(\mathbf{x}; \mathbf{y}) - I_{t=1}(\mathbf{z}; \mathbf{y}), \end{aligned} \quad (36)$$

in which $I_{t=1}(\mathbf{z}; \mathbf{y} | \mathbf{x}) = 0$ since \mathbf{z} depends only on \mathbf{x} .

Since $I(\mathbf{x}; \mathbf{y})$ is constant in $q(\mathbf{z}|\mathbf{x})$:

$$\min_{q(\mathbf{z}|\mathbf{x})} I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z}) = I(\mathbf{x}; \mathbf{y}) - \max_{q(\mathbf{z}|\mathbf{x})} I_{t=1}(\mathbf{y}; \mathbf{z}). \quad (37)$$

The upper-bound is derived considering $q(\mathbf{y}|\mathbf{z})$ as a variational distribution. For any $q(\mathbf{y}|\mathbf{z})$:

$$\begin{aligned} I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z}) &= D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} \| p_{\mathbf{y}|\mathbf{z}}^{t=1}) \\ &= D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}) - D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}) \\ &\leq D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}), \end{aligned} \quad (38)$$

and in particular

$$I_{t=1}(\mathbf{x}; \mathbf{y}|\mathbf{z}) \leq \min_{q(\mathbf{y}|\mathbf{z})} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{x}}^{t=1} \| q_{\mathbf{y}|\mathbf{z}}), \quad (39)$$

in which equality holds for $q(\mathbf{y}|\mathbf{z}) = p(\mathbf{y}|\mathbf{z}, t = 1)$. \square

B.6 Proof for Proposition 4

Proof. Consider the latent concept shift $I(\mathbf{y}; t|\mathbf{z})$. Using the chain rule of mutual information, we obtain:

$$\begin{aligned} I(\mathbf{y}; t|\mathbf{z}) &= I(\mathbf{y}; \mathbf{e}t|\mathbf{z}) - I(\mathbf{y}; \mathbf{e}|\mathbf{z}t) \\ &\leq I(\mathbf{y}; \mathbf{e}t|\mathbf{z}) \\ &= I(\mathbf{y}; \mathbf{e}|\mathbf{z}) + I(\mathbf{y}; t|\mathbf{z}\mathbf{e}). \end{aligned} \quad (40)$$

Whenever t can be expressed as a function of \mathbf{e} and some independent noise ϵ , we have:

$$\begin{aligned} t &= f(\mathbf{e}, \epsilon) \\ \implies p(t|\mathbf{e}) &= p(t|\mathbf{e}, \mathbf{z}) = p(t|\mathbf{e}, \mathbf{z}, \mathbf{y}) \\ \implies D_{\text{KL}}(p(t|\mathbf{e}, \mathbf{z}) \| p(t|\mathbf{e}, \mathbf{z}, \mathbf{y})) &= 0 \\ \iff I(\mathbf{y}; t|\mathbf{z}\mathbf{e}) &= 0. \end{aligned} \quad (41)$$

Therefore, we conclude:

$$I(\mathbf{y}; t|\mathbf{z}) \leq I(\mathbf{y}; \mathbf{e}|\mathbf{z}) \quad (42)$$

\square

B.7 Proof for Proposition 5

Proof. Consider the latent concept shift $I(\mathbf{y}; t|\mathbf{z})$. Using the chain rule of mutual information, we obtain:

$$\begin{aligned} I(\mathbf{y}; t|\mathbf{z}) &= I(\mathbf{y}\mathbf{z}; t) - I(t; \mathbf{z}) \\ &\leq I(\mathbf{y}\mathbf{z}; t) \\ &= I(\mathbf{y}; t) + I(t; \mathbf{z}|\mathbf{y}) \end{aligned} \quad (43)$$

Then, we express $I(t; \mathbf{z}|\mathbf{y})$ as a function of $I(\mathbf{e}; \mathbf{z}|\mathbf{y})$:

$$\begin{aligned} I(t; \mathbf{z}|\mathbf{y}) &= I(\mathbf{e}t; \mathbf{z}|\mathbf{y}) - I(\mathbf{e}; \mathbf{z}|\mathbf{y}t) \\ &\leq I(\mathbf{e}t; \mathbf{z}|\mathbf{y}) \\ &= I(\mathbf{e}; \mathbf{z}|\mathbf{y}) + I(t; \mathbf{z}|\mathbf{y}\mathbf{e}). \end{aligned} \quad (44)$$

