

Achieving Robustness in the Wild via Adversarial Mixing with Disentangled Representations

Sven Gowal* Chongli Qin* Po-Sen Huang Taylan Cemgil
 DeepMind chongliqin@google.com posenhuang@google.com taylancemgil@google.com
 sgowal@google.com

Krishnamurthy (Dj) Dvijotham Timothy Mann Pushmeet Kohli
 dvij@google.com timothymann@google.com pushmeet@google.com

Abstract

Recent research has made the surprising finding that state-of-the-art deep learning models sometimes fail to generalize to small variations of the input. Adversarial training has been shown to be an effective approach to overcome this problem. However, its application has been limited to enforcing invariance to analytically defined transformations like ℓ_p -norm bounded perturbations. Such perturbations do not necessarily cover plausible real-world variations that preserve the semantics of the input (such as a change in lighting conditions). In this paper, we propose a novel approach to express and formalize robustness to these kinds of real-world transformations of the input. The two key ideas underlying our formulation are (1) leveraging disentangled representations of the input to define different factors of variations, and (2) generating new input images by adversarially composing the representations of different images. We use a StyleGAN model to demonstrate the efficacy of this framework. Specifically, we leverage the disentangled latent representations computed by a StyleGAN model to generate perturbations of an image that are similar to real-world variations (like adding make-up, or changing the skin-tone of a person) and train models to be invariant to these perturbations. Extensive experiments show that our method improves generalization and reduces the effect of spurious correlations (reducing the error rate of a “smile” detector by 21% for example).

1. Introduction

The principle by which neural networks are trained to minimize their average error on the training data is known as Empirical Risk Minimization (ERM) [1]. ERM has, for the most part, enabled breakthroughs in a wide variety of fields [2–4], and this success has led to the usage of neural networks in applications that are safety-critical [5]. ERM, however, is only guaranteed to produce meaningful models




Figure 1. Variations of the same faces. A model obtained through classical training classifies the same face as both “smiling” and “not smiling” (depending on the variations). Our model remains consistent in terms of classification. Note that these persons “do not exist” and have been generated using a StyleGAN model.

when the data encountered during training and deployment is drawn independently from the same distribution. When a mismatch between training and testing data occurs, models can fail in catastrophic ways; and, unfortunately, such occurrence is commonplace: training data is often collected through a biased process that highlights confounding factors and spurious correlations [6, 7], which can lead to undesirable consequences (e.g., <http://gendershades.org>).

The effects of such data shifts are largely detailed in the literature. For example, both Recht et al. [8] and Hendrycks et al. [9] show that the accuracy of IMAGENET models is severely impacted by changes in the data collection process. Methods to counteract such effect, which mainly consist of *data augmentation* techniques, also struggle. Training against corrupted data only forces the memorization of such corruptions and, as a result, these models fail to generalize to new corruptions [10, 11]. Works such as *mixup* [12] or *AutoAugment* [13] pave the way to further improvements,

but still require intricate fine-tuning to succeed in practice.

Another parallel and important line of work uncovered that the addition of small but carefully chosen deviations to the input, called adversarial perturbations, can cause the neural network to make incorrect predictions with high confidence [14–18]. Techniques to build models that are robust to adversarially perturbed examples, such as adversarial training [19], have received a significant amount of attention in the recent years [16, 20–22]. The existence of imperceptible perturbations that alter a model’s output demonstrates that supervised learning algorithms still fail to capture the true causal relationships between signal and label. The degradation of performance occurred when shifting between training and adversarial (or otherwise corrupted) distributions indicates that neural networks pick up on correlations that are not necessarily robust to small input perturbations [23]. The existence of imperceptible adversarial perturbations highlights just one form of spurious correlation that causes undesirable behaviors in the networks we train.

This paper focuses on training models that are robust to plausible real-world perturbations that preserve semantic content (such as those presented in Figure 1). We go beyond conventional *data augmentation* and *adversarial training* on l_p -norm bounded perturbations by leveraging high-quality generative models that can describe such perturbations. In particular, we address the question: “Given a generative model with a sufficiently good disentangled representation that aligns well with the perturbations of interest, can we train neural networks that are resistant to bias and spurious correlations present in the training data?” More specifically, we consider *StyleGAN* [24] as our underlying generative model. Our contributions are as follows:

1. We develop a framework dubbed *Adversarial Mixing with Disentangled Representations (AdvMix)* which leverages the disentangled latents of a generative model to train networks that are robust to real-world variations.
2. We demonstrate how to leverage *StyleGAN’s mixing* property to systematically transfer image attributes likely to be misclassified across image instances, thus allowing us to generate realistic worst-case semantic variations. This enables us to define semantic perturbations in a purely data-driven fashion, as opposed to methods that require data collection under different conditions [25].
3. We conduct extensive experiments on a controlled Color-MNIST dataset that compare *Adversarial Mixing with Disentangled Representations* with *random data augmentation* and demonstrate under which conditions *AdvMix* achieves higher accuracy.
4. Finally, we demonstrate empirically on CELEBA that accuracy is not necessarily at odds with robustness [26], once we consider semantic variations other than l_p -norm bounded variations.



Figure 2. Comparison of different data augmentation techniques. These transformations tend to destroy the image semantics.

2. Related work

Robustness to l_p -norm perturbations. Generating pixel-level adversarial perturbations has been and remains extensively studied [16, 18–20, 27, 28]. Most works focus the robustness of classifiers under l_p -norm bounded perturbations. In particular, it is expected that a *robust* classifier be invariant to small perturbations in the pixel space (as defined by the l_p -norm). Goodfellow et al. [16] and Madry et al. [19] laid down foundational principles to train robust networks, and recent works [29, 30] continue to find novel approaches to enhance robustness. While existing work is able to train models that are robust to imperceptible pixel-level variations, the study of robustness against semantically meaningful perturbations is largely under-explored.

