Invariant Representations through Adversarial Forgetting

Ayush Jaiswal, Daniel Moyer, Greg Ver Steeg, Wael AbdAlmageed, Premkumar Natarajan

Information Sciences Institute, University of Southern California
 {ajaiswal, gregv, wamageed, pnataraj}@isi.edu, moyerd@usc.edu

Abstract

We propose a novel approach to achieving invariance for deep neural networks in the form of inducing amnesia to unwanted factors of data through a new adversarial forgetting mechanism. We show that the forgetting mechanism serves as an information-bottleneck, which is manipulated by the adversarial training to learn invariance to unwanted factors. Empirical results show that the proposed framework achieves state-of-the-art performance at learning invariance in both nuisance and bias settings on a diverse collection of datasets and tasks.

1 Introduction

Supervised machine learning models learn to associate a target y with underlying factors of data x that are informative of y . However, trained models often learn to associate irrelevant factors of x with y (Domingos 2012), e.g., stroke-width of text in optical character recognition. Learning incorrect associations between y and *nuisance* factors leads to overfitting. Models might also learn associations between the target y and factors that are correlated with y in collected training data due to external (sometimes historical) reasons, e.g., age, gender, and race in socio-economical prediction tasks. These *biasing* factors, which are sometimes known *a priori*, make trained models unfair to under-represented groups, posing ethical and legal challenges to the usage of such models. It is, therefore, imperative to develop models that are invariant to both nuisance and biasing factors.

Popular approaches for achieving invariance to undesired factors s have involved training with data augmentation through small perturbations to real x with respect to s (Bengio, Courville, and Vincent 2013). While these methods have been utilized for training deep neural networks (DNNs) (Ko et al. 2015; Krizhevsky, Sutskever, and Hinton 2012), in recent years methods have emerged that remove s from the latent representations $\tilde{z}$, penalizing the model for the presence of s (Jaiswal et al. 2018; 2019a; Li, Swersky, and Zemel 2014; Lopez et al. 2018; Louizos et al. 2016; Moyer et al. 2018; Xie et al. 2017; Zemel et al. 2013). These methods perform better than data

augmentation, due to their approach of invariance by *exclusion* rather than *inclusion* (Jaiswal et al. 2018; 2019b; Hsu, Jaiswal, and Natarajan 2019).

We propose a novel framework for invariance in DNNs that promotes removal of information about s from $\tilde{z}$ through an adversarial forgetting mechanism. The working principle of the model is to map data x to an embedding z that encodes everything about x and use a forget-mask to filter out s information from z while retaining information about y to produce the invariant $\tilde{z}$. Specifically, an encoder network generates a latent code z from x , which is used to reconstruct x through a decoder. At the same time, a forget-gate network generates a mask m from x , which is multiplied elementwise with z to produce $\tilde{z}$. The encoding $\tilde{z}$ is used by a predictor to infer the target y . These components of the framework are trained adversarially with a discriminator that aims to predict s from $\tilde{z}$. However, during training, gradients from the discriminator are allowed to only affect the training of the forget-gate. Finally, the framework is augmented with a regularizer that pushes the components m_i of the forget-mask to be close to either 0 or 1, inducing disentanglement within the components of z and effective masking of s to produce an invariant representation $\tilde{z}$.

We show that the forgetting mechanism is equivalent to a bound on the mutual information $I(\tilde{z} : z)$ and that coupled with the y -prediction task, can be interpreted as an information bottleneck. Further, by the data processing inequality, the forgetting mechanism bounds $I(\tilde{z} : s)$. The generated mask can be manipulated through adversarial training to remove information about s from z . Empirical results show that the proposed framework exhibits state-of-the-art performance at inducing invariance to both nuisance and biasing factors across a diverse collection of datasets and tasks.

2 Related Work

Recent work (Achille and Soatto 2018b; Alemi et al. 2016; Moyer et al. 2018) has modeled invariance in supervised DNNs through information bottleneck (Tishby, Pereira, and Bialek 1999), wherein representations minimize the mutual information $I(x : \tilde{z})$ while maximizing $I(\tilde{z} : y)$. For nuisance variables ($s \perp y$), these methods bring about compression in the latent space, which removes information about

s and indirectly minimizes $I(\tilde{z} : s)$. Under optimality, the information bottleneck objective will remove all such nuisance s completely from $\tilde{z}$ (Achille and Soatto 2018b). The bottleneck objective is, however, difficult to optimize and has, hence, been approximated using variational inference in prior work (Alemi et al. 2016) or implemented as Information Dropout (Achille and Soatto 2018a).