Whenever t can be expressed as a function of \mathbf{e} , \mathbf{y} and some independent noise ϵ , we have:

$$\begin{aligned} \exists f : t &= f(\mathbf{e}, \mathbf{y}, \epsilon) \\ \implies p(t|\mathbf{e}, \mathbf{y}) &= p(t|\mathbf{e}, \mathbf{y}, \mathbf{z}) \\ \implies D_{\text{KL}}(p(t|\mathbf{e}, \mathbf{y}, \mathbf{z}) \| p(t|\mathbf{e}, \mathbf{y})) &= 0 \\ \iff I(t; \mathbf{z}|\mathbf{y}\mathbf{e}) &= 0. \end{aligned} \quad (45)$$

Using the results from equations 43, 44 and 45:

$$\begin{aligned} I(\mathbf{y}; t|\mathbf{z}) &\leq I(\mathbf{y}; t) + I(t; \mathbf{z}|\mathbf{y}) \\ &\leq I(\mathbf{y}; t) + I(\mathbf{e}; \mathbf{z}|\mathbf{y}). \end{aligned} \quad (46)$$

\square

B.8 Correcting for prior shift

When the separation constraint $I(\mathbf{e}; \mathbf{z}|\mathbf{y}) = 0$ is enforced the *reverse concept shift* $I(\mathbf{t}; \mathbf{z}|\mathbf{y})$, which represents how much the distribution $p(\mathbf{z}|\mathbf{y})$ changes as a result of the selection, is zero (equations 44 and 45). Considering the stability of $p(\mathbf{z}|\mathbf{y})$, we can express the test predictive distribution $p(\mathbf{y}|\mathbf{z}, \mathbf{t} = 0)$ as a function of the train one using Bayes rule:

$$\begin{aligned}
p(\mathbf{y}|\mathbf{z}, \mathbf{t} = 0) &= \frac{p(\mathbf{z}|\mathbf{y}, \mathbf{t} = 0)p(\mathbf{y}|\mathbf{t} = 0)}{p(\mathbf{z}|\mathbf{t} = 0)} \\
&\stackrel{*}{=} \frac{p(\mathbf{z}|\mathbf{y}, \mathbf{t} = 1)p(\mathbf{y}|\mathbf{t} = 0)}{p(\mathbf{z}|\mathbf{t} = 0)} \\
&= \frac{\frac{p(\mathbf{y}|\mathbf{z}, \mathbf{t} = 1)p(\mathbf{z}|\mathbf{t} = 1)}{p(\mathbf{y}|\mathbf{t} = 1)}p(\mathbf{y}|\mathbf{t} = 0)}{p(\mathbf{z}|\mathbf{t} = 0)} \\
&= p(\mathbf{y}|\mathbf{z}, \mathbf{t} = 1) \underbrace{\frac{p(\mathbf{y}|\mathbf{t} = 0)}{p(\mathbf{y}|\mathbf{t} = 1)}}_{r(\mathbf{y})} \underbrace{\frac{p(\mathbf{z}|\mathbf{t} = 1)}{p(\mathbf{z}|\mathbf{t} = 0)}}_{1/Z}, \tag{47}
\end{aligned}$$

where * uses $p(\mathbf{z}|\mathbf{y}, \mathbf{t} = 0) = p(\mathbf{z}|\mathbf{y}, \mathbf{t} = 1)$. The ratio $r(\mathbf{y})$ represents the fraction of the marginal probabilities of \mathbf{y} , while Z is a normalization constant.

We define the corrected model $q'(\mathbf{y}|\mathbf{z})$ following the result from equation 47:

$$\hat{q}'(\mathbf{y}|\mathbf{z}) := \frac{1}{Z}q(\mathbf{y}|\mathbf{z})r(\mathbf{y}). \tag{48}$$

The test error for the model $q'(\mathbf{y}|\mathbf{z})$ for a representation \mathbf{z} that has positive support in both train and test ($p(\mathbf{z}|\mathbf{t} = 1) > 0, p(\mathbf{z}|\mathbf{t} = 0) > 0$) is determined by the train latent test error:

