

Sequential training of GANs against GAN-classifiers reveals correlated “knowledge gaps” present among independently trained GAN instances

Arkanath Pathak Nicholas Dufour

Google Research

{arkanath, ndufour}@google.com

Abstract

Modern Generative Adversarial Networks (GANs) generate realistic images remarkably well. Previous work has demonstrated the feasibility of “GAN-classifiers” that are distinct from the co-trained discriminator, and operate on images generated from a frozen GAN. That such classifiers work at all affirms the existence of “knowledge gaps” (out-of-distribution artifacts across samples) present in GAN training. We iteratively train GAN-classifiers and train GANs that “fool” the classifiers (in an attempt to fill the knowledge gaps), and examine the effect on GAN training dynamics, output quality, and GAN-classifier generalization. We investigate two settings, a small DCGAN architecture trained on low dimensional images (MNIST), and StyleGAN2, a SOTA GAN architecture trained on high dimensional images (FFHQ). We find that the DCGAN is unable to effectively fool a held-out GAN-classifier without compromising the output quality. However, the StyleGAN2 can fool held-out classifiers with no change in output quality, and this effect persists over multiple rounds of GAN/classifier training which appears to reveal an ordering over optima in the generator parameter space. Finally, we study different classifier architectures and show that the architecture of the GAN-classifier has a strong influence on the set of its learned artifacts.

1. Introduction

GAN [8] architectures like StyleGAN2 [17] generate high-resolution images that appear largely indistinguishable from real images to the untrained eye [14, 18, 24]. While there are many positive applications, the ability to generate large amounts of realistic images is also a source of concern given its potential application in scaled abuse and misinformation. In particular, GAN-generated human faces are widely available (e.g., thispersondoesnotexist.com) and have been used for creating fake identities on the internet [12].

Detection of GAN-generated images is an active research area (see [9] for a survey of approaches), with some us-

ing custom methods and others using generic CNN-based classifiers. Such *classifiers* are distinct from the discriminator networks that are trained alongside the generator in the archetypal GAN setup. Given the adversarial nature of the training loss for GANs, the existence of the GAN-classifiers suggest consistent generator *knowledge gaps* (i.e., artifacts present across samples that distinguish generated images from those of the underlying distribution) left by discriminators during training. Specialized classifiers [31] are able to detect images sampled from held-out GAN instances and even from held-out GAN architectures. These generalization capabilities imply that the knowledge gaps are consistent not only across samples from a GAN generator but across independent GAN generator instances.

In this work we modify the GAN training loss in order to fool a GAN-classifier in addition to the co-trained discriminator, and examine the effect on training dynamics and output quality. We conduct multiple rounds of training independent pools (initialized differently) of GANs followed by GAN-classifiers, and gain new insights into the GAN optimization process. We investigate two different settings: in the first setting, we choose the low-dimensional domain of handwritten digits (MNIST [19]), using a small DCGAN [25] architecture and a vanilla GAN-classifier architecture. For the second setting, we choose a high-dimensional domain of human faces (FFHQ [16]) with StyleGAN2 (SG2) as a SOTA GAN architecture, and three different GAN-classifier architectures (ResNet-50 [10], Inception-v3 [28], and MobileNetV2 [27]). Our findings in this paper are as follows:

- Samples drawn from a GAN instance exhibit a space of “artifacts” that are exploited by the classifiers, and this space is strongly correlated with those of other GAN generator instances. This effect is present in both the DCGAN and SG2 settings.
- Upon introducing the need to fool held-out classifiers, the DCGAN is unable to generate high quality outputs.
- In the high dimensional setting, however, SG2 gener-

ators can easily fool held-out trained classifiers, and move to a new artifact space. Strikingly, we find that the artifact space is correlated among the new population of generators as it was in the original population. This correlation appears to persist in subsequent rounds as new classifiers are introduced that are adapted to the new artifact spaces.

- MobileNetV2 classifier instances in the SG2 setting appear unable to learn all of the artifacts available for them to exploit. Instead, MobileNetV2 instances form clusters based on the subset of artifacts learned. We hypothesize this being an effect of classifier capacity.
- An SG2 generator trained to reliably fool unseen classifier instances from a given architecture is not guaranteed to fool classifiers from another architecture. Therefore, the artifacts learned by a given classifier depends strongly on the classifier’s architecture.