The bottleneck objective cannot remove s from $\tilde{z}$ that are correlated with y (biasing factors). In such cases, it is necessary to force the exclusion of these biasing s from $\tilde{z}$. Methods that employ such mechanisms additionally benefit the removal of nuisance s for which annotations are available. Moyer *et al.* (2018) achieve this by augmenting the bottleneck objective with $I(\tilde{z} : s)$ and optimizing its variational bound, which we call the Conditional Variational Information Bottleneck (CVIB). The Variational Fair Autoencoder (VFAE) (Louizos et al. 2016) optimizes this objective indirectly as a Variational Autoencoder (VAE) (Kingma and Welling 2014) and uses Maximum Mean Discrepancy (MMD) (Gretton et al. 2007) to boost the removal of undesired s from the latent embedding. The Hilbert-Schmidt Independence Criterion (HSIC) (Gretton et al. 2005) constrained VAE (HCV) (Lopez et al. 2018) uses HSIC to enforce independence between the latent embedding and s . NN+MMD (Li, Swersky, and Zemel 2014) directly minimizes MMD as a regularizer for DNNs.

Invariance to nuisance variables has also been modeled as an implicit feature selection method within DNNs in the Unsupervised Adversarial Invariance framework (UAI) (Jaiswal et al. 2018; 2019a) that splits data representation into an invariant embedding, which is used for predicting y , and a nuisance embedding through competitive training between prediction and reconstruction tasks combined with disentanglement between the two embeddings.

The Domain-Adversarial Neural Network (DANN) (Ganin et al. 2016) and the Controllable Adversarial Invariance (CAI) model (Xie et al. 2017) indirectly optimize an invariance objective using the gradient-reversal trick (Ganin et al. 2016) that penalizes a DNN if it encodes s . In the case of invariant representation learning for a supervised prediction task (which is the subject of this work), CAI, DANN and Fader Networks (Lample et al. 2017) reduce to the same model. In contrast, we propose a new forgetting mechanism that more explicitly “masks” s out of z , which encodes everything about x , to generate a new representation $\tilde{z}$ that is maximally informative of y but invariant to s . Inspired by (1) the discovery of richer features in DNN classifiers upon augmentation with reconstruction objectives (Sabour, Frosst, and Hinton 2017; Zhang, Lee, and Lee 2016), and (2) the forgetting operation in Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber 1997) cells, the proposed framework adopts the idea of “discovery and separation of information” for invariance, which is *foundationally different* from the direct removal of undesired factors as in DANN and CAI.

3 Invariance through Adversarial Forgetting

Invariant representation learning aims to produce a mapping of data (x) to code ($\tilde{z}$) that is minimally informative of unde-

sired factors of data (s) but maximally discriminative for the prediction task (y). There are two cases that arise from this formulation – (1) s is nuisance, i.e., there is little or no information shared between s and y asymptotically (e.g., pose in face recognition), and (2) s contains biasing factors, i.e., there is correlation between s and y , but for outside reasons, it is necessary to exclude these biases from the prediction process (e.g., gender, race, etc. in socio-economical prediction tasks using historical data). Invariance to s leads to robust models that generalize better on test data for case (1) (Jaiswal et al. 2018), and produces fair models that do not incorporate s while making predictions for case (2).

We induce invariance to s within a DNN using a novel approach of adversarial forgetting, which is inspired by forget-gates in Long Short-Term Memory (LSTM) cells in recurrent neural networks (Hochreiter and Schmidhuber 1997). Within the proposed framework, the model learns to embed everything about a data sample into an intermediate representation, which is transformed into an invariant representation through multiplication with a mask generated by an adversarial forgetting mechanism. Figure 1 shows the complete framework design. Data samples x are encoded into an intermediate representation z using an encoder E , while a forget-mask m ($m_i \in (0, 1)$) is simultaneously produced from x through a forget-gate network F . The invariant representation $\tilde{z}$ is then computed as the element-wise multiplication $\tilde{z} = z \odot m$. This is similar to how forget gates are used in LSTMs to “forget” certain information learned from past data. A decoder R is used to reconstruct x from z , such that E learns to encode everything about x into z . A predictor P infers y from $\tilde{z}$, while an adversarial discriminator D tries to predict s from $\tilde{z}$. Hence, the combined objectives of P and D aim to allow only factors of x that are predictive of y but not of s to pass from z to $\tilde{z}$.