$$\begin{aligned}
D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}=\mathbf{z}}^{\mathbf{t}=0} || q'_{\mathbf{y}|\mathbf{z}=\mathbf{z}}) &= \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{z}=\mathbf{z}, \mathbf{t}=0)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, \mathbf{t} = 0)}{q'(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \right] \\
&= \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{z}=\mathbf{z}, \mathbf{t}=0)} \left[\log \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, \mathbf{t} = 0)}{q(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{t} = 1)}{p(\mathbf{y} = \mathbf{y}|\mathbf{t} = 0)} \frac{p(\mathbf{z} = \mathbf{z}|\mathbf{t} = 0)}{p(\mathbf{z} = \mathbf{z}|\mathbf{t} = 1)} \right] \\
&= \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{z}=\mathbf{z}, \mathbf{t}=0)} \left[\log \frac{\frac{p(\mathbf{y}=\mathbf{y}|\mathbf{z}=\mathbf{z}, \mathbf{t}=0)p(\mathbf{z}=\mathbf{z}|\mathbf{t}=0)}{p(\mathbf{y}=\mathbf{y}|\mathbf{t}=0)}}{\frac{p(\mathbf{y}=\mathbf{y}|\mathbf{z}=\mathbf{z}, \mathbf{t}=1)p(\mathbf{z}=\mathbf{z}|\mathbf{t}=1)}{p(\mathbf{y}=\mathbf{y}|\mathbf{t}=1)}} \frac{p(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z}, \mathbf{t} = 1)}{q(\mathbf{y} = \mathbf{y}|\mathbf{z} = \mathbf{z})} \right] \\
&= \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{z}=\mathbf{z}, \mathbf{t}=0)} \left[\log \frac{p(\mathbf{z} = \mathbf{z}|\mathbf{y} = \mathbf{y}, \mathbf{t} = 0)}{p(\mathbf{z} = \mathbf{z}|\mathbf{y} = \mathbf{y}, \mathbf{t} = 1)} \right] + D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{\mathbf{t}=1} || q_{\mathbf{y}|\mathbf{z}}) \\
&\stackrel{*}{=} D_{\text{KL}}(p_{\mathbf{y}|\mathbf{z}}^{\mathbf{t}=1} || q_{\mathbf{y}|\mathbf{z}}), \tag{49}
\end{aligned}$$

where * uses $p(\mathbf{z}|\mathbf{y}, \mathbf{t} = 0) = p(\mathbf{z}|\mathbf{y}, \mathbf{t} = 1)$.

B.9 Relation between the different criteria

The different criteria mentioned in section 3 are related by the following expression, which follows from the chain rule of mutual information:

$$\underbrace{I(\mathbf{e}; \mathbf{z}|\mathbf{y})}_{\text{separation}} + I(\mathbf{e}; \mathbf{y}) = \underbrace{I(\mathbf{e}; \mathbf{y}|\mathbf{z})}_{\text{sufficiency}} + \underbrace{I(\mathbf{e}; \mathbf{z})}_{\text{independence}}$$

From which we can derive the following:

1. if \mathbf{y} and \mathbf{e} are independent ($I(\mathbf{y}; \mathbf{e}) = 0$) enforcing separation is equivalent to enforcing both sufficiency and independence:

$$\begin{aligned}
I(\mathbf{e}; \mathbf{z}|\mathbf{y}) + I(\mathbf{e}; \mathbf{y}) &= 0 \\
\iff I(\mathbf{y}; \mathbf{e}|\mathbf{z}) + I(\mathbf{e}; \mathbf{z}) &= 0 \tag{50}
\end{aligned}$$

2. if \mathbf{y} and \mathbf{e} are dependent ($I(\mathbf{e}; \mathbf{y}) > 0$), sufficiency and independence are mutually exclusive conditions (Barocas et al., 2018):

$$I(\mathbf{e}; \mathbf{y}|\mathbf{z}) + I(\mathbf{e}; \mathbf{z}) \geq I(\mathbf{e}; \mathbf{y}) \tag{51}$$

C Datasets

Here we report the conditional probability distributions used to produce the d-CMNIST and y-CMNIST datasets, underlining the commonalities and differences from the original CMNIST distribution.

Across the three datasets:

- The first two environments are selected for training.

$$\begin{array}{c|ccc}
 & \multicolumn{3}{c}{p(t = t|e = e)} \\
 & e = 0 & e = 1 & e = 2 \\
 \hline
 t = 1 & 1 & 1 & 0 \\
 t = 0 & 0 & 0 & 1
 \end{array}$$

- The correlation between color c and label y is positive in the first two environments and negative on the last.

$$\begin{array}{c|ccc}
 & \multicolumn{3}{c}{p(c = 1|y = y, e = e)} \\
 & e = 0 & e = 1 & e = 2 \\
 \hline
 y = 0 & 9/10 & 4/5 & 1/10 \\
 y = 1 & 1/10 & 1/5 & 9/10
 \end{array}$$

- For a given color c and digit d , the corresponding picture is a MNIST digit with label d and red ($c = 1$) or green ($c = 0$) color. In the discrete settings we assume that pictures contain only color and digit information (ignoring the style) and we construct \mathbf{x} as a concatenation of the color c and digit d : $\mathbf{x} := [c, d]$.