2. Related Work

Research into detection of GAN-generated media has largely tracked the increasing prominence and output quality of GANs themselves. Several studies [3, 4, 5, 9, 13, 21, 23, 30, 31] focus on detection of GAN-generated images using CNNs, and their robustness to data augmentation at test time. Of particular interest to us is [31], who train a ResNet-50 classifier (pre-trained using ImageNet [26]) on images generated using just one modern GAN architecture, ProGAN [15]. They show that the classifier generalizes to unseen GAN architectures, concluding that the task of general GAN detection is fairly straightforward, at least in the absence of image augmentations. Later studies dispute this [6, 9], demonstrating that the test performance of the classifier is decreased if GAN architectures used in training predate those used during test.

Previous research [9, 31] train classifiers on samples from multiple generators but train each generator on separate data domains or datasets, meaning it is not possible to discern the variation due to the generators themselves. We also note these classifiers are not robust to perturbations and can be fooled with specialized targeted attacks [1, 7], as is characteristic of classifiers trained using neural networks [29]. We consider this out of scope for this work relative to the questions we seek to address.

[32] study GAN attribution, a related problem where the architecture of the source generator for a given sample is inferred. They show that multi-class classification works well to distinguish different GANs, where the learned latent embeddings and weights are used as the image and GAN model fingerprints respectively. [22] attribute fingerprints to distinguish different GAN model architectures and datasets. These studies show that GAN-generated images contain architecture- and instance-specific artifacts.

There is limited research on the behavior of GAN generators trained to fool such classifiers. [2] train a specialized patch-based classifier then finetune the GAN generator to fool the classifier, which results in a significant drop in the classifier’s accuracy. A second classifier trained using images from the finetuned generator is able to recover in accuracy. [33] study a related problem of automatically eliminating generator artifacts, training a lightweight CNN generator that adds minimal perturbations to GAN-generated images, allowing them to fool even unseen classifiers. [20] build another such “trace removal network” which learns to remove several types of traces left in various types of “DeepFakes.” For our study, we are principally interested in the consistency of these artifacts across GAN generators and how the GAN generators adjust themselves when their training loss is modified to include a classifier, to better study the space of artifacts present in GAN-generated images.

To our knowledge, there has not been any work evaluating the effect of classifier model architecture and capacity on the learned artifacts, which we have also studied in this paper.

3. Approach

We study the phenomena outlined in the introduction by creating and measuring the performance of classifiers trained to detect images sampled from *unseen* generators and subsequently training new generators to fool them, in sequential rounds, forming a chain of generators and classifiers. We do this in one of two settings, first with low dimensional images (MNIST), a simplistic DCGAN, and a basic classifier architecture. In the second setting, we use higher dimensional images (FFHQ), and perform experiments using the unmodified StyleGAN2 (SG2) architecture. Seeking to minimize sources of variance as much as possible, we limit to a single GAN architecture and a fixed dataset in both settings. We also do not use the “truncation” trick [16], a sample-time heuristic commonly used with the SG2 architecture to improve the output visual quality at the expense of diversity (*see* Supplement Sec. 2 for more discussion on this). In the SG2 setting, we test three different widely-used classifier architectures: ResNet-50, Inception-v3, and MobileNetV2. These architectures were chosen for their architectural diversity. All classifiers and generators are trained from scratch, without any pre-training. Supplement Sec. 1 provides details about the model architectures and training parameters.

3.1. A note on terminology

Because our procedure involves both GANs and classifiers, there is potential ambiguity in terminology as GANs themselves are trained with a subnetwork designed to distinguish generated images from natural images, which is commonly called the “discriminator”, “adversarial network”, or “critic”, among others. To keep the text clear, we will refer to subnetworks co-trained with a generator which to-

gether comprise a GAN as “**discriminators**”, denoted D . The networks trained on samples from multiple, independently trained generators are referred to as “**classifiers**”, C . Each sequential round of training a pool of GANs followed by training classifiers is an “**iteration**” (detailed in Sec. 3.2, and Figs. 1 and 2) and is indexed with a superscript. Iterations are distinct from training steps: during a single iteration, GANs are fully trained, then classifiers are fully trained using those GAN generators. Broadly speaking an “**artifact**” is any property of a generated image that distinguishes it from a real image. By “**knowledge gaps**”, we are referring to a specific class of artifacts that reliably occur *across* samples from a generator. Since this class of artifacts is the only one studied in this work, we use artifact and knowledge gap interchangeably.