The complete framework is trained with an adversarial objective such that the discriminator is pitted against all the other modules, as depicted with colors in Figure 1. The discriminator is allowed a more active role in the development of forget-masks m by allowing adversarial gradients from the discriminator to only flow to the forget-gate network and not to the encoder during training. This is illustrated in the figure with the break on the arrow between z and the multiplication operation. In order to further encourage the development of masks that truly filter out some information from z but retain everything else, a mask regularizer, in the form of $m^T(1 - m)$, is added to push components of m to either 0 or 1. We found the mask regularizer to always improve results in our experiments. The complete training objective can be written as shown in Equation 1.

$$\min_{E, F, P, R} \max_D J(E, F, P, R, D); \text{ where:}$$

$$J(E, F, P, R, D) = L_y(y, P(\tilde{z})) + \rho L_x(x, R(z)) + \delta L_s(s, D(\tilde{z})) + \lambda m^T(1 - m) \quad (1)$$

The losses L_y and L_s are implemented as cross-entropy while mean squared error is used for L_x in our experiments. The proposed model is trained using a scheduled update scheme similar to the training mechanism of the UAI

Figure 1: Adversarial forgetting framework for invariant representation learning

model (Jaiswal et al. 2019a). Hence, the weights of the discriminator are frozen when updating the rest of the model and vice versa. The adversarial training of the proposed model would benefit from training the discriminator to convergence before any update to the other modules (Jaiswal et al. 2019a). However, in practice, this is infeasible and training the discriminator much more frequently than the rest of the model (depending on the nature of the prediction task and the dataset) is sufficient to achieve good performance. This is especially true because the training of the discriminator is resumed from its previous state rather than starting from scratch after every update to the other modules. Therefore, the weights of the discriminator and the rest of the model are updated in the frequency ratio of $k : 1$. We found $k = 10$ to work well in our experiments.

The model training does not incorporate the popular approach of gradient-reversal and instead follows (Jaiswal et al. 2019a). The targets of the discriminator D are set to the ground-truth s labels while updating D , but to random s values (sampled from the empirical s -distribution) when the parameters of the rest of the model are updated. Hence, D tries to predict the correct s during its training phase, but the rest of the model is updated to elicit random-chance performance at s -prediction, leading to the desired invariance to s . The model was implemented in Keras with TensorFlow backend. The Adam optimizer was used with 10^{-4} learning rate and 10^{-4} decay. The hyperparameters ρ , λ , and δ were tuned through grid search in powers of 10.

The proposed framework can also be extended to multi-task settings, with z treated as the common encoding for tasks involving prediction of targets $\{y^{(1)}, y^{(2)}, \dots, y^{(n)}\}$ with corresponding undesired factors $\{s^{(1)}, s^{(2)}, \dots, s^{(n)}\}$. Forget-gates $F^{(j)}$ are added to the framework, one for each prediction task $y^{(j)}$, to generate associated masks $m^{(j)}$ and invariant representations $\tilde{z}^{(j)}$ from z through adversarial training with discriminators $D^{(j)}$, each of which tries to predict $s^{(j)}$ from $\tilde{z}^{(j)}$. Thus, the multi-task extension of the proposed model is intuitive and straightforward.

4 Characterizing Forgetting with Forget-gate

Forget-gates were introduced as components of LSTMs, where they cause “forgetting” of information from the past in the recurrence formulation conditioned on the input at a given step as well as the existing state. A forget-gate typically produces a mask m with $m_i \in (0, 1)$ that is multi-

plied elementwise with a latent encoding within an LSTM cell. Thus, the forget-gate can scale or remove information in each dimension of the encoding *but not add information to it*. Inspired by this formulation, we employ forget-gates to induce invariance to s , i.e., to “forget” s -related information. In this section, we characterize the erasure properties of the forget-gate in the proposed framework. Intuitively, if a mask element $m_i = 0$, the information passed from z_i to $\tilde{z}_i$ is also zero; likewise, if $m_i = 1$, the information passed is complete. We characterize here the behavior for $m_i \in (0, 1)$, showing that under reasonable assumptions, there is a non-trivial forget regime besides zero. We show that the forget-gate *acts as* an information bottleneck, which can be manipulated for invariance to specific s .