C.1 CMNIST

- All 10 digits d the same probability:

$$\forall d \in 0, \dots, 9 : p(d = d) = 1/10$$

- The 3 environments have the same probability to occur:

$$\forall e \in 0, 1, 2 : p(e = e) = 1/3$$

Note that the first two environments will have probability 0.5 on the selected training set, while the last one will be drawn with probability 1 from the test distribution.

- The label $y = 1$ is assigned with probability 0.75 for digits 5-9 and 0.25 for digits 0-4 across all environments.

$$\begin{array}{c|cc}
 & \multicolumn{2}{c}{p(y = y|d = d)} \\
 & d < 5 & d \geq 5 \\
 \hline
 y = 0 & 3/4 & 1/4 \\
 y = 1 & 1/4 & 3/4
 \end{array}$$

C.2 d-CMNIST

- Digits from 0-4 are more likely to occur on the first environment, less likely on the second, while the probability is uniform on the third environment.

$$\begin{array}{c|ccc}
 & \multicolumn{3}{c}{p(d \in D|e = e)} \\
 & e = 0 & e = 1 & e = 2 \\
 \hline
 D = [0, 4] & 3/5 & 1/5 & 1/2 \\
 D = [5, 9] & 2/5 & 4/5 & 1/2
 \end{array}$$

Digits from 0 to 4 (and from 5 to 9) have the same probability:

$$\forall d \in [0, 4] : p(\mathbf{d} = d) = \frac{\sum_{d' \in [0,4]} p(\mathbf{d} = d')}{5}$$

$$\forall d \in [5, 9] : p(\mathbf{d} = d) = \frac{\sum_{d' \in [5,9]} p(\mathbf{d} = d')}{5}$$

- The first environment is more likely than the second:

$$p(\mathbf{e} = 0) = 1/2$$

$$p(\mathbf{e} = 1) = 1/6$$

$$p(\mathbf{e} = 2) = 1/3$$

Note that this assignments is designed to ensure that marginally the digits d have uniform distribution on both train and test.

- Labels y are assigned in the same way as the original CMNIST distribution, based on the digits d .

C.3 y-CMNIST

- Digits 0-4 are more likely to occur for $y = 0$ and less likely for $y = 1$:

$$\begin{array}{c|cc} p(\mathbf{d} \in D | y = y) & y = 0 & y = 1 \\ \hline D = [0, 4] & 3/4 & 1/4 \\ D = [5, 9] & 1/4 & 3/4 \end{array}$$

Similarly to d-CMNIST, digits in the same group ($[0,4]$, $[5,9]$) have equal likelihood:

$$\forall d \in [0, 4] : p(\mathbf{d} = d) = \frac{\sum_{d' \in [0,4]} p(\mathbf{d} = d')}{5}$$

$$\forall d \in [5, 9] : p(\mathbf{d} = d) = \frac{\sum_{d' \in [5,9]} p(\mathbf{d} = d')}{5}$$

Note that the conditional distribution $p(\mathbf{d}|y)$ is the same as the one of the original CMNIST distribution.

- The label $y = 0$ is more likely in the first environment and less on the second. The two labels have the same probability on the third environment:

$$\begin{array}{c|ccc} p(y = y | \mathbf{e} = e) & e = 0 & e = 1 & e = 2 \\ \hline y = 0 & 3/5 & 1/5 & 1/2 \\ y = 1 & 2/5 & 4/5 & 1/2 \end{array}$$

- The marginal environment distribution is the same as the one described in the d-CMNIST dataset. This assignment ensures that both labels y and digits d have marginal uniform distribution on train $t = 1$ and test $t = 0$ splits.

C.4 Dataset properties

Table 1 reports some of the mutual information measurements for the three CMNIST, d-CMNIST, and y-CMNIST datasets. Note that only representations that contain digit or no information result in zero concept shift.