3.2. Overview of setup and iterations

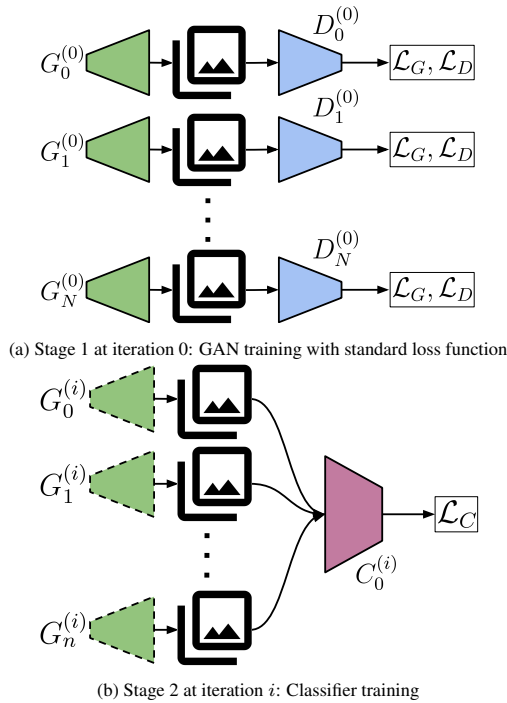


Figure 1. **Experimental setup & training classifiers.** Generators G are green, co-trained discriminators D are blue and classifiers C trained using multiple, frozen generators are purple. Dashed borders indicate that the subnetwork is not being updated during this stage of the iteration. (a) Generators trained in iteration 0 are trained in the typical way. (b) Classifiers are trained in the second stage of all iterations, on samples drawn from subsets of the generators trained in the first stage.

Our experiments consist of sequential rounds (“iterations”), each with two stages: first, a pool of GAN generators initialized randomly is trained, then classifiers are trained to detect samples from the generators trained in the first

stage. In the first stage of the first iteration ($i = 0$), a number of GANs (DCGAN in the first setting, SG2 in the second setting) are trained independently on the training images (MNIST in the first setting, FFHQ in the second setting), as shown in Fig. 1a. This setup is modified slightly in later iterations (*see* Fig. 2) as detailed below. Classifier training follows in the second stage (Fig. 1b) as a standard classification task where each classifier is trained on a balanced dataset of real images and images sampled from a subset of generators trained in the first stage. The second stage is the same in every iteration, always sampling images from generators trained in the first stage of the iteration. The first stage of subsequent iterations ($i > 0$) proceeds like the first stage of the first iteration but with a modified generator loss function: generators are trained to fool not only the discriminator they are co-trained with, but also frozen classifiers from preceding iterations. To do this we modify the “classical” GAN generator loss function \mathcal{L} :

$$\mathcal{L}_{G^{(i)}} = -\log(D^{(i)}(G^{(i)}(w))) \quad (1)$$

in one of two ways. In the first, $\mathcal{L}_{G^{(i)}}^{\Sigma}$, generators must fool a classifier from *every* preceding iteration:

$$\mathcal{L}_{G^{(i)}}^{\Sigma} = -[\log(D^{(i)}(G^{(i)}(w))) + \phi \sum_{j=0}^{i-1} \log(C_0^{(j)}(G^{(i)}(w)))] \quad (2)$$

A graphical depiction of a single generator using this loss function is shown in Fig. 2a. ϕ is used to weight the relative influence of classifiers. Because a classifier from each previous iteration must be fooled in order to minimize this function, we refer to it as the “fool-all” loss function.


