

Ensembling with Deep Generative Views

Lucy Chai^{1,2} Jun-Yan Zhu^{2,3} Eli Shechtman² Phillip Isola¹ Richard Zhang²
¹MIT ²Adobe Research ³CMU
 {lrchai, phillipi}@mit.edu {elishe, rizhang}@adobe.com junyanz@cs.cmu.edu

Abstract

Recent generative models can synthesize “views” of artificial images that mimic real-world variations, such as changes in color or pose, simply by learning from unlabeled image collections. Here, we investigate whether such views can be applied to real images to benefit downstream analysis tasks such as image classification. Using a pre-trained generator, we first find the latent code corresponding to a given real input image. Applying perturbations to the code creates natural variations of the image, which can then be ensembled together at test-time. We use StyleGAN2 as the source of generative augmentations and investigate this setup on classification tasks involving facial attributes, cat faces, and cars. Critically, we find that several design decisions are required towards making this process work; the perturbation procedure, weighting between the augmentations and original image, and training the classifier on synthesized images can all impact the result. Currently, we find that while test-time ensembling with GAN-based augmentations can offer some small improvements, the remaining bottlenecks are the efficiency and accuracy of the GAN reconstructions, coupled with classifier sensitivities to artifacts in GAN-generated images.

1. Introduction

Image datasets are the backbone of learning-based vision problems, but images are only sparsely-sampled discrete snapshots of the underlying continuous world. However, recent generative adversarial networks (GANs) [18] have shown promise in learning to imitate the real-image manifold, mapping random samples from a latent distribution to realistic image outputs. A heavily exploited property of these models is that the latent space is locally smooth: samples nearby in latent space will appear perceptually similar in image space [42]. Therefore, GANs can be viewed as a type of “interpolating mechanism” that can blend and recombine images in a continuous manner. From individual image samples, can we use a GAN to generate nearby alternatives on the image manifold, or “views,” giving us unlimited variants of a given image?

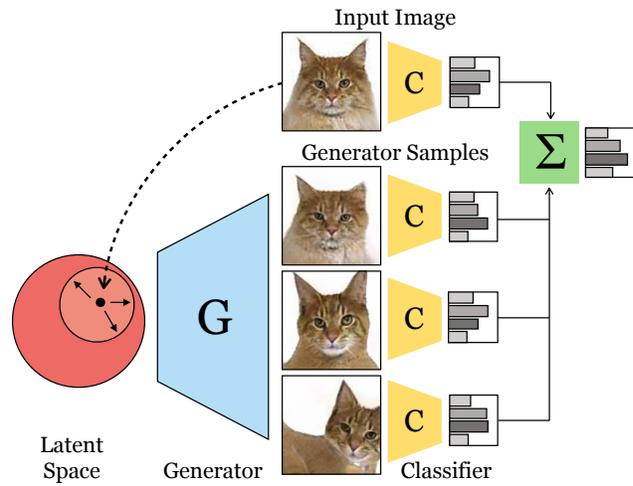


Figure 1: We project an input image into the latent space of a pre-trained GAN and perturb it slightly to obtain modifications of the input image. These alternative views from the GAN are ensembled at test-time, together with the original image, in a downstream classification task.

Here, we investigate using GAN outputs as test-time augmentation for classification tasks. In the standard classification pipeline, passing an image through a trained classifier yields predictions of the image belonging to one of several classes. However, performance is often improved using more than one sample – if we have multiple views of an image, we can use the classifier to obtain predictions for each view and average the results together as an ensemble. The classic approach to generating additional views has been to crop the image at different locations before classification. Using a GAN, we have an orthogonal, data-driven method of generating additional views of a given image, such as altering its pose, shape, or color based on the directions of variation that a GAN learns.

A secondary advantage of unconditional GANs is that they can be trained on image collections without requiring image labels. As data labeling is often vastly more expensive than data collection, GANs learn from much larger datasets compared to tasks involving manual annotation, such as classification. Training on large datasets allows sev-

eral interesting properties to emerge, where the generator learns meaningful variations in data without requiring an explicit training objective to do so [21, 53, 12].

A challenge in using GANs to generate augmented samples is the potential domain gap between real images and GAN outputs – the generated samples must be of sufficient quality to be used in classification tasks, and must adequately reconstruct the target image sample while preserving the relevant visual patterns for accurate discrimination. If this condition is not met, the classifier may behave differently on the generated samples than the natural images, which is undesirable for data augmentation. To reconstruct image samples using the GAN, we use a hybrid encoder and optimization approach [69]: the encoder network initializes the latent code, which is further optimized to improve the similarity between the target image and the reconstruction. In addition to high-quality reconstruction, we also desire that the generated variations of an image do not cross classifier boundaries; *i.e.*, it cannot modify visual appearance that affect its classification. To this end, we experiment with a variety of possible image modifications using the GAN generator, ranging from coarse pose and shape changes to finer-grained color changes.

Using the recent StyleGAN2 generator [28], we apply our method on several classification tasks involving facial attributes [33, 26], cat faces [39], and cars [29]. Given the relative simplicity of the face domain, we find test-time ensembling with GAN samples helps even when the classifier is trained only on real images; however, training the classifier on generated samples offers further improvements, particularly for the more difficult car and cat domains. Code is available on our website: <https://chail.github.io/gan-ensembling/>.