C.5 Sampling

Since all the three dataset have uniform digit distribution on both train and test splits ($p(\mathbf{d} = d | t = 1) = p(\mathbf{d} = d | t = 0) = 1/10$), the training dataset are produced using the following sampling procedure:

dataset	Concept shift				Independence			Sufficiency				Separation		
	$I(y; t \mathbf{x})$	$I(y; t c)$	$I(y; t d)$	$I(y; t)$	$I_{t=1}(e; \mathbf{x})$	$I_{t=1}(e; c)$	$I_{t=1}(e; d)$	$I_{t=1}(y; e)$	$I_{t=1}(y; e \mathbf{x})$	$I_{t=1}(y; e c)$	$I_{t=1}(y; e d)$	$I_{t=1}(e; \mathbf{x} y)$	$I_{t=1}(e; c y)$	$I_{t=1}(e; d y)$
CMNIST	0.219	0.283	0.0	0.0	0.00143	0.0	0.0	0.0	0.00854	0.00997	0.0	0.00997	0.00997	0.0
d-CMNIST	0.238	0.306	0.0	0.0	0.0642	0.000636	0.0633	0.0152	0.00588	0.0164	0.0	0.0549	0.00760	0.0481
y-CMNIST	0.238	0.306	0.0	0.0	0.0331	0.0258	0.0152	0.0633	0.0371	0.0444	0.0481	0.00683	0.00683	0.0

Table 1: Mutual information quantities for the three CMNIST datasets. Note that only a representation containing digit or no information results in zero concept shift. Both color and digit satisfy the independence criterion on CMNIST, while digit information is not satisfying sufficiency or separation for the y-CMNIST and d-CMNIST datasets.

1. Determine $p(c, e, y|d, t = 1)$ using Bayes rule:

$$p(c, e, y|d, t = 1) = \frac{p(c, e, y, d|t = 1)}{p(d|t = 1)}.$$

2. Sample a MNIST picture $\tilde{\mathbf{x}}$ with label d from the MNIST dataset.
3. Sample from $p(c, e, y|d = d, t = 1)$ to determine color c , environment e and label y .
4. Color the picture $\tilde{\mathbf{x}}$ in green or red depending on the value of c to obtain \mathbf{x} .
5. Return the triple (\mathbf{x}, y, e)

Note that in the reported experiments, step 3) is performed at run-time. Therefore, the same picture $\tilde{\mathbf{x}}$ can appear in different environments with different colors if sampled multiple times.

Digits used for the train $t = 1$ and test $t = 0$ splits are disjoint, following the traditional train-test MNIST splits. The code used to generate the three datasets can be found at <https://github.com/mfederici/dsit>.

D Direct Optimization

In order to optimize for the objective prescribed by the different criteria directly, we define a discrete random variable $\hat{\mathbf{x}}$ obtained by concatenating the random variables representing color c and digit d . All the extra information in \mathbf{x} (e.g. style, position, small rotations, ..) does not interact with the environment e or label y , therefore we can safely assert that \mathbf{x} is conditionally independent of y and e once $\hat{\mathbf{x}}$ is observed:

$$I(\mathbf{x}; ye|\hat{\mathbf{x}}) = 0. \quad (52)$$

At the same time, we can safely assume that color and digit can be identified when a picture \mathbf{x} is observed:

$$\exists g \text{ s.t. } \hat{\mathbf{x}} = g(\mathbf{x}) \implies I(\hat{\mathbf{x}}; ye|\mathbf{x}) = 0. \quad (53)$$

Using the assumptions in 52 and 53, we can infer that $\hat{\mathbf{x}}$ is a sufficient statistic for \mathbf{x} . In particular, we can show that any (conditional or not) mutual information measurement involving \mathbf{x} , y and e is unchanged if $\hat{\mathbf{x}}$ is used instead of \mathbf{x} .

For example:

