



# Conditional variance penalties and domain shift robustness

Christina Heinze-Deml<sup>1</sup> · Nicolai Meinshausen<sup>1</sup>

Received: 15 July 2019 / Revised: 1 October 2020 / Accepted: 15 October 2020 /  
Published online: 23 November 2020  
© The Author(s) 2020

## Abstract

When training a deep neural network for image classification, one can broadly distinguish between two types of latent features of images that will drive the classification. We can divide latent features into (i) ‘core’ or ‘conditionally invariant’ features  $C$  whose distribution  $C|Y$ , conditional on the class  $Y$ , does not change substantially across domains and (ii) ‘style’ features  $S$  whose distribution  $S|Y$  can change substantially across domains. Examples for style features include position, rotation, image quality or brightness but also more complex ones like hair color, image quality or posture for images of persons. Our goal is to minimize a loss that is robust under changes in the distribution of these style features. In contrast to previous work, we assume that the domain itself is not observed and hence a latent variable. We do assume that we can sometimes observe a typically discrete identifier or “ID variable”. In some applications we know, for example, that two images show the same person, and ID then refers to the identity of the person. The proposed method requires only a small fraction of images to have ID information. We group observations if they share the same class and identifier  $(Y, \text{ID}) = (y, \text{id})$  and penalize the conditional variance of the prediction or the loss if we condition on  $(Y, \text{ID})$ . Using a causal framework, this conditional variance regularization (CoRE) is shown to protect asymptotically against shifts in the distribution of the style variables in a partially linear structural equation model. Empirically, we show that the CoRE penalty improves predictive accuracy substantially in settings where domain changes occur in terms of image quality, brightness and color while we also look at more complex changes such as changes in movement and posture.

**Keywords** Domain shift · Dataset shift · Causal models · Distributional robustness · Anti-causal prediction · Image classification

---

Editor: Paolo Frasconi.

---

✉ Christina Heinze-Deml  
heinzdeml@stat.math.ethz.ch

Nicolai Meinshausen  
meinshausen@stat.math.ethz.ch

<sup>1</sup> Seminar for Statistics, ETH Zurich, Zurich, Switzerland

## 1 Introduction

Deep neural networks (DNNs) have achieved outstanding performance on prediction tasks like visual object and speech recognition (Krizhevsky et al. 2012; Szegedy et al. 2015; He et al. 2015). Issues can arise when the learned representations rely on dependencies that vanish in test distributions (see for example Quionero-Candela et al. (2009), Torralba and Efros (2011), Csurka (2017) and references therein). Such domain shifts can be caused by changing conditions such as color, background or location changes. Predictive performance is then likely to degrade. For example, consider the analysis presented in Kuehlkamp et al. (2017) which is concerned with the problem of predicting a person's gender based on images of their iris. The results indicate that this problem is more difficult than previous studies have suggested due to the remaining effect of cosmetics after segmenting the iris from the whole image.<sup>1</sup> Previous analyses obtained good predictive performance on certain datasets but when testing on a dataset only including images without cosmetics accuracy dropped. In other words, the high predictive performance previously reported relied to a significant extent on exploiting the confounding effect of mascara on the iris segmentation which is highly predictive for gender. Rather than the desired ability of discriminating based on the iris' texture the systems would mostly learn to detect the presence of cosmetics.

More generally, existing biases in datasets used for training machine learning algorithms tend to be replicated in the estimated models (Bolukbasi et al. 2016). For an example involving Google's photo app, see Crawford (2016) and Emspak (2016). In Sect. 5 we show many examples where unwanted biases in the training data are picked up by the trained model. As any bias in the training data is in general used to discriminate between classes, these biases will persist in future classifications, raising also considerations of fairness and discrimination (Barocas and Selbst 2016).

Addressing the issues outlined above, we propose Conditional variance Regularization (CoRE) to give differential weight to different latent features. Conceptually, we take a causal view of the data generating process and categorize the latent data generating factors into 'conditionally invariant' (*core*) and 'orthogonal' (*style*) features, as in Gong et al. (2016). The core and style features are unobserved and can in general be highly nonlinear transformations of the observed input data. It is desirable that a classifier only extracts the latent core features from the input data as they pertain to the target of interest in a stable and coherent fashion. Basing a prediction on the core features alone yields stable predictive accuracy even if the style features are altered. Under suitable assumptions, CoRE yields an estimator which is approximately invariant under changes in the conditional distribution of the style features (conditional on the class labels) and it is asymptotically robust with respect to domain shifts, arising through interventions on the style features. CoRE relies on the fact that for certain datasets we can observe *grouped observations* in the sense that we observe the same object under different conditions. For instance, such grouping information is available

- (i) in natural image data when several pictures of the same person are taken;
- (ii) in medical imaging when several images belonging to the same patient are made;

---

<sup>1</sup> Segmenting eyelashes from the iris is not entirely accurate which implies that the iris images can still contain parts of eyelashes, occluding the iris. As mascara causes the eyelashes to be thicker and darker, it is difficult to entirely remove the presence of cosmetics from the iris images.

- (iii) in speech recognition when multiple recordings from the same speaker are available;
- (iv) in video data where nearby frames showing the same objects can be exploited to group observations;
- (v) in data augmentation where a transformed data point can be grouped together with the original one.

We will show examples for the first and last category. For the last category, we will show that pairing the augmented data with the original image they were generated from helps to improve accuracy and robustness with respect to the chosen transformation.

Rather than pooling over all examples, CoRE exploits knowledge about this grouping, i.e., that a number of instances relate to the same object. By penalizing between-object variation of the prediction less than variation of the prediction for the same object, we can steer the prediction to be based more on the latent core features and less on the latent style features. While the proposed methodology can be motivated from the desire to achieve representational invariance with respect to the style features, the causal framework we use throughout this work allows to precisely formulate the distribution shifts we aim to protect against.

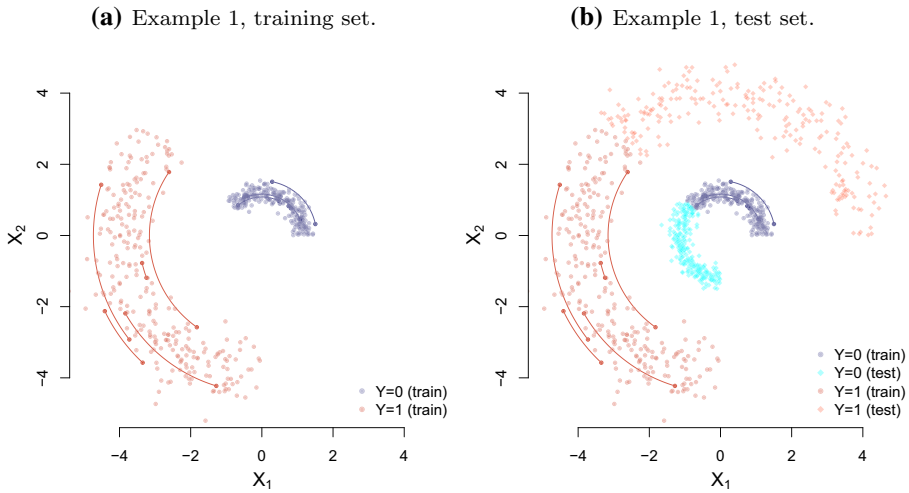
The remainder of this manuscript is structured as follows: Sect. 1.1 starts with a few motivating examples, showing simple settings where the style features change in the test distribution such that standard empirical risk minimization approaches would fail. In Sect. 1.2 we review related work, introduce notation in Sect. 2 and in Sect. 3 we formally introduce conditional variance regularization CoRE. In Sect. 4, CoRE is shown to be asymptotically equivalent to minimizing the risk under a suitable class of strong interventions in a partially linear classification setting, provided one chooses sufficiently strong CoRE penalties. We also show that the population CoRE penalty induces domain shift robustness for general loss functions to first order in the intervention strength. The size of the conditional variance penalty can be shown to determine the size of the distribution class over which we can expect distributional robustness. In Sect. 5 we evaluate the performance of CoRE in a variety of experiments.

To summarize, our contributions are the following:

- (i) **Causal framework and distributional robustness** We build on the causal framework from Gong et al. (2016) to define distributional shifts for style variables. This allows us to formulate the objective of interest in terms of distributional robust inference. Specifically, the distribution class, on which the estimator should achieve a guaranteed performance bound, consists of those distributions that are generated by interventions on the latent style variables in a causal model. Our framework allows that the domain variable itself is latent.
- (ii) **Conditional variance penalties** We introduce conditional variance penalties and show two robustness properties in Theorems 1 and 2.
- (iii) **Software** We illustrate our ideas using synthetic and real-data experiments. A TensorFlow implementation of CoRE as well as code to reproduce some of the experimental results are available at <https://github.com/christinaheinz/core>.

## 1.1 Motivating examples

To motivate the methodology we propose, consider the examples shown in Figs. 1 and 2. Example 1 shows a setting where a nonlinear decision boundary is required. Here, the core



**Fig. 1** Motivating example 1: The distributions are shifted in the test data by style interventions where style is the polar angle. Standard estimators achieve error rates of 0% on the training data and test data drawn from the same distribution as the training data (panel a). On the shown test set where the distribution of the style conditional on  $Y$  has changed the error rates are > 50% (panel b)



**Fig. 2** Motivating example 2: The goal is to predict whether a person is wearing glasses. The distributions are shifted in test data by style interventions where style is the image quality. A 5-layer CNN achieves 0% training error and 2% test error for images that are sampled from the same distribution as the training images a, but a 65% error rate on images where the confounding between image quality and glasses is changed b. See Sect. 5.3 for more details

feature corresponds to the distance from the origin while the style feature corresponds to the angle between the  $x_1$ -axis and the vector from the origin to  $(x_1, x_2)$ . Panel (a) shows a subsample of the training data where class 1 is associated with red points, dark blue points correspond to class 0. Panel (b) additionally shows a subsample of the test data where the style—i.e. the distribution of the angle—is intervened upon: class 1 is associated with orange squares, cyan squares correspond to class 0. Clearly, a circular decision boundary yields optimal performance on both training and test set but is unlikely to be found by a standard classification algorithm when only using the training set for the estimation. We will return to these examples in Sect. 3.4.

Secondly, we introduce a strong dependence between the class label and the style feature “image quality” in the third example by manipulating the face images from the CelebA dataset (Liu et al. 2015): in the training set images of class “wearing glasses” are associated with a lower image quality than images of class “not wearing glasses”. Examples are shown in Fig. 2a. In the test set, this relation is reversed, i.e. images showing persons wearing glasses are of higher quality than images of persons without glasses, with examples in Fig. 2b. We will return to this example in Sect. 5.3 and show that training a convolutional

neural network to distinguish between people wearing glasses or not works well on test data that are drawn from the same distribution (with error rates below 2%) but fails entirely on the shown test data, with error rates worse than 65%.

## 1.2 Related work

For general distributional robustness, the aim is to learn

$$\operatorname{argmin}_{\theta} \sup_{F \in \mathcal{F}} E_F(\ell(Y, f_{\theta}(X))) \quad (1)$$

for a given set  $\mathcal{F}$  of distributions, twice differentiable and convex loss  $\ell$ , and prediction  $f_{\theta}(x)$ . The set  $\mathcal{F}$  is the set of distributions on which one would like the estimator to achieve a guaranteed performance bound.

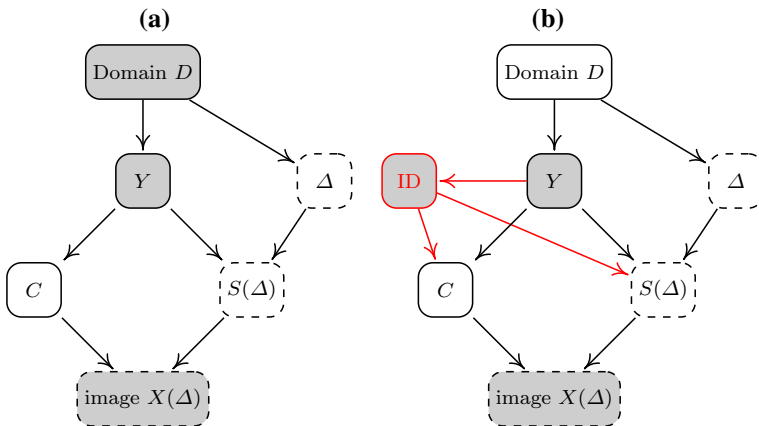
Causal inference can be seen to be a specific instance of distributional robustness, where we take  $\mathcal{F}$  to be the class of all distributions generated under do-interventions on the predictors  $X$  (Meinshausen 2018; Rothenhäusler et al. 2018). Causal models thus have the defining advantage that the predictions will be valid even under arbitrarily large interventions on all predictor variables (Haavelmo 1944; Aldrich 1989; Pearl 2009; Schölkopf et al. 2012; Peters et al. 2016; Zhang et al. 2013, 2015; Yu et al. 2017; Rojas-Carulla et al. 2018; Magliacane et al. 2018). There are two difficulties in transferring these results to the setting of domain shifts in image classification. The first hurdle is that the classification task is typically anti-causal since the image we use as a predictor is a descendant of the true class of the object we are interested in rather than the other way around. The second challenge is that the input data consists of pixel intensities and we do not want (or could) guard against arbitrary interventions on any or all variables but only would like to guard against a shift of the unobserved style features. It is hence not immediately obvious how standard causal inference can be used to guard against large domain shifts.

Another line of work uses a class of distributions of the form  $\mathcal{F} = \mathcal{F}_{\epsilon}(F_0)$  with

$$\mathcal{F}_{\epsilon}(F_0) := \{\text{distributions } F \text{ such that } D(F, F_0) \leq \epsilon\}, \quad (2)$$

with  $\epsilon > 0$  a small constant and  $D(F, F_0)$  being, for example, a  $\phi$ -divergence (Namkoong and Duchi 2017; Ben-Tal et al. 2013; Bagnell 2005; Volpi et al. 2018) or a Wasserstein distance (Shafieezadeh-Abadeh et al. 2017; Sinha et al. 2018; Gao et al. 2017). The distribution  $F_0$  can be the true (but generally unknown) population distribution  $P$  from which the data were drawn or its empirical counterpart  $P_n$ . The distributionally robust targets in Eq. (2) can often be expressed in penalized form (Gao et al. 2017; Sinha et al. 2018; Xu et al. 2009). A Wasserstein-ball is a suitable class of distributions for example in the context of adversarial examples (Sinha et al. 2018; Szegedy et al. 2014; Goodfellow et al. 2015).

In this work, we do not try to achieve robustness with respect to a set of distributions that are pre-defined by a Kullback-Leibler divergence or a Wasserstein metric as in Eq. (2). Instead, we try to achieve robustness against a set of distributions that are generated by interventions on latent style variables in a causal model (we will make this precise in Sect. 2). We will formulate the class of distributions over which we try to achieve robustness as in Eq. (1) but with the class of distributions in Eq. (2) now replaced with the class of distributions  $\mathcal{F}_{\xi}$  defined as  $\{F : D_{\text{style}}(F, F_0) \leq \xi\}$ , where  $F_0$  is again the distribution the training data are drawn from. The difference to standard distributional robustness approaches listed below Eq. (2) is now that the metric  $D_{\text{style}}$  measures the shift of the



**Fig. 3** Observed quantities are shown as shaded nodes; nodes of latent quantities are transparent. Left: data generating process for the considered model as in Gong et al. (2016), where the effect of the domain on the orthogonal features  $S$  is mediated via unobserved noise  $\Delta$ . The style interventions and all its descendants are shown as nodes with dashed borders to highlight variables that are affected by style interventions. Right: our setting. The domain itself is unobserved but we can now observe the (typically discrete) ID variable we use for grouping. The arrow between ID and  $Y$  can be reversed, depending on the sampling scheme

orthogonal style features. We do not know a priori which features are prone to distributional shifts and which features have a stable (conditional) distribution. The metric is hence not known a priori and needs to be inferred in a suitable sense from the data.

Similar to this work in terms of their goals are the work of Gong et al. (2016) and Domain-Adversarial Neural Networks (DANN) proposed in Ganin et al. (2016), an approach motivated by the work of Ben-David et al. (2007). The main idea of Ganin et al. (2016) is to learn a representation that contains no discriminative information about the origin of the input (source or target domain). This is achieved by an adversarial training procedure: the loss on domain classification is maximized while the loss of the target prediction task is minimized simultaneously. The data generating process assumed in Gong et al. (2016) is similar to our model, introduced in Sect. 2.1, where we detail the similarities and differences between the models (cf. Fig. 3). Gong et al. (2016) identify the conditionally independent features by adjusting a transformation of the variables to minimize the squared MMD distance between distributions in different domains.<sup>2</sup> The fundamental difference between these very promising methods and our approach is that we use a different data basis. The domain identifier is explicitly observable in Gong et al. (2016) and Ganin et al. (2016), while it is latent in our approach. In contrast, we exploit the presence of an identifier variable ID that relates to the identity of an object (for example identifying a person). In other words, we do *not* assume that we have data from different domains but just different realizations of the same object under different interventions. This also differentiates this work from *latent* domain adaptation papers from the computer vision literature (Hoffman et al. 2012; Gong et al. 2013). Further related work is discussed in Sect. 6.

<sup>2</sup> The distinction between ‘conditionally independent’ features and ‘conditionally transferable’ (which is the former modulo location and scale transformations) is for our purposes not relevant as we do not make a linearity assumption in general.

## 2 Setting

We introduce the assumed underlying causal graph and some notation before discussing notions of domain shift robustness.

### 2.1 Causal graph

Let  $Y \in \mathcal{Y}$  be a target of interest. Typically  $\mathcal{Y} = \mathbb{R}$  for regression or  $\mathcal{Y} = \{1, \dots, K\}$  in classification with  $K$  classes. Let  $X \in \mathbb{R}^p$  be predictor variables, for example the  $p$  pixels of an image. The causal structural model for all variables is shown in panel (b) of Fig. 3. The domain variable  $D$  is latent, in contrast to Gong et al. (2016) whose model is shown in panel (a) of Fig. 3. We add the ID variable to the graph. In Fig. 3,  $Y \rightarrow \text{ID}$  but in some settings it might be more plausible to consider  $\text{ID} \rightarrow Y$ . For the proposed method both options are possible. Together with  $Y$ , the ID variable is used to group observations. It is typically discrete and relates to the identity of the underlying object. The variable can be assumed to be latent in the setting of Gong et al. (2016).