Adversarial robustness beyond l_p -norm. Engstrom et al. [31] and Kanbak et al. [32] explored geometric transformations such as rotations and translation of images. Early works (e.g., Baluja and Fischer [33]) also demonstrated that it is possible to go beyond analytically defined variations by using generative models to create perturbations. Song et al. [34] and Xiao et al. [35] used a pre-trained AC-GAN [36] to generate perturbations; and they demonstrated that it is possible to generate semantically relevant perturbations for tasks such as MNIST, SVHN and CELEBA. Lastly, Qiu et al. [37] have attempted to generate adversarial examples by interpolating through the attribute space defined by a generative model. With the exception of [38], in which the authors strongly limit semantic variations by keeping the perturbed image close to its original counterpart, there has been little to no work demonstrating robustness to large semantically plausible variations. As such the effect of training models robust to such variations is unclear. To the best of our knowledge, this paper is the first to analyze the difference between *adversarial training* and *data augmentation* in the space of semantically meaningful variations.

Data augmentation Data augmentation can reduce generalization error. For image classification tasks, random flips, rotations and crops are commonly used [39]. More sophisticated techniques such as *Cutout* [40] (which produces random occlusions), *CutMix* [41] (which replaces parts of

an image with another) and *mixup* [12] (which linearly interpolates between two images) all demonstrate extremely compelling and surprising results. Indeed, while these methods often result in images that are visibly corrupted and void of semantic meaning (even to the human eye), the resulting models often achieve state-of-the-art accuracy across a wide range of datasets. Figure 2 shows a comparison of these different techniques. Some of these data augmentation techniques have been applied to latent representations of the input (rather than the input itself) [42]. However, these do not focus on the effect of data bias.

Causal reasoning using additional data. Heinze-Deml and Meinshausen [43] use grouped observations (e.g., the same object under different conditions) to discover variations that should not explain the classification label. More recently Arjovsky et al. [25] developed a method called Invariant Risk Minimization (IRM) which tries to find an invariant predictor across different environments (or groups of data points). Both methods were able to build classifiers that were less sensitive to spurious correlations, which, in turn, lead to classifiers that were less biased than classifiers trained purely on an original biased training set. However, they require explicitly annotated data collected under different environmental conditions.

3. Adversarial Mixing with Disentangled Representations

In this paper, we consider a model f_θ parametrized by θ . We would like our model to be robust or invariant to a set of transformations \mathcal{T} . Formally, our goal is to find the model parameters θ that minimize the *semantic adversarial risk*

$$\mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\max_{t \in \mathcal{T}} L(f_\theta(t(x)), y) \right], \quad (1)$$

where $\mathcal{D} \subset \mathcal{X} \times \mathcal{Y}$ is a data distribution over pairs of examples x and corresponding labels y , and L is a suitable loss function (such as the 0 – 1 loss in the context of classification tasks). The set of semantic transformations \mathcal{T} contains functions of the form $t : \mathcal{X} \rightarrow \mathcal{X}$. Each element $t \in \mathcal{T}$ is irreducible and, crucially, for the optimal classifier $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, we would like that $f_\theta(t(x)) = f_\theta(x)$ for all $t \in \mathcal{T}$. For example, an MNIST classifier should not be affected by changes in the digit color. In the following, we define a set of transformations \mathcal{T} via a decoder that leverages a disentangled latent representation and explain how to evaluate the resulting risk in Equation (1).

Invariant latent factors. Disentanglement is perceived as a desirable property of representations. Often, one hopes to obtain a representation of the observed data $x \in \mathcal{X}$ in terms of separate and conditionally independent factors $z \in$

\mathcal{Z} given x under a certain class of input transformations [44]. In our particular setting, we will assume a task-specific disentangled representation. Formally, we assume that we have an *ideal* generator (or decoder), $\text{dec} : \mathcal{Z} \rightarrow \mathcal{X}$, where the latent space \mathcal{Z} is a product space of the form $\mathcal{Z} = \mathcal{Z}_\parallel \times \mathcal{Z}_\perp$. For a given classification task that predicts the label y , only the coordinates corresponding to \mathcal{Z}_\parallel are relevant, while \mathcal{Z}_\perp is irrelevant. We formalize the above notions using conditional independence: given an example $x = \text{dec}(z_\parallel, z_\perp)$ with $z_\perp \in \mathcal{Z}_\perp$, $z_\parallel \in \mathcal{Z}_\parallel$ and corresponding label $y \in \mathcal{Y}$, we have

$$\mathbb{P}(y|z_\parallel, z_\perp) = \mathbb{P}(y|z_\parallel). \quad (2)$$

Hence, the ideal invariant classifier f^* that outputs a probability distribution over \mathcal{Y} should be consistent with the invariance assumption

$$f^*(\text{dec}(z_\parallel, z_\perp)) = f^*(\text{dec}(z_\parallel, \tilde{z}_\perp)) \quad (3)$$

for all $\tilde{z}_\perp \in \mathcal{Z}_\perp$, and should output the correct label:

$$\operatorname{argmax}_{y' \in \mathcal{Y}} f^*(\text{dec}(z_\parallel, z_\perp)) = y. \quad (4)$$

Finally, referring back to Equation (1), we define the set of transforms \mathcal{T} that induce semantically irrelevant perturbations as:

$$\mathcal{T} = \{t \mid t(x) = \text{dec}(z_\parallel, \tilde{z}_\perp) \text{ with } \tilde{z}_\perp \in \mathcal{Z}_\perp \\ \text{s.t. } \exists z_\perp x = \text{dec}(z_\parallel, z_\perp)\}. \quad (5)$$

Adversarial training. Given a model f_θ with enough capacity, minimizing the *semantic adversarial risk* in Equation (1) results in parameters θ^*

$$\theta^* = \operatorname{argmin}_\theta \mathbb{E}_{\substack{(x,y) \sim \mathcal{D} \\ x = \text{dec}(z_\parallel, z_\perp)}} \left[\max_{\tilde{z}_\perp \in \mathcal{Z}_\perp} L(f_\theta(\text{dec}(z_\parallel, \tilde{z}_\perp)), y) \right] \quad (6)$$

that satisfy Equations (3) and (4). In other words, there exists no transformation $t \in \mathcal{T}$ that, when applied to x , would result in a misclassification of the optimal classifier $f^* = f_{\theta^*}$. Solving the saddle point problem in Equation (6) requires solving the corresponding inner-maximization problem

$$\tilde{z}_\perp^* = \operatorname{argmax}_{\tilde{z}_\perp \in \mathcal{Z}_\perp} L(f_\theta(\text{dec}(z_\parallel, \tilde{z}_\perp)), y). \quad (7)$$