$$\begin{aligned}
I(\mathbf{y}; \mathbf{e}|\mathbf{x}) &= I(\mathbf{xy}; \mathbf{e}) - I(\mathbf{x}; \mathbf{y}) \\
&= I(\mathbf{xy}; \mathbf{e}) - I(\mathbf{x}\hat{\mathbf{x}}; \mathbf{y}) + \underbrace{I(\hat{\mathbf{x}}; \mathbf{y}|\mathbf{x})}_{=0 \text{ using eq.53}} \\
&= I(\mathbf{xy}; \mathbf{e}) - I(\hat{\mathbf{x}}; \mathbf{y}) - \underbrace{I(\mathbf{x}; \mathbf{y}|\hat{\mathbf{x}})}_{=0 \text{ using eq.52}} \\
&= I(\mathbf{y}; \mathbf{e}) + I(\mathbf{x}; \mathbf{e}|\mathbf{y}) - I(\hat{\mathbf{x}}; \mathbf{y}) \\
&= I(\mathbf{y}; \mathbf{e}) + I(\mathbf{x}\hat{\mathbf{x}}; \mathbf{e}|\mathbf{y}) - \underbrace{I(\hat{\mathbf{x}}; \mathbf{e}|\mathbf{xy})}_{=0 \text{ using eq.53}} - I(\hat{\mathbf{x}}; \mathbf{y}) \\
&= I(\mathbf{y}; \mathbf{e}) + I(\hat{\mathbf{x}}; \mathbf{e}|\mathbf{y}) + \underbrace{I(\mathbf{x}; \mathbf{e}|\hat{\mathbf{x}}\mathbf{y})}_{=0 \text{ using eq.52}} - I(\hat{\mathbf{x}}; \mathbf{y}) \\
&= I(\hat{\mathbf{x}}\mathbf{y}; \mathbf{e}) - I(\hat{\mathbf{x}}; \mathbf{y}) \\
&= I(\mathbf{y}; \mathbf{e}|\hat{\mathbf{x}}).
\end{aligned}$$

Note that in this proof we also show $I(\mathbf{x}; \mathbf{y}) = I(\hat{\mathbf{x}}; \mathbf{y})$ and $I(\mathbf{x}; \mathbf{e}|\mathbf{y}) = I(\hat{\mathbf{x}}; \mathbf{e}|\mathbf{y})$. The proof for the other reported quantity is analogous.

The dashed trajectories in figure 3 are computed by directly optimizing the objectives reported in section 3 using $\hat{\mathbf{x}}$ instead of \mathbf{x} . We use a $[(10 \times 2) \times 64]$ normalized matrix \mathbf{W} to parametrize the encoder $q(\mathbf{z}|\mathbf{x})$. Each column of \mathbf{W} sums to 1 to represent a valid conditional probability distribution. This matrix maps each combination of digit (10) and color (2) into a representation \mathbf{z} consisting of 64 different options using a dot product:

$$\mathbf{z} = \mathbf{W}^T \text{flatten}(\hat{\mathbf{x}}), \quad (54)$$

in which $\text{flatten}(\cdot)$ represent the operation flattening the 2×10 matrix into a 20-dimensional vector. Since $\hat{\mathbf{x}}$ is a categorical variable, this encoder can represents any possible function from $[2 \times 10]$ to the latent space of size [64]. We expect the family of encoders parametrized by \mathbf{W} to be sufficiently flexible since the dimensionality of the representation is bigger when compared to the number of possible inputs.

All mutual information terms in the objective can be computed and differentiated directly, and we can progressively update \mathbf{W} using stochastic gradient descent until convergence. The matrix \mathbf{W} is initialized randomly and updated using the ADAM optimizer (Kingma & Ba, 2015) with a learning rate of 10^{-3} . The optimization stops when the total variation of train and test error is less than 10^{-4} for at least 1000 iterations.

The results for different values of the hyper-parameter λ are obtained by training a new matrix \mathbf{W} from scratch for each considered value of λ . The parameter λ varies from 0 to 10^6 for models trained using the independence, sufficiency and separation criteria, while a maximum value of 10 is used for the information bottleneck experiments since $\lambda = 2$ is usually sufficient to obtain a constant representation. Further details on the procedure can be found at <https://github.com/mfederici/dsit>.

E Training details

Each model in section 4 is composed of the same encoder $q(\mathbf{z}|\mathbf{x})$ and decoder $q(\mathbf{y}|\mathbf{z})$ neural network consisting of multi-layer perceptrons:

- Encoder:

$$\begin{aligned}
&\text{Flatten()} \rightarrow \text{Linear}(\text{in_size}= 768, \text{out_size}= 1024) \rightarrow \\
&\quad \text{Linear}(\text{in_size}= 1024, \text{out_size}= 128) \rightarrow \\
&\quad \text{Linear}(\text{in_size}= 128, \text{out_size}= 64(\times 2))
\end{aligned}$$

with 0.25 dropout probability and ReLU activation in-between each layer. Note that a final size of $128 = 64 + 64$ is used for the VIB experiments to separately parametrize mean and variance of the stochastic encoder $q(\mathbf{z}|\mathbf{x})$.