2. Related Work

Latent manipulation in GANs. Advances in generative adversarial networks have allowed them to create increasingly realistic images [14, 42, 26, 7, 63, 27, 28], and furthermore, the generated outputs mimic the variations found in their training data. For example, Radford et al. [42] demonstrate linear separability in latent space and use it to modify attributes of interest in GAN-generated samples. This linear separability has been exploited in a number of subsequent works, including for faces [52], for camera attributes [24, 41], and for quantifying disentanglement of the latent space [27]. While these manipulations require a supervised or self-supervised objective, another direction of interest involves learning meaningful editing directions without direct supervision. One set of approaches aims to uncover primary directions of variation in an intermediate latent space [21, 53, 12], while another enforces distinctness of optimized directions during training [58]. Apart from linear edit directions, recent works [67, 36] learn latent editing

operations that control the 3D appearance of generated images. Interestingly, the architecture of a GAN itself may lend itself to natural edits; for example, StyleGAN, which we use [27, 28], provides a hierarchical latent code that controls visual patterns at various scales.

GANs for real image editing. While GANs are able to generate interesting variations of synthetically generated images, one is often more interested in editing a given *real* image. To do this, we first need to find the latent code that best reconstructs the image [69, 8], a challenging problem due to the generator’s inherent limitations and the complex optimization landscape through a deep network. Optimizing in an intermediate layer or an expanded latent space allows the generator to reconstruct a greater variety of real images [1, 2, 5], which relaxes the former restriction. However, one runs the risk of overfitting or drifting off the manifold of natural latent codes; as such, works [68, 61] propose regularizers to constrain the optimized code to the latent manifold. To speed up the optimization process, several methods [69, 68, 6, 46] propose to train an encoder to initialize latent optimization. To further improve the generator’s flexibility and reconstruction capability, Pan et al. [37] and Bau et al. [4] finetune the generator weights towards a particular target image, while Huh et al. [23] propose to spatially transform the image due to the generator’s inherent biases. As a fast and accurate image projection is critical to our pipeline, we directly draw upon the insights from these works in our method.

Data augmentation with generative models. Generative models’ ability to create realistic variations of images opens up the possibility of using them as data augmentation in downstream tasks, as acquiring images from a generative model is vastly cheaper than collecting additional data. Several works investigate using generative samples as data augmentation in an image translation setting, in which an autoencoder-style network is trained to produce variations of an input image, on a variety of domains such as face attribute editing [57], gaze estimation [54], emotion classification [70], and honeybee tracking [55]. Rather than taking an image-to-image translation approach, we use *projection into a pretrained generator*, which allows us to use emergent variations from this generator rather than directly training for a specified transformation. Instead of a generative model for images, Ratner et al. [43] formulates a generator to select from a sequence of predefined image augmentations used to train a downstream classifier. Specifically in the medical community, synthesizing additional images to augment limited-size datasets has been beneficial for classification [40, 60, 17, 22, 48] or segmentation tasks [32, 11, 20]. However, using GANs for data augmentation is limited by the possible domain gap between generated samples and dataset images [45, 44]. Studies show

that despite their rapid improvement, the generative models do, in fact, still exhibit artifacts, exploitable for reliable detection of GAN-generated imagery [59, 10, 62, 66, 16]. Outputs from a generative model can also be used in adversarial settings: Samangouei et al. [49] demonstrate projecting adversarial images through a generator on MNIST and Fashion-MNIST, while other works [51, 25, 19] train models using the generator to improve robustness. Augmentation using GANs can be applied at training time for robustness or test-time for ensembling; while we primarily focus on the latter, concurrent works [35, 50] investigate the benefits of the former approach while [56] investigates intermediate GAN representations for few-shot segmentation.

3. Method

Our goal is to leverage a generative model to synthesize useful variations of a given image. As summarized in Fig. 1, the first step is to “project” the image to the latent code of a generator. From there, we explore different methods for producing variations. Finally, we show how these variations can be effectively used by a downstream classifier.

3.1. GAN preliminaries

A generative network maps a low-dimensional code z to an image x . Specifically, we use the StyleGAN2 generator [28]. A useful property of the architecture is the intermediate feature space $w \in \mathbb{R}^{B \times D}$, which contains B blocks of dimension D , designed to control the “style” of the image [28].¹ Earlier and later blocks affect the image on coarser and finer scales, respectively. Combining the mapping network M and backbone generator G produces image $x = G(w) = G(M(z))$. Previous work [1] finds that the intermediate w space is better able to represent images than the original code z , while moving in this space offers controllable and interesting effects [21]. As such, work on StyleGAN2 inversion [2, 52, 68] typically uses this intermediate space, and we refer to it as the “latent code”. Following StyleGAN [27], we subdivide $w \in \mathbb{R}^{B \times D}$ into “coarse” styles (first four style codes), “middle” styles (next six style codes), and “fine” styles (the remaining style codes). As we find that changing the “middle” layers can alter object identity, we focus on modifying the “coarse” and “fine” styles to produce image variations.

3.2. Projecting Images into GAN Latent Space

To edit a real image x , we must first find the latent code w that generates the image. As an exact match cannot usually be found, this problem is relaxed to finding the closest

¹Previous work [1] refers to this as the w^+ space, with a separate, more constrained w space. We do not need to make this distinction, so we refer to it as w for simplicity.

image by solving the following optimization problem over an image distance metric:

$$\begin{aligned} w^* &= \arg \min_w \mathcal{L}_{\text{img}}(G(w), x) \\ &= \arg \min_w \ell_1(G(w), x) + \mathcal{L}_{\text{percep}}(G(w), x). \end{aligned} \quad (1)$$

We use ℓ_1 and LPIPS [65] distances, denoted as $\mathcal{L}_{\text{percep}}$. Solving this projection problem directly via optimization is challenging due to the difficult optimization landscape, and tends to heavily depend on the initialization. This has been the subject of active research [69, 1, 6, 2, 23]. Furthermore, optimization can be slow to converge to a reasonable reconstruction of the target image x . We follow best practices, with the needs of our problem in mind. Specifically, a projection algorithm that is very computationally-intensive is not tenable for our use-case, as our approach must be feasible across a full dataset. To balance reconstruction quality and optimization time, we use a two-step approach to project images into the GAN latent space.