4.1 Forget-gate

The proposed model generates a d -dimensional encoding z of x and a forgetting mask m with components $m_i \in (0, 1)$. These are multiplied element-wise to produce $\tilde{z} = z \odot m$. We consider the multiplication as a noisy operation $\tilde{z} = z \odot m + \varepsilon$ with a small ε in order to facilitate a theoretical analysis of “forgetting”. We discuss the practicalities of ε later in Section 4.2. Assuming $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon I)$, we get $P(\tilde{z}|z) \sim \mathcal{N}(z \odot m, \sigma_\varepsilon I)$. In each dimension, we get:

$$\begin{aligned} I(\tilde{z}_i : z_i) &= H(\tilde{z}_i) - H(\tilde{z}_i|z_i) = H(\tilde{z}_i) - H(\varepsilon_i) \\ &= H(\tilde{z}_i) - \frac{1}{2} \log(\text{Var}(\varepsilon_i)) - \frac{1}{2} \log(2\pi e) \end{aligned} \quad (2)$$

where $H(\varepsilon_i)$ is constant with respect to z_i and m_i . Thus, the information passed from z_i to $\tilde{z}_i$ is proportional to $H(\tilde{z}_i)$. Assuming that $\text{Var}(z_i)$ is defined, the max-entropy Gaussian upper bound gives us the following in each dimension of the embedding (Cover and Thomas 2012):

$$H(\tilde{z}_i) = H(m_i z_i + \varepsilon_i) \quad (3)$$

$$\leq \frac{1}{2} \log(\text{Var}(m_i z_i) + \text{Var}(\varepsilon_i)) + \frac{1}{2} \log(2\pi e) \quad (4)$$

If m is non-random, the mutual information is:

$$I(\tilde{z}_i : z_i) = H(\tilde{z}_i) - H(\varepsilon_i) \quad (5)$$

$$\leq \frac{1}{2} \log(m_i^2 \text{Var}(z_i) + \text{Var}(\varepsilon_i)) - \frac{1}{2} \log(\text{Var}(\varepsilon_i)) \quad (6)$$

When $m_i \rightarrow 1$, assuming that $\text{Var}(\varepsilon_i) \ll \text{Var}(z_i)$, $H(\tilde{z}_i) \approx H(z_i)$. As $m_i \rightarrow 0$, $\tilde{z}_i = \varepsilon_i$, $H(\tilde{z}_i) \rightarrow H(\varepsilon_i)$ and $I(\tilde{z}_i :$

Figure 2: Chairs Dataset – t-SNE visualization of z and $\tilde{z}$ labeled with orientation class (s). Visualization with chair-type (y) annotations are not shown because there are 1,393 y classes. The invariant encoding $\tilde{z}$ shows no clustering by orientation as s is masked out of z , which exhibits s -grouping.

$z_i) \rightarrow 0$. Importantly, when $m_i < \sqrt{\text{Var}(\varepsilon_i)}$ yet away from zero, the information loss is still non-trivial. The result in Equation 6 was derived assuming that m is fixed. We can extend this bound to random m , including those dependent on x as described in Section 3, by using the following identity (shown in Appendix A):

$$\text{Var}(m_i z_i) \leq 2\text{Var}((m_i - \mathbb{E}[m_i])z_i) + 2\mathbb{E}[m_i]^2 \text{Var}(z_i) \quad (7)$$

This makes the bound on $I(\tilde{z}_i : z_i)$:

$$\begin{aligned} I(\tilde{z}_i : z_i) &\leq \frac{1}{2} \log(2\text{Var}((m_i - \mathbb{E}[m_i])z_i) + \text{Var}(\varepsilon_i)) \\ &\quad + 2\mathbb{E}[m_i]^2 \text{Var}(z_i) - \frac{1}{2} \log(\text{Var}(\varepsilon_i)) \end{aligned} \quad (8)$$

Though more difficult to interpret, this has approximately the same characteristic as the fixed m_i case. In order to extend this to the multivariate case (from a bound on $I(\tilde{z}_i : z_i)$ to that on $I(\tilde{z} : z)$), we first note that the max-entropy bound still holds for multivariate Gaussians, and by Hadamard’s inequality, we can bound that distribution by its diagonal as:

$$H(\tilde{z}) \leq \log \det(\Sigma_{\tilde{z}} + \sigma_\varepsilon I) + \frac{d}{2} \log(2\pi e) \quad (9)$$

$$\leq \log \det(\Sigma_{\tilde{z}}^{diag} + \sigma_\varepsilon I) + \frac{d}{2} \log(2\pi e) \quad (10)$$

where $\Sigma_{\tilde{z}}^{diag}$ has only the diagonal elements of Σ_z and zero elsewhere. This gives us the bound:

$$I(\tilde{z} : z) \leq \sum_i \log(\text{Var}(m_i z_i) + \text{Var}(\varepsilon_i)) \quad (11)$$

The max-entropy bound fails if there are degenerate elements of z , i.e. completely duplicate channels, but still holds on subsets of channels without duplicates. We have somewhat abused notation in this section; really, our bound is on $I(\tilde{z} : (z, m))$. While this is less intuitive, the distinction is necessary for data processing inequalities showing, e.g., that $I(\tilde{z} : s)$ is controlled by the forget gate.