Figure 4: Train and test error on the CMNIST, d-CMNIST and y-CMNIST datasets obtained by training a convolutional architecture using the same hyper-parameters described in section E. The results are consistent with the ones obtained with a multi-layer perceptron (MLP), which are reported in figure 3.

- Classifier:

Linear(in_size= 64, out_size= 128) →
 Linear(in_size= 128, out_size= 2)

with ReLU activation and softmax output normalization.

Each model is trained for a total of 25000 iterations with batches of size 256 (with the exception of the IRM and VREx model that use a batch size of 4096) using the ADAM optimizer (Kingma & Ba, 2015) with a learning rate of 10^{-4} . The neural networks are trained for a total of 3000 iterations without any regularization, then λ is linearly scaled by a constant value after each iteration (5×10^{-3} for VREx, 2×10^{-2} for DANN, CDANN and IRM, 5×10^{-5} for VIB). The increment policy has been chosen to maintain stability while exploring the full range of latent representation for each one of the models. The train and test cross-entropy error are measured and reported every 2000 training steps across 3 different seeded runs.

The DANN and CDANN models also require an auxiliary discriminator architecture. In both cases we use multi-layer perceptrons:

Linear(in_size= 64(+2), out_size= 1024) →
 Linear(in_size= 1024, out_size= 128) →
 Linear(in_size= 128, out_size= 2)

with ReLU activations, softmax output and spectral normalization (Miyato et al., 2018) for additional stability. The +2 in parenthesis accounts for the extra dimensions that are required for the label concatenation in the CDANN model. In between each update of the encoder and classifier neural network, the discriminator model is trained for a total of 64 iterations with a learning rate of 10^{-4} . This procedure is required to ensure stability of the adversarial training. The entirety of the experiments reported in this work required approximately a total of 250 compute hours on Nvidia 1080 GPUs. The code used to train the architectures is publicly available at <https://www.github.com/mfederici/dsit>.

E.1 CNN results

In order to validate the result obtained with the different regularization strategies, we perform some additional experiments using a convolutional encoder architecture:

Conv(in_channels= 2, out_channels= 32, kernel_size= 5, stride= 2)→
 Conv(in_channels= 32, out_channels= 64, kernel_size= 5, stride= 2)→
 Conv(in_channels= 64, out_channels= 128, kernel_size= 4, stride= 2)→
 Flatten()→ Linear(in_size= 128, out_size= 64)




Figure 5: Measurements of train and test cross-entropy for varying number of discriminator training iterations (columns) on the three CMNIST variations (rows) on three different runs. The model convergence and stability is strongly affected by the number of discriminator iterations, especially on the d-CMNIST and y-CMNIST datasets.

with ReLU activations and dropout.

The corresponding results reported in Figure 4 for the three CMNIST variations are consistent with the ones reported in figure 3, underlying that the observed behaviour depends on the model objectives more than the chosen architectures.

F Hyper-parameters and their effect

F.1 On the number of adversarial steps

The results reported in figure 5 clearly show that the effectiveness of the DANN and CDANN model strongly depends on the number of adversarial iterations used for each training step of the encoder model. The plots in the figure 5 display that an increased number of adversarial iterations helps with both stability and performance of the model. This indicates that the optimality of the adversarial predictors $q(e|z)$ and $q(e|z, y)$ plays a fundamental role in the training procedure of the model.

Note that the training cost and time scales up linearly with the number of adversarial iterations. This becomes especially problematic with bigger models and datasets.

F.2 On the size of the training batch

We observed that both the IRM and VREx models are sensitive to the choice of the size of the training batch used for optimization. We hypothesize that this is due to the fact that both method rely on the computation of a higher order batch statistic (variance and gradient norm).

Figure 6 reports train and test cross-entropy obtained for different values of batch size. We observe that increasing batch size is required to stabilize the training procedure and improve generalization results at the cost of increased memory and computational requirements.

F.3 Practical limitations of Invariant Risk Minimization

We observed that the IRM model is unable to minimize the test error on the d-CMNIST dataset despite the optimality of the sufficiency criterion. We identify two main reasons that can explain this issue:




Figure 6: Visualization of the effect of different choices of batch size (different columns) for the IRM and VREx models on the three CMNIST variations (rows) using the MLP architectures described in appendix E. Increasing the batch size generally results in better stability and performance at the cost of extra computation time. The models in the pictures are trained for the same amount of iterations.