Preprocessing by alignment. As GANs have a tendency to accentuate spatial biases and generate centered objects [23], we shift the image to center the target object prior to image projection (Fig. 2 left). As this may cause some missing pixels around the edges of the image, we use spatial masks in the projection objective, so that unknown edge pixels do not contribute to the loss. Please refer to the supplement for additional details.

Encoder initialization. As a first step, we train an encoder network E , to initialize the latent code with a single forward pass. The encoder network is trained using the objective:

$$\mathcal{L}(x, w, E) = \mathcal{L}_{\text{img}}(x, G(E(x))) + \lambda \mathcal{L}_{\text{latent}}(w, E(x)). \quad (2)$$

The $\mathcal{L}_{\text{latent}}$ is an ℓ_2 reconstruction loss for supervision. For this stage, we set $\lambda = 1.0$. The encoder E is trained across randomly drawn, corresponding latent codes and images:

$$E^* = \arg \min_E \mathbb{E}_{w,x} \mathcal{L}(x, w, E). \quad (3)$$

Iterative optimization. Given a specific image x , a forward pass through the trained encoder yields an initialization. A closer match can be obtained by further optimization:

$$w^* = \arg \min_w \mathcal{L}(x, w, E^*). \quad (4)$$

Note that nearly the same objective can be used as for training the encoder. We reduce λ to 0.5, as here, the latent loss serves as a regularizer to the encoder-provided initialization, rather than as ground truth supervision. We optimize for 500 iterations using an L-BFGS optimizer [31] taking 30-45 seconds per image on a V100 GPU, depending on the generator resolution. As this optimization is performed for each image, we limit the optimization time to be short enough so that computation is tractable over the dataset, yet long enough to reasonably match the target image.



Figure 2: Synthesizing deep generative views. We first align (Aligned Input) and reconstruct an image by finding the corresponding latent code in StyleGAN2 [27] (GAN Reconstruction). We then investigate different approaches to produce image variations using the GAN, such as style-mixing on fine layers (Style-mix Fine), which predominantly changes color, or coarse layers (Style-mix Coarse), which changes pose. We show additional perturbations in supplementary material.

3.3. Image Augmentations using Pretrained GANs

Once we optimize for a latent code w^* matching a given target image, we then perturb it to obtain image variations. We try local perturbations using an isotropic Gaussian and PCA directions, along with a “style” swapping operation.

Isotropic Gaussian. One approach is to simply sample from an isotropic Gaussian ball centered at the optimized latent:

$$\tilde{w} \sim \mathcal{N}(w^*, \sigma I), \quad (5)$$

where σ scales variance of the sampled points. We obtain hyperparameter σ by cross-validating over the validation set, and ultimately report on the test set. Note that using $\sigma = 0$ yields the reconstructed image, while a large σ corresponds to randomly drawing a generated image with little regard to the original image x . We conduct experiments adding noise to either the “coarse” or “fine” style codes.

Principal directions. Alternatively, we also experiment with sampling according to principal directions; these directions were found to correspond well with interpretable controls in GANspace [21]. We perform PCA on random samples in the latent space to obtain principal directions \hat{v}_d , which are unit vectors, and eigenvalues λ_d . To produce perturbations, we randomly sample a principal component direction d uniformly among n principal components, $d \sim \mathcal{U}[1, n]$, and a perturbation factor β from a uniform distribution $\beta \sim \mathcal{U}[-\sigma, \sigma]$. In practice, as only the top directions have a visible effect, we restrict ourselves to the top $n = 20$ principal components:

$$\tilde{w} = w^* + \beta \lambda_d \hat{v}_d. \quad (6)$$

We modify either the “coarse” or “fine” style codes while holding the remaining layers fixed.

Style-mixing. Recall that a property of StyleGAN2 is that $w^* \in \mathbb{R}^{B \times D}$ where the early layers correspond to the coarse visual patterns, while later layers correspond to fine details. The third method of perturbation corresponds to “style-mixing”, which swaps in a randomly generated latent code w at some granularity (e.g., the “fine” styles), while preserving the remaining layers from the optimized w^* ([28]; Fig. 2-right). Geometrically, rather than jittering locally, this corresponds to jumping onto a vertex on a hypercube defined by the two latent codes. Visually, this enables us to achieve greater changes in the appearance of the perturbed output, compared to local jittering.

3.4. Ensembling deep generative views

In typical classification, feeding an image x through a trained classifier C yields the prediction logits for that image $y = C(x) \in \mathbb{R}^L$, where L is the number of classes.

Ensembling. By perturbing the optimized latent codes, we obtain additional samples, which can be ensembled at test-time. We jitter the optimized latents to obtain a series of latent codes $\tilde{w}_1 \cdots \tilde{w}_N$ and then run these codes through the generator to obtain variations of the target image: $G(\tilde{w}_1) \cdots G(\tilde{w}_N)$ as inputs to a trained classifier C . To ensemble the classifier predictions, we find that an appropriate weighting between the original image x and the generated samples $\{G(\tilde{w}_n)\}$ improves results. Therefore, we ensemble the end classifier decision using:

$$y_{\text{ens}} = (1 - \alpha)C(x) + \frac{\alpha}{N} \sum_{n=1}^N C(G(\tilde{w}_n)). \quad (7)$$

The hyperparameter $\alpha \in [0, 1]$ is selected using the validation data, and the optimal value in validation is applied to the test partition. Note that $\alpha = 0$ means that the perturbations actually reduce validation performance. This could

be due to (a) a poor reconstruction of the input image or (b) overaggressive perturbations that do not preserve the property of interest in the image. On the other hand, $\alpha = 1$ means that the ensembled variations are entirely able to capture necessary properties for the classification task.