The rest of the graph is in analogy to Gong et al. (2016). The prediction is anti-causal, that is the predictor variables  $X$  that we use for  $\hat{Y}$  are non-ancestral to  $Y$ . In other words, the class label is here seen to be causal for the image and not the other way around.<sup>3</sup> The causal effect from the class label  $Y$  on the image  $X$  is mediated via two types of latent variables: the so-called *core* or ‘conditionally invariant’ features  $C$  and the orthogonal or *style* features  $S$ . The distinguishing factor between the two is that external interventions  $\Delta$  are possible on the style features but not on the core features. If the interventions  $\Delta$  have different distributions in different domains, then the conditional distributions  $C|Y = y, \text{ID} = \text{id}$  are invariant for all  $(y, \text{id})$  while  $S|Y = y, \text{ID} = \text{id}$  can change. The style variable can include point of view, image quality, resolution, rotations, color changes, body posture, movement etc. and will in general be context-dependent.<sup>4</sup> The style intervention variable  $\Delta$  influences both the latent style  $S$ , and hence also the image  $X$ . In potential outcome notation, we let  $S(\Delta = \delta)$  be the style under intervention  $\Delta = \delta$  and  $X(Y, \text{ID}, \Delta = \delta)$  the image for class  $Y$ , identity  $\text{ID}$  and style intervention  $\Delta$ . The latter is sometimes abbreviated as  $X(\Delta = \delta)$  for notational simplicity. Finally,  $f_\theta(X(\Delta = \delta))$  is the prediction under the style intervention  $\Delta = \delta$ . For a formal justification of using a causal graph and potential outcome notation simultaneously see Richardson and Robins (2013).

To be specific, if not mentioned otherwise we will assume a causal graph as follows. For independent  $\varepsilon_Y, \varepsilon_{\text{ID}}, \varepsilon_{\text{style}}$  in  $\mathbb{R}, \mathbb{R}, \mathbb{R}^q$  respectively with positive density on their support and continuously differentiable functions  $k_Y, k_{\text{id}},$  and  $k_{\text{style}}, k_{\text{core}}, k_X,$

<sup>3</sup> If an existing image is classified by a human, then the image is certainly ancestral for the attached label. If the label  $Y$  refers, however, to the underlying true object (say if you generate images by asking people to take pictures of objects or if you record the status of cells after performing a gene knockout), then the more fitting model is the one where  $Y$  is ancestral for  $X$ . We here focus on modeling the underlying true object since ultimately the goal is to deploy the trained system in the real world.

<sup>4</sup> The type of features we regard as style and which ones we regard as core features can conceivably change depending on the circumstances—for instance, is the color “gray” an integral part of the object “elephant” or can it be changed so that a colored elephant is still considered to be an elephant?

$$\begin{aligned}
 Y &\leftarrow k_y(D, \varepsilon_Y) \\
 \text{identifier ID} &\leftarrow k_{\text{id}}(Y, \varepsilon_{\text{ID}}) \\
 \text{core or conditionally invariant features } C &\leftarrow k_{\text{core}}(Y, \text{ID}) \\
 \text{style or orthogonal features } S &\leftarrow k_{\text{style}}(Y, \text{ID}, \varepsilon_{\text{style}}) + \Delta \\
 \text{image } X &\leftarrow k_x(C, S).
 \end{aligned} \tag{3}$$

The distribution of  $\varepsilon_{\text{style}}$  is assumed to be identical across domains, while  $\Delta$  can change. In more generality, one could discard  $\Delta$  and instead allow the distribution of  $\varepsilon_{\text{style}}$  to change. Here the assumption is slightly more restrictive because of additivity of  $\Delta$  in the structural equation for the style variable. The core features are here assumed to be a deterministic function of  $Y$  and  $\text{ID}$  to allow for theoretical analysis. In more generality (and as indicated in the graph), these would also be non-deterministic relations. The theoretical results will also require positive density for the style features in an  $\varepsilon$ -ball around the origin, as made precise in assumption (A1) later.

The prediction  $\hat{y}$  for  $y$ , given  $X = x$ , is of the form  $f_{\theta}(x)$  for a suitable function  $f_{\theta}$  with parameters  $\theta \in \mathbb{R}^d$ , where the parameters  $\theta$  correspond to the weights in a DNN, for example.

We would like to stress that the above model is fairly general and subsumes many simpler ones as special cases. To give a concrete example, consider the task of classifying a health condition  $Y$  from medical images  $X$ . Style features could be, for example, technical noise, orientation or resolution. The unobserved domain  $D$  could correspond to different hospitals or doctors. Due to the usage of different measuring devices in each of these locations, the conditional distribution  $S|Y$  will change substantially across different domains. In contrast, the core features  $C$ , i.e. those image features that carry the actual signal, will remain invariant conditional on the underlying health condition  $Y$ .

## 2.2 Data

We assume we have  $n$  data points  $(x_i, y_i, \text{id}_i)$  for  $i = 1, \dots, n$ , where the observations  $\text{id}_i$  with  $i = 1, \dots, n$  of variable  $\text{ID}$  can also contain unobserved values. Let  $m \leq n$  be the number of unique realizations of  $(Y, \text{ID})$  and let  $G_1, \dots, G_m$  be a partition of  $\{1, \dots, n\}$  such that, for each  $j \in \{1, \dots, m\}$ , the realizations  $(y_i, \text{id}_i)$  are identical<sup>5</sup> for all  $i \in G_j$ . While our prime application is classification, regression settings with continuous  $Y$  can be approximated in this framework by slicing the range of the response variable into distinct bins in analogy to the approach in sliced inverse regression (Li 1991). The cardinality of  $G_j$  is denoted by  $n_j := |G_j| \geq 1$ . Then  $n = \sum_i n_i$  is again the total number of samples and  $c = n - m$  is the total number of grouped observations in the following sense: if we count all samples in a group except the first one we have, if summing over all groups, a total of  $c = n - m = \sum_{i=1}^m (n_i - 1)$  observations left that are ‘grouped’ with the first example in their corresponding group.

Typically  $n_i = 1$  for most samples and occasionally  $n_i \geq 2$  but one can also envisage scenarios with larger groups of the same identifier  $(y, \text{id})$ .

<sup>5</sup> Observations where the  $\text{ID}$  variable is unobserved are not grouped, that is each such observation is counted as a unique observation of  $(Y, \text{ID})$ .



## 2.3 Domain shift robustness

In this section, we clarify against which classes of distributions we hope to achieve robustness. Let  $\ell$  be a suitable loss that maps  $y$  and  $\hat{y} = f_\theta(x)$  to  $\mathbb{R}^+$ . The risk under distribution  $F$  and parameter  $\theta$  is given by

$$E_F \left[ \ell(Y, f_\theta(X)) \right].$$

Let  $F_0$  be the joint distribution of  $(ID, Y, S)$  in the training distribution. A new domain and explicit interventions on the style features can now shift the distribution of  $(ID, Y, \tilde{S})$  to  $F$ . We can measure the distance between distributions  $F_0$  and  $F$  in different ways. Below we will define the distance considered in this work and denote it by  $D_{\text{style}}(F, F_0)$ . Once defined, we get a class of distributions

$$\mathcal{F}_\xi = \{F : D_{\text{style}}(F_0, F) \leq \xi\} \quad (4)$$

and the goal will be to optimize a worst-case loss over this distribution class in the sense of Eq. (1), where larger values of  $\xi$  afford protection against larger distributional changes. The relevant loss for distribution class  $\mathcal{F}_\xi$  is then

$$L_\xi(\theta) = \sup_{F \in \mathcal{F}_\xi} E_F \left[ \ell(Y, f_\theta(X)) \right]. \quad (5)$$

In the limit of arbitrarily strong interventions on the style features  $S$ , the loss is given by

$$L_\infty(\theta) = \lim_{\xi \rightarrow \infty} \sup_{F \in \mathcal{F}_\xi} E_F \left[ \ell(Y, f_\theta(X)) \right]. \quad (6)$$

Minimizing the loss  $L_\infty(\theta)$  with respect to  $\theta$  guarantees an accuracy in prediction which will work well across arbitrarily large shifts in the conditional distribution of the style features.

A natural choice to define  $D_{\text{style}}$  is to use a Wasserstein-type distance (see e.g. Villani 2003). We will first define a distance  $D_{y,\text{id}}$  for the conditional distributions

$$S|Y = y, \text{ID} = \text{id} \quad \text{and} \quad \tilde{S}|Y = y, \text{ID} = \text{id},$$

and then set  $D(F_0, F) = E(D_{Y,\text{ID}})$ , where the expectation is with respect to random ID and labels  $Y$ . The distance  $D_{y,\text{id}}$  between the two conditional distributions of  $S$  will be defined as a Wasserstein  $W_2^2(F_0, \tilde{F})$ -distance for a suitable cost function  $c(x, \tilde{x})$ . Specifically, let  $\Pi_{y,\text{id}}$  be the couplings between the conditional distributions of  $S$  and  $\tilde{S}$ , meaning measures supported on  $\mathbb{R}^q \times \mathbb{R}^q$  such that the marginal distribution over the first  $q$  components is equal to the distribution of  $S$  and the marginal distribution over the remaining  $q$  components equal to the distribution of  $\tilde{S}$ . Then the distance between the conditional distributions is defined as

$$D_{y,\text{id}} = \min_{M \in \Pi_{y,\text{id}}} E \left[ c(x, \tilde{x}) \right],$$

where  $c : \mathbb{R}^q \times \mathbb{R}^q \mapsto \mathbb{R}^+$  is a nonnegative, lower semi-continuous cost function. Here, we focus on a Mahalanobis distance as cost

$$c^2(x, \tilde{x}) = (x - \tilde{x})^t \Sigma_{y,\text{id}}^{-1} (x - \tilde{x}).$$

The cost of a shift is hence measured against the variability under the distribution  $F_0$ ,  $\Sigma_{y,\text{id}} = \text{Var}(S|Y, \text{ID})$ .<sup>6</sup>

Clearly, since the core and style features are unobserved, we cannot directly optimize the loss (5) but need to infer the metric  $D_{\text{style}}$  from the input data  $X$ . In the next section, we will show how this can be achieved. Intuitively, the model in Fig. 3 implies that the variance conditional on  $Y, \text{ID}$  stems from the difference in the style features. Hence, we would like to minimize the conditional variance of the prediction or loss when we condition on  $Y, \text{ID}$ . This enforces the desired invariance with respect to the style features.

### 3 Conditional variance regularization

#### 3.1 Pooled estimator

Let  $(x_i, y_i)$  for  $i = 1, \dots, n$  be the observations that constitute the training data and  $\hat{y}_i = f_\theta(x_i)$  the prediction for  $y_i$ . The standard approach is to simply pool over all available observations, ignoring any grouping information that might be available. The pooled estimator thus treats all examples identically by summing over the empirical loss as

$$\hat{\theta}^{\text{pool}} = \underset{\theta}{\text{argmin}} \hat{E} \left[ \ell(Y, f_\theta(X)) \right] + \gamma \cdot \text{pen}(\theta), \tag{7}$$

where the first part is simply the empirical loss over the training data,

$$\hat{E} \left[ \ell(Y, f_\theta(X)) \right] = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i)).$$

In the second part,  $\text{pen}(\theta)$  is a complexity penalty, for example a squared  $\ell_2$ -norm of the weights  $\theta$  in a convolutional neural network as a ridge penalty.

#### 3.2 CoRE estimator

The CoRE estimator is defined in Lagrangian form for penalty  $\lambda \geq 0$  as

$$\hat{\theta}^{\text{core}}(\lambda) = \underset{\theta}{\text{argmin}} \hat{E} \left[ \ell(Y, f_\theta(X)) \right] + \lambda \cdot \hat{C}_\theta. \tag{8}$$

The penalty  $\hat{C}_\theta$  is a conditional variance penalty of the form

$$\text{conditional-variance-of-prediction:} \quad \hat{C}_{f, v, \theta} := \hat{E} \left[ \widehat{\text{Var}}(f_\theta(X) | Y, \text{ID})^v \right] \tag{9}$$

<sup>6</sup> As an example, if the change in distribution for  $S$  is caused by random shift-interventions  $\Delta$ , then  $\tilde{S} \leftarrow S + \Delta$ , and the distance  $D_{\text{style}}$  induced in the distributions is

$$D_{\text{style}}(F_0, F) \leq E \left[ E(\Delta^t \Sigma_{y,\text{id}}^{-1} \Delta | Y = y, \text{ID} = \text{id}) \right],$$

ensuring that the strength of the shifts is measured against the natural variability  $\Sigma_{y,\text{id}}$  of the style features.

$$\text{conditional-variance-of-loss: } \hat{C}_{\ell, \nu, \theta} := \hat{E}[\widehat{\text{Var}}(\ell(Y, f_{\theta}(X)) | Y, \text{ID})^{\nu}], \quad (10)$$

where typically  $\nu \in \{1/2, 1\}$ . For  $\nu = 1/2$ , we also refer to the respective penalties as “conditional-standard-deviation” penalties. In practice in the context of classification and DNNs, we apply the penalty (9) to the predicted logits. The conditional-variance-of-loss penalty (10) takes a similar form to Namkoong and Duchi (2017). The crucial difference of our approach to Namkoong and Duchi (2017) is that we penalize with the expected *conditional* variance or standard deviation. The fact that we take a conditional variance is here important as we try to achieve distributional robustness with respect to interventions on the style variables. Conditioning on ID allows to guard specifically against these interventions. An unconditional variance penalty, in contrast, can achieve robustness against a pre-defined class of distributions such as a ball of distributions defined in a Kullback-Leibler or Wasserstein metric. The population CoRE estimator is defined as in Eq. (8) where empirical estimates are replaced by their respective population quantities.

Before showing numerical examples, we discuss the estimation of the expected conditional variance in Sect. 3.3 and return to the simple examples of Sect. 1.1 in Sect. 3.4. Domain shift robustness in a classification setting for a partially linear version of the structural equation model (3) is shown in Sect. 4.1. Furthermore, we discuss the population limit of  $\hat{\theta}^{\text{core}}(\lambda)$  in Sect. 4.2, where we show that the regularization parameter  $\lambda \geq 0$  is proportional to the size of the future style interventions that we want to guard against for future test data.

### 3.3 Estimating the expected conditional variance

Recall that  $G_j \subseteq \{1, \dots, n\}$  contains samples with identical realizations of  $(Y, \text{ID})$  for  $j \in \{1, \dots, m\}$ . For each  $j \in \{1, \dots, m\}$ , define  $\hat{\mu}_{\theta, j}$  as the arithmetic mean across all  $f_{\theta}(x_i), i \in G_j$ . The canonical estimator of the conditional variance  $\hat{C}_{f, 1, \theta}$  is then

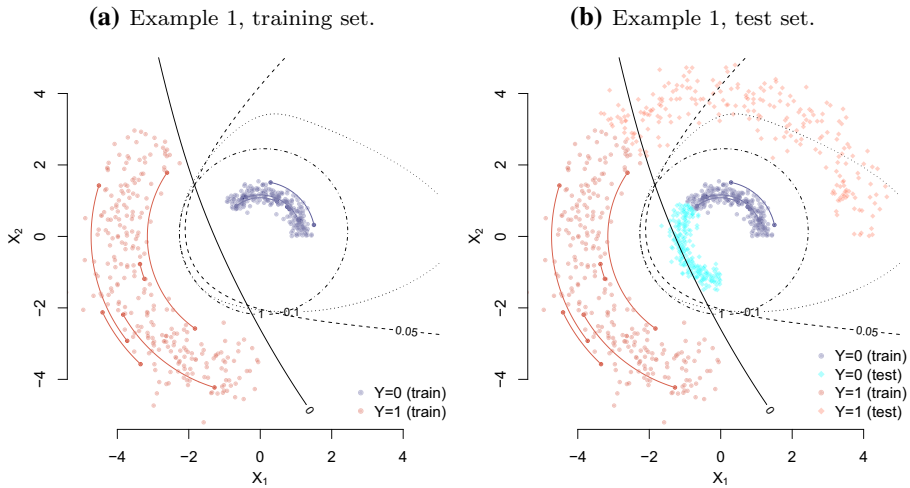
$$\hat{C}_{f, 1, \theta} := \frac{1}{m} \sum_{j=1}^m \frac{1}{|G_j|} \sum_{i \in G_j} (f_{\theta}(x_i) - \hat{\mu}_{\theta, j})^2, \quad \text{where } \hat{\mu}_{\theta, j} = \frac{1}{|G_j|} \sum_{i \in G_j} f_{\theta}(x_i)$$

and analogously for the conditional-variance-of-loss, defined in Eq. (10)<sup>7</sup>. If there are no groups of samples that share the same identifier  $(y, \text{id})$ , we define  $\hat{C}_{f, 1, \theta}$  to vanish. The CoRE estimator is then identical to pooled estimation in this special case.

### 3.4 Motivating examples (continued)

We revisit the first example from Sect. 1.1. Figure 4 shows subsamples of the training and test set with the estimated decision boundaries for different values of the penalty parameter  $\lambda$  when using a 2-layer fully connected neural network. Here,  $n = 20,000$  and  $c = 500$ . Additionally, grouped examples that share the same  $(y, \text{id})$  are visualized: two grouped observations are connected by a line or curve, respectively. Ten such groups are shown. Panel (a) shows the decision boundaries for  $\lambda = 0$ , equivalent to the pooled estimator, and for CoRE with  $\lambda \in \{0, 0.05, 0.1, 1\}$ . The pooled estimator misclassifies a large number of

<sup>7</sup> The right hand side can also be interpreted as the graph Laplacian (Belkin et al. 2006) of an appropriately weighted graph that fully connects all observations  $i \in G_j$  for each  $j \in \{1, \dots, m\}$ .



**Fig. 4** The decision boundary as function of the penalty parameters  $\lambda$  for Example 1 from Fig. 1. There are ten pairs of samples visualized that share the same identifier ( $y, id$ ) and these are connected by a curve in the figures. The decision boundary associated with a solid line corresponds to  $\lambda = 0$ , the standard pooled estimator that ignores the groupings. The broken lines are decision boundaries for increasingly strong penalties, taking into account the groupings in the data. Here, we only show a subsample of the data to avoid overplotting

test points as can be seen in panel (b), suffering from a test error of  $\approx 58\%$ . In contrast, the decision boundary of the CoRE estimator with  $\lambda = 1$  aligns with the direction along which the grouped observations vary, classifying the test set with almost perfect accuracy (test error is  $\approx 0\%$ ).