1. The Invariant Risk Minimization introduces a relaxation of the sufficiency criterion:

$$\text{Sufficiency} \implies \text{IRM-optimality}, \quad \text{IRM-optimality} \not\Rightarrow \text{Sufficiency}.$$

As a consequence, the model can converge to an IRM-optimal solution which is not enforcing sufficiency. Further discussion on the relation between sufficiency and IRM is discussed in Huszár (2019).

2. The gradient penalized by the regularization term $\mathbb{E}_e \|\nabla \ell_e(q_{y|z} \circ q_{z|x})\|^2$ becomes too small to produce meaningful encoder updates.

Figure 7 shows that the gradient on the different datasets quickly goes to zero as the regularization strength λ is increased. Increasing the batch size results in a decrease in the variance of the gradient estimate, but it does not increase its magnitude. This suggests that either the model already achieved IRM-optimality, or that the gradient is so small that the variance of the regularization term needs to be decreased by orders of magnitudes to result in meaningful updates. Although possible, the latter option would imply that drastically bigger batch size are needed, reducing the applicability of the IRM model to larger settings.

G Generalizing constraints from the train distribution

In section 3, we mention that different criteria implicitly rely on the assumption that enforcing a constraint (or maximizing predictive information) on the training distribution is sufficient to ensure that the constraint holds for the whole distribution, even outside of the selection. Here we report a simple example in which these assumption are clearly violated to underline the importance of environment variety at training time and pinpoint the limitations of the approaches explored in this work.




Figure 7: Gradient penalty measured during training of the IRM models using different batch size on the three variations of the CMNIST dataset. Increasing the size of the training batches reduces the variance of the gradient but it does not increase its magnitude.

Consider the problem of constructing a classifier to distinguish pictures of cows from camels. To simplify our reasoning, we will assume that each picture contains only animal features $a \in \{\text{patches, horns, humps, ..}\}$ and background features $b \in \{\text{sand, grass}\}$. The data is collected from different locations e , which represent the environmental factors. If the training distribution does not contain any example of cows on sandy terrains or camels on pastures, we can not expect any of the described criteria to successfully classify pictures. This happens for two main reasons:

1. On the training distribution animal and background feature are equally predictive of the label:

$$I_{t=1}(y; a) = I_{t=1}(y; b) = I_{t=1}(y; ab).$$

As a consequence, we can not expect the encoder $q(z|x)$ to preferentially extract animal features a , especially when the background b is much easier to identify.

Note that the amount of predictive information for the background b on the distribution that has not been selected (camels on pastures, cows on sand for $t = 0$) is strictly less than the amount of information carried by a :

$$I_{t=0}(y; b) < I_{t=0}(y; a) = I_{t=0}(y; ab).$$

This is because not all camels are necessarily in sandy locations, but all of them have humps ($H_{t=0}(y|b) > H(y|a) = 0$). In other words, the lack of diversity in the selection makes the background information b as predictive as the animal features a on the training distribution, although background is not more predictive than the animal features in general.

2. Background information is stable across all the training locations e :

$$p(y, b|e, t = 1) = p(y, b|t = 1).$$

This happens since, within the training locations, all cows are on green patches and all camels are on sandy terrains. Since the joint distribution of background and label is the same across different training locations, the constraint imposed by both sufficiency and separation

criteria cannot remove the background information \mathbf{b} from the representation \mathbf{z} . Since both $p(\mathbf{y}|\mathbf{b}, t = 1)$ and $p(\mathbf{b}|\mathbf{y}, t = 1)$ appear stable we have:

$$I_{t=1}(\mathbf{y}; \mathbf{e}|\mathbf{b}) = 0, I_{t=1}(\mathbf{e}; \mathbf{b}|\mathbf{y}) = 0.$$

In other words, no regularization strength λ can force a model trained using the sufficiency or separation (or independence) criterion to remove background information from the representation.

Note that even if background appears to be a stable feature on the selection, it is not stable in general and we have:

$$I(\mathbf{y}; \mathbf{e}|\mathbf{b}) > 0, I(\mathbf{e}; \mathbf{b}|\mathbf{y}) > 0,$$

since both $p(\mathbf{y}|\mathbf{b})$ and $p(\mathbf{b}|\mathbf{y})$ can change for some environmental conditions that have not been selected. This is because in some (unobserved) locations it is possible to find cows on sandy beaches or camels on grasslands.

To summarize, lack of diversity in the training environment may compromise the effectiveness of the criteria mentioned in this analysis. This is because models can rely on features that appear to be stable on the training distribution to make predictions. Whenever the stability is due to lack of diversity in the data selection, we cannot expect optimal out-of-distribution performance.