Classifier training. Due to the domain gap between real images and their generated counterparts, we find that training the classifier on generated images often improves results. Given tuples of images, optimized latents, and labels (x, w^*, y) , we train the classifier using images generated from the latents and standard cross entropy loss \mathcal{L}_{cls} :

$$C^* = \arg \min_C \mathcal{L}_{cls}(C(G(w^*), y)). \quad (8)$$

We also experiment with training on perturbed \tilde{w} using the methods in Sec. 3.3, thus replacing $C(G(w^*))$ with $C(G(\tilde{w}))$ in Eqn. 8. We provide details on classifier architecture and training parameters in supplementary material.

4. Experiments

We use the previously described approach of projecting images into a GAN’s latent space and then modifying the latent code to create alternative views. We conduct experiments on facial attribute, car, and cat face classification.

4.1. Facial Attribute Classification

Human faces are one of the most successful domains for recent GAN architectures, as aligned faces reduce the amount of variation that the GAN needs to model. Furthermore, large datasets exist of labeled facial attributes, which can be subsequently used for classification. We start by investigating this setting by pairing the StyleGAN2 [28] face model with the CelebA-HQ dataset [33, 26], containing 30,000 images with 40 labeled binary attributes.

Ensemble weighting. As described in Sec. 3.4, a soft weighting between the real image and its GAN-generated variants is necessary. Since the GAN reconstructions may not exactly approximate the ground-truth image, this causes the classifier to behave differently on the real images compared to their GAN-generated counterparts. We visualize this effect in Fig. 3, where we plot accuracy as a function of the ensemble weight α and find that an intermediate weighting yields the highest accuracy.

GAN perturbation variations. Next, we investigate the effect of the different types of perturbations – isotropic Gaussian perturbations (Eqn. 5), principal component directions (Eqn. 6), or swapping layers of the optimized latent code (style-mixing; Sec. 3.3). We start with a subset of 12 out of 40 facial attributes, and investigate each type of perturbation on the resulting ensemble classification accuracy. In Fig. 2, we show qualitative examples of the latent code perturbations for the style-mixing operation on fine and coarse

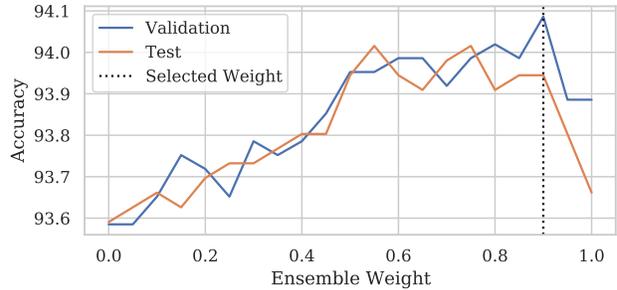


Figure 3: Selecting ensemble weight α by cross-validation; $\alpha = 0$ corresponds to the standard test accuracy, while $\alpha = 1$ corresponds to discarding the original image. We find that often using an intermediate value of α yields the highest accuracy, trading off the additional views provided by the GAN with potential imperfection in reconstruction; we show validation accuracy as a function of α in blue, and test accuracy in orange for the Smiling attribute. We select the value of α based on the validation split and apply the same value to the test split (shown in the dotted black line); note that this may not be the same as the optimal α for the test split.

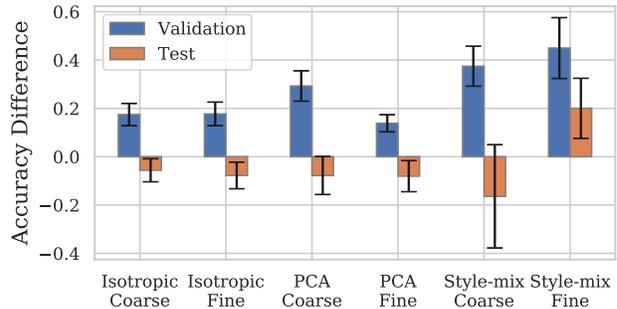


Figure 4: Investigating GAN perturbation methods. On a subset of 12 face attributes, we first investigate test-time ensembling using each type of latent code perturbation. The ensemble weight hyperparameter α is selected from the validation split. Therefore, we find that all types of perturbation can increase validation accuracy. However, the style-mixing perturbation on fine layers performs the best at both validation and test time, and only the fine style-mixing operation yields robust improvements at test time. Error bars indicate standard error over the 12 attributes.

layers, and visualize the remaining types of perturbations in the supplementary material. While perturbations in the coarse layers preserve color attributes, they do not always preserve the facial identity. We find that style-mixing in the fine layers offers the most robust improvements at test time; while all types of perturbations offer improvements on the validation split (using the previous ensemble weight α selected on validation data), only the fine style-mixing method offers reliable improvements at test time (Fig. 4). Therefore, we primarily focus on this type of perturbation for the remaining face experiments.

Comparing to image augmentations. Standard image augmentations used during training, such as random flipping or cropping, can also be ensembled together at test

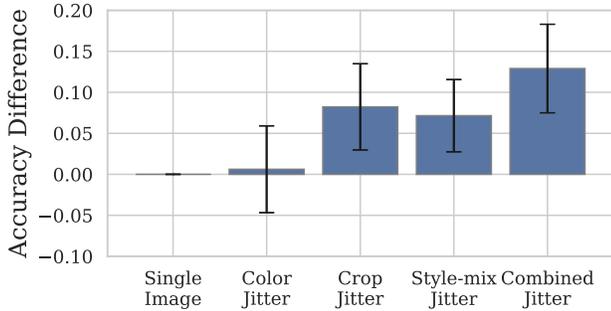


Figure 5: Traditional vs. generative ensembling. Averaged over 40 facial attributes, we plot the accuracy increase from test-time ensembling with various forms of augmentation. While ensembling with image augmentations (Color, Crop) or deep generative views (Style-mix) improves over no ensembling (Single Image), more importantly, they provide *orthogonal benefits*; combining them works best. Errorbars show standard error over 40 attributes.

time to obtain a more robust classification result. We compare the GAN-based augmentations to these image-based augmentation techniques, including (1) color jitter with horizontal flip or (2) spatial jitter with horizontal flip (as faces are aligned, we use small spatial perturbation by resizing to 1056px before taking a random 1024px crop). Averaged across all 40 facial attributes, we find that ensembling with color jitter offers small improvements over single-image classification. Adding the style-mixing perturbation provides similar benefits to spatial jittering. However, classification improves even more when combining all three augmentation methods: we show the difference between each type of test-time augmentation and single-image classification accuracy in Fig. 5.