As enumerating all possible latents $\tilde{z}_\perp \in \mathcal{Z}_\perp$ is often intractable, we resort to a technique popularized by Madry et al. [19] in the context of adversarial training, which consists of using projected gradient ascent on a differentiable surrogate loss. For a classification task, the 0 – 1 loss is replaced with the cross-entropy loss:

$$\hat{L}(f_\theta(x), y) = -\log([f_\theta(x)]_y) \quad (8)$$




Figure 3. Illustration of the maximization process in Equation (9).

where $[a]_i$ returns the i -th coordinate of a . Gradient ascent steps are then interleaved with projection steps for a given number of iterations K . Formally, we find an estimate $\hat{z}_\perp^{(K)}$ of \hat{z}_\perp^* using the following recursion:

$$\hat{z}_\perp^{(k+1)} = \text{proj}_{\mathcal{Z}_\perp} \left(\hat{z}_\perp^{(k)} + \alpha \nabla_{\hat{z}_\perp^{(k)}} \hat{L}(f_\theta(\text{dec}(z_\parallel, \hat{z}_\perp^{(k)}), y)) \right) \quad (9)$$

where $\hat{z}_\perp^{(0)}$ is chosen at random within \mathcal{Z}_\perp , α is a constant step-size and $\text{proj}_{\mathcal{A}}(a)$ is a projection operator that project a onto \mathcal{A} . Figure 3 illustrates the process.

Ultimately, *Adversarial Mixing with Disentangled Representations* (shortened as *AdvMix*) tries to find parameters that minimize the worst-case loss that could arise from altering the input examples through plausible transformations. It guarantees that transformations of the input are meaningful by using a disentangled latent representation that encodes independent controllable factors, where some of these factors are known to be independent from the label. Finding such a disentangled representation is rarely possible, as it is not always known which variations of the input should or should not affect the label. In some cases, however, it is possible to train generative models such that we expect some subset of the latents to not affect the label. Section 4 implements *AdvMix* using a *StyleGAN* model.

Data with low density regions. The motivation behind *AdvMix* stems from the *manifold hypothesis* [45]. It states that high dimensional data present in the real-world, such as images, often lies on a low-dimensional manifold. As a consequence, there exists large regions in the input space that are outside the support of the data distribution. Hence, for maximal efficiency, *data augmentation* and *adversarial training* should be done carefully to make sure that the augmented data is still within the support of the original data distribution. Data augmentation techniques presented in Figure 2 clearly violate this condition, and despite their success, we cannot expect that they perform well across all datasets (in fact, *mixup* performs poorly on Color-MNIST). Similarly,




Figure 4. Comparison of *mixup* and *AdvMix* on a toy example. In this example, we are given 200 datapoints. Each data point (x_1, x_2) is sampled according to $x_1 \sim \mathcal{N}(z_\perp, \sqrt{3})$ where $z_\perp \in \mathcal{Z}_\perp = \{0., 10.\}$ and $x_2 \sim \mathcal{N}(z_\parallel, 1)$ where $z_\parallel \in \mathcal{Z}_\parallel = \{0., 20.\}$. The colors represent the label. Note that the latent variable $z_\parallel = 20y$ is dependent on the label while z_\perp is independent of the label. Panel (a) shows the original set of 200 datapoints; panel (b) shows the effect of sampling additional data using *AdvMix*; and panel (c) shows the effect of *mixup*. Of course, we should point out that our method, *AdvMix*, is aware of the underlying latent representation, while *mixup* is not.

adversarial training targeting ℓ_p -norm bounded perturbations tend to trade-off accuracy for robustness [23]. Figure 4 compares *mixup* and *AdvMix* on a toy example. In this example, we artificially construct a dataset with two classes and an underlying disentangled latent representation. We observe that by exploiting the knowledge of the disentangled latent representation, *AdvMix* is capable of generating additional datapoints that are consistent with the original dataset, while *mixup* generates additional datapoints that are unlikely.

Relationship to *mixup*. *mixup* augments data with respect to the input space. Given two pairs of inputs (x_A, y_A) , (x_B, y_B) and a linear interpolation factor sampled from a β -distribution $\lambda \sim \beta(\alpha, \alpha)$, *mixup* generate a new input pair as follows:

$$\begin{aligned} \tilde{x} &= \lambda x_A + (1 - \lambda) x_B \\ \tilde{y} &= \lambda y_A + (1 - \lambda) y_B. \end{aligned} \quad (10)$$

Our methodology combines inputs (x_A, y_A) and (x_B, y_B) in the latent space. If $x_A = \text{dec}(z_{A\parallel}, z_{A\perp})$ and $x_B = \text{dec}(z_{B\parallel}, z_{B\perp})$, we obtain

$$\begin{aligned} \tilde{x} &= \text{dec}(z_{A\parallel}, z_{B\perp}) \\ \tilde{y} &= y_A. \end{aligned} \quad (11)$$

Crucially, this combination only affects the latent sub-space that is independent from the label, thus the label remains unchanged. We also note that, unlike [42], no interpolation occurs in the latent space (i.e., $\lambda z_{A\perp} + (1 - \lambda) z_{B\perp}$) as this could result in points that are outside \mathcal{Z}_\perp when \mathcal{Z}_\perp is not convex.