Table 1: Chairs results (random chance of $s = 0.25$).

Model	A_y	A_s
NN+MMD	0.73 ± 0.02	0.46 ± 0.04
VFAE	0.72 ± 0.04	0.37 ± 0.02
CAI	0.68	0.69
CVIB	0.67 ± 0.01	0.52 ± 0.01
UAI	0.74	0.34
Ours	0.84 ± 0.01	0.25 ± 0.00
Δ over UAI	38.5%	100%

4.2 Practicalities of ε -noise and bottleneck-mask

With exact computation, information is only lost when $m_i = 0$ because scaling by $m_i \in (0, 1)$ is an isomorphic map. In real computation, however, multiplication operations are not isomorphic due to imprecision in floating point arithmetic. These imprecisions, alongside commonly undertaken computational procedures (e.g., clipping) induce a non-trivial forgetting region, under which z_i with “reasonable” variance may be forgotten. Thus, practically speaking, we *do not need to add ε -noise artificially* to erase information even when $m_i \neq 0$. Hence, we do not need m_i to be exactly zero to lose information. The background noise of computation is sufficient for $I(\tilde{z} : z)$ to be smoothly controlled outside of zero m_i . The proposed framework implements exactly this control on $I(\tilde{z} : z)$, and thereby $I(\tilde{z} : x)$. Further, since it also optimizes a distortion measure between $\tilde{z}$ and y , this forms an information bottleneck (Alemi et al. 2016). For categorical y with cross-entropy loss, minimization of m_i coupled with the prediction task is equivalent to the bottleneck objective from (Tishby, Pereira, and Bialek 1999).

The goal of this work, however, is to generate invariant representations and not necessarily optimal bottleneck embeddings. In order to induce invariance to specific s , we learn the parameters of the bottleneck (forget-gate) so that it filters this information out of z (mechanism described in Section 3). This encourages the encoder and the forget-gate

Figure 3: Extended Yale-B – t-SNE visualizations labeled with lighting direction (s) and subject-identity (y). Subject-identities are marked with numbers and not colors. The invariant $\tilde{z}$ shows no clustering by lighting but groups by subject identity.

Table 2: Extended Yale-B results (random chance $s = 0.2$)

Model	A_y	A_s
NN+MMD	0.82	–
VFAE	0.85	0.57
CAI	0.89	0.57
CVIB	0.82 ± 0.01	0.45 ± 0.03
UAI	0.95	0.24
Ours	0.95 ± 0.01	0.20 ± 0.01
Δ over UAI	–	100%

to generate $\tilde{z}$ with minimal $I(\tilde{z} : s)$. The adversary operates only on m and thus can be thought of as optimizing element-wise the channel between z and $\tilde{z}$, minimizing $I(\tilde{z} : s)$ (or the equivalent general co-dependence term).

The masks generated by the forget-gate have an intuitive interpretation: for each component, it either allows information to pass from z_i to $\tilde{z}_i$ or does not. The overall design of the proposed framework causes separation of factors of x that are correlated with s from those that are not, so that they occur in different components of z , allowing a component-wise mask to effectively include or exclude factors from the final representation $\tilde{z}$ without cross-factor considerations. Removal of nuisance s does not penalize the training objective (Equation 1). However, for biasing s correlated with y , the forget-gate will choose whether to allow their inclusion in $\tilde{z}$ based on the loss-weights in the objective.

5 Experimental Evaluation

The proposed framework is compared with NN+MMD, VFAE, CAI, CVIB and UAI. Performance is evaluated on two metrics: accuracy of predicting y from $\tilde{z}$ (A_y) using the jointly trained predictor and that of predicting s from $\tilde{z}$ (A_s) using a two-layer neural network trained *post hoc*. While a high A_y is desired, for true invariance A_s should be random chance for nuisances and the share of the majority s -class for biasing s . Mean and standard deviation are reported based on five runs, except when results are quoted from previous works. Relative improvements in error-rate (Δ) are also reported with error-rate for A_s defined as the gap between the observed A_s and its optimal value. We further report evaluation results of the framework in a multi-task setting.

5.1 Robustness through invariance to nuisances

Invariance to nuisances is evaluated on the Chairs (Jaiswal et al. 2018), Extended Yale-B (Georgiades, Belhumeur, and Kriegman 2001), and MNIST-ROT (Jaiswal et al. 2018) datasets. The network architectures for the forget-gate and the encoder are kept the same in all experiments. Besides quantitative results, we show t-SNE plots of z and $\tilde{z}$ to visualize the transformation of the latent space due to $\tilde{z} = z \odot m$. These plots show that invariance is indeed brought about by the adversarial forgetting operation.