Training with GAN-generated views. Due to the relative simplicity of the face domain and the GAN’s ability to reasonably reconstruct aligned facial images, adding GAN-generated augmentations at test time is beneficial even when the classifier is not trained on the GAN domain. In the supplementary material, we investigate training the classifiers on GAN reconstructions (Eqn. 8) and reconstructed latent code perturbations as additional augmentation.

Investigating dependency on ensemble-size. While we use an ensemble size of 32 images (1 dataset image and 31 GAN views) for all prior experiments, here we investigate how the number of samples used in the ensemble impacts classification accuracy. Using the Smiling classifier trained on dataset images (rather than GAN images), we plot accuracy as a function of the number of GAN samples in the ensemble (zero indicates only the real image is used for classification, corresponding to standard test accuracy) in Fig. 6. We find that accuracy increases as the number of views used for classification increases but plateaus after a certain point, after which increasing the number of GAN samples has limited returns.

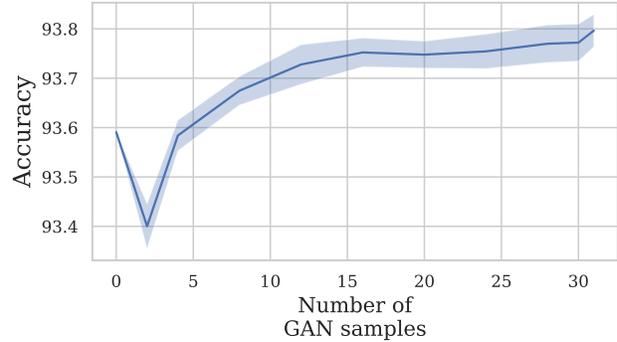


Figure 6: Classification accuracy as a function of the number of ensembled deep generative views on the Smiling attribute. Zero corresponds to using the original input image. Even adding a couple of views increases accuracy. Generally, adding more views further increases accuracy. We use ensemble size of 32 in our experiments, as performance saturates. We show additional attributes in the supplement.

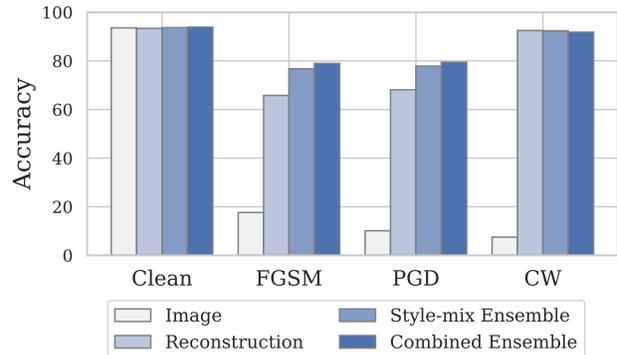


Figure 7: Robustness to adversarial attacks. We show accuracy on a corrupted image (Image), the GAN reconstruction (Reconstruction), ensembling with GAN style-mixing (Style-mix Ensemble), and ensembling over both traditional and GAN views (Combined Ensemble), on the Smiling attribute. Adversarial attacks (FGSM, PGD, CW) greatly reduce accuracy. On all cases, just GAN reconstruction recovers significant performance, and for FGSM and PGD attacks, ensembling with GAN views further improves accuracy. We show additional attributes and experiments on untargeted corruptions in the supplementary material.

Ensembling on corrupted images The GAN generator is trained only on clean images, and a projection step through the latent code can potentially map off-manifold images back to the image manifold, serving as a useful intermediary prior to a classification task. Here, we investigate these scenarios using a set of targeted image corruptions, which adversarially change the output prediction of the classifier. We test against FGSM [18], PGD [34], and CW attacks [9]. Our initial step, projecting the corrupted images through the latent code of the generator, corresponds to the DefenseGAN method [49]. Classifying the corrupted images results in a large decrease in classification accuracy, which the projection step can partially recover. Next, we perturb the opti-