The other generator loss function variation, $\mathcal{L}_{G^{(i)}}^*$, relies purely on a classifier from the iteration immediately preceding the current one, rather than all preceding iterations:

$$\mathcal{L}_{G^{(i)}}^* = -[\log(D^{(i)}(G^{(i)}(w))) + \phi \log(C_0^{(i-1)}(G^{(i)}(w)))] \quad (3)$$


This is depicted in Fig. 2b. Because $\mathcal{L}_{G^{(i)}}^*$ depends only on the current iteration and the preceding iteration, we refer to this as the “memoryless” loss function.

The two modifications result in markedly different training dynamics. Reported results will generally be for the “fool-all” \mathcal{L}^{Σ} variation (Fig. 2a). When results are based on experiments using the “memoryless” variation \mathcal{L}^* (Fig. 2b), they will be explicitly noted as such. Classifiers are frozen (i.e., their weights are not updated) during the first stage of every iteration.

The classifier subscript 0, used in Figs. 2a and 2b (e.g., $C_0^{(i-1)}$), is purely to distinguish classifiers within the same iteration. In each iteration, multiple classifiers are trained that



(a) Stage 1 at iteration i : GAN training with “fool-all” modified loss function



(b) Stage 1 at iteration i : GAN training with “memoryless” modified loss function

Figure 2. GANs trained in higher iterations. In subsequent iterations ($i > 0$), stage 1 GAN training is modified from the first iteration ($i = 0$, see Fig. 1a) such that the generator $G_k^{(i)}$ learns to fool not only its co-trained discriminator $D_k^{(i)}$ but also (a) i classifiers, one from each preceding iteration (Eq. (2)) or (b) a single classifier from the immediately preceding iteration (Eq. (3)). At $i = 1$, these two approaches are equivalent.

are initialized randomly and trained independently. When testing a GAN trained to fool the previous iteration’s classifiers, classifiers used for training and testing are trained on disjoint subsets of generators, to measure generalization. For example, if $G_k^{(i)}$ is trained to fool $C_0^{(i-1)}$, and is evaluated against $C_1^{(i-1)}$, then $C_0^{(i-1)}$ and $C_1^{(i-1)}$ are trained on disjoint subsets of iteration $i - 1$ generators.

4. Results

We study the interaction between $G^{(0)}$ and $C^{(0)}$ in Sec. 4.1, $C^{(0)}$ and $G^{(1)}$ in Secs. 4.2 and 4.3, and compare generators and classifiers of multiple iterations ($C^{(n)}, G^{(n)}, C^{(n+1)}, \dots$) in Sec. 4.4.

4.1. Classifiers generalize when sampling training data from multiple generator instances

Using our pool of generators, we can profile the number of independent generators necessary to train a classifier that can reliably generalize to samples from generators unseen during training. After the first stage of the first iteration, suppose we have a sufficiently large set of trained GAN

generators. We split them into two subsets:

$$S_0 = \{G_0^{(0)}, \dots, G_n^{(0)}\} \quad \text{and} \quad S_1 = \{G_{n+1}^{(0)}, \dots, G_N^{(0)}\}$$

We then train a classifier $C_0^{(0)}$ on samples drawn from generators in S_0 and measure its performance by testing it on samples from generators in S_1 (while S_1 can vary in size, to keep results comparable we fix S_1 for a given experiment). By measuring the effect of n on the performance of $C_0^{(0)}$ in this way (Fig. 3), we can implicitly measure the distinctiveness of artifacts produced by different generators.

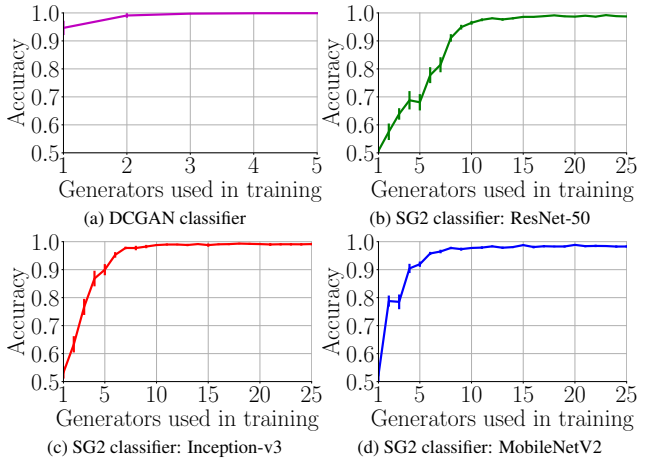


Figure 3. Classifier generalization requires samples from several generators. For all classifier architectures tested, the ability to distinguish between images sampled from unseen generators and real images (y-axis) depends on the number of generators used to produce training data (x-axis). This effect is present in the DCGAN setting (a) and is amplified in the SG2 setting (b), (c) & (d). The performance of each classifier is reported as accuracy on a balanced dataset of unseen natural images and samples drawn from a pool of 5 held-out generators in the DCGAN setting, and 25 held-out generators in the SG2 setting; 0.5 represents random classification.