Chairs. This is a dataset of 86,366 images of 1,393 types of chairs at 31 yaw and two pitch angles. We use the same version of this dataset as prior works, which is split into training and testing sets by picking alternate yaw angles,

Figure 4: MNIST-ROT – t-SNE visualization of z and $\tilde{z}$ labeled with rotation-angle (s) and digit-class (y). The invariant $\tilde{z}$ shows no clustering by s while z shows clear s -subgroups within each y -cluster.

such that there is no overlap of angles between the two sets. The chair type is treated as y and the yaw binned into four classes (front, left, right and back) as s . We use the same architecture for the encoder, the predictor, and the decoder as UAI (Jaiswal et al. 2018), i.e., two-layer networks. The discriminator is modeled as a two-layer network. Results of our experiments are summarized in Table 1. The proposed model achieves large improvements over the prior state-of-the-art (UAI) on both A_y and A_s , with A_s that is exactly random chance. Figure 2 shows the t-SNE visualization of z and $\tilde{z}$, exhibiting that z groups by orientation but $\tilde{z}$ does not.

Extended Yale-B. This is a dataset of face-images of 38 subjects captured under various lighting conditions. The prediction target y is the subject-ID, while the nuisance s is the lighting condition binned into five classes (four corners and frontal). For each subject, one image from each s -category is used for training and the rest of the dataset is used for testing (Jaiswal et al. 2018; Louizos et al. 2016; Xie et al. 2017). We use the same architecture for the encoder and the predictor as previous works (Jaiswal et al. 2018; Xie et al. 2017), i.e., one layer for each of these modules. The decoder and the discriminator are modeled with two layers each. Results of our experiments are shown in Table 2. The proposed model exhibits state-of-the-art performance on both A_y and A_s , with A_s being exactly random chance. This shows that the proposed model is able to completely remove information about the lighting direction from the latent embedding while retaining high A_y . Figure 3 shows the t-SNE visualizations of z and $\tilde{z}$ and validates this

claim as the clustering of the latent embedding completely changes from grouping by s for z to grouping by y for $\tilde{z}$.

MNIST-ROT. This is a variant of the MNIST (LeCun et al. 1998) dataset, which is augmented with digits rotated at $\pm 22.5^\circ$ and $\pm 45^\circ$. The digit class is treated as y and the rotation angle as categorical s . We use the same architecture for the encoder, the predictor, and the decoder as UAI, i.e., two layers for the encoder and the predictor each and three layers for the decoder. The discriminator is modeled with two layers. Table 3 summarizes the results of our experiments on test images with rotation angles both seen and unseen ($\pm 55^\circ$ and $\pm 65^\circ$) during training. Our model achieves not only the best A_y but also A_s that is exactly random chance, showing that it is able to filter out s while retaining more information about y , leading to more accurate y -predictions. The t-SNE visualizations of z and $\tilde{z}$ shown in Figure 4 further validate this as $\tilde{z}$ is clustered by y but has uniformly distributed s , while z shows distinct groups of s in each digit cluster.

5.2 Fairness through invariance to biasing factors

The Adult (Dheeru and Karra Taniskidou 2017) and German (Dheeru and Karra Taniskidou 2017) datasets have been popularly employed in fairness settings (Li, Swersky, and Zemel 2014; Louizos et al. 2016; Moyer et al. 2018; Xie et al. 2017). The former is an income dataset for classifying whether a person has more than \$50,000 savings and the biasing s is age. The latter is used for predicting whether a person has a good credit rating, and gender is the biasing s . We use architectures similar to VFAE for encoders (two lay-

Table 3: MNIST-ROT results (random chance of $s = 0.2$). Θ represents the angles seen during training, i.e., $\{0, \pm 22.5^\circ, \pm 45^\circ\}$.

Model	A_y			A_s
	Θ	$\pm 55^\circ$	$\pm 65^\circ$	Θ
NN+MMD	0.970 ± 0.001	0.831 ± 0.001	0.665 ± 0.002	0.380 ± 0.011
VFAE	0.953 ± 0.004	$\times$	$\times$	0.389 ± 0.076
CAI	0.958	0.829	0.663	0.384
CVIB	0.960 ± 0.008	0.819 ± 0.007	0.674 ± 0.009	0.382 ± 0.005
UAI	0.977	0.856	0.696	0.338
Ours	0.991 ± 0.001	0.863 ± 0.001	0.730 ± 0.001	0.201 ± 0.001
Δ over UAI	60.9%	4.86%	11.18%	99.3%

Table 4: Adult results (majority class of $s = 0.67$).