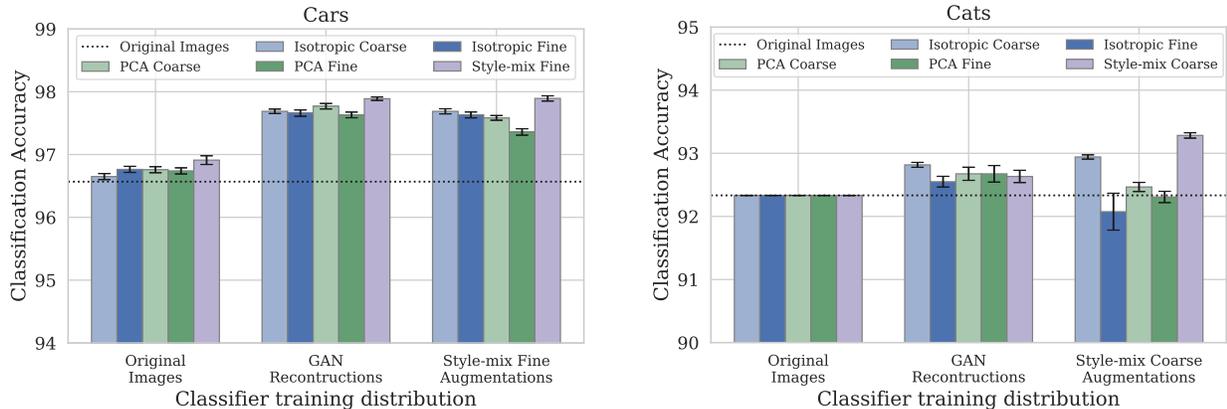


Figure 8: Effect of ensembling deep generative views. Classification accuracy for cars (left) and cats (right), as a function of training distribution (original images, GAN reconstructions, and GAN style-mixing, on x-axis) and test procedure (original images, and isotropic, PCA, and mixing-based GAN augmentations, as colored bars). The dotted line is the baseline, trained and tested on real image crops. When the classifier is trained on real images, GAN-generated views ensembled at test offers improvements only in the Car dataset. Fine-tuning the classifier on GAN-generated views improves the effects of test-time ensembling, and outperforms the baseline (error bars indicate standard error over bootstrapped samples from the ensemble). Fine layer style-mixing at test time is the best perturbation type on the cars domain, and coarse layer style-mixing is the best on the cats domain. The benefit of GAN augmentations on cat classification is weaker compared to cars; this is perhaps because cat classification is a harder problem (12-way vs. 3-way), resulting in a more constrained space of useful perturbations. We show results on additional classifier training distributions in the supplement.

mized latent code by applying style-mixing, and ensemble together these alternative views for classification. For the FGSM and PGD attacks, ensembled images from the GAN can boost classification beyond the single reconstructed images, and applying standard image augmentations in conjunction with the GAN-generated augmentations can offer an additional increase. In supplementary material we investigate targeted attacks on additional facial attributes and also untargeted corruptions, although we find that ensembling with GAN views offers stronger benefits in the targeted case compared to the untargeted scenario.

4.2. Classifying Cars

Due to its binary nature, the face classification problem is a relatively simple task, and given the success of recent GANs at imitating aligned faces, we find that adding views from the GAN at test time, without training on GAN-generated images, offers improvements. We next aim to investigate a slightly more difficult scenario of classifying types of cars, namely the ‘SUV’, ‘Sedan’, and ‘Cab’ superclasses derived from [29] (note that we use these superclasses, rather than the original 196 fine-grained classes, as the GAN reconstruction cannot recover fine details such as make and model of a car). Following best practices, we center the car using a bounding box prior to GAN projection, as GAN reconstruction tends to be better on centered objects [23]. After projecting the images into the GAN latent space, we try isotropic, PCA, or style-mixing perturbations (Sec. 3.3). Fig 2 shows qualitative examples of fine and coarse style-mixing perturbations; the remaining perturba-

tion methods are shown in supplementary material, along with additional dataset and classifier training details.

With a classifier trained on car images, ensembling the predicted classes over multiple random crops at test time boosts accuracy over a single crop (95.8% vs. 96.6% accuracy). Adding additional GAN augmentations at test time offers a small benefit, with the fine style-mixing augmentation increasing accuracy to 96.9%. However, due to imperfect reconstructions there can be a domain gap between the training domain of the classifier and the reconstructed images from the GAN. If we finetune this initial classifier using the GAN reconstructions or perturbed GAN reconstructions, the addition of GAN augmentations offers further benefits, increasing accuracy to 97.9%. In Fig. 8 we report classification accuracy on these training variations, and show results on additional training variations in supplementary material. For this domain, since using random image crops improves classification, we combine 16 random image crops and 16 random crops of the GAN-generated views at test time to form the ensemble; we average the classifier predictions on the random image crops with weight $1 - \alpha$ and average the predictions on GAN outputs with weight α (Eqn. 7).

4.3. Classifying Cats

We next investigate a 12-way classification task on aligned cat faces from [39]. Since the cat images are aligned, we find that ensembling with random image crops performs similarly to using a single center-cropped image, so we combine the real image with 31 random GAN-

generated variants (32 images in total) following Eqn. 7. As this task is more difficult than the previous two scenarios, we find that adding GAN augmentations does not improve accuracy when the classifier is only trained on real images. However, ensembling with GAN-generated views helps when the classifier is finetuned on the GAN domain. We note that training with the perturbed GAN reconstructions also benefits standard image classification as well, increasing from 92.3% when trained only with images to 93.2% when trained with GAN perturbations on the image dataset (shown in supplementary material), to which adding GAN augmentations at test time further increases to 93.3%. We show additional training variations in the supplement.

5. Discussion

While GANs show promise in synthesizing alternative views of a given image, several challenges remain in using them for downstream classification. Firstly, the GAN must be able to reconstruct the salient features used for classification. Here, we choose the StyleGAN2 generator [27], which focus on modeling a single object category (such as faces, cars, and cats), rather than class-conditional models like BigGAN [7] on 1000 ImageNet categories [47]. Due to the higher variation in multi-class data, image projection (a critical part of generating deep augmentations) is still not sufficiently fast and reliable across the full dataset [23, 37, 35]. Even within a single-class StyleGAN2 generator, some aspects of the original image cannot be accurately recovered, such as fine textures, ornate backgrounds, or non-canonical poses (shown in supplementary material), impacting downstream classification.

Related to GAN reconstruction *quality* is the ability to *efficiently* find a latent code that corresponds to a target image. Longer optimization can better reconstruct the image, but becomes intractable over a large dataset. In supplementary material, we investigate classification accuracy as a function of the number of optimization steps, and also perform experiments using an alternative inversion method [68] on a smaller face GAN, obtaining similar results. Moreover, recent alternative architectures trained specifically for efficiency [30] or invertibility [38] may help further reduce the computational cost of image reconstruction.