### 4 Domain shift robustness for the CoRE estimator

We show two properties of the CoRE estimator. First, consistency is shown under the risk definition (6) for an infinitely large conditional variance penalty and the logistic loss in a partially linear structural equation model. Second, the population CoRE estimator is shown to achieve distributional robustness against shift interventions in a first order expansion.

#### 4.1 Asymptotic domain shift robustness under strong interventions

We analyze the loss under strong domain shifts, as given in Eq. (6), for the pooled and the CoRE estimator in a one-layer network for binary classification (logistic regression) in an asymptotic setting of large sample size and strong interventions.

Assume the structural equation for the image  $X \in \mathbb{R}^p$  is linear in the style features  $S \in \mathbb{R}^q$  (with generally  $p \gg q$ ) and we use logistic regression to predict the class label  $Y \in \{-1, 1\}$ . Let the interventions  $\Delta \in \mathbb{R}^q$  act additively on the style features  $S$  (this is only for notational convenience) and let the style features  $S$  act in a linear way on the image  $X$  via a matrix  $W \in \mathbb{R}^{p \times q}$  (this is an important assumption without which results are more involved). The core or ‘conditionally invariant’ features are  $C \in \mathbb{R}^r$ , where in general  $r \leq p$  but this is not important for the following. For independent  $\epsilon_Y, \epsilon_{ID}, \epsilon_{style}$  in  $\mathbb{R}, \mathbb{R}, \mathbb{R}^q$

respectively with positive density on their support and continuously differentiable functions  $k_y, k_{id}, k_{style}, k_{core}, k_x$ ,

$$\begin{aligned}
 \text{class } Y &\leftarrow k_y(D, \epsilon_Y) \\
 \text{identifier ID} &\leftarrow k_{id}(Y, \epsilon_{ID}) \\
 \text{core or conditionally invariant features } C &\leftarrow k_{core}(Y, ID) \\
 \text{style or orthogonal features } S &\leftarrow k_{style}(Y, ID, \epsilon_{style}) + \Delta \\
 \text{image } X &\leftarrow k_x(C) + WS.
 \end{aligned} \tag{11}$$

We assume a logistic regression as a prediction of  $Y$  from the image data  $X$ :

$$f_\theta(x) := \frac{\exp(x^t \theta)}{1 + \exp(x^t \theta)}.$$

Given training data with  $n$  samples, we estimate  $\theta$  with  $\hat{\theta}$  and use here a logistic loss  $\ell_\theta(y_i, x_i) = \log(1 + \exp(-y_i(x_i^t \theta)))$ .

The formulation of Theorem 1 relies on the following assumptions.

**Assumption 1** We require the following conditions:

- (A1) Assume the conditional distribution  $S|Y = y, ID = id$  under the training distribution  $F_0$  has positive density (with respect to the Lebesgue measure) in an  $\epsilon$ -ball in  $\ell_2$ -norm around the origin for some  $\epsilon > 0$  for all  $y \in \mathcal{Y}$  and  $id \in \mathcal{I}$ .
- (A2) Assume the matrix  $W$  has full rank  $q$ .
- (A3) Let  $M \leq n$  be the number of unique realizations among  $n$  iid samples of  $(Y, ID)$  and let  $p_n := P(M \leq n - q)$ . Assume that  $p_n \rightarrow 1$  for  $n \rightarrow \infty$ .

Assumption (A1) is a key assumption about the style variations we observe in the training set. It requires that we observe some variance in those directions that we expect to be subject to domain shifts in the future. If, on the other hand, the conditional variance in a particular direction is vanishing, we also expect it to vanish in the future. A violation of this assumption would imply that the guarantee of the CoRE regularization no longer holds. Assumption (A3) guarantees that the number  $c = n - m$  of grouped examples is at least as large as the dimension of the style variables. If we have too few or no grouped examples (small  $c$ ), we cannot estimate the conditional variance accurately. Under these assumptions we can prove domain shift robustness.

**Theorem 1** (Asymptotic domain shift robustness under strong interventions) *Under model (11) and Assumption 1, with probability 1, the pooled estimator (7) has infinite loss (6) under arbitrarily large shifts in the distribution of the style features,*

$$L_\infty(\hat{\theta}^{pool}) = \infty.$$

The CoRE estimator (8)  $\hat{\theta}^{core}$  with  $\lambda \rightarrow \infty$  is domain shift robust under strong interventions in the sense that for  $n \rightarrow \infty$ ,

$$L_\infty(\hat{\theta}^{core}) \rightarrow_p \inf_\theta L_\infty(\theta).$$

A proof is given in “Appendix A”. The respective ridge penalties in both estimators (7) and (8) are assumed to be zero for the proof, but the proof can easily be generalized to include ridge penalties that vanish sufficiently fast for large sample sizes. The Lagrangian regularizer  $\lambda$  is assumed to be infinite for the CoRE estimator to achieve domain shift robustness under these strong interventions. The next section considers the population CoRE estimator in a setting with weak interventions and finite values of the penalty parameter.

### 4.2 Population domain shift robustness under weak interventions

The previous theorem states that the CoRE estimator can achieve domain shift robustness under strong interventions for an infinitely strong penalty in an asymptotic setting. An open question is how the loss (5),

$$L_\xi(\theta) = \sup_{F \in \mathcal{F}_\xi} E_F \left[ \ell(Y, f_\theta(X)) \right]$$

behaves under interventions of small to medium size and correspondingly smaller values of the penalty. Here, we aim to minimize this loss for a given value of  $\xi$  and show that domain shift robustness can be achieved to first order with the population CoRE estimator using the conditional-standard-deviation-of-loss penalty, i.e., Eq. (10) with  $\nu = 1/2$ , by choosing an appropriate value of the penalty  $\lambda$ . Below we will show this appropriate choice of the penalty weight is  $\lambda = \sqrt{\xi}$ .

**Assumption 2** We require the following conditions:

(B1) Define the loss under a deterministic shift  $\delta$  as

$$h_\theta(\delta) := E_{F_\theta} [\ell(Y, f_\theta(X))],$$

where the expectation is with respect to random  $(ID, Y, \tilde{S}) \sim F_\theta$ , with  $F_\theta$  defined by the deterministic shift intervention  $\tilde{S} = S + \delta$  and  $(ID, Y, \tilde{S}) \sim F_0$ . Assume that for all  $\theta \in \Theta$ ,  $h_\theta(\delta)$  is twice continuously differentiable with bounded second derivative for a deterministic shift  $\delta \in \mathbb{R}^q$ .

(B2) The spectral norm of the conditional variance  $\Sigma_{y, id}$  of  $S|Y, ID$  under  $F_0$  is assumed to be smaller or equal to some  $\zeta \in \mathbb{R}$  for all  $y \in \mathcal{Y}$  and  $id \in \mathcal{I}$ .

The first assumption (B1) ensures that the loss is well behaved under interventions on the style variables. The second assumption (B2) allows to take the limit of small conditional variances in the style variables.

If setting  $\lambda = \sqrt{\xi}$  and using the conditional-standard-deviation-of-loss penalty, the CoRE estimator optimizes according to

$$\hat{\theta}^{core}(\sqrt{\xi}) = \underset{\theta}{\operatorname{argmin}} \hat{E}_{F_0} [\ell(Y, f_\theta(X))] + \sqrt{\xi} \cdot \hat{C}_{\ell, 1/2, \theta}.$$

The next theorem shows that this is to first order equivalent to minimizing the worst-case loss over the distribution class  $\mathcal{F}_\xi$ . The following result holds for the population CoRE estimator, see below for a discussion about consistency.

**Theorem 2** *The supremum of the loss over the class of distribution  $F_{\xi}$  is to first-order given by the expected loss under distribution  $F_0$  with an additional conditional-standard-deviation-of-loss penalty  $C_{\ell,1/2,\theta}$*

$$\sup_{F \in \mathcal{F}_{\xi}} E_F[\ell(Y, f_{\theta}(X))] = E_{F_0}[\ell(Y, f_{\theta}(X))] + \sqrt{\xi} \cdot C_{\ell,1/2,\theta} + O(\max\{\xi, \zeta\}). \quad (12)$$

A proof is given in “Appendix B”. The objective of the population CoRE estimator matches thus to first order the loss under domain shifts if we set the penalty weight  $\lambda = \sqrt{\xi}$ . Larger anticipated domain shifts thus require naturally a larger penalty  $\lambda$  in the CoRE estimation. The result is possible as we have chosen the Mahalanobis distance to measure shifts in the style variable and define  $\mathcal{F}_{\xi}$ , ensuring that the strength of shifts on style variables are measured against the natural variance on the training distribution  $F_0$ .

In practice, the choice of  $\lambda$  involves a somewhat subjective choice about the strength of the distributional robustness guarantee. A stronger distributional robustness property is traded off against a loss in predictive accuracy if the distribution is not changing in the future. One option for choosing  $\lambda$  is to choose the largest penalty weight before the validation loss increases considerably. This approach would provide the best distributional robustness guarantee that keeps the loss of predictive accuracy in the training distribution within a pre-specified bound.<sup>8</sup>

As a caveat, the result takes the limit of small conditional variance of  $S$  in the training distribution and small additional interventions. Under larger interventions higher-order terms could start to dominate, depending on the geometry of the loss function and  $f_{\theta}$ . A further caveat is that the result looks at the population CoRE estimator. For finite sample sizes, we would optimize a noisy version on the rhs of (12). To show domain shift robustness in an asymptotic sense, we would need additional uniform convergence (in  $\theta$ ) of both the empirical loss and the conditional variance in that for  $n \rightarrow \infty$ ,

$$\begin{aligned} \sup_{\theta} |\hat{E}_{F_0}[\ell(Y, f_{\theta}(X))] - E_{F_0}[\ell(Y, f_{\theta}(X))]| &\rightarrow_p 0, \quad \text{and} \\ \sup_{\theta} |\hat{C}_{\ell,1/2,\theta} - C_{\ell,1/2,\theta}| &\rightarrow_p 0. \end{aligned}$$

While this is in general a reasonable assumption to make, the validity of the assumption will depend on the specific function class and on the chosen estimator of the conditional variance.

## 5 Experiments

We perform an array of different experiments, showing the applicability and advantage of the conditional variance penalty for two broad settings:

1. Settings where we **do not know** what the style variables correspond to but still want to protect against a change in their distribution in the future. In the examples we show

<sup>8</sup> If some labeled test data are available and our goal is to perform optimally for samples that come from the same distribution as those test samples, then we can and should adjust the parameter to minimize estimated test error.

cases where the style variable ranges from fashion (Sect. 5.2), image quality (Sect. 5.3), movement (Sect. 5.4) and brightness (“Appendix D.1”), which are all not known explicitly to the method. We also include genuinely unknown style variables in Sect. 5.1 (in the sense that they are unknown not only to the methods but also to us as we did not explicitly create the style interventions).

2. Settings where we **do know** what type of style interventions we would like to protect against. This is usually dealt with by data augmentation (adding images which are, say, rotated or shifted compared to the training data if we want to protect against rotations or translations in the test data; see for example Schölkopf et al. (1996)). The conditional variance penalty is here exploiting that some augmented samples were generated from the same original sample and we use as ID variable the index of the original image. We show that this approach generalizes better than simply pooling the augmented data, in the sense that we need fewer augmented samples to achieve the same test error. This setting is shown in Sect. 5.5.

We compare against the pooled estimator which has the same architecture as the network to which we add the CoRE penalty. For both the pooled and the CoRE estimator we apply an  $\ell_2$  penalty as regularization. We would like to stress that the related work discussed in Sects. 1.2 and 6 cannot be directly compared to the CoRE estimator as these approaches cannot exploit the ID information but rely on having data from different domains available at training time instead. Since this is a different problem setting, we can only compare against the pooled estimator which is a standard approach to classification. As a downside, our approach requires availability of an ID variable, which might not always be available<sup>9</sup> To further understand the behavior of the CoRE penalty, we perform a number of analyses and ablation studies to show

- (i) How sensitive the performance of CoRE is to the value of the penalty weight  $\lambda$  (Sects. 5.1.1, 5.2);
- (ii) How the CoRE penalty differs from a standard  $\ell_2$  penalty (Sect. 5.1.1);
- (iii) How the value of the CoRE penalty can be used as a qualitative measure for the presence of sample bias (Sects. 5.1.1, 5.2);
- (iv) How sensitive the performance of both the CoRE and the pooled estimator is to label shift in the grouped observations (Sect. 5.2.1);
- (v) How the relative performance of both estimators is affected when using pre-trained InceptionV3 features (Sect. 5.2.2);
- (vi) How sensitive the performance is to different grouping strategies (Sects. 5.3, “Appendix D.1”, D.3.1, D.4);
- (vii) how sensitive the performance is as a function of the strength of the domain shift and the number of grouped observations (“Appendices D.1, D.4, D.5, D.6”).

<sup>9</sup> For example, the home office dataset in (Venkateswara et al. 2017) contains images of various objects (kettle, clock etc) from different domains (real world photographs, cliparts, art images and product images). If we aim to predict the category of an object, it seems difficult to identify an appropriate ID variable in the dataset. However, if we aim instead to predict properties of an object (does it use electricity?), then we can use as ID the object category (kettle, clock etc.)





**Fig. 5** Eyeglass detection for CelebA dataset with small sample size. The goal is to predict whether a person wears glasses or not. Random samples from training and test data are shown. Groups of observations in the training data that have common  $(Y, ID)$  here correspond to pictures of the same person with either glasses on or off. These are labelled by red boxes in the training data and the conditional variance penalty is calculated across these groups of pictures

Details of the network architectures can be found in Appendix “Appendix C”. All reported error rates are averaged over five runs of the respective method. A TensorFlow (Abadi et al. 2015) implementation of CoRE can be found at <https://github.com/christinaheinze/core>.

## 5.1 Eyeglasses detection with small sample size

In this example, we explore a setting where training and test data are drawn from the same distribution, so we might not expect a distributional shift between the two. However, we consider a small training sample size which gives rise to statistical fluctuations between training and test data. We assess to which extent the conditional variance penalty can help to improve test accuracies in this setting.

Specifically, we use a subsample of the CelebA dataset (Liu et al. 2015) and try to classify images according to whether or not the person in the image wears glasses. For construction of the ID variable, we exploit the fact that several photos of the same person are available and set ID to be the identifier of the person in the dataset. Figure 5 shows examples from both the training and the test dataset. The conditional variance penalty is estimated across groups of observations that share a common  $(Y, ID)$ . Here, this corresponds to pictures of the same person where all pictures show the person either with glasses (if  $Y = 1$ ) or all pictures show the person without glasses ( $Y = 0$ ). Statistical fluctuations between training and test set could for instance arise if by chance the background of eyeglass wearers is darker in the training sample than in test samples, the eyeglass wearers happen to be outdoors more often or might be more often female than male etc.

Below, we present the following analyses. First, we look at five different datasets and analyze the effect of adding the CoRE penalty (using conditional-variance-of-prediction) to the cross-entropy loss. Second, we focus on one dataset and compare the four different variants of the CoRE penalty in Eqs. (9) and (10) with  $\nu \in \{1/2, 1\}$ .

### 5.1.1 CoRE penalty using the conditional variance of the predicted logits

We consider five different training sets which are created as follows. For each person in the standard CelebA training data we count the number of available images and select the 50 identities for which most images are available individually. We partition these 50 identities into 5 disjoint subsets of size 10 and consider the resulting 5 datasets,

**Table 1** Eyeglass detection, trained on small subsets (DS1–DS5) of the CelebA dataset with disjoint identities

	Method	Error		Penalty value	
		Training	Test	Training	Test
DS 1	5-layer CNN	0.0% (0.00%)	18.08% (0.24%)	19.14 (1.70)	18.86 (1.87)
	+ CoRE	0.0% (0.00%)	<b>15.08% (0.43%)</b>	0.01 (0.01)	0.70 (0.05)
DS 2	5-layer CNN	0.0% (0.00%)	23.81% (0.51%)	6.20 (0.35)	6.97 (0.46)
	+ CoRE	0.0% (0.00%)	<b>17.00% (0.75%)</b>	0.00 (0.00)	0.41 (0.04)
DS 3	5-layer CNN	0.0% (0.00%)	18.61% (0.52%)	7.33 (1.40)	7.91 (1.13)
	+ CoRE	0.0% (0.00%)	<b>14.79% (0.89%)</b>	0.00 (0.00)	0.26 (0.03)
DS 4	5-layer CNN	0.0% (0.00%)	25.97% (0.24%)	6.19 (0.43)	7.13 (0.54)
	+ CoRE	0.0% (0.00%)	<b>21.12% (0.40%)</b>	0.00 (0.00)	0.63 (0.04)
DS 5	5-layer CNN	0.0% (0.00%)	23.64% (0.64%)	20.20 (2.46)	24.85 (3.56)
	+ CoRE	0.0% (0.00%)	<b>21.49% (1.27%)</b>	0.00 (0.00)	0.59 (0.10)

For each experimental setting, bold marks the best performing method

We report training and test error as well as the value of the CoRE penalty  $\hat{C}_{f,1,\theta}$  on the training and the test set after training, evaluated for both the pooled estimator (“5-layer CNN”) and the CoRE estimator, applied to training the same architecture (“+ CoRE”). The weights of the ridge and the CoRE penalty were chosen based on their performance on the validation set

containing the images of 10 unique identities each. The resulting 5 datasets have sizes {289, 296, 292, 287, 287}. For the validation and the test set, we consider the usual CelebA validation and test split but balance these with respect to the target variable “Eyeglasses”. The balanced validation set consists of 2766 observations; the balanced test set contains 2578 images. The identities in the validation and test sets are disjoint from the identities in the training sets.

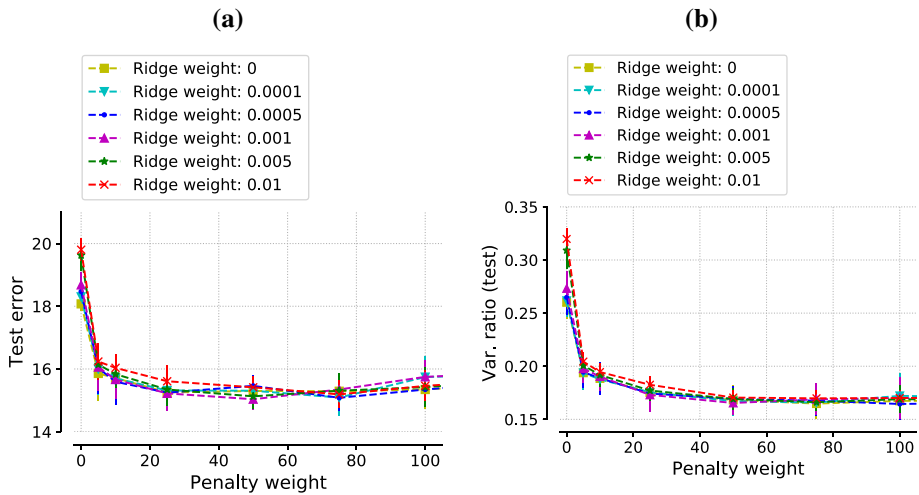
Given a training dataset, the standard approach would be to pool all examples. The only additional information we exploit is that some observations can be grouped. If using a 5-layer convolutional neural network with a standard ridge penalty (details can be found in Table 5) and pooling all data, the test error on unseen images ranges from 18.08 to 25.97%. Exploiting the group structure with the CoRE penalty (in addition to a ridge penalty) results in test errors ranging from 14.79 to 21.49%, see Table 1. The relative improvements when using the CoRE penalty range from 9 to 28.6%.