Relationship to Invariant Risk Minimization. Arjovsky et al. [25] consider the case where we have multiple datasets

$D_e = \{x_i, y_i\}_{i=1}^n$ drawn from different training environments $e \in \mathcal{E}$. As explained in [25], the motivation behind IRM is to minimize the worst-case risk

$$\max_{e \in \mathcal{E}} \mathbb{E}_{(x,y) \in D_e} [L(f_\theta(x), y)]. \quad (12)$$

In this paper, the environments are defined by the different instances of $z_\perp \in \mathcal{Z}_\perp$. Given a dataset $\{\text{dec}(z_{i\parallel}, z_{i\perp}), y_i\}_{i=1}^n$, we can rewrite the *semantic adversarial risk* shown in Equation (1) as Equation (12) by setting the environment set \mathcal{E} to

$$\mathcal{E} = \{\{\text{dec}(z_{i\parallel}, z_\perp), y_i\}_{i=1}^n \mid z_\perp \in \mathcal{Z}_\perp\}. \quad (13)$$

This effectively create an ensemble of datasets for all possible combinations of $z_\perp \in \mathcal{Z}_\perp$ for all examples.

The crucial difference between IRM and *AdvMix* is in the formulation of the risk. While IRM computes the risk by enumerating over a countable set of environments and picking the worst-case, *AdvMix* attempts to compute the worst-case risk by finding the combination of variations that maximize the risk over all examples.

4. Implementation using *StyleGAN*

So far, we have assumed the presence of a generator (or decoder) that is capable of using a perfectly disentangled latent representation: we have assumed that this representation is partitioned into two subsets, one of which is known to be independent from the target label. In practice, the methodology is often reversed: generative models are trained in the hope of obtaining some level of disentanglement. If a partition of the trained latent space does not influence the label, we can use the corresponding trained generator within *AdvMix*. This section explains why *StyleGAN* is a good candidate and details how to implement *AdvMix* using *StyleGAN*. In particular, as we rely on *StyleGAN*'s *mixing* property to enforce a partitioning of the latents, only three elements are needed: (i) a transformation set \mathcal{Z}_\perp from which label-independent variants \tilde{z}_\perp can be chosen; (ii) a dataset $\mathcal{D} = \{z_{i\parallel}, y_i\}_{i=1}^n$ of latents and labels; and (iii) an efficient method to find a worst-case variation $\tilde{z}_\perp \in \mathcal{Z}_\perp$.

StyleGAN. *StyleGAN* is a generator architecture for generative adversarial networks proposed by Karras et al. [24]. It borrows interesting properties from the style transfer literature [46]. In this work, we rely on the style mixing property. Formally, the *StyleGAN* architecture is composed of two stages. The first stage takes a latent variable $z \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ that is not necessarily disentangled and projects it into a disentangled latent space $z = \text{map}(z)$. The second stage synthesizes an image x from the disentangled latents z using a decoder $x = \text{dec}(z)$. Overall, the process of generating an image x using a *StyleGAN* network is defined as

$$x = \text{dec} \circ \text{map}(z) \quad \text{where } z \sim \mathcal{N}(\mathbf{0}, \mathbf{1}). \quad (14)$$

The intermediate latent variables z provide some level of disentanglement that affects image generation at different spatial resolutions which allows us to control the synthesis of an image. Particularly, we can apply the ‘‘style’’ of an image to another by mixing the disentangled latents of these images together. In the context of face generation, the styles corresponding to coarse spatial resolutions affect high-level aspects such as pose, and styles of fine resolutions affect mainly the color scheme. In the rest of this manuscript, we focus on variations of the finer style.¹ Concretely, our experiments in Section 5 assume that the fine attributes z_\perp are label-independent, while the coarse attributes z_\parallel may be label-dependent. Consequently, the finer style $z_{B\perp}$ of an image x_B can be applied to another image $x_A = \text{dec}(z_{A\parallel}, z_{A\perp})$ via $\text{dec}(z_{A\parallel}, z_{B\perp})$. Figure 5b shows a nominal image and two variations of that image obtained by mixing the finer style of two other images.

Definition of the transformation set. For completeness, we now define the set of transforms \mathcal{T} in Equation (5) by defining \mathcal{Z}_\perp . While the formulation of *StyleGAN* allows z to be sampled within an infinite support, our formulation requires \mathcal{Z}_\perp to be bounded. Additionally, as explained by Nalisnick et al. [47], due to concentration of measure, a generative model usually draws samples from its typical set [48] (a subset of the model’s full support) rather than regions of high probability density.² As such, if $z \in \mathbb{R}^d$, we wish to define \mathcal{Z}_\perp as follows:

$$\mathcal{Z}_\perp = \left\{ \text{map}(z)_\perp \mid \sqrt{d} - \delta d^{\frac{1}{4}} \leq \|z\|_2 \leq \sqrt{d} + \delta d^{\frac{1}{4}} \right\} \quad (15)$$

where δ is a small tunable positive constant. In practice, however, we do not want to backpropagate through the *map* operation as it is inefficient. Instead, a small collection of latents is sampled, passed through the *map* operation, and \mathcal{Z}_\perp is limited to a neighborhood of the points in this collection. This collection is re-sampled for each example and in expectation covers the typical set well (more details are given in Algorithm 2).

Construction of a dataset of disentangled latents. Constructing a dataset of labelled latents $\mathcal{D} = \{z_{i\parallel}, y_i\}_{i=1}^n$ requires finding the latents z_i that decode into each example x_i of an original labelled dataset $\{x_i, y_i\}_{i=1}^n$. Hence, we need to find a mapping between the image space and the latent space. This mapping, which can be computed offline, is used to construct the dataset \mathcal{D} , and is only required once for each new dataset. Specifically, this mapping is denoted as $\text{enc} : \mathcal{X} \mapsto \mathcal{Z}$ and finds z_i such that $x_i \approx \text{dec}(z_i)$.

¹Other variations are possible as long as we have the certainty that they do not affect the label of interest.

²For d -dimensional isotropic Gaussian with standard deviation σ , the typical set resides at a distance of $\sigma\sqrt{d}$ from the mode [49].



Figure 5. Panel **a** shows how the latents are progressively able to match a target image (on the far right). Panel **b** shows two different variations of the obtained image.