Model	A_y	A_s
NN+MMD	0.75 ± 0.00	0.67 ± 0.01
VFAE	0.76 ± 0.01	0.67 ± 0.01
CAI	0.83	0.89
CVIB	0.69 ± 0.01	0.68 ± 0.01
Ours	0.85 ± 0.00	0.67 ± 0.00
Δ over VFAE	37.5%	—

ers), predictors (one layer), and decoders (two layers), along with a two-layer discriminator. Results of experiments on these datasets are shown in Tables 4 and 5. The proposed model completely removes information about s , as reflected by A_s being the same as the population share of the majority s -class for both the datasets. Previous works have also achieved perfect score on A_s but the proposed model outperforms them on A_y on both the datasets, showing that it is able to retain more information for predicting y while being fair by successfully filtering only biasing factors.

5.3 Invariance in multi-task learning

The proposed framework is evaluated on the dSprites (Matthey et al. 2017) dataset of shapes with independent factors: color, shape, scale, orientation, and position. The dataset was preprocessed following (Higgins et al. 2018), resulting in two classes for scale and four for position and orientation each. Shape ($y^{(1)}$) and scale ($y^{(2)}$) are treated as the prediction tasks, where shape is desired to be invariant to position ($s^{(1)}$) and scale to orientation ($s^{(2)}$). We use the same component networks that we used for Extended Yale-B. We compare results with a version of our model without the decoder, the maskers and the discriminators, i.e., both $y^{(1)}$ and $y^{(2)}$ are predicted from z . Evaluation could not be conducted on NN+MMD, VFAE, CAI and CVIB because they have one $\tilde{z}$ and all tasks have to be invariant to the same s , and on UAI because it works only for a single y . Table 6 presents the results. The proposed framework achieves the same accuracies for $y^{(1)}$ and $y^{(2)}$ as the baseline, while maintaining random chance accuracies for $s^{(1)}$ (0.5) and $s^{(2)}$ (0.25) as opposed to significantly

Table 5: German results (majority class of $s = 0.8$).

Model	A_y	A_s
NN+MMD	0.74 ± 0.01	0.80 ± 0.00
VFAE	0.70 ± 0.00	0.80 ± 0.00
CAI	0.70	0.81
CVIB	0.74 ± 0.00	0.80 ± 0.00
Ours	0.76 ± 0.00	0.80 ± 0.00
Δ over CVIB	7.7%	—

Table 6: Results on dSprites in multi-task setting. Random chance of s for task #1 is 0.5 and for task #2 is 0.25.

Acc	Task #1		Task #2	
	Baseline	Ours	Baseline	Ours
A_y	0.99	0.99	0.99	0.99
A_s	0.94	0.50	0.40	0.25

higher corresponding scores for the baseline. Hence, our framework works effectively in multi-task settings.

6 Conclusion

We have presented a novel framework for invariance induction in neural networks through “forgetting” of information related to unwanted factors. We showed that the forget-gate used in the proposed framework acts as an information bottleneck and that adversarial training encourages generation of forget-masks that remove unwanted factors. Empirical results show that the proposed model exhibits state-of-the-art performance in both nuisance and bias settings.

Acknowledgements. This work is based on research sponsored by the Defense Advanced Research Projects Agency (agreement number FA8750-16-2-0204). The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

A Variance Inequality

For two random variables A and B , possibly dependent,

$$\text{Cov}(A, B) \leq \sqrt{\text{Var}(A)\text{Var}(B)} \leq \frac{1}{2}(\text{Var}(A) + \text{Var}(B)) \quad (12)$$

$$\text{Var}(A + B) = \text{Var}(A) + \text{Var}(B) + 2\text{Cov}(A, B) \leq 2(\text{Var}(A) + \text{Var}(B)) \quad (13)$$

Equation 12 holds due to the geometric mean–arithmetic mean inequality. Let $A = (m_i - \mathbb{E}[m_i])z_i$ and $B = \mathbb{E}[m_i]z_i$. Then, $\text{Var}(m_i z_i) = \text{Var}(A + B)$ gives:

$$\text{Var}(m_i z_i) \leq 2\text{Var}((m_i - \mathbb{E}[m_i])z_i) + 2\mathbb{E}[m_i]^2 \text{Var}(z_i) \quad (14)$$