On the classification side, we note that classifiers are sensitive to GAN reconstructions; accuracy on the reconstructions tends to be lower than that of the dataset images, requiring an ensemble weighting hyperparameter to merge the predictions of the image and the GAN outputs at test time. In most cases, we find that GAN transformations that modify style tend to be more beneficial than those that modify poses. This is in line with previous works that note the benefits of style-based *training* augmentations for image classification [35, 13] and related positional sensitivities of classifiers [64, 3, 15]. In the more difficult Imagenet classification

problem, we found performance degrades substantially during image projection, and therefore GAN perturbations offer limited benefits. In supplementary material, we show results on using the class-conditional CIFAR10 StyleGAN2, where we also find the GAN reconstructions are more difficult to classify than the original images, such that adding GAN-reconstructed views does not benefit classification at both training and test time. When training classifiers, we use standard random flip, resize, and crop transformations on the image, but we further find that the alternative image augmentation strategies *during training* can slightly outperform using the GAN-based augmentations *at test time*, also shown in the supplement.

In our experiments, we find that ensembling images with GAN views helps in simple classification settings, but gains are small and are impacted by (1) the quality and efficiency of GAN reconstructions and (2) classifier sensitivities to the imperfect GAN outputs. Accordingly, as generative modeling technology advances in the future to create better reconstructions, similar GAN augmentation and ensembling strategies may yield even greater improvements.

6. Conclusion

We investigate the capability of StyleGAN2 to generate alternative views of an image for use in downstream classification tasks. We first project the image into GAN latent space using a hybrid encoder and optimization setup to balance reconstruction speed and image similarity. Next, we investigate several types of perturbations in the StyleGAN2 latent space, such as isotropic Gaussian noise, Principal Component directions, or style-mixing the optimized latent code with a random latent code at certain generator layers. We find that two adjustments to naive ensembling are beneficial: (1) appropriately weighting the real image and the GAN outputs in the ensembled predictions, and (2) finetuning the classifier on reprojected GAN samples to account for the domain gap between real images and GAN-generated variants. We conduct experiments on ensembling with GAN-generated perturbations on face attribute, cat, and car classification tasks, and investigate the impacts of ensemble size and corruptions in the input image. Due to current limitations in GAN reconstructions and classifier sensitivities, we are constrained to relatively simple tasks with small datasets, which may be mitigated with future improvements in generative modeling technology.

Acknowledgments. We would like to thank Jonas Wulff, David Bau, Minyoung Huh, Matt Fisher, Aaron Hertzmann, Connelly Barnes, and Evan Shelhamer for helpful discussions. LC is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. This work was started while LC was an intern at Adobe Research.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Int. Conf. Comput. Vis.*, 2019. 2, 3
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 3
- [3] Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? 2020. 8
- [4] David Bau, Hendrik Strobelt, William Peebles, Bolei Zhou, Jun-Yan Zhu, Antonio Torralba, et al. Semantic photo manipulation with a generative image prior. *ACM Trans. Graph.*, 2020. 2
- [5] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B Tenenbaum, William T Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *Int. Conf. Learn. Represent.*, 2019. 2
- [6] David Bau, Jun-Yan Zhu, Jonas Wulff, William Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a gan cannot generate. In *Int. Conf. Comput. Vis.*, 2019. 2, 3
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *Int. Conf. Learn. Represent.*, 2018. 2, 8
- [8] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016. 2
- [9] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017. 6
- [10] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. In *Eur. Conf. Comput. Vis.*, 2020. 3
- [11] Krishna Chaitanya, Neerav Karani, Christian F Baumgartner, Anton Becker, Olivio Donati, and Ender Konukoglu. Semi-supervised task-driven data augmentation for medical image segmentation. In *Information Processing in Medical Imaging*, 2020. 2
- [12] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in style: Uncovering the local semantics of gans. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2
- [13] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 113–123, 2019. 8
- [14] Emily Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In *Adv. Neural Inform. Process. Syst.*, 2015. 2
- [15] Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. Exploring the landscape of spatial robustness. In *Int. Conf. Machine Learning*, 2019. 8
- [16] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *Int. Conf. Machine Learning*, 2020. 3
- [17] Maayan Frid-Adar, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. Synthetic data augmentation using gan for improved liver lesion classification. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 2018. 2
- [18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Int. Conf. Learn. Represent.*, 2015. 1, 6
- [19] Sven Gowal, Chongli Qin, Po-Sen Huang, Taylan Cemgil, Krishnamurthy Dvijotham, Timothy Mann, and Pushmeet Kohli. Achieving robustness in the wild via adversarial mixing with disentangled representations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1211–1220, 2020. 3
- [20] John T Guibas, Tejpal S Virdi, and Peter S Li. Synthetic medical images from dual generative adversarial networks. In *Accepted to NeurIPS 2017 Workshops on Medical Imaging and Machine Learning for Health*, 2017. 2
- [21] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *Adv. Neural Inform. Process. Syst.*, 2020. 2, 3, 4
- [22] Le Hou, Ayush Agarwal, Dimitris Samaras, Tahsin M Kurc, Rajarsi R Gupta, and Joel H Saltz. Unsupervised histopathology image synthesis. *arXiv preprint arXiv:1712.05021*, 2017. 2
- [23] Minyoung Huh, Richard Zhang, Jun-Yan Zhu, Sylvain Paris, and Aaron Hertzmann. Transforming and projecting images into class-conditional generative networks. In *Eur. Conf. Comput. Vis.*, 2020. 2, 3, 7, 8
- [24] Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. In *Int. Conf. Learn. Represent.*, 2020. 2
- [25] Ajil Jalal, Andrew Ilyas, Constantinos Daskalakis, and Alexandros G Dimakis. The robust manifold defense: Adversarial training using generative models. In *arXiv preprint arXiv:1712.09196*, 2017. 3
- [26] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *Int. Conf. Learn. Represent.*, 2018. 2, 5
- [27] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 2, 3, 4, 8
- [28] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 3, 4, 5
- [29] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 2, 7
- [30] Ji Lin, Richard Zhang, Frieder Ganz, Song Han, and Jun-Yan Zhu. Anycost gans for interactive image synthesis and editing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 8
- [31] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 1989. 3