Algorithm 1 Encoder enc

Input: Target image x , trained *StyleGAN* model $\text{dec} \circ \text{map}$, and trained VGG network vgg . α_i and β_i are hyperparameters all set to 1 and 1/5 respectively. $\gamma^{(k)}$ is a step-size schedule.

Output: Disentangled latents \hat{z} such that $\text{dec}(\hat{z}) \approx x$

- 1: $\hat{z} \leftarrow \frac{1}{M} \sum_{i=1}^M \text{map}(z^{(i)})$ with $z^{(i)} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ \triangleright Average latents
- 2: **for** $k \in \{1, \dots, N\}$ **do** $\triangleright N$ is the number of iterations
- 3: $\hat{x} = \text{dec}(\hat{z})$
- 4: $\hat{\mathcal{A}} = \text{vgg}(\hat{x})$ $\triangleright \hat{\mathcal{A}}$ is a list of activations (after the 2nd convolution of 1st, 2nd and 3rd blocks)
- 5: $\mathcal{A} = \text{vgg}(x)$
- 6: $\mathcal{A}_{\text{mix}} = \text{vgg}(\text{dec}(\hat{z}_{\parallel}, \text{map}(z_{\perp})))$ with $z \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$
- 7: $L_{\text{reconstruct}} = \alpha_0 \|\hat{x} - x\|_2^2 + \sum_{i=1}^{|\mathcal{A}|} \alpha_i \|\hat{\mathcal{A}}_i - \mathcal{A}_i\|_2^2$ \triangleright Reconstruction loss
- 8: $L_{\text{mix}} = \sum_{i=1}^{|\mathcal{A}|} \beta_i \|\mathcal{A}_{\text{mix}, i} - \mathcal{A}_i\|_2^2$ \triangleright Mixing loss
- 9: $\hat{z} \leftarrow \hat{z} - \gamma^{(k)} \nabla_{\hat{z}} (L_{\text{reconstruct}} + L_{\text{mix}})$
- 10: **end for**

Algorithm 1 defines this mapping through an optimization process. Inspired by [50], and rather than relying solely on the distance between pixel values to define the loss of that optimization, we use the perceptual loss [51, 52] – which helps steer the optimization process. The perceptual loss is defined on the intermediate activations of a trained VGG-16 network [53] (see line 7). We also found that the *StyleGAN* generator, dec , is a surjective mapping between its disentangled latent space and the image space (i.e., multiple latents can decode into the same image). Hence, since we heavily rely on the mixing property of *StyleGAN*, and to the contrary of [50], we propose to add an additional component to the loss that steers the latents towards a subset of latents that can be mixed. In particular, we add a perceptual loss between the synthesized image and a mixed version of the same image (see lines 6 and 8). Figure 5 shows the evolution of the optimization process as well as mixed variants of the resulting image.

Generating worst-case examples to train robust models.

As explained in Section 3, minimizing the *semantic adversarial risk* requires solving an inner-maximization problem. We rely on projected gradient ascent on the cross-entropy loss \hat{L} to efficiently find perturbed latents $\tilde{z}_{\perp} \in \mathcal{Z}_{\perp}$ such that, when mixed with z_{\parallel} , make the classifier output a label other than the true label. Algorithm 2 illustrates the process. This algorithm approximates the typical set in Equation (15) by randomly sampling initial latents $\tilde{z}_{\perp}^{(0)}$ N_r times and project-

Algorithm 2 Solution to Equation (7)

Input: A nominal input x and label y , a model f_{θ} , a *StyleGAN* model $\text{dec} \circ \text{map}$ and an encoder enc . L is the 0 – 1 loss and \hat{L} is the cross-entropy loss.

Output: Possible misclassified example \tilde{x}

- 1: $\tilde{x} \leftarrow x$
- 2: $[z_{\parallel}, z_{\perp}] = \text{enc}(x)$ \triangleright See Algorithm 1
- 3: **for** $r \in \{1, \dots, N_r\}$ **do** \triangleright Repeat N_r times
- 4: $\tilde{z}_{\perp}^{(0)} \leftarrow \text{map}(z_{\perp})$ with $z \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ \triangleright Initial latents
- 5: $\tilde{x}^{(0)} = \text{dec}(z_{\parallel}, \tilde{z}_{\perp}^{(0)})$
- 6: **for** $k \in \{1, \dots, K\}$ **do** $\triangleright K$ is the number of optimization steps
- 7: $\tilde{z}_{\perp}^{(k)} \leftarrow \text{proj} \left(\tilde{z}_{\perp}^{(k-1)} + \alpha \nabla_{\tilde{z}_{\perp}^{(k-1)}} \hat{L}(f_{\theta}(\tilde{x}^{(0)}), y) \right)$
- 8: $\tilde{x}^{(k)} = \text{dec}(z_{\parallel}, \tilde{z}_{\perp}^{(k)})$
- 9: **if** $L(f_{\theta}(\tilde{x}^{(k)}), y) > L(f_{\theta}(\tilde{x}), y)$ **then**
- 10: $\tilde{x} \leftarrow \tilde{x}^{(k)}$
- 11: **return** \tilde{x} \triangleright Since L is the 0 – 1 loss, the procedure can terminate early
- 12: **end if**
- 13: **end for**
- 14: **end for**

ing intermediate solutions $\tilde{z}_{\perp}^{(k)}$ back onto a neighborhood of $\tilde{z}_{\perp}^{(0)}$.³ It refines the initial latents using gradient ascent with the goal of finding latents $\tilde{z}_{\perp}^{(K)}$ that, when mixed with the original image latents z_{\parallel} , generate an image $\text{dec}(z_{\parallel}, \tilde{z}_{\perp}^{(K)})$ that is misclassified. Figure 1 shows the result of this optimization procedure where the original image (on the top-left) is classified as “not smiling” and the optimized image (on the bottom-left) is classified as “smiling”. Once perturbed latents $\tilde{z}_{\perp} = \tilde{z}_{\perp}^{(K)}$ are found, we can compute the cross-entropy loss on the image generated by $\text{dec}(z_{\parallel}, \tilde{z}_{\perp})$. Formally, for a classifier f_{θ} and a dataset $\mathcal{D} = \{z_{i\parallel}, y_i\}_{i=1}^n$, we want to solve

$$\text{argmin}_{\theta} \mathbb{E}_{z_{i\parallel}, y_i \sim \mathcal{D}} [L(f_{\theta}(\text{dec}(z_{i\parallel}, \tilde{z}_{\perp})), y_i)] \quad (16)$$

$$\text{and } \tilde{z}_{\perp} = \text{argmax}_{z_{\perp} \in \mathcal{Z}_{\perp}} L(f_{\theta}(\text{dec}(z_{i\parallel}, z_{\perp})), y_i).$$

Random mixing with disentangled representations.