References

- Achille, A., and Soatto, S. 2018a. Information dropout: Learning optimal representations through noisy computation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40(12):2897–2905.
- Achille, A., and Soatto, S. 2018b. Emergence of invariance and disentanglement in deep representations. *Journal of Machine Learning Research* 19(50):1–34.
- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35(8):1798–1828.
- Cover, T. M., and Thomas, J. A. 2012. *Elements of Information Theory*. John Wiley & Sons.
- Dheeru, D., and Karra Taniskidou, E. 2017. UCI machine learning repository.
- Domingos, P. 2012. A few useful things to know about machine learning. *Commun. ACM* 55(10):78–87.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17(1):2096–2030.
- Georghiades, A. S.; Belhumeur, P. N.; and Kriegman, D. J. 2001. From few to many: illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(6):643–660.
- Gretton, A.; Bousquet, O.; Smola, A.; and Schölkopf, B. 2005. Measuring statistical dependence with hilbert-schmidt norms. In Jain, S.; Simon, H. U.; and Tomita, E., eds., *Algorithmic Learning Theory*, 63–77. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Gretton, A.; Borgwardt, K. M.; Rasch, M.; Schölkopf, B.; and Smola, A. J. 2007. A kernel method for the two-sample-problem. In Schölkopf, B.; Platt, J. C.; and Hoffman, T., eds., *Advances in Neural Information Processing Systems 19*. MIT Press. 513–520.
- Higgins, I.; Sonnerat, N.; Matthey, L.; Pal, A.; Burgess, C. P.; Bošnjak, M.; Shanahan, M.; Botvinick, M.; Hassabis, D.; and Lerchner, A. 2018. Scan: Learning hierarchical compositional visual concepts.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Comput.* 9(8):1735–1780.
- Hsu, I.-H.; Jaiswal, A.; and Natarajan, P. 2019. Niesr: Nuisance invariant end-to-end speech recognition. *Proceedings of Interspeech 2019* 456–460.
- Jaiswal, A.; Wu, R. Y.; Abd-Almageed, W.; and Natarajan, P. 2018. Unsupervised adversarial invariance. In *Advances in Neural Information Processing Systems 31*. 5097–5107.
- Jaiswal, A.; Wu, Y.; AbdAlmageed, W.; and Natarajan, P. 2019a. Unified Adversarial Invariance. *arXiv preprint arXiv:1905.03629*.
- Jaiswal, A.; Xia, S.; Masi, I.; and AbdAlmageed, W. 2019b. RoPAD: Robust Presentation Attack Detection through Unsupervised Adversarial Invariance. In *12th IAPR International Conference on Biometrics (ICB)*.
- Kingma, D. P., and Welling, M. 2014. Auto-encoding Variational Bayes. In *International Conference on Learning Representations*.
- Ko, T.; Peddinti, V.; Povey, D.; and Khudanpur, S. 2015. Audio augmentation for speech recognition. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Lample, G.; Zeghidour, N.; Usunier, N.; Bordes, A.; Denoyer, L.; et al. 2017. Fader networks: Manipulating images by sliding attributes. In *Advances in Neural Information Processing Systems*, 5967–5976.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based Learning Applied to Document Recognition. *Proceedings of the IEEE* 86(11):2278–2324.
- Li, Y.; Swersky, K.; and Zemel, R. 2014. Learning unbiased features. *arXiv preprint arXiv:1412.5244*.
- Lopez, R.; Regier, J.; Jordan, M. I.; and Yosef, N. 2018. Information constraints on auto-encoding variational bayes. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*. 6117–6128.
- Louizos, C.; Swersky, K.; Li, Y.; Welling, M.; and Zeme, R. 2016. The variational fair autoencoder. In *Proceedings of International Conference on Learning Representations*.
- Matthey, L.; Higgins, I.; Hassabis, D.; and Lerchner, A. 2017. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>.
- Moyer, D.; Gao, S.; Brekelmans, R.; Galstyan, A.; and Ver Steeg, G. 2018. Invariant representations without adversarial training. In *Advances in Neural Information Processing Systems 31*. 9102–9111.
- Sabour, S.; Frosst, N.; and Hinton, G. E. 2017. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, 3859–3869.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 1999. The information bottleneck method. In *37th Annual Allerton Conference on Communication, Control and Computing*, 368–377.
- Xie, Q.; Dai, Z.; Du, Y.; Hovy, E.; and Neubig, G. 2017. Controllable invariance through adversarial feature learning. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30*. 585–596.
- Zemel, R.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning fair representations. In Dasgupta, S., and McAllester, D., eds., *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, 325–333. Atlanta, Georgia, USA: PMLR.
- Zhang, Y.; Lee, K.; and Lee, H. 2016. Augmenting supervised neural networks with unsupervised objectives for large-scale image classification. In *International Conference on Machine Learning*, 612–621.