The test error is not very sensitive to the weight of the CoRE penalty as shown in Fig. 6a: for a large range of penalty weights, adding the CoRE penalty decreases the test error compared to the pooled estimator (identical to a CoRE penalty weight of 0). This holds true for various ridge penalty weights.

While the test error rates shown in Fig. 6 suggest already that the CoRE penalty differentiates itself clearly from a standard ridge penalty, we examine next the differential effect of the CoRE penalty on the between- and within-group variances. Concretely, the variance of the predictions can be decomposed as

$$\text{Var}(f_\theta(X)) = E[\text{Var}(f_\theta(X)|Y, \text{ID})] + \text{Var}[E(f_\theta(X)|Y, \text{ID})],$$

where the first term on the rhs is the within-group variance that CoRE penalizes, while a ridge penalty would penalize both the within- and also the between-group variance (the



**Fig. 6** Eye-glass detection, trained on a small subset (DS1) of the CelebA dataset with disjoint identities. **a** Average test error as a function of both the CoRE penalty on  $x$ -axis and various levels of the ridge penalty. The results can be seen to be fairly insensitive to the ridge penalty. **b** The variance ratio (13) on test data as a function of both the CoRE and ridge penalty weights. The CoRE penalty can be seen to penalize the within-group variance selectively, whereas a strong ridge penalty decreases both the within- and between-group variance

second term on the rhs above). In Fig. 6b we show the ratio between the CoRE penalty and the between-group variance where groups are defined by conditioning on  $(Y, ID)$ . Specifically, the ratio is computed as

$$\hat{E}[\widehat{\text{Var}}(f_{\theta}(X)|Y, ID)] / \widehat{\text{Var}}[\hat{E}(f_{\theta}(X)|Y, ID)]. \tag{13}$$

The results shown in Fig. 6b are computed on dataset 1 (DS 1). While increasing ridge penalty weights do lead to a smaller value of the CoRE penalty, the between-group variance is also reduced such that the ratio between the two terms does not decrease with larger weights of the ridge penalty.<sup>10</sup> With increasing weight of the CoRE penalty, the variance ratio decreases, showing that the CoRE penalty indeed penalizes the within-group variance more than the between-group variance.

Table 1 also reports the value of the CoRE penalty after training when evaluated for the pooled and the CoRE estimator on the training and the test set. As a qualitative measure to assess the presence of sample bias in the data (provided the model assumptions hold), we can compare the value the CoRE penalty takes after training when evaluated for the pooled estimator and the CoRE estimator. The difference yields a measure for the extent the respective estimators are functions of  $\Delta$ . If the respective hold-out values are both small, this would indicate that the style features are not very predictive for the target variable. If, on the other hand, the CoRE penalty evaluated for the pooled estimator takes a much larger value than for the CoRE estimator (as in this case), this would indicate the presence of sample bias.

<sup>10</sup> In Fig. 18 in the Appendix, the numerator and the denominator are plotted separately as a function of the CoRE penalty weight.

**Table 2** Eyeglass detection, trained on a small subset (DS1) of the CelebA dataset with disjoint identities

Method	Error Test	Penalty value	
		Training	Test
5-layer CNN	18.08% (0.24%)	19.14 (1.70)	18.86 (1.87)
5-layer CNN + CoRE w/ $\hat{C}_{f,1,\theta}$	15.08% (0.43%)	0.01 (0.01)	0.70 (0.05)
5-layer CNN + CoRE w/ $\hat{C}_{f,1/2,\theta}$	15.34% (0.83%)	0.03 (0.01)	0.89 (0.03)
5-layer CNN + CoRE w/ $\hat{C}_{l,1,\theta}$	15.12% (0.27%)	0.00 (0.00)	0.38 (0.03)
5-layer CNN + CoRE w/ $\hat{C}_{l,1/2,\theta}$	15.59% (0.36%)	0.00 (0.00)	0.35 (0.02)

We report training and test error as well as the value of the CoRE penalties  $\hat{C}_{f,1,\theta}$ ,  $\hat{C}_{f,1/2,\theta}$ ,  $\hat{C}_{l,1,\theta}$  and  $\hat{C}_{l,1/2,\theta}$  on the training and the test set after training, evaluated for both the pooled estimator and the CoRE estimator. The weights of the ridge and the CoRE penalty were chosen based on their performance on the validation set. The four CoRE penalty variants' performance differences are not statistically significant

### 5.1.2 Other CoRE penalty types


We now compare all CoRE penalty types, i.e., penalizing with (i) the conditional variance of the predicted logits  $\hat{C}_{f,1,\theta}$ , (ii) the conditional standard deviation of the predicted logits  $\hat{C}_{f,1/2,\theta}$ , (iii) the conditional variance of the loss  $\hat{C}_{l,1,\theta}$  and (iv) the conditional standard deviation of the loss  $\hat{C}_{l,1/2,\theta}$ . For this comparison, we use the training dataset 1 (DS 1) from above. Table 2 contains the test error (training error was 0% for all methods) as well as the value the respective CoRE penalty took after training on the training set and the test set. The four CoRE penalty variants' performance differences are not statistically significant. Hence, we mostly focus on the conditional variance of the predicted logits  $\hat{C}_{f,1,\theta}$  in the other experiments.

### 5.1.3 Discussion

While the distributional shift in this example arises due to statistical fluctuations which will diminish as the sample size grows, the following examples are more concerned with biases that will persist even if the number of training and test samples is very large. A second difference to the subsequent examples is the grouping structure—in this example, we consider only a few identities, namely  $m = 10$ , with a relatively large number  $n_i$  of associated observations (about thirty observations per individual). In the following examples,  $m$  is much larger while  $n_i$  is typically smaller than five.

## 5.2 Gender classification with unknown confounding

In the following set of experiments, we work again with the CelebA dataset and the 5-layer convolutional neural network architecture described in Table 5. This time we consider the problem of classifying whether the person shown in the image is male or female. We create a confounding in training and test set I by including mostly images of men wearing glasses and women not wearing glasses. In test set 2 the association between gender and glasses is flipped: women always wear glasses while men never




**Fig. 7** Classification for  $Y \in \{\text{woman, man}\}$ . There is an unknown confounding here as men are very likely to wear glasses in training and test set 1 data, while it is women that are likely to wear glasses in test set 2. Estimators that pool all observations are making use of this confounding and hence fail for test set 2. The conditional variance penalty for the CoRE estimator is computed over groups of images of the same person (and consequently same class label), such as the images in the red box on the left. The number of grouped examples  $c$  is 500. We vary the proportion of males in the grouped examples between 50 and 100% (cf. Sect. 5.2.1)

wear glasses. Examples from the training and test sets 1 and 2 are shown in Fig. 7. The training set, test set 1 and 2 are subsampled such that they are balanced with respect to  $Y$ , resulting in 16,982, 4224 and 1120 observations, respectively.

To compute the conditional variance penalty, we use again images of the same person. The ID variable is, in other words, the identity of the person and gender  $Y$  is constant across all examples with the same ID. Conditioning on  $(Y, \text{ID})$  is hence identical to conditioning on ID alone. Another difference to the other experiments is that we consider a binary style feature here.

### 5.2.1 Label shift in grouped observations

We compare six different datasets that vary with respect to the distribution of  $Y$  in the grouped observations. In all training datasets, the total number of observations is 16982 and the total number of grouped observations is 500. In the first dataset, 50% of the grouped observations correspond to males and 50% correspond to females. In the remaining 5 datasets, we increase the number of grouped observations with  $Y = \text{“man”}$ , denoted by  $\kappa$ , to 75%, 90%, 95%, 99% and 100%, respectively. Table 3 shows the performance obtained for these datasets when using the pooled estimator compared to the CoRE estimator with  $\hat{C}_{f,1,\theta}$ . The results show that both the pooled estimator as well as the CoRE estimator perform better if the distribution of  $Y$  in the grouped observations is more balanced. The CoRE estimator improves the error rate of the pooled estimator by  $\approx 28 - 39\%$  on a relative scale. Figure 8 shows the performance for  $\kappa = 50\%$  as a function of the CoRE penalty weight. Significant improvements can be obtained across a large range of values for the CoRE penalty and the ridge penalty. Test errors become more sensitive to the chosen value of the CoRE penalty for very large values of the ridge penalty weight as the overall amount of regularization is already large.



**Fig. 8** Classification for  $Y \in \{\text{woman, man}\}$  with  $\kappa = 0.5$ . Panels **a** and **b** show the test error on test data sets 1 and 2 respectively as a function of the CoRE and ridge penalty. Panels **c** and **d** show the variance ratio (13) (comparing within- and between- group variances) for females and males separately

### 5.2.2 Using pre-trained Inception V3 features

To verify that the above conclusions do not change when using more powerful features, we here compare  $\ell_2$ -regularized logistic regression using pre-trained Inception V3 features<sup>11</sup> with and without the CoRE penalty. Table 4 shows the results for  $\kappa = 0.5$ . While the results show that both the pooled estimator as well as the CoRE estimator perform better using pre-trained Inception features, the relative improvement with the CoRE penalty is still 28% on test set 2.

### 5.2.3 Ablation experiments

In Sect. D.3.1, we report results for the following two additional baselines: (i) we group all examples sharing the same class label and penalize with the conditional variance of the

<sup>11</sup> Retrieved from [https://tfhub.dev/google/imagenet/inception\\_v3/feature\\_vector/1](https://tfhub.dev/google/imagenet/inception_v3/feature_vector/1).

**Table 3** Classification for  $Y \in \{\text{woman, man}\}$ 

	Method	Error			Penalty value		
		Train (%)	Test 1 (%)	Test 2 (%)	Train	Test: Females	Test: Males
$\kappa = .5$	5-layer CNN	<b>0.00</b>	<b>2.00</b>	38.54	22.77	74.05	30.67
	+ CoRE	6.43	5.85	<b>24.07</b>	0.01	1.61	0.93
$\kappa = .75$	5-layer CNN	<b>0.00</b>	<b>1.98</b>	43.41	8.23	32.98	11.76
	+ CoRE	7.61	6.99	<b>27.05</b>	0.00	1.44	0.62
$\kappa = .9$	5-layer CNN	<b>0.00</b>	<b>2.00</b>	47.64	9.47	40.51	14.37
	+ CoRE	8.76	7.74%	<b>30.63</b>	0.00	1.26	0.42
$\kappa = .95$	5-layer CNN	<b>0.00</b>	<b>1.89</b>	48.96	13.62	61.01	21.26
	+ CoRE	10.45	9.35	<b>29.57</b>	0.00	0.42	0.16
$\kappa = .99$	5-layer CNN	<b>0.00</b>	<b>1.70</b>	50.11	20.66	70.80	27.80
	+ CoRE	11.10	10.51	<b>32.91</b>	0.00	0.00	0.00
$\kappa = 1$	5-layer CNN	<b>0.00</b>	<b>1.93</b>	49.41	821.32	2524.77	1253.21
	+ CoRE	11.12	10.11	<b>35.68</b>	0.00	0.02	0.01

For each experimental setting, bold marks the best performing method

We compare six different datasets that vary with respect to the distribution of  $Y$  in the grouped observations. Specifically, we vary the proportion of images showing men between  $\kappa = 0.5$  and  $\kappa = 1$ . In all training datasets, the total number of observations is 16,982 and the total number of grouped observations is 500. Both the pooled estimator as well as the CoRE estimator perform better if the distribution of  $Y$  in the grouped observations is more balanced. The CoRE estimator improves the error rate of the pooled estimator by  $\approx 28 - 39\%$  on a relative scale. Table 6 in the Appendix additionally contains the standard error of all shown results

**Table 4** Classification for  $Y \in \{\text{woman, man}\}$  with  $\kappa = 0.5$  Here, we compared  $\ell_2$ -regularized logistic regression based on Inception V3 features with and without the CoRE penalty

Method	Error		
	Train (%)	Test 1 (%)	Test 2 (%)
Inception V3	<b>5.74</b>	<b>5.53</b>	30.29
Inception V3 + CoRE	6.15	5.85	<b>21.70</b>

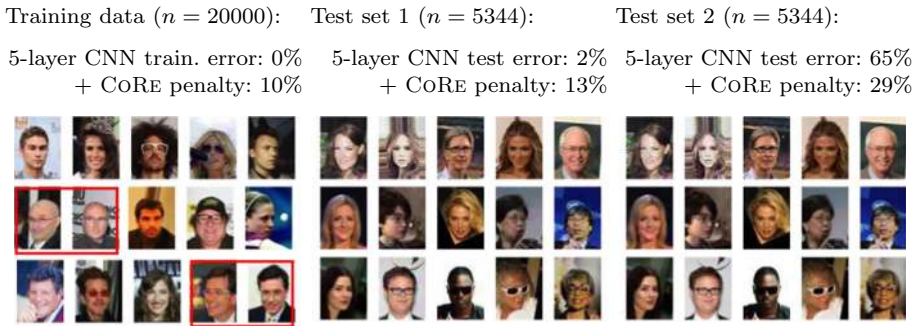
For each experimental setting, bold marks the best performing method  
The CoRE estimator improves the performance of the pooled estimator by  $\approx 28\%$  on a relative scale

predicted logits, computed over these two groups; (ii) we penalize the overall variance of the predicted logits, i.e., a form of unconditional variance regularization.

### 5.3 Eyeglasses detection with known and unknown image quality intervention

We now revisit the second example from Sect. 1.1. We again use the CelebA dataset and consider the problem of classifying whether the person in the image is wearing eyeglasses. Here, we modify the images in the following way: in the training set and in test set 1, we sample the image quality<sup>12</sup> for all samples  $\{i : y_i = 1\}$  (all samples that show glasses) from

<sup>12</sup> We use ImageMagick (<https://www.imagemagick.org>) to change the level of the JPEG compression through `convert -quality qi input.jpg output.jpg` where  $q_{i,j} \sim \mathcal{N}(30, 100)$ .



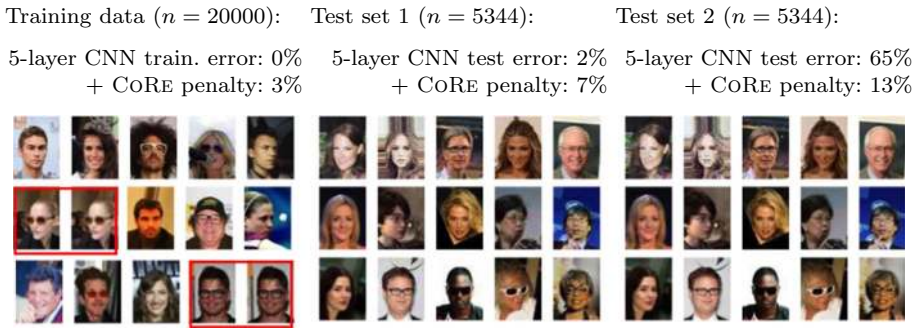
**Fig. 9** Eyeglass detection for CelebA dataset with image quality interventions (which are unknown to any procedure used). The JPEG compression level is lowered for  $Y = 1$  (glasses) samples on training data and test set 1 and lowered for  $Y = 0$  (no glasses) samples for test set 2. To the human eye, these interventions are barely visible but the CNN that uses pooled data without CoRE penalty has exploited the correlation between image quality and outcome  $Y$  to achieve a (arguably spurious) low test error of 2% on test set 1. However, if the correlation between image quality and  $Y$  breaks down, as in test set 2, the CNN that uses pooled data without a CoRE penalty has a 65% misclassification rate. The training data on the left show paired observations in two red boxes: these observations share the same label  $Y$  and show the same person ID. They are used to compute the conditional variance penalty for the CoRE estimator that does not suffer from the same degradation in performance for test set 2

a Gaussian distribution with mean  $\mu = 30$  and standard deviation  $\sigma = 10$ . Samples with  $y_i = 0$  (no glasses) are unmodified. In other words, if the image shows a person wearing glasses, the image quality tends to be lower. In test set 2, the quality is reduced in the same way for  $y_i = 0$  samples (no glasses), while images with  $y_i = 1$  are not changed. Figure 9 shows examples from the training set and test sets 1 and 2. For the CoRE penalty, we calculate the conditional variance across images that share the same ID if  $Y = 1$ , that is across images that show the same person wearing glasses on all images. Observations with  $Y = 0$  (not wearing glasses) are not grouped. Two examples are shown in the red box of Fig. 9. Here, we have  $c = 5000$  grouped observations among a total sample size of  $n = 20,000$ .

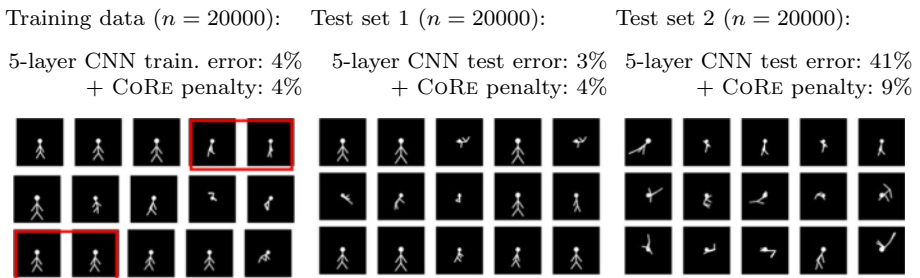
Figure 9 shows misclassification rates for CoRE and the pooled estimator on test sets 1 and 2. The pooled estimator (only penalized with an  $\ell_2$  penalty) achieves low error rates of 2% on test set 1, but suffers from a 65% misclassification error on test set 2, as now the relation between  $Y$  and the implicit  $S$  variable (image quality) has been flipped. The CoRE estimator has a larger error of 13% on test set 1 as image quality as a feature is penalized by CoRE implicitly and the signal is less strong if image quality has been removed as a dimension. However, in test set 2 the performance of the CoRE estimator is 28% and improves substantially on the 65% error of the pooled estimator. The reason is again the same: the CoRE penalty ensures that image quality is not used as a feature to the same extent as for the pooled estimator. This increases the test error slightly if the samples are generated from the same distribution as training data (as here for test set 1) but substantially improves the test error if the distribution of image quality, conditional on the class label, is changed on test data (as here for test set 2).