While this section describes an instantiation of *AdvMix* using *StyleGAN*, it is possible to formulate an equivalent random *data augmentation* baseline. For an input x , we generate a

³The actual implementation used in the experimental section projects back onto a ℓ_{∞} -bounded neighborhood around $\tilde{z}_{\perp}^{(0)}$: $\{z_{\perp} \mid \|\tilde{z}_{\perp}^{(0)} - z_{\perp}\|_{\infty} < \epsilon\}$ where ϵ is set to 0.03.



Figure 6. Mean colors given to each digit in the training set of our Color-MNIST case-study.

random variation as follows:

$$\tilde{x} = \text{dec}(\text{enc}(x)_{\parallel}, \text{map}(z)_{\perp}) \text{ with } z \sim \mathcal{N}(\mathbf{0}, \mathbf{1}) \quad (17)$$

5. Results

In this section, we compare *AdvMix* to (i) nominal training which minimizes the *empirical risk*, (ii) Adversarial Training (AT) which minimizes the adversarial risk over ℓ_{∞} -norm bounded perturbations of size ϵ in input space [19], and (iii) Random Mixing with Disentangled Representations (*RandMix*) which minimizes the *vicinal risk* by randomly sampling latents from \mathcal{Z}_{\perp} (rather than systematically finding the worst-case variations). We perform two experiments to assess the generalization abilities of *AdvMix*. The first experiment is done on an artificially constructed dataset called Color-MNIST (it bares resemblance to the Color-MNIST experiments present in [25]). The second experiment uses CELEBA. Both experiment demonstrate that methods using semantic variations as expressed by a trained *StyleGAN* model achieve higher accuracy. It also demonstrates that, when the distribution of variations is skewed (i.e., some variations z_{\perp} appear more often than others in the dataset used to train the *StyleGAN* model), *AdvMix* obtains higher accuracy than *RandMix*. For both experiments, we train a truncated VGG network with 5 layers using 5 epochs on Color-MNIST and 20 epochs on CELEBA. We use the Adam [54] optimizer with a learning rate of 10^{-3} . *AdvMix* is trained with N_r set to 5.

5.1. Color-MNIST

Color-MNIST consists of a dataset of MNIST [55] digits that are artificially colored to emphasize bias. On the training set, we color each pair (x, y) of the original MNIST dataset with a color drawn randomly from a normal distribution with mean μ_y and standard deviation σ (means μ_y for $y \in \{0, \dots, 9\}$ are shown in Figure 6). On the test set, we color digits uniformly at random. In other words, the colors present in the training set spuriously correlate with the label. We can use σ to affect this correlation: by progressively increasing σ the dataset becomes less biased. For all techniques (including *mixup*), we vary the level of bias and train models using 5 epochs. The *StyleGAN* model is trained on the training set only, once for each setting of σ . The disentangled latents defining the finer style correspond to the final resolution of 32×32 .⁴

⁴32 corresponds to the next power of two after 28 which is the size of an image in the original MNIST dataset.




Figure 7. Accuracy of different training methods on images from our unbiased Color-MNIST test set. The training set is progressively debiased by increasing the standard deviation of the colors present.

Table 1. Effect of bias when training a *StyleGAN* model on our Color-MNIST dataset.

Method	Test accuracy on clean images		
	Unbiased	Less biased	More biased
<i>RandMix</i>	99.11%	98.87%	97.63%
<i>AdvMix</i>	99.19%	99.07%	98.79%

Figure 7 shows the results. Across all settings, *RandMix* and *AdvMix* outperform the other methods. As expected, the gap between all methods decreases as the training set becomes less biased. It is also worth noting that AT is useful (compared to nominal training and *mixup*) as on this dataset ℓ_{∞} -norm bounded perturbations allow the exploration of slight variations in colors. *RandMix* and *AdvMix* are both expected to do well as all variations z_{\perp} (that correspond to applications of different colors) are equally likely to be drawn from the *StyleGAN* model (since they are uniformly distributed in the training set).

To further emphasize the difference between *RandMix* and *AdvMix*, we purposefully bias the training of the *StyleGAN* model. We create two additional datasets (with $\sigma = 0$). With the first dataset (named “more biased”), the *StyleGAN* model is trained on a large fraction of zeros (and few other digits), while on the second dataset (named “less biased”, the *StyleGAN* model is trained on a large fraction of zeros and ones. As a result, rarely occurring variations (colors of digits from 1 to 9 for the first dataset and colors of digits from 2 to 9 for the second) are less likely to be randomly selected by *RandMix*. Table 1 shows the results. We observe that *AdvMix* performs better. However, we note that the gap is not large, as all color variations all contain red, green and blue components (which allows the network to implicitly learn about other color combinations).

Finally, to create a stronger effect, we limit digits to the

Table 2. Effect of bias when training a *StyleGAN* model on our RGB Color-MNIST dataset (limited to red, blue or green colors). The classifier is a linear model (instead of a convolutional network).