- [32] Shuangting Liu, Jiaqi Zhang, Yuxin Chen, Yifan Liu, Zengchang Qin, and Tao Wan. Pixel level data augmentation for semantic image segmentation using generative adversarial networks. In *ICASSP*, 2019. 2
- [33] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Int. Conf. Comput. Vis.*, 2015. 2, 5
- [34] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Int. Conf. Learn. Represent.*, 2018. 6
- [35] Chengzhi Mao, Amogh Gupta, Augustine Cha, Hao Wang, Junfeng Yang, and Carl Vondrick. Generative interventions for causal learning. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 3, 8
- [36] Xingang Pan, Bo Dai, Ziwei Liu, Chen Change Loy, and Ping Luo. Do 2d gans know 3d shape? unsupervised 3d shape reconstruction from 2d image gans. In *Int. Conf. Learn. Represent.*, 2021. 2
- [37] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. In *Eur. Conf. Comput. Vis.*, 2020. 2, 8
- [38] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A. Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. In *Adv. Neural Inform. Process. Syst.*, 2020. 8
- [39] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2012. 2, 7
- [40] Fábio Perez, Cristina Vasconcelos, Sandra Avila, and Eduardo Valle. Data augmentation for skin lesion analysis. In *OR 2.0 Context-Aware Operating Theaters, Computer Assisted Robotic Endoscopy, Clinical Image-Based Procedures, and Skin Image Analysis*. Springer, 2018. 2
- [41] Antoine Plumerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuous factors of variations. In *Int. Conf. Learn. Represent.*, 2020. 2
- [42] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 1, 2
- [43] Alexander J Ratner, Henry Ehrenberg, Zeshan Hussain, Jared Dunnmon, and Christopher Ré. Learning to compose domain-specific transformations for data augmentation. In *Adv. Neural Inform. Process. Syst.*, 2017. 2
- [44] Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. In *Adv. Neural Inform. Process. Syst.*, 2019. 2
- [45] Suman Ravuri and Oriol Vinyals. Seeing is not necessarily believing: Limitations of biggans for data augmentation. In *ICLR 2019 Workshop on Learning from Limited Labeled Data*, 2019. 2
- [46] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 2
- [47] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 2015. 8
- [48] Hojjat Salehinejad, Shahrokh Valaei, Tim Dowdell, Errol Colak, and Joseph Barfett. Generalization of deep neural networks for chest pathology classification in x-rays using generative adversarial networks. In *ICASSP*, 2018. 2
- [49] Pouya Samangouei, Maya Kabkab, and Rama Chellappa. Defense-gan: Protecting classifiers against adversarial attacks using generative models. In *Int. Conf. Learn. Represent.*, 2018. 3, 6
- [50] Axel Sauer and Andreas Geiger. Counterfactual generative networks. *Int. Conf. Learn. Represent.*, 2021. 3
- [51] Shiwei Shen, Guoqing Jin, Ke Gao, and Yongdong Zhang. Ape-gan: Adversarial perturbation elimination with gan. In *arXiv preprint arXiv:1707.05474*, 2017. 3
- [52] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 3
- [53] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *arXiv preprint arXiv:2007.06600*, 2020. 2
- [54] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. 2
- [55] Leon Sixt, Benjamin Wild, and Tim Landgraf. Rendergan: Generating realistic labeled data. *Frontiers in Robotics and AI*, 2018. 2
- [56] Nontawat Tritrong, Pitchaporn Rewatbowornwong, and Supasorn Suwajanakorn. Repurposing gans for one-shot semantic part segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 3
- [57] Yuri Viazovetskyi, Vladimir Ivashkin, and Evgeny Kashin. Stylegan2 distillation for feed-forward image manipulation. *arXiv preprint arXiv:2003.03581*, 2020. 2
- [58] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *arXiv preprint arXiv:2002.03754*, 2020. 2
- [59] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 3
- [60] Eric Wu, Kevin Wu, David Cox, and William Lotter. Conditional infilling gans for data augmentation in mammogram classification. In *Image Analysis for Moving Organ, Breast, and Thoracic Images*. Springer, 2018. 2
- [61] Jonas Wulff and Antonio Torralba. Improving inversion and generation diversity in stylegan using a gaussianized latent space. In *arXiv preprint arXiv:2009.06529*, 2020. 2
- [62] Ning Yu, Larry S Davis, and Mario Fritz. Attributing fake images to gans: Learning and analyzing gan fingerprints. In *Int. Conf. Comput. Vis.*, 2019. 3

- [63] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *Int. Conf. Machine Learning*, 2019. [2](#)
- [64] Richard Zhang. Making convolutional networks shift-invariant again. In *Int. Conf. Machine Learning*, 2019. [8](#)
- [65] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. [3](#)
- [66] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2019. [3](#)
- [67] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yinan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. *arXiv preprint arXiv:2010.09125*, 2020. [2](#)
- [68] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *Eur. Conf. Comput. Vis.*, 2020. [2](#), [3](#), [8](#)
- [69] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *Eur. Conf. Comput. Vis.*, 2016. [2](#), [3](#)
- [70] Xinyue Zhu, Yifan Liu, Jiahong Li, Tao Wan, and Zengchang Qin. Emotion classification with data augmentation using generative adversarial networks. In *Pacific-Asia conference on knowledge discovery and data mining*, 2018. [2](#)