*Eyeglasses detection with known image quality intervention* To compare to the above results, we repeat the experiment by changing the grouped observations as follows. Above, we grouped images that had the same person ID when  $Y = 1$ . We refer to this scheme of grouping observations with the same  $(Y, ID)$  as ‘Grouping setting 2’. Here, we use an explicit augmentation scheme and augment  $c = 5000$  images with  $Y = 1$  in





**Fig. 10** Eyeglass detection for CelebA dataset with image quality interventions. The only difference to Fig. 9 is in the training data where the paired images now use the same underlying image in two different JPEG compressions. The compression level is drawn from the same distribution. The CoRE penalty performs better than for the experiment in Fig. 9 since we could explicitly control that only  $S \equiv \text{image quality}$  varies between grouped examples. On the other hand, the performance of the pooled estimator is not changed in a noticeable way if we add augmented images as the (spurious) correlation between image quality and outcome  $Y$  still persists in the presence of the extra augmented images. Thus, the pooled estimator continues to be susceptible to image quality interventions



**Fig. 11** Classification into {adult, child} based on stickmen images, where children tend to be smaller and adults taller. In training and test set 1 data, children tend to have stronger movement whereas adults tend to stand still. In test set 2 data, adults show stronger movement. The two red boxes in the panel with the training data show two out of the  $c = 50$  pairs of examples over which the conditional variance is calculated. The CoRE penalty leads to a network that generalizes better for test set 2 data, where the spurious correlation between age and movement is reversed, if compared to the training data

the following way: each image is paired with a copy of itself and the image quality is adjusted as described above. In other words, the only difference between the two images is that image quality differs slightly, depending on the value that was drawn from the Gaussian distribution with mean  $\mu = 30$  and standard deviation  $\sigma = 10$ , determining the strength of the image quality intervention. Both the original and the copy get the same value of identifier variable ID. We call this grouping scheme ‘Grouping setting 1’. Compare the left panels of Figs. 9 and 10 for examples.

While we used explicit changes in image quality in both above and here, we referred to grouping setting 2 as ‘unknown image quality interventions’ as the training sample as in the left panel of Fig. 9 does not immediately reveal that image quality is the important style variable. In contrast, the augmented data samples (grouping setting 1) we use here differ only in their image quality for a constant ( $Y, ID$ ).

Figure 10 shows examples and results. The pooled estimator performs more or less identical to the previous dataset. The explicit augmentation did not help as the association between image quality and whether eyeglasses are worn is not changed in the pooled data after including the augmented data samples. The misclassification error of the CoRE estimator is substantially better than the error rate of the pooled estimator. The error rate on test set 2 of 13% is also improving on the rate of 28% of the CoRE estimator in grouping setting 2. We see that using grouping setting 1 works best since we could explicitly control that only  $S \equiv \textit{image quality}$  varies between grouped examples. In grouping setting 2, different images of the same person can vary in many factors, making it more challenging to isolate image quality as the factor to be invariant against.

A similar example where  $S \equiv \textit{brightness}$  is summarized in “Appendix D.1”.

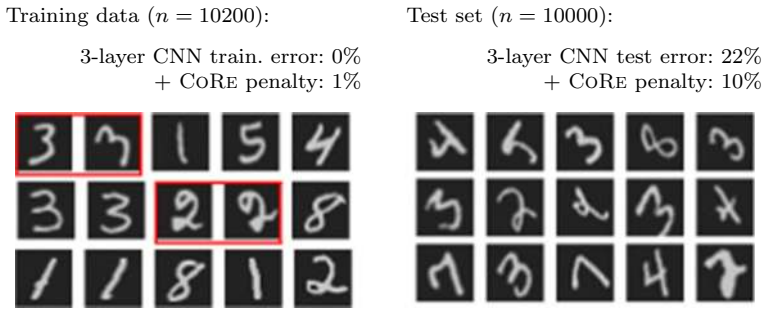
## 5.4 Stickmen image-based age classification with unknown movement interventions

In this example we consider synthetically generated stickmen images; see Fig. 11 for some examples. The target of interest is  $Y \in \{\textit{adult}, \textit{child}\}$ . The core feature  $C$  is here the height of each person. The class  $Y$  is causal for height and height cannot be easily intervened on or change in different domains. Height is thus a robust predictor for differentiating between children and adults. As style feature we have here the movement of a person (distribution of angles between body, arms and legs). For the training data we created a dependence between age and the style feature ‘movement’, which can be thought to arise through a hidden common cause  $D$ , namely the place of observation. For instance, the images of children might mostly show children playing while the images of adults typically show them in more “static” postures. The left panel of Fig. 11 shows examples from the training set where large movements are associated with children and small movements are associated with adults. Test set 1 follows the same distribution, as shown in the middle panel. A standard CNN will exploit this relationship between movement and the label  $Y$  of interest, whereas this is discouraged by the conditional variance penalty of CoRE. The latter is pairing images of the same person in slightly different movements as shown by the red boxes in the leftmost panel of Fig. 11. If the learned model exploits this dependence between movement and age for predicting  $Y$ , it will fail when presented images of, say, dancing adults. The right panel of Fig. 11 shows such examples (test set 2). The standard CNN suffers in this case from a 41% misclassification rate, as opposed to the 3% on test set 1 data. For as few as  $c = 50$  paired observations, the network with an added CoRE penalty, in contrast, achieves also 4% on test set 1 data and succeeds in achieving a 9% performance on test set 2, whereas the pooled estimator fails on this dataset with a test error of 41%.

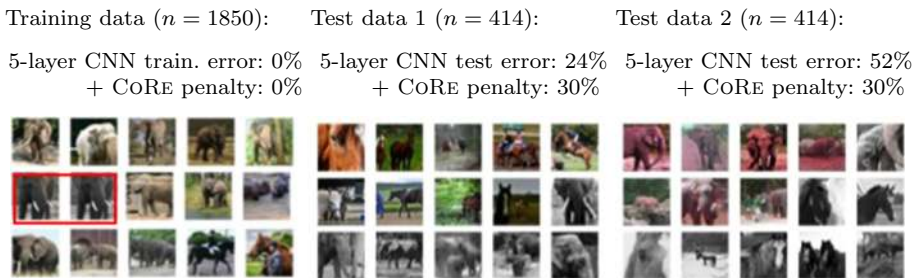
These results suggest that the learned representation of the pooled estimator uses movement as a predictor for age while CoRE does not use this feature due to the conditional variance regularization. Importantly, including more grouped examples would not improve the performance of the pooled estimator as these would be subject to the same bias and hence also predominantly have examples of heavily moving children and “static” adults (also see Fig. 23 which shows results for  $c \in \{20, 500, 2000\}$ ).

## 5.5 MNIST: more sample efficient data augmentation

The goal of using CoRE in this example is to make data augmentation more efficient in terms of the required samples. In data augmentation, one creates additional samples



**Fig. 12** Data augmentation for MNIST images. The left shows training data with a few rotated images. Evaluating on only rotated images from the test set, a standard network achieves only 22% accuracy. We can add the CoRE penalty by computing the conditional variance over images that were generated from the same original image. The test error is then lowered to 10% on the test data of rotated images



**Fig. 13** Elmer-the-Elephant dataset. The left panel shows training data with a few additional grayscale elephants. The pooled estimator learns that color is predictive for the animal class and achieves test error of 24% on test set 1 where this association is still true but suffers a misclassification error of 53% on test set 2 where this association breaks down. By adding the CoRE penalty, the test error is consistently around 30%, irrespective of the color distribution of horses and elephants

by modifying the original inputs, e.g. by rotating, translating, or flipping the images (Schölkopf et al. 1996). In other words, additional samples are generated by interventions on style features. Using this augmented data set for training results in invariance of the estimator with respect to the transformations (style features) of interest. For CoRE we can use the grouping information that the original and the augmented samples belong to the same object. This enforces the invariance with respect to the style features more strongly compared to normal data augmentation which just pools all samples. We assess this for the style feature ‘rotation’ on MNIST (LeCun et al. 1998) and only include  $c = 200$  augmented training examples for  $m = 10,000$  original samples, resulting in a total sample size of  $n = 10200$ . The degree of the rotations is sampled uniformly at random from  $[35, 70]$ . Figure 12 shows examples from the training set. By using CoRE the average test error on rotated examples is reduced from 22% to 10%. Very few augmented sample are thus sufficient to lead to stronger rotational invariance. The standard approach of creating augmented data and pooling all images requires, in contrast, many more samples to achieve the same effect. Additional results for  $m \in \{1000, 10,000\}$  and  $c$  ranging from 100 to 5000 can be found in Fig. 22 in Appendix Sect. D.5.

**Table 5** Details of the model architectures used

Dataset	Optimizer		Architecture
MNIST	Adam	Input	$28 \times 28 \times 1$
		CNN	Conv $5 \times 5 \times 16$ , $5 \times 5 \times 32$ (same padding, strides = 2, ReLu activation), fully connected, softmax layer
Stickmen	Adam	Input	$64 \times 64 \times 1$
		CNN	Conv $5 \times 5 \times 16$ , $5 \times 5 \times 32$ , $5 \times 5 \times 64$ , $5 \times 5 \times 128$ (same padding, strides = 2, leaky ReLu activation), fully connected, softmax layer
CelebA (all experiments using CelebA)	Adam	Input	$64 \times 48 \times 3$
		CNN	Conv $5 \times 5 \times 16$ , $5 \times 5 \times 32$ , $5 \times 5 \times 64$ , $5 \times 5 \times 128$ (same padding, strides = 2, leaky ReLu activation), fully connected, softmax layer
AwA2	Adam	Input	$32 \times 32 \times 3$
		CNN	Conv $5 \times 5 \times 16$ , $5 \times 5 \times 32$ , $5 \times 5 \times 64$ , $5 \times 5 \times 128$ (same padding, strides = 2, leaky ReLu activation), fully connected, softmax layer

## 5.6 Elmer the Elephant

In this example, we want to assess whether invariance with respect to the style feature ‘color’ can be achieved. In the children’s book ‘Elmer the elephant’<sup>13</sup> one instance of a colored elephant suffices to recognize it as being an elephant, making the color ‘gray’ no longer an integral part of the object ‘elephant’. Motivated by this process of concept formation, we would like to assess whether CoRE can exclude ‘color’ from its learned representation by penalizing conditional variance appropriately.

We work with the ‘Animals with attributes 2’ (AwA2) dataset (Xian et al. 2017) and consider classifying images of horses and elephants. We include additional examples by adding grayscale images for  $c = 250$  images of elephants. These additional examples do not distinguish themselves strongly from the original training data as the elephant images are already close to grayscale images. The total training sample size is 1850.

Figure 13 shows examples and misclassification rates from the training set and test sets for CoRE and the pooled estimator on different test sets. Examples from these and more test sets can be found in Fig. 24. Test set 1 contains original, colored images only. In test set 2 images of horses are in grayscale and the colorspace of elephant images is modified, effectively changing the color gray to red-brown. We observe that the pooled estimator does not perform well on test set 2 as its learned representation seems to exploit the fact that ‘gray’ is predictive for ‘elephant’ in the training set. This association is no longer valid for test set 2. In contrast, the predictive performance of CoRE is hardly affected by the changing color distributions. More details can be found in “Appendix D.7”.

It is noteworthy that a colored elephant can be recognized as an elephant by adding a few examples of a grayscale elephant to the very lightly colored pictures of natural

<sup>13</sup> [https://en.wikipedia.org/wiki/Elmer\\_the\\_Patchwork\\_Elephant](https://en.wikipedia.org/wiki/Elmer_the_Patchwork_Elephant).

elephants. If we just pool over these examples, there is still a strong bias that elephants are gray. The CoRE estimator, in contrast, demands invariance of the prediction for instances of the same elephant and we can learn color invariance with a few added grayscale images.

## 6 Further related work

Encoding certain invariances in estimators is a well-studied area in computer vision and machine learning with an extensive body of literature. While a large part of this work assumes the desired invariance to be known, fewer approaches aim to learn the required invariances from data and the focus often lies on geometric transformations of the input data or explicitly creating augmented observations (Sohn and Lee 2012; Khasanova and Frossard 2017; Hashimoto et al. 2017; Devries and Taylor 2017). The main difference between this line of work and CoRE is that we do not require to know the style feature explicitly, the set of possible style features is not restricted to a particular class of transformations and we do not aim to create augmented observations in a generative framework.

Recently, various approaches have been proposed that leverage causal motivations for deep learning or use deep learning for causal inference, related to e.g. the problems of cause-effect inference and generative adversarial networks (Chalupka et al. 2014; Lopez-Paz et al. 2017; Lopez-Paz and Oquab 2017; Goudet et al. 2017; Bahadori et al. 2017; Besserve et al. 2018; Kocaoglu et al. 2018).

Kilbertus et al. (2017) exploit causal reasoning to characterize fairness considerations in machine learning. Distinguishing between the protected attribute and its proxies, they derive causal non-discrimination criteria. The resulting algorithms avoiding proxy discrimination require classifiers to be constant as a function of the proxy variables in the causal graph, thereby bearing some structural similarity to our style features.

Distinguishing between core and style features can be seen as some form of disentangling factors of variation. Estimating disentangled factors of variation has gathered a lot of interest in the context of generative modeling. As in CoRE, Bouchacourt et al. (2018) exploit grouped observations. In a variational autoencoder framework, they aim to separate style and content—they assume that samples within a group share a common but unknown value for one of the factors of variation while the style can differ. Denton and Birodkar (2017) propose an autoencoder framework to disentangle style and content in videos using an adversarial loss term where the grouping structure induced by clip identity is exploited. Here we try to solve a classification task directly without estimating the latent factors explicitly as in a generative framework.

In the computer vision literature, various works have used identity information to achieve pose invariance in the context of face recognition (Bartlett and Sejnowski 1996; Tran et al. 2017). More generally, the idea of exploiting various observations of the same underlying object is related to multi-view learning (Xu et al. 2013). In the context of adversarial examples, Kannan et al. (2018) recently proposed the defense “Adversarial logit pairing” which is methodologically equivalent to the CoRE penalty  $C_{f,1,\theta}$  when using the squared error loss. Several empirical studies have shown mixed results regarding the performance on  $\ell_\infty$  perturbations (Engstrom et al. 2018; Mosbach et al. 2018), so far this setting has not been analyzed theoretically and hence it is an open question whether a CoRE-type penalty constitutes an effective defense against adversarial examples.

## 7 Conclusion

Distinguishing the latent features in an image into CoRE and *style* features, we have proposed conditional variance regularization (CoRE) to achieve robustness with respect to interventions on the style or “orthogonal” features. The main idea of the CoRE estimator is to exploit the fact that we often have instances of the same object in the training data. By demanding invariance of the classifier amongst a group of instances that relate to the same object, we can achieve invariance of the classification performance with respect to interventions on style features such as image quality, fashion type, color, or body posture. The training also works despite sampling biases in the data.

There are two main application areas:

1. If the style features are known explicitly, we can achieve the same classification performance as standard data augmentation approaches with substantially fewer augmented samples, as shown for example in Sect. 5.5.
2. Perhaps more interesting are settings in which it is unknown what the style features are, with examples in Sects. 5.1, 5.2, 5.3, 5.4 and “Appendix D.1”. CoRE regularization forces predictions to be based on features that do not vary strongly between instances of the same object. We could show in the examples and in Theorems 1 and 2 that this regularization achieves distributional robustness with respect to changes in the distribution of the (unknown) style variables.

An interesting line of work would be to use larger models such as Inception or large ResNet architectures (Szegedy et al. 2015; He et al. 2016). These models have been trained to be invariant to an array of explicitly defined style features. In Sect. 5.2.2 we include results which show that using Inception V3 features does not guard against interventions on more implicit style features. We would thus like to assess what benefits CoRE can bring for training Inception-style models end-to-end, both in terms of sample efficiency and in terms of generalization performance. While we showed some examples where the necessary grouping information is available, an interesting possible future direction would be to use video data since objects display temporal constancy and the temporal information can hence be used for grouping and conditional variance regularization. Beyond that our results show that it can be worthwhile to collect ID information when new datasets are created. As CoRE only requires a subset of the observations to have ID annotations, in many cases this information might be cheap to collect while it can improve performance substantially when future test data is subject to domain shifts.

**Acknowledgements** We thank Brian McWilliams, Jonas Peters, and Martin Arjovsky for helpful comments and discussions and CSCS for provision of computational resources. A preliminary version of this work was presented at the NIPS 2017 Interpretable ML Symposium and we thank participants of the symposium for very helpful discussions.

**Funding** Open access funding provided by Swiss Federal Institute of Technology Zurich.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

### Proof of Theorem 1

**First part** To show the first part, namely that with probability 1,

$$L_\infty(\hat{\theta}^{pool}) = \infty,$$

we need to show that  $W^t \hat{\theta}^{pool} \neq 0$  with probability 1. The reason this is sufficient is as follows: if  $W^t \theta \neq 0$ , then  $L_\infty(\theta) = \infty$  as we can then find a  $v \in \mathbb{R}^q$  such that  $\gamma := \theta^t W v \neq 0$ . Assume without limitation of generality that  $v$  is normed such that  $E(E(v^t \sum_{y, \text{id}}^{-1} v | Y = y, \text{ID} = \text{id})) = 1$ . Setting  $\Delta_\xi = \xi v$  for  $\xi \in \mathbb{R}$ , we have that  $(\text{ID}, Y, S + \Delta_\xi)$  is in the class  $F_{|\xi|}$  if the distribution of  $(\text{ID}, Y, S)$  is equal to  $F_0$ . Furthermore,  $x(\Delta_\xi)^t \theta = x(\Delta = 0)^t \theta + \xi \gamma$ . Hence  $\log(1 + \exp(-y \cdot x(\Delta_\xi)^t \theta)) \rightarrow \infty$  for either  $\xi \rightarrow \infty$  or  $\xi \rightarrow -\infty$ .