Method	Test accuracy on clean images		
	Unbiased	99% red Less biased	99.9% red More biased
<i>RandMix</i>	88.55%	83.18%	53.56%
<i>AdvMix</i>	85.07%	85.02%	85.00%

red, green and blue colors only (resulting in new datasets), and use a linear classifier (instead of a truncated VGG network). Table 2 demonstrates that, when the *StyleGAN* model is trained with a significant proportion of red digits, *AdvMix* does much better. Indeed, *AdvMix* is able to systematically find the corner cases (i.e., green and blue variations) that are currently misclassified rather than relying on the random sampling of such cases. We note that adversarial training can result in unstable learning, which can explain why *RandMix* does slightly better when the *StyleGAN* model is unbiased.

5.2. CELEBA

CELEBA [56] is a large-scale public dataset with forty different face attribute annotations including whether a person smiles or wears a hat. We make no modifications to the dataset and use a pretrained *StyleGAN* model⁵. For all techniques, we train models using 20 epochs. We evaluate all methods on their ability to classify the “smiling” attribute, as well as three other attributes.⁶ In this experiment, the disentangled latents defining the finer style correspond to resolutions ranging from 128×128 to 1024×1024 .⁷

In Table 3, we observe that *AdvMix* is the only method that systematically achieves high accuracy. This clearly demonstrates *AdvMix* can lead to a lower generalization error. It is also interesting to see that *RandMix* does not always improve on nominal training and that AT consistently trades off clean accuracy for ℓ_∞ -robustness (as seen in [23]). Finally, Figure 8 shows qualitative examples of images that are all correctly classified by the nominal model, but for which we can find plausible variants that are misclassified. Appendix B shows more results and includes other data augmentation schemes.

Overall, these results are confirming the observations made on the Color-MNIST dataset. They seem to indicate that there is a slightly distributional shift between CELEBA’s train and test sets (at least when it comes to the finer image style). By systematically probing variations that are difficult to classify, *AdvMix* is able to overcome this shift and reach

⁵<https://github.com/NVLabs/stylegan>

⁶Due to their sensitive connotation, we purposefully anonymized other attribute names and picked them from easier to harder classification tasks.

⁷We always re-scale the resulting images to a 64×64 resolution.

Table 3. Test accuracy on different classification tasks of the CELEBA dataset.

Method	Test accuracy on attribute			
	#1	#2 (smiling)	#3	#4
Nominal	96.49%	90.22%	83.52%	78.05%
AT ($\epsilon = 4/255$)	95.34%	91.11%	81.43%	76.61%
AT ($\epsilon = 8/255$)	95.22%	89.29%	79.46%	74.39%
<i>RandMix</i>	96.70%	90.36%	84.49%	76.41%
<i>AdvMix</i>	97.56%	92.29%	85.65%	79.47%



Figure 8. The top row shows examples of clean images from CELEBA that are all classified correctly by the nominal model. The bottom row shows semantically plausible variants of these images that are all misclassified.

better classification accuracy (to the contrary of *RandMix* which can only stumble on difficult variants by chance).

6. Conclusion

We have demonstrated a novel approach to achieving robustness to input variations encountered in the real world by generating adversarial instances that compose disentangled representations. We have shown how this framework can be realized by leveraging the *StyleGAN* architecture – resulting in models that are not only robust to systematic evaluation of insensitivity to variations but also exhibit better generalization, demonstrating that that accuracy is not necessarily at odds with robustness. Our formulation relies on good generative models that can learn a disentangled representation from which some directions are orthogonal to the label we are trying to predict. Methods such as *AdvMix* are intended to be used to reduce the effect of bias and spurious correlations on classifiers.⁸ We hope the promising results shown in this paper encourage the development of more effective disentangled representations that cover most factors of variations encountered in the real world. Finally, we hope this work leads to the exploration of this paradigm in the context of other Computer Vision applications and leads to the development of robust perception systems that can be safely used in the real world.

⁸It may be possible to use *AdvMix* to increase classification bias. However, other approaches are likely more effective to achieve this.