To show that  $W^t \hat{\theta}^{pool} \neq 0$  with probability 1, let  $\hat{\theta}^*$  be the oracle estimator that is constrained to be orthogonal to the column space of  $W$ :

$$\hat{\theta}^* = \underset{\theta: W^t \theta = 0}{\operatorname{argmin}} L_n(\theta) \quad \text{with} \quad L_n(\theta) := \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i)). \tag{14}$$

We show  $W^t \hat{\theta}^{pool} \neq 0$  by contradiction. Assume hence that  $W^t \hat{\theta}^{pool} = 0$ . If this is indeed the case, then the constraint  $W^t \theta = 0$  in (14) becomes non-active and we have  $\hat{\theta}^{pool} = \hat{\theta}^*$ . This would imply that taking the directional derivative of the training loss with respect to any  $\delta \in \mathbb{R}^p$  in the column space of  $W$  should vanish at the solution  $\hat{\theta}^*$ . In other words, define the gradient as  $g(\theta) = \nabla_\theta L_n(\theta) \in \mathbb{R}^p$ . The implication is then that for all  $\delta$  in the column-space of  $W$ ,

$$\delta^t g(\hat{\theta}^*) = 0 \tag{15}$$

and we will show the latter condition is violated almost surely.

As we work with the logistic loss and  $\mathcal{Y} \in \{-1, 1\}$ , the loss is given by  $\ell(y_i, f_\theta(x_i)) = \log(1 + \exp(-y_i x_i^t \theta))$ . Define  $r_i(\theta) := y_i / (1 + \exp(y_i x_i^t \theta))$ . For all  $i = 1, \dots, n$  we have  $r_i \neq 0$ . Then

$$g(\hat{\theta}^*) = \frac{1}{n} \sum_{i=1}^n r_i(\hat{\theta}^*) x_i. \tag{16}$$

The training images can be written according to the model as  $x_i = x_i^0 + W s_i$ , where  $X^0 := k_x(C, \epsilon_X)$  are the images in absence of any style variation. Since the style features only have an effect on the column space of  $W$  in  $X$ , the oracle estimator  $\hat{\theta}^*$  is identical under the true training data and the (hypothetical) training data  $x_i^0, i = 1, \dots, n$  in absence of style variation. As  $X - X^0 = WS$ , Eq. (16) can also be written as

$$\delta^t g(\hat{\theta}^*) = \frac{1}{n} \sum_{i=1}^n r_i(\hat{\theta}^*) (x_i^0)^t \delta + \frac{1}{n} \sum_{i=1}^n r_i(\hat{\theta}^*) (s_i)^t W^t \delta. \tag{17}$$

Since  $\delta$  is in the column-space of  $W$ , there exists  $u \in \mathbb{R}^q$  such that  $\delta = Wu$  and we can write (17) as

$$\delta^t g(\hat{\theta}^*) = \frac{1}{n} \sum_{i=1}^n r_i(\hat{\theta}^*) (x_i^0)^t Wu + \frac{1}{n} \sum_{i=1}^n r_i(\hat{\theta}^*) (s_i)^t W^t Wu. \tag{18}$$

From (A2) we have that the eigenvalues of  $W^tW$  are all positive. Also  $r_i(\hat{\theta}^*)$  is not a function of the interventions  $s_i, i = 1, \dots, n$  since, as above, the estimator  $\hat{\theta}^*$  is identical whether trained on the original data  $x_i$  or on the intervention-free data  $x_i^0, i = 1, \dots, n$ . If we condition on everything except for the random interventions by conditioning on  $(x_i^0, y_i)$  for  $i = 1, \dots, n$ , then the rhs of (18) can be written as

$$a^t u + B^t u,$$

where  $a \in \mathbb{R}^q$  is fixed (conditionally) and  $B = \frac{1}{n} \sum_{i=1}^n r_i(\hat{\theta}^*)(s_i)^t W^t W \in \mathbb{R}^q$  is a random vector and  $B \neq -a \in \mathbb{R}^q$  with probability 1 by (A1) and (A2). Hence the left hand side of (18) is not identically 0 with probability 1 for any given  $\delta$  in the column-space of  $W$ . This shows that the implication (15) is incorrect with probability 1 and hence completes the proof of the first part by contradiction.

**Invariant parameter space** Before continuing with the second part of the proof, some definitions. Let  $I$  be the *invariant parameter space*

$$I := \{ \theta : f_\theta(x(\Delta)) \text{ is constant as function of } \Delta \in \mathbb{R}^q \text{ for all } x \in \mathbb{R}^p \}.$$

For all  $\theta \in I$ , the loss (6) for any  $F \in F_\xi$  is identical to the loss under  $F_0$ . That is for all  $\xi \geq 0$ ,

$$\text{if } \theta \in I, \text{ then } \sup_{F \in F_\xi} E_F \left[ \ell(Y, f_\theta(X)) \right] = E_{F_0} \left[ \ell(Y, f_\theta(X)) \right].$$

The optimal predictor in the invariant space  $I$  is

$$\theta^* = \operatorname{argmin}_\theta E_{F_0} \left[ \ell(Y, f_\theta(X)) \right] \text{ such that } \theta \in I. \tag{19}$$

If  $f_\theta$  is only a function of the core features  $C$ , then  $\theta \in I$ . The challenge is that the core features are not directly observable and we have to infer the invariant space  $I$  from data.

**Second part** For the second part, we first show that with probability at least  $p_n$ , as defined in (A3),  $\hat{\theta}^{core} = \hat{\theta}^*$  with  $\hat{\theta}^*$  defined as in (14). The invariant space for this model is the linear subspace  $I = \{ \theta : W^t \theta = 0 \}$  and by their respective definitions,

$$\begin{aligned} \hat{\theta}^* &= \operatorname{argmin}_\theta \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i)) \text{ such that } \theta \in I, \\ \hat{\theta}^{core} &= \operatorname{argmin}_\theta \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i)) \text{ such that } \theta \in I_n. \end{aligned}$$

Since we use  $I_n = I_n(\tau)$  with  $\tau = 0$ ,

$$I_n = \{ \theta : \hat{E}(\hat{\text{Var}}(f_\theta(X)|Y, \text{ID})) = 0 \}.$$

This implies that for  $\theta \in I_n, f_\theta(x_i) = f_\theta(x_{i'})$  if  $i, i' \in G_j$  for some  $j \in \{1, \dots, m\}$ .<sup>14</sup> Since  $f_\theta(x) = f_\theta(x')$  implies  $(x - x')^t \theta = 0$ , it follows that  $(x_i - x_{i'})^t \theta = 0$  if  $i, i' \in G_j$  for some  $j \in \{1, \dots, m\}$  and hence

<sup>14</sup> Recall that  $(y_i, \text{id}_i) = (y_{i'}, \text{id}_{i'})$  if  $i, i' \in G_j$  as the subsets  $G_j, j = 1, \dots, m$ , collect all observations that have a unique realization of  $(Y, \text{ID})$



$$I_n \subseteq \{ \theta : (x_i - x_{i'})^t \theta = 0 \text{ if } i, i' \in G_j \text{ for some } j \in \{1, \dots, m\} \}.$$

Since  $S$  has a linear influence on  $X$  in (11),  $x_i - x_{i'} = W(\Delta_i - \Delta_{i'})$  if  $i, i'$  are in the same group  $G_j$  of observations for some  $j \in \{1, \dots, m\}$ . Note that the number of grouped examples  $n - m$  is equal to or exceeds the rank  $q$  of  $W$  with probability  $p_n$ , using (A3), and  $p_n \rightarrow 1$  for  $n \rightarrow \infty$ . By (A2), it follows then with probability at least  $p_n$  that  $I_n \subseteq \{ \theta : W^t \theta = 0 \} = I$ . As, by definition,  $I \subseteq I_n$  is always true, we have with probability  $p_n$  that  $I = I_n$ . Hence, with probability  $p_n$  (and  $p_n \rightarrow 1$  for  $n \rightarrow \infty$ ),  $\hat{\theta}^{core} = \hat{\theta}^*$ . It thus remains to be shown that

$$L_\infty(\hat{\theta}^*) \rightarrow_p \inf_\theta L_\infty(\theta). \tag{20}$$

Since  $\hat{\theta}^*$  is in  $I$ , we have  $\ell(y, x(\Delta)) = \ell(y, x^0)$ , where  $x^0$  are the previously defined data in absence of any style variance. Hence

$$\hat{\theta}^* = \operatorname{argmin}_\theta \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i^0)) \text{ such that } \theta \in I, \tag{21}$$

that is the estimator is unchanged if we use the (hypothetical) data  $x_i^0, i = 1, \dots, n$  as training data. The population optimal parameter vector defined in (19) as

$$\theta^* = \operatorname{argmin}_\theta E_{F_0} \left[ \ell(Y, f_\theta(X)) \right] \text{ such that } \theta \in I. \tag{22}$$

is for all  $\xi \geq 0$  identical to

$$\operatorname{argmin}_\theta \sup_{F \in \mathcal{F}_\xi} E_F \left[ \ell(Y, f_\theta(X)) \right] \text{ such that } \theta \in I.$$

Hence (21) and (22) can be written as

$$\begin{aligned} \hat{\theta}^* &= \operatorname{argmin}_{\theta: \theta \in I} L_n^{(0)}(\theta) \text{ with } L_n^{(0)}(\theta) := \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i^0)) \\ \theta^* &= \operatorname{argmin}_{\theta: \theta \in I} L^{(0)}(\theta) \text{ with } L^{(0)}(\theta) := E[\ell(Y, f_\theta(X))]. \end{aligned}$$

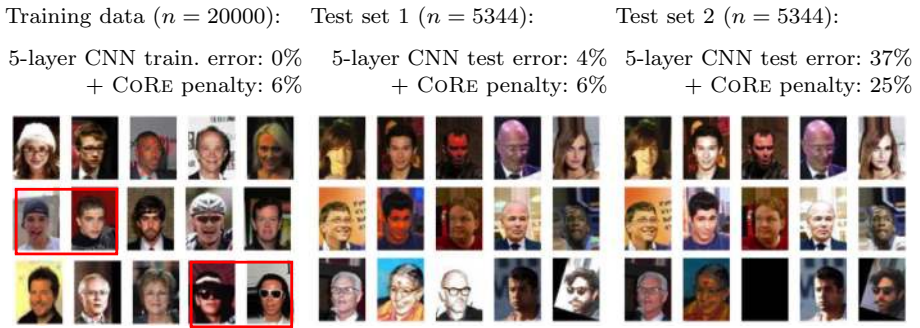
By uniform convergence of  $L_n^{(0)}$  to the population loss  $L^{(0)}$ , we have  $L^{(0)}(\hat{\theta}^*) \rightarrow_p L^{(0)}(\theta^*)$ . By definition of  $I$  and  $\theta^*$ , we have  $L_\infty^* = L_\infty(\theta^*) = L^{(0)}(\theta^*)$ . As  $\hat{\theta}^*$  is in  $I$ , we also have  $L_\infty(\hat{\theta}^*) = L^{(0)}(\hat{\theta}^*)$ . Since, from above,  $L^{(0)}(\hat{\theta}^*) \rightarrow_p L^{(0)}(\theta^*)$ , this also implies  $L_\infty(\hat{\theta}^*) \rightarrow_p L_\infty(\theta^*) = L_\infty^*$ . Using the previously established result that  $\hat{\theta}^{core} = \hat{\theta}^*$  with probability at least  $p_n$  and  $p_n \rightarrow 1$  for  $n \rightarrow \infty$ , this completes the proof.

### Proof of Theorem 2

Let  $F_0$  be the training distribution of  $(ID, Y, S)$  and  $F$  a distribution for  $(ID, Y, \tilde{S})$  in  $\mathcal{F}_\xi$ . By definition of  $\mathcal{F}_\xi$ , we can write  $\tilde{S} = S + \Delta$  for a suitable random variable  $\Delta \in \mathbb{R}^q$  with

$$\Delta \in \mathcal{U}_\xi, \quad \text{where } \mathcal{U}_\xi = \{ \Delta : E(E(\Delta^t \Sigma_{Y, ID}^{-1} \Delta | Y, ID)) \leq \xi \}.$$

Vice versa: if we can write  $\tilde{S} = S + \Delta$  with  $\Delta \in \mathcal{U}_\xi$ , then the distribution is in  $F_\xi$ . While  $X$  under  $F_0$  can be written as  $X(\Delta = 0)$ , the distribution of  $X$  under  $F$  is of the form  $X(\Delta)$



**Fig. 14** Eyeglass detection for CelebA dataset with brightness interventions (which are unknown to any procedure used). On training data and test set 1 data, images where people wear glasses tend to be brighter whereas on test set 2 images where people do not wear glasses tend to be brighter

or, alternatively,  $X(\sqrt{\xi}U)$  with  $U \in \mathcal{U}_1$ . Adopting from now on the latter constraint that  $U \in \mathcal{U}_1$ , and using (B2),

$$E_F \left[ \ell(Y, f_\theta(X)) \right] = E_{F_0} \left[ h_\theta(0) \right] + \sqrt{\xi} E_{F_0} \left[ (\nabla h_\theta)^t U \right] + o(\xi),$$

where  $\nabla h_\theta$  is the gradient of  $h_\theta(\delta)$  with respect to  $\delta$ , evaluated at  $\delta \equiv 0$ . Hence

$$\sup_{F \in \mathcal{F}_\xi} E_F \left[ h_\theta(\Delta) \right] = E_{F_0} \left[ h_\theta(0) \right] + \sqrt{\xi} \sup_{U \in \mathcal{U}_1} E_{F_0} \left[ (\nabla h_\theta)^t U \right] + o(\xi).$$

The proof is complete if we can show that

$$C_{\ell,1/2,\theta} = \sup_{U \in \mathcal{U}_1} E_{F_0} \left[ (\nabla h_\theta)^t U \right] + O(\zeta).$$

On the one hand,

$$\sup_{U \in \mathcal{U}_1} E_{F_0} \left[ (\nabla h_\theta)^t U \right] = E_{F_0} \left[ \sqrt{(\nabla h_\theta)^t \Sigma_{Y, \text{ID}} (\nabla h_\theta)} \right].$$

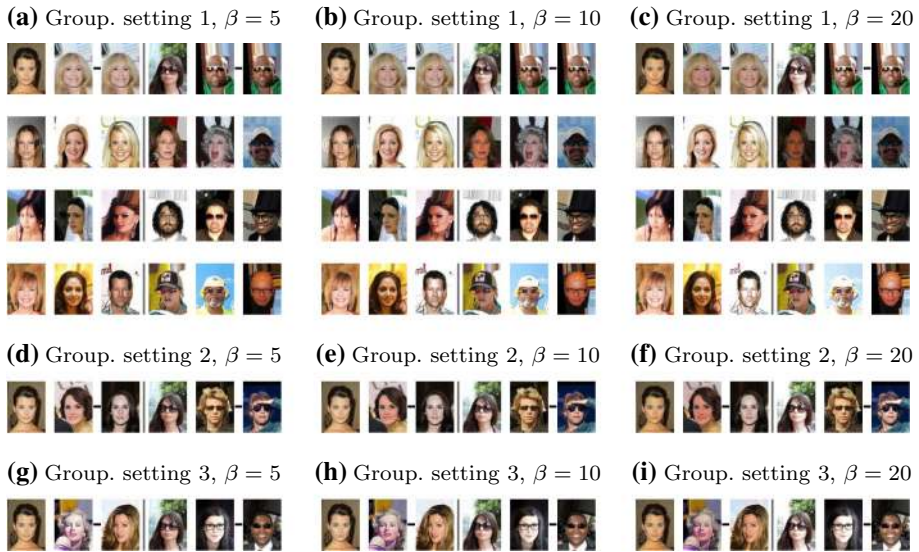
This follows for a matrix  $\Sigma$  with Cholesky decomposition  $\Sigma = V^t V$ ,

$$\begin{aligned} \max_{u: u^t \Sigma^{-1} u \leq 1} (\nabla h_\theta)^t u &= \max_{w: \|w\|_2 \leq 1} (\nabla h_\theta)^t V^t w \\ &= \|V(\nabla h)\|_2 = \sqrt{(\nabla h)^t \Sigma (\nabla h)}. \end{aligned}$$

On the other hand, the conditional-variance-of-loss can be expanded as

$$C_{\ell,1/2,\theta} = E_{F_0} \left[ \sqrt{\text{Var}(\ell(Y, f_\theta(X)) | Y, \text{ID})} \right] = E_{F_0} \left[ \sqrt{(\nabla h_\theta)^t \Sigma_{Y, \text{ID}} (\nabla h_\theta)} \right] + O(\zeta),$$

which completes the proof.



**Fig. 15** Examples from the CelebA eyeglasses detection with brightness interventions, grouping settings 1–3 with  $\beta \in \{5, 10, 20\}$ . In all rows, the first three images from the left have  $y \equiv \text{no glasses}$ ; the remaining three images have  $y \equiv \text{glasses}$ . Connected images are grouped examples. In panels **a–c**, row 1 shows examples from the training set, rows 2–4 contain examples from test sets 2–4, respectively. Panels **d–i** show examples from the respective training sets

## Network architectures

We implemented the considered models in TensorFlow (Abadi et al. 2015). The model architectures used are detailed in Table 5. CoRECoRE and the pooled estimator use the same network architecture and training procedure; merely the loss function differs by the CoRE regularization term. In all experiments we use the Adam optimizer (Kingma and Ba 2015). All experimental results are based on training the respective model five times (using the same data) to assess the variance due to the randomness in the training procedure. In each epoch of the training, the training data  $x_i, i = 1, \dots, n$  are randomly shuffled, keeping the grouped observations  $(x_i)_{i \in I_j}$  for  $j \in \{1, \dots, m\}$  together to ensure that mini batches will contain grouped observations. In all experiments the mini batch size is set to 120. For small  $c$  this implies that not all mini batches contain grouped observations, making the optimization more challenging.

## Additional experiments

### Eyeglasses detection: known and unknown brightness interventions

As in Sect. 5.3 we work with the CelebA dataset and try to classify whether the person in the image is wearing eyeglasses. Here we analyze a confounded setting that could arise as follows. Say the hidden common cause  $D$  of  $Y$  and  $S$  is a binary variable and indicates

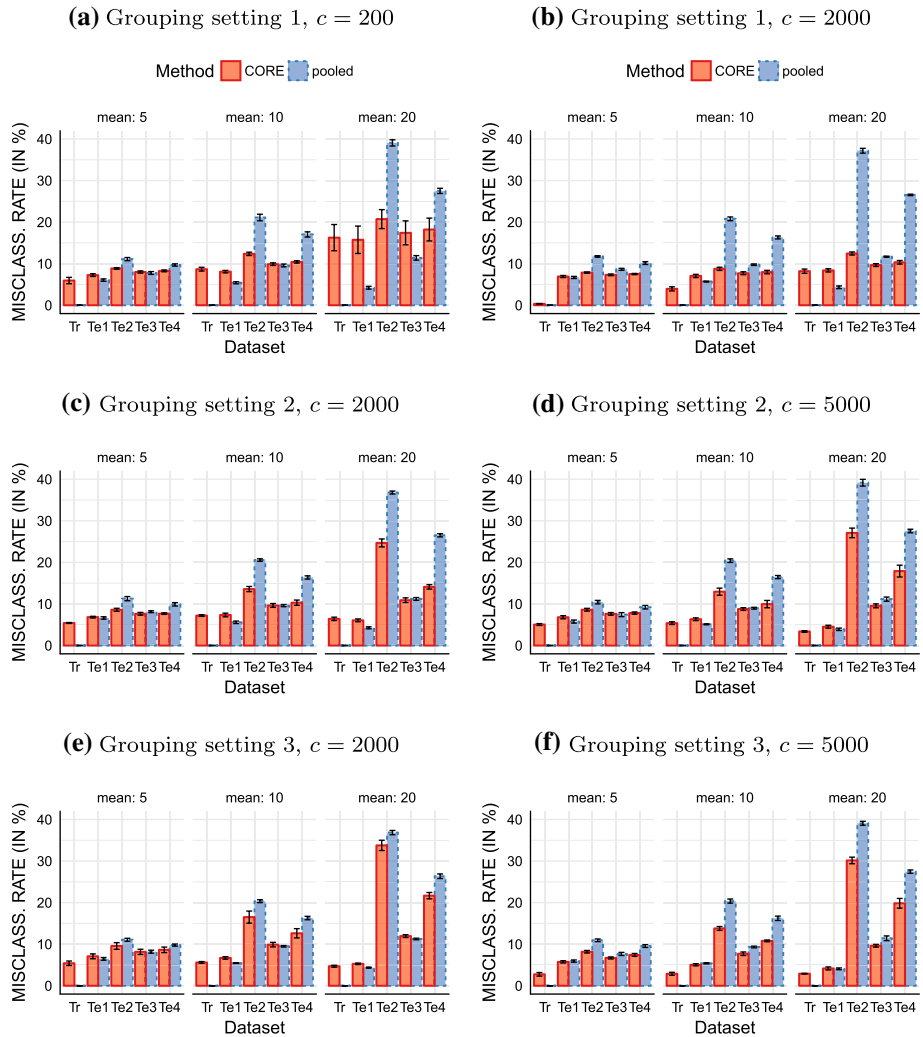
whether the image was taken outdoors or indoors. If it was taken outdoors, then the person tends to wear (sun-)glasses more often and the image tends to be brighter. If the image was taken indoors, then the person tends not to wear (sun-)glasses and the image tends to be darker. In other words, the style variable  $S$  is here equivalent to brightness and the structure of the data generating process is equivalent to the one shown in Fig. 3. Figure 14 shows examples from the training set and test sets. As previously, we compute the conditional variance over images of the same person, sharing the same class label (and the CoRE estimator is hence not using the knowledge that brightness is important). Two alternatives for constructing grouped observations in this setting are discussed further below. We use  $c = 2000$  and  $n = 20,000$ . For the brightness intervention, we sample the value for the magnitude of the brightness increase resp. decrease from an exponential distribution with mean  $\beta = 20$ . In the training set and test set 1, we sample the brightness value as  $b_{ij} = [100 + y_i e_{ij}]_+$  where  $e_{ij} \sim \text{Exp}(\beta^{-1})$  and  $y_i \in \{-1, 1\}$ , where  $y_i = 1$  indicates presence of glasses and  $y_i = -1$  indicates absence.<sup>15</sup> For test set 2, we use instead  $b_{ij} = [100 - y_i e_{ij}]_+$ , so that the relation between brightness and glasses is flipped.

Figure 14 shows misclassification rates for CoRE and the pooled estimator on different test sets. Examples from all test sets can be found in Fig. 15. First, we notice that the pooled estimator performs better than CoRE on test set 1. This can be explained by the fact that it can exploit the predictive information contained in the brightness of an image while CoRE is restricted not to do so. Second, we observe that the pooled estimator does not perform well on test set 2 as its learned representation seems to use the image's brightness as a predictor for the response which fails when the brightness distribution in the test set differs significantly from the training set. In contrast, the predictive performance of CoRE is hardly affected by the changing brightness distributions.

We now discuss two alternatives for constructing different test sets and we vary the number of grouped observations in  $c \in \{200, 2000, 5000\}$  as well as the strength of the brightness interventions in  $\beta \in \{5, 10, 20\}$ , all with sample size  $n = 20,000$ . Generation of training and test sets 1 and 2 were already described above. Here, we consider additionally test set 3 where all images are left unchanged (no brightness interventions at all) and in test set 4 the brightness of all images is increased. Furthermore, we consider three different ways of grouping images. Above, we used images of the same person to create a grouped observation by sampling a different value for the brightness intervention. We refer to this as 'Grouping setting 2' here. An alternative is to use the same image of the same person in different brightnesses (drawn from the same distribution) as a group over which the conditional variance is calculated. We call this 'Grouping setting 1' and it can be useful if we know that we want to protect against brightness interventions in the future. For comparison, we also evaluate grouping with an image of a different person (but sharing the same class label) as a baseline ('Grouping setting 3'). Examples from the training sets using grouping settings 1, 2 and 3 can be found in Fig. 15.

Results for all grouping settings,  $\beta \in \{5, 10, 20\}$  and  $c \in \{200, 5000\}$  can be found in Fig. 16. We see that using grouping setting 1 works best since we could explicitly control that only  $S \equiv \textit{brightness}$  varies between grouping examples. In grouping setting 2, different images of the same person can vary in many factors, making it more challenging to isolate brightness as the factor to be invariant against. Lastly, we see that if we group images of

<sup>15</sup> Specifically, we use ImageMagick (<https://www.imagemagick.org>) and modify the brightness of each image by applying the command `convert -modulate b_ij,100,100 input.jpg output.jpg` to the image.



**Fig. 16** Misclassification rates for the CelebA eyeglasses detection with brightness interventions, grouping settings 1–3 with  $c \in \{200, 2000, 5000\}$  and the mean of the exponential distribution  $\beta \in \{5, 10, 20\}$

different persons (‘Grouping setting 3’), the difference between CoRE estimator and the pooled estimator becomes much smaller than in the previous settings. Figure 17 shows some examples of misclassified observations for Grouping setting 1.

**Eyeglasses detection with small sample size**

Figure 18 shows the numerator and the denominator of the variance ratio defined in Eq. (13) separately as a function of the CoRE penalty weight. In conjunction with Fig. 6b, we observe that a ridge penalty decreases both the within- and between-group variance while the CoRE penalty penalizes the within-group variance selectively.

(a) Examples of misclassified observations.

$y \equiv \textit{glasses}$	$y \equiv \textit{no glasses}$	$y \equiv \textit{glasses}$
$\hat{P}^{core}(gl.) = 1.00$	$\hat{P}^{core}(no\ gl.) = 0.84$	$\hat{P}^{core}(gl.) = 0.90$
$\hat{P}^{pool}(gl.) = 0.21$	$\hat{P}^{pool}(no\ gl.) = 0.13$	$\hat{P}^{pool}(gl.) = 0.14$



Fig. 17 CelebA eyeglasses detection with brightness interventions, grouping setting 1. Examples of misclassified observations from the test sets

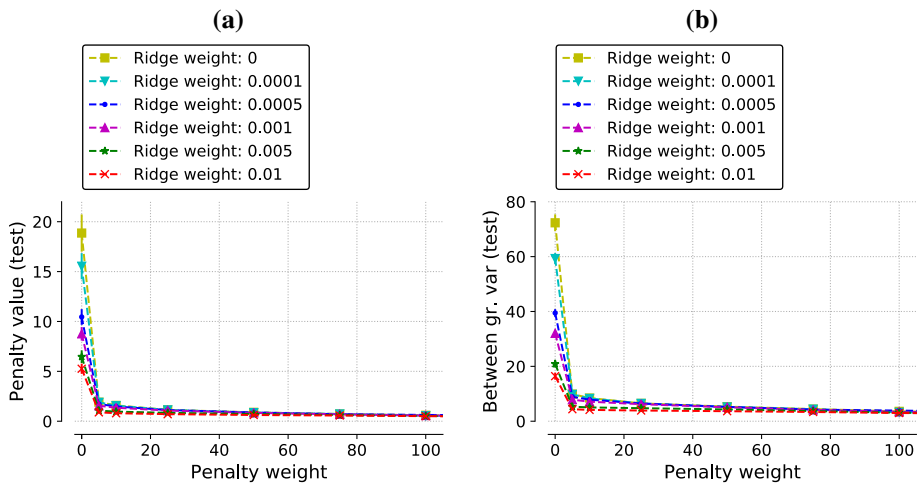
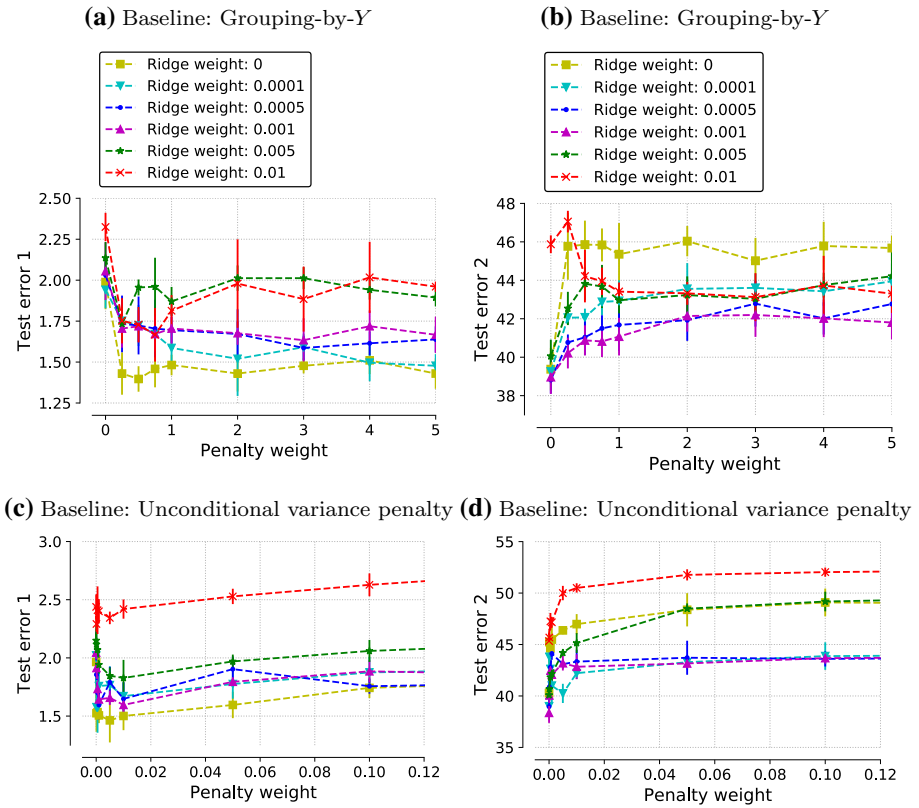


Fig. 18 Eyeglass detection, trained on a small subset (DS1) of the CelebA dataset with disjoint identities. Panel a shows the numerator of the variance ratio defined in Eq. (13) on test data as a function of both the CoRE and ridge penalty weights. Panel b shows the equivalent plot for the denominator. A ridge penalty decreases both the within- and between-group variance while the CoRE penalty penalizes the within-group variance selectively (the latter can be seen more clearly in Fig. 6b)

### Gender classification

#### Additional baselines: Unconditional variance regularization and grouping by class label

As additional baselines, we consider the following two schemes: (i) we group all examples sharing the same class label and penalize with the conditional variance of the



**Fig. 19** Classification for  $Y \in \{\text{woman, man}\}$  with  $\kappa = 0.5$ , using the baselines which (i) penalize the variance of the predicted logits conditional on the class label  $Y$  only; and (ii) penalize the overall variance of the predicted logits (cf. Sect. D.3.1). For baseline (i), panels (a) and (b) show the test error on test data sets 1 and 2 respectively as a function of the “baseline penalty weight” for various ridge penalty strengths. For baseline (ii), the equivalent plots are shown in panels (c) and (d). In contrast to the CoRE penalty, regularizing with these two baselines does not yield performance improvements on test set 2, compared to the pooled estimator (corresponding to a penalty weight of 0)

predicted logits, computed over these two groups; (ii) we penalize the *overall* variance of the predicted logits, i.e., a form of unconditional variance regularization. Figure 19 shows the performance of these two approaches. In contrast to the CoRE penalty, regularizing with the variance of the predicted logits conditional on  $Y$  only does not yield performance improvements on test set 2, compared to the pooled estimator (corresponding to a penalty weight of 0). Interestingly, using baseline (i) without a ridge penalty does yield an improvement on test set I, compared to the pooled estimator with various strengths of the ridge penalty.

**Additional results**

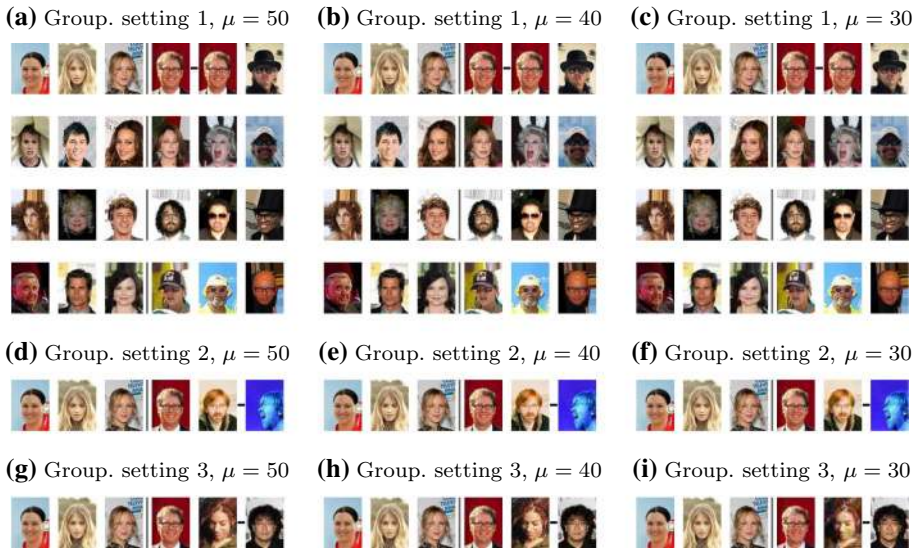
Table 6 additionally reports the standard errors for the results discussed in Sect. 5.2.

**Table 6** Classification for  $Y \in \{\text{woman, man}\}$

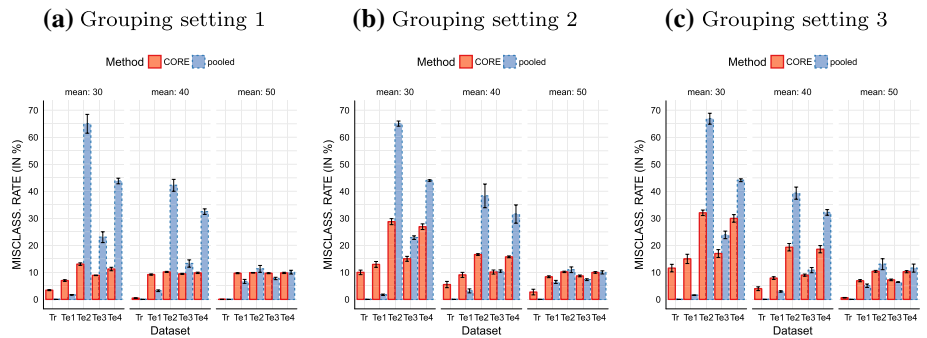
	Method	Error						Penalty value		
		Error		Test 1		Test 2		Train	Test: Females	Test: Males
		Train		Test 1		Test 2				
$\kappa = .5$	5-layer CNN + CoRE	0.00% (0.00%) 6.43% (0.25%)		2.00% (0.17%) 5.85% (0.17%)		38.54% (0.81%) 24.07% (0.80%)	22.77 (0.27) 0.01 (0.01)	74.05 (2.17) 1.61 (0.14)	30.67 (0.88) 0.93 (0.09)	
$\kappa = .75$	5-layer CNN + CoRE	0.00% (0.00%) 7.61% (0.44%)		1.98% (0.09%) 6.99% (0.42%)		43.41% (0.44%) 27.05% (1.32%)	8.23 (0.33) 0.00 (0.00)	32.98 (1.44) 1.44 (0.20)	11.76 (0.59) 0.62 (0.03)	
$\kappa = .9$	5-layer CNN + CoRE	0.00% (0.00%) 8.76% (0.59%)		2.00% (0.11%) 7.74% (0.68%)		47.64% (1.11%) 30.63% (1.73%)	9.47 (0.75) 0.00 (0.01)	40.51 (1.62) 1.26 (0.31)	14.37 (1.12) 0.42 (0.08)	
$\kappa = .95$	5-layer CNN + CoRE	0.00% (0.00%) 10.45% (1.42%)		1.89% (0.16%) 9.35% (1.79%)		48.96% (1.02%) 29.57% (2.43%)	13.62 (1.64) 0.00 (0.00)	61.01 (1.22) 0.42 (0.30)	21.26 (1.66) 0.16 (0.11)	
$\kappa = .99$	5-layer CNN + CoRE	0.00% (0.00%) 11.10% (1.17%)		1.70% (0.10%) 10.51% (1.26%)		50.11% (0.65%) 32.91% (0.81%)	20.66 (1.25) 0.00 (0.00)	70.80 (1.81) 0.00 (0.00)	27.80 (1.12) 0.00 (0.00)	
$\kappa = 1$	5-layer CNN + CoRE	0.00% (0.00%) 11.12% (0.34%)		1.93% (0.05%) 10.11% (0.32%)		49.41% (1.50%) 35.68% (2.03%)	821.32 (144.68) 0.00 (0.00)	2524.77 (219.96) 0.02 (0.02)	1253.21 (184.84) 0.01 (0.01)	

We compare six different datasets that vary with respect to the distribution of  $Y$  in the grouped observations. Specifically, we vary the proportion of images showing men between  $\kappa = 0.5$  and  $\kappa = 1$ . In all training datasets, the total number of observations is 16982 and the total number of grouped observations is 500. Both the pooled estimator as well as the CoRE estimator perform better if the distribution of  $Y$  in the grouped observations is more balanced. The CoRE estimator improves the error rate of the pooled estimator by  $\approx 28 - 39\%$  on a relative scale