References

- [1] V. Vapnik, “Statistical learning theory,” 1998. [1](#)
- [2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org> [1](#)
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [4] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and others, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012. [1](#)
- [5] K. D. Julian, J. Lopez, J. S. Brush, M. P. Owen, and M. J. Kochenderfer, “Policy compression for aircraft collision avoidance systems,” in *IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*. IEEE, 2016, pp. 1–10. [1](#)
- [6] A. Torralba, A. A. Efros *et al.*, “Unbiased look at dataset bias,” in *CVPR*, vol. 1, no. 2. Citeseer, 2011, p. 7. [1](#)
- [7] A. Kuehlikamp, B. Becker, and K. Bowyer, “Gender-from-iris or gender-from-mascara?” in *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2017, pp. 1151–1159. [1](#)
- [8] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, “Do imagenet classifiers generalize to imagenet?” *arXiv preprint arXiv:1902.10811*, 2019. [1](#)
- [9] D. Hendrycks, K. Zhao, S. Basart, J. Steinhardt, and D. Song, “Natural adversarial examples,” *arXiv preprint arXiv:1907.07174*, 2019. [1](#)
- [10] I. Vasiljevic, A. Chakrabarti, and G. Shakhnarovich, “Examining the impact of blur on recognition by convolutional networks,” *arXiv preprint arXiv:1611.05760*, 2016. [1](#)
- [11] R. Geirhos, C. R. Temme, J. Rauber, H. H. Schütt, M. Bethge, and F. A. Wichmann, “Generalisation in humans and deep neural networks,” in *Advances in Neural Information Processing Systems*, 2018, pp. 7538–7550. [1](#)
- [12] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, 2017. [1](#), [3](#)
- [13] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le, “Autoaugment: Learning augmentation policies from data,” *arXiv preprint arXiv:1805.09501*, 2018. [1](#)
- [14] N. Carlini and D. Wagner, “Adversarial examples are not easily detected: Bypassing ten detection methods,” in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*. ACM, 2017, pp. 3–14. [2](#)
- [15] —, “Towards evaluating the robustness of neural networks,” in *2017 IEEE Symposium on Security and Privacy*. IEEE, 2017, pp. 39–57.
- [16] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014. [2](#)
- [17] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial machine learning at scale,” *arXiv preprint arXiv:1611.01236*, 2016.
- [18] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2013. [2](#)
- [19] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017. [2](#), [3](#), [7](#)
- [20] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, “Distillation as a defense to adversarial perturbations against deep neural networks,” *arXiv preprint arXiv:1511.04508*, 2015. [2](#)
- [21] H. Kannan, A. Kurakin, and I. Goodfellow, “Adversarial Logit Pairing,” *arXiv preprint arXiv:1803.06373*, 2018.
- [22] C. Xie, Y. Wu, L. van der Maaten, A. Yuille, and K. He, “Feature denoising for improving adversarial robustness,” *arXiv preprint arXiv:1812.03411*, 2018. [2](#)
- [23] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Adversarial examples are not bugs, they are features,” *arXiv preprint arXiv:1905.02175*, 2019. [2](#), [4](#), [8](#)
- [24] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410. [2](#), [5](#)
- [25] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz, “Invariant risk minimization,” *arXiv preprint arXiv:1907.02893*, 2019. [2](#), [3](#), [4](#), [5](#), [7](#)
- [26] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, “Robustness may be at odds with accuracy,” *arXiv preprint arXiv:1805.12152*, 2018. [2](#)
- [27] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *arXiv preprint arXiv:1607.02533*, 2016. [2](#)
- [28] S.-M. Moosavi-Dezfooli, A. Fawzi, J. Uesato, and P. Frossard, “Robustness via curvature regularization, and vice versa,” *arXiv preprint arXiv:1811.09716*, 2018. [2](#)
- [29] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, “Theoretically principled trade-off between robustness and accuracy,” *arXiv preprint arXiv:1901.08573*, 2019. [2](#)
- [30] C. Qin, J. Martens, S. Gowal, D. Krishnan, A. Fawzi, S. De, R. Stanforth, P. Kohli *et al.*, “Adversarial robustness through local linearization,” *arXiv preprint arXiv:1907.02610*, 2019. [2](#)
- [31] L. Engstrom, B. Tran, D. Tsipras, L. Schmidt, and A. Madry, “A rotation and a translation suffice: Fooling cnns with simple transformations,” *arXiv preprint arXiv:1712.02779*, 2017. [2](#)
- [32] C. Kanbak, S.-M. Moosavi-Dezfooli, and P. Frossard, “Geometric robustness of deep networks: analysis and improvement,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4441–4449. [2](#)
- [33] S. Baluja and I. Fischer, “Adversarial transformation networks: Learning to generate adversarial examples,” *arXiv preprint arXiv:1703.09387*, 2017. [2](#)
- [34] Y. Song, R. Shu, N. Kushman, and S. Ermon, “Constructing unrestricted adversarial examples with generative models,” in *Advances in Neural Information Processing Systems*, 2018, pp. 8312–8323. [2](#)
- [35] C. Xiao, B. Li, J.-Y. Zhu, W. He, M. Liu, and D. Song, “Generating adversarial examples with adversarial networks,” *arXiv preprint arXiv:1801.02610*, 2018. [2](#)

- [36] A. Odena, C. Olah, and J. Shlens, “Conditional image synthesis with auxiliary classifier gans,” in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2642–2651. 2
- [37] H. Qiu, C. Xiao, L. Yang, X. Yan, H. Lee, and B. Li, “Semanticadv: Generating adversarial examples via attribute-conditional image editing,” *arXiv preprint arXiv:1906.07927*, 2019. 2
- [38] A. Jalal, A. Ilyas, C. Daskalakis, and A. G. Dimakis, “The robust manifold defense: Adversarial training using generative models,” *arXiv preprint arXiv:1712.09196*, 2017. 2
- [39] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778. 2
- [40] T. DeVries and G. W. Taylor, “Improved regularization of convolutional neural networks with cutout,” *arXiv preprint arXiv:1708.04552*, 2017. 2
- [41] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, “Cutmix: Regularization strategy to train strong classifiers with localizable features,” *arXiv preprint arXiv:1905.04899*, 2019. 2
- [42] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, A. Courville, D. Lopez-Paz, and Y. Bengio, “Manifold mixup: Better representations by interpolating hidden states,” *arXiv preprint arXiv:1806.05236*, 2018. 3, 4
- [43] C. Heinze-Deml and N. Meinshausen, “Conditional variance penalties and domain shift robustness,” *arXiv preprint arXiv:1710.11469*, 2017. 3
- [44] I. Higgins, D. Amos, D. Pfau, S. Racaniere, L. Matthey, D. Rezende, and A. Lerchner, “Towards a Definition of Disentangled Representations,” *arXiv e-prints*, p. arXiv:1812.02230, Dec 2018. 3
- [45] C. Fefferman, S. Mitter, and H. Narayanan, “Testing the manifold hypothesis,” *Journal of the American Mathematical Society*, vol. 29, no. 4, pp. 983–1049, 2016. 4
- [46] X. Huang and S. Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1501–1510. 5
- [47] E. Nalisnick, A. Matsukawa, Y. W. Teh, and B. Lakshminarayanan, “Detecting out-of-distribution inputs to deep generative models using a test for typicality,” *arXiv preprint arXiv:1906.02994*, 2019. 5
- [48] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012. 5
- [49] R. Vershynin, *High-dimensional probability: An introduction with applications in data science*. Cambridge University Press, 2018, vol. 47. 5
- [50] R. Abdal, Y. Qin, and P. Wonka, “Image2stylegan: How to embed images into the stylegan latent space?” *arXiv preprint arXiv:1904.03189*, 2019. 6
- [51] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *European conference on computer vision*. Springer, 2016, pp. 694–711. 6
- [52] A. Dosovitskiy and T. Brox, “Generating images with perceptual similarity metrics based on deep networks,” in *Advances in neural information processing systems*, 2016, pp. 658–666. 6
- [53] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014. 6
- [54] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014. 7
- [55] Y. LeCun and C. Cortes, “MNIST handwritten digit database,” 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/> 7
- [56] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 3730–3738. 8
- [57] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, and others, “Tensorflow: a system for large-scale machine learning,” in *OSDI*, vol. 16, 2016, pp. 265–283. 15