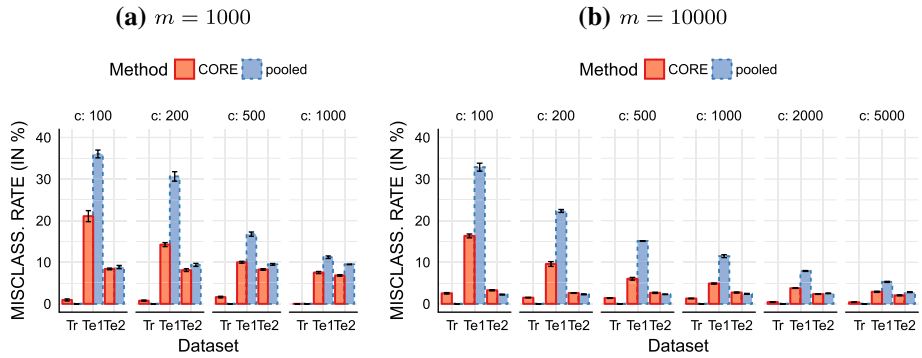
**Fig. 20** Examples from the CelebA image quality datasets, grouping settings 1–3 with  $\mu \in \{30, 40, 50\}$ . In all rows, the first three images from the left have  $y \equiv no\ glasses$ ; the remaining three images have  $y \equiv glasses$ . Connected images are grouped observations over which we calculate the conditional variance. In panels **a–c**, row 1 shows examples from the training set, rows 2–4 contain examples from test sets 2–4, respectively. Panels **d–i** show examples from the respective training sets



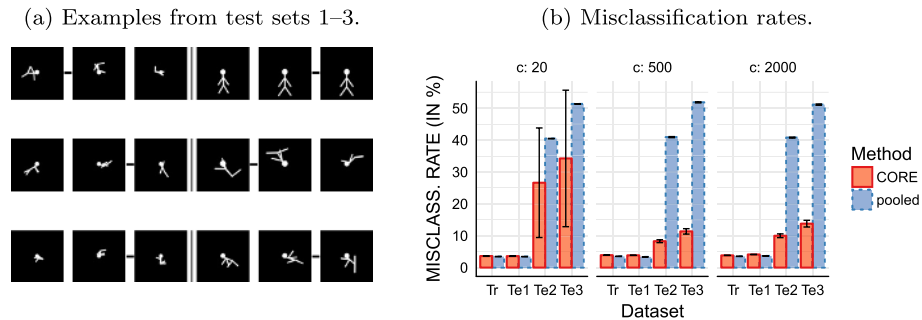
**Fig. 21** Misclassification rates for the CelebA eyeglasses detection with image quality interventions, grouping settings 1–3 with  $c = 5000$  and the mean of the Gaussian distribution  $\mu \in \{30, 40, 50\}$

### Eyeglasses detection: image quality intervention

Here, we show further results for the experiments introduced in Sect. 5.3. Specifically, we consider interventions of different strengths by varying the mean of the quality intervention in  $\mu \in \{30, 40, 50\}$ . Recall that we use ImageMagick to modify the image quality. In the training set and in test set 1, we sample the image quality value as  $q_{ij} \sim \mathcal{N}(\mu, \sigma = 10)$  and apply the command `convert -quality q_ij input`.



**Fig. 22** Data augmentation setting: Misclassification rates for MNIST and  $S \equiv rotation$ . In test set 1 all digits are rotated by a degree randomly sampled from  $[35, 70]$ . Test set 2 is the usual MNIST test set



**Fig. 23** **a** Examples from the stickmen test set 1 (row 1), test set 2 (row 2) and test sets 3 (row 3). In each row, the first three images from the left have  $y \equiv child$ ; the remaining three images have  $y \equiv adult$ . Connected images are grouped examples. **b** Misclassification rates for different numbers of grouped examples


`jpg output.jpg` if  $y_i \equiv glasses$ . If  $y_i \equiv no\ glasses$ , the image is not modified. In test set 2, the above command is applied if  $y_i \equiv no\ glasses$  while images with  $y_i \equiv glasses$  are not changed. In test set 3 all images are left unchanged and in test set 4 the command is applied to all images, i.e. the quality of all images is reduced.

We run experiments for grouping settings 1–3 and for  $c = 5000$ , where the definition of the grouping settings 1–3 is identical to “Appendix D.1”. Figure 20 shows examples from the respective training and test sets and Fig. 21 shows the corresponding misclassification rates. Again, we observe that grouping setting 1 works best, followed by grouping setting 2. Interestingly, there is a large performance difference between  $\mu = 40$  and  $\mu = 50$  for the pooled estimator. Possibly, with  $\mu = 50$  the image quality is not sufficiently predictive for the target.


**MNIST: more sample efficient data augmentation**

Here, we show further results for the experiment introduced in Sect. 5.5. We vary the number of augmented training examples  $c$  from 100 to 5000 for  $m = 10,000$  and  $c \in \{100, 200, 500, 1000\}$  for  $m = 1000$ . The degree of the rotations is sampled uniformly


**Fig. 24** Examples from the sub-sampled and augmented AwA2 dataset (Elmer-the-Elephant dataset). Row 1 shows examples from the training set, rows 2–5 show examples from test sets 1–4, respectively



**(a)** Examples of misclassified observations.



**(b)** Misclassification rates.



**Fig. 25** Elmer-the-Elephant dataset. **a** Misclassified examples from the test sets. **b** Misclassification rates on test sets 1 to 4

at random from [35, 70]. Figure 22 shows the misclassification rates. Test set 1 contains rotated digits only, test set 2 is the usual MNIST test set. We see that the misclassification rates of CoRE are always lower on test set 1, showing that it makes data augmentation more efficient. For  $m = 1000$ , it even turns out to be beneficial for performance on test set 2.

### Stickmen image-based age classification

Here, we show further results for the experiment introduced in Sect. 5.4. Recall that test set 1 follows the same distribution as the training set. In test sets 2 and 3 large movements are associated with both children and adults, while the movements are heavier in test set 3 than in test set 2. Figure D.10b shows results for different numbers of grouping examples. For  $c = 20$  the misclassification rate of CoRE estimator has a large variance. For  $c \in \{50, 500, 2000\}$ , the CoRE estimator shows similar results. Its performance is thus not sensitive to the number of grouped examples, once there are sufficiently many grouped observations in the training set. The pooled estimator fails to achieve good predictive performance on test sets 2 and 3 as it seems to use “movement” as a predictor for “age” (Fig. 23).

## Elmer the Elephant

The color interventions for the experiment introduced in Sect. 5.6 were created as follows. In the training set, if  $y_i \equiv \textit{elephant}$  we apply the following ImageMagick command for the grouped examples `convert -modulate 100,0,100 input.jpg output.jpg`. Test sets 1 and 2 were already discussed in Sect. 5.6: in test set 1, all images are left unchanged. In test set 2, the above command is applied if  $y_i \equiv \textit{horse}$ . If  $y_i \equiv \textit{elephant}$ , we sample  $c_{ij} \sim \mathcal{N}(\mu = 20, \sigma = 1)$  and apply `convert -modulate 100,100,100-c_ij input.jpg output.jpg` to the image. Here, we consider again some more test sets than in Sect. 5.6. In test set 4, the latter command is applied to all images. It rotates the colors of the image, in a cyclic manner.<sup>16</sup> In test set 3, all images are changed to grayscale. Examples from all four test sets are shown in Fig. 24 and classification results are shown in Fig. 25.

## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., & Zheng, X. (2015). TensorFlow: Large-scale machinelearning on heterogeneous systems. <https://www.tensorflow.org/>, software available from tensorflow.org.
- Aldrich, J. (1989). Autonomy. *Oxford Economic Papers*, 41, 15–34.
- Bagnell, J. (2005). Robust supervised learning. In *Proceedings of the national conference on artificial intelligence*, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; (1999) Vol. 20, p. 714.
- Bahadori, M. T., Chalupka, K., Choi, E., Chen, R., Stewart, W. F., & Sun, J. (2017). Causal regularization. [arXiv:170202604](https://arxiv.org/abs/170202604).
- Barocas, S., & Selbst, A.D. (2016). Big Data’s Disparate Impact. 104 California Law Review 671.
- Bartlett, M.S., Sejnowski, T. J. (1996). Viewpoint invariant face recognition using independent component analysis and attractor networks. In *Proceedings of the 9th International Conference on Neural Information Processing Systems* (pp. 817–823), MIT Press, Cambridge, MA, USA, NIPS’96.
- Belkin, M., Niyogi, P., & Sindhvani, V. (2006). Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7(Nov), 2399–2434.
- Ben-David, S., Blitzer, J., Crammer, K., & Pereira, F. (2007). Analysis of representations for domain adaptation. In *Advances in Neural Information Processing Systems*, Vol. 19.
- Ben-Tal, A., Den Hertog, D., De Waegenaere, A., Melenberg, B., & Rennen, G. (2013). Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2), 341–357.
- Besserve, M., Shajarisales, N., Schölkopf, B., & Janzing, D. (2018). Group invariance principles for causal generative models. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR, *Proceedings of Machine Learning Research* (pp. 557–565), Vol. 84.
- Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems*, Vol. 29.
- Bouchacourt, D., Tomioka, R., & Nowozin, S. (2018). Multi-level variational autoencoder: Learning disentangled representations from grouped observations. In *AAAI Conference on Artificial Intelligence*.
- Chalupka, K., Perona, P., & Eberhardt, F. (2014). Visual Causal Feature Learning. *Uncertainty in Artificial Intelligence*.
- Crawford, K. (2016). Artificial intelligence’s white guy problem. The New York Times, June 25 2016 <https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html>.

<sup>16</sup> For more details, see [http://www.imagemagick.org/Usage/color\\_mods/#color\\_mods](http://www.imagemagick.org/Usage/color_mods/#color_mods).

- Csurka, G. (2017). A comprehensive survey on domain adaptation for visual applications. In *Domain Adaptation in Computer Vision Applications* (pp. 1–35).
- Denton, E. L., & Birodkar, V. (2017). Unsupervised learning of disentangled representations from video. In *Advances in Neural Information Processing Systems*, Vol. 30.
- Devries, T., & Taylor, G. W. (2017). Dataset augmentation in feature space. ICLR Workshop Track.
- Empak, J. (2016). How a machine learns prejudice. Scientific American, December 29 2016 <https://www.scientificamerican.com/article/how-a-machine-learns-prejudice/>.
- Engstrom, L., Ilyas, A., & Athalye, A. (2018). Evaluating and understanding the robustness of adversarial logit pairing. [arXiv:1807.10272](https://arxiv.org/abs/1807.10272).
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., et al. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1), 2096–2030.
- Gao, R., Chen, X., & Kleywegt, A. (2017). [arXiv:1712.06050](https://arxiv.org/abs/1712.06050).
- Gong, B., Grauman, K., & Sha, F. (2013). Reshaping visual datasets for domain adaptation. In *Advances in Neural Information Processing Systems* (Vol. 26, pp. 1286–1294), Curran Associates, Inc.
- Gong, M., Zhang, K., Liu, T., Tao, D., Glymour, C., & Schölkopf, B. (2016). Domain adaptation with conditional transferable components. In *International Conference on Machine Learning*.
- Goodfellow, I., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *International conference on learning representations*.
- Goudet, O., Kalainathan, D., Caillou, P., Lopez-Paz, D., Guyon, I., Sebag, M., Tritas, A., & Tubaro, P. (2017). Learning functional causal models with generative neural networks. [arXiv:1709.05321](https://arxiv.org/abs/1709.05321).
- Haavelmo, T. (1944). The probability approach in econometrics. *Econometrica* 12:S1–S115 (supplement).
- Hashimoto, T. B., Liang, P. S., & Duchi, J. C. (2017). Unsupervised transformation learning via convex relaxations. In *Advances in neural information processing systems* (Vol. 30, pp. 6875–6883), Curran Associates, Inc.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the 2015 IEEE international conference on computer vision (ICCV)* (pp. 1026–1034).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 770–778).
- Hoffman, J., Kulis, B., Darrell, T., & Saenko, K. (2012). Discovering latent domains for multisource domain adaptation. In *Computer Vision - ECCV 2012* (pp. 702–715). Berlin Heidelberg, Berlin, Heidelberg: Springer.
- Kannan, H., Kurakin, A., & Goodfellow, I. J. (2018). Adversarial logit pairing. [arXiv:1803.06373](https://arxiv.org/abs/1803.06373).
- Khasanova, R., & Frossard, P. (2017). Graph-based isometry invariant representation learning. In *Proceedings of the 34th international conference on machine learning* (Vol. 70, pp. 1847–1856).
- Kilbertus, N., Rojas Carulla, M., Parascandolo, G., Hardt, M., Janzing, D., & Schölkopf, B. (2017). Avoiding discrimination through causal reasoning. *Advances in Neural Information Processing Systems*, 30, 656–666.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. International Conference on Learning Representations (ICLR).
- Kocaoglu, M., Snyder, C., Dimakis, A., & Vishwanath, S. (2018). CausalGAN: Learning causal implicit generative models with adversarial training. In *International conference on learning representations*.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, Vol. 25.
- Kuehlikamp, A., Becker, B., & Bowyer, K. (2017). Gender-from-iris or gender-from-mascara? In *2017 IEEE winter conference on applications of computer vision (WACV)*.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*.
- Li, K. C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414), 316–327.
- Liu, Z., Luo, P., Wang, X., & Tang, X. (2015). Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Lopez-Paz, D., & Oquab, M. (2017). Revisiting classifier two-sample tests. In *International conference on learning representations (ICLR)*.
- Lopez-Paz, D., Nishihara, R., Chintala, S., Schölkopf, B., & Bottou, L. (2017). Discovering causal signals in images. In *The IEEE conference on computer vision and pattern recognition (CVPR 2017)*.
- Magliacane, S., van Ommen, T., Claassen, T., Bongers, S., Versteeg, P., & Mooij, J. (2018). Domain adaptation by using causal inference to predict invariant conditional distributions. In *Advances in neural information processing systems*.

- Meinshausen, N. (2018). Causality from a distributional robustness point of view. In *2018 IEEE Data Science Workshop (DSW)* (pp. 6–10).
- Mosbach, M., Andriushchenko, M., Trost, T., Hein, M., & Klakow, D. (2018). Logit pairing methods can fool gradient-based attacks. [arXiv:1810.2042](https://arxiv.org/abs/1810.2042).
- Namkoong, H., & Duchi, J. (2017). Variance-based regularization with convex objectives. In *Advances in Neural Information Processing Systems* (pp. 2975–2984).
- Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). New York: Cambridge University Press.
- Peters, J., Bühlmann, P., & Meinshausen, N. (2016). Causal inference using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society, Series B*, 78, 947–1012.
- Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., & Lawrence, N. D. (2009). *Dataset shift in machine learning*. Cambridge: The MIT Press.
- Richardson, T., & Robins, J. M. (2013). Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. Center for the Statistics and the Social Sciences, University of Washington Series Working Paper 128, 30 April 2013.
- Rojas-Carulla, M., Schölkopf, B., Turner, R., & Peters, J. (2018). Causal transfer in machine learning. To appear in *Journal of Machine Learning Research*.
- Rothenhäusler, D., Bühlmann, P., Meinshausen, N., & Peters, J. (2018). Anchor regression: heterogeneous data meets causality. [arXiv:1801.06229](https://arxiv.org/abs/1801.06229).
- Schölkopf, B., Burges, C., & Vapnik, V. (1996). Incorporating invariances in support vector learning machines. *Artificial Neural Networks – ICANN 96* (pp. 47–52). Berlin Heidelberg, Berlin, Heidelberg: Springer.
- Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., & Mooij, J. (2012). On causal and anti-causal learning. In *Proceedings of the 29th international conference on machine learning (ICML)* (pp. 1255–1262).
- Shafieezadeh-Abadeh, S., Kuhn, D., & Esfahani, P. (2017). Regularization via mass transportation. [arXiv:1710.10016](https://arxiv.org/abs/1710.10016).
- Sinha, A., Namkoong, H., & Duchi, J. (2018). Certifiable distributional robustness with principled adversarial training. In *International conference on learning representations*.
- Sohn, K., & Lee, H. (2012). Learning invariant representations with local transformations. In *Proceedings of the 29th international conference on machine learning*, Omnipress, USA, ICML'12, pp. 1339–1346.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. In *Computer vision and pattern recognition (CVPR)*.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2014). Intriguing properties of neural networks. In *International conference on learning representations*.
- Torralba, A., & Efros, A. A. (2011). Unbiased look at dataset bias. In *Computer Vision and Pattern Recognition (CVPR)*.
- Tran, L., Yin, X., & Liu, X. (2017). Disentangled representation learning gan for pose-invariant face recognition. In *Proceeding of IEEE computer vision and pattern recognition*, Honolulu, HI
- Venkateswara, H., Eusebio, J., Chakraborty, S., & Panchanathan, S. (2017). Deep hashing network for unsupervised domain adaptation. In (IEEE) Conference on Computer Vision and Pattern Recognition (CVPR).
- Villani, C. (2003). *Topics in optimal transportation* (Vol. 58). Providence: American Mathematical Society.
- Volpi, R., Namkoong, H., Sener, O., Duchi, J., Murino, V., & Savarese, S. (2018). Generalizing to unseen domains via adversarial data augmentation. [arXiv:1805.12018](https://arxiv.org/abs/1805.12018).
- Xian, Y., Lampert, C. H., Schiele, B., & Akata, Z. (2017). Zero-shot learning—A comprehensive evaluation of the good, the bad and the ugly. [arXiv:1707.00600](https://arxiv.org/abs/1707.00600).
- Xu, H., Caramanis, C., & Mannor, S. (2009). Robust regression and lasso. In *Advances in Neural Information Processing Systems* (pp. 1801–1808).
- Xu, C., Tao, D., & Xu, C. (2013). A survey on multi-view learning. [arXiv:1304.5634](https://arxiv.org/abs/1304.5634).
- Yu, X., Liu, T., Gong, M., Zhang, K., & Tao, D. (2017). Transfer learning with label noise. [arXiv:1707.0724](https://arxiv.org/abs/1707.0724).
- Zhang, K., Gong, M., & Schölkopf, B. (2015). Multi-source domain adaptation: A causal view. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Zhang, K., Schölkopf, B., Muandet, K., & Wang, Z. (2013). Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*.