We find that a DCGAN classifier generalizes well using a single generator, and almost perfectly when trained using more than one different generators. However, in the second setting, we find that SG2 generators produce sufficient diversity between generator instances that a classifier requires samples from several generators to produce reliable generalization. With sufficient generators to sample from, however, all classifiers become extremely accurate. By contrast, when the so-called “truncation trick” [16] is used to generate samples (see Supplement Sec. 2), a single generator is sufficient to achieve nearly-perfect classification accuracy on unseen generators. Based on this finding, we use 3 generators in the DCGAN setting and 15 generators in the SG2 setting when training classifiers for the rest of our experiments.

4.2. DCGAN generators fail to fool classifiers in a generalizable way

When using Eq. (2) to train a DCGAN of the second iteration $i = 1$, we find that, surprisingly, the GAN struggles to fool a held-out classifier. This effect is shown in Fig. 4 where the DCGAN learns to fool the classifier included in Eq. (2) at higher values of ϕ , but it fails to fool a held-out classifier of iteration $i = 0$. As we increase the value of ϕ to large values, we see the output quality degrades, as shown in Tab. 1. Note that because we experiment with higher values of ϕ in this section, we normalize the coefficients as:

$$\mathcal{L}_G^* = -\left[\frac{1}{1+\phi} \log(D) + \frac{\phi}{1+\phi} \log(C)\right] \quad (4)$$




Figure 4. **DCGAN fails to fool held-out classifiers.** Here, the accuracy is reported on GAN-sampled images *only*, where an accuracy of 0 represents the GAN completely fooling the classifier. We report the mean accuracy of 5 held-out classifiers.

ϕ	DCGAN Image Samples
10^{-2}	
10^{-1}	
10^0	
10^1	
10^2	
10^3	
10^4	
10^5	
10^6	

Table 1. **DCGAN output collapses as ϕ takes large values.** We show 10 random image samples from each GAN trained with a different value of ϕ .

4.3. SG2 generators can be trained to fool classifiers in a generalizable way, with caveats

We observe a different behavior in the SG2 setting than in the DCGAN. In particular, very low values of ϕ are sufficient

to cause the generator to learn to fool the classifiers. When SG2 generators are trained to fool the classifiers by using the modified loss described in Eq. (2) (or Eq. (3) where noted), they learn to do so early in their training, as shown in Fig. 5.




Figure 5. **SG2 generators rapidly learn to fool classifiers during training.** Accuracy of first iteration $i = 0$ classifiers included in the modified loss with $\phi = 0.001$ as a second iteration $i = 1$ SG2 generator learns to fool them (at iteration $i = 1$, the Eq. (2) “fool-all” and Eq. (3) “memoryless” loss functions are equivalent). We plot 10 runs in each graph, corresponding to 10 different first-iteration classifier instances against 10 different second-iteration generators. Here, classifiers are evaluated against a balanced set of generated and real images; accuracy of 0.5 is random classification.

When using ResNet-50 classifiers, SG2 generators showed a striking ability to generalize: learning to fool one classifier conferred the ability to fool any other ResNet-50 classifier (see Tab. 2). The reliability of this finding, which persists over multiple iterations (see Sec. 4.4), suggests that all ResNet-50 classifier instances $C^{(i)}$ learn strongly overlapping subsets of the artifacts exhibited by the $G^{(i)}$ generators.

However, this is not the case for all classifier architectures tested. The effect was weakened slightly but statistically significantly ($p < 0.01$) in Inception-v3, though generators that fool one Inception-v3 classifier will still fool almost all others. For MobileNetV2, however, the attenuation of the effect was substantial (and statistically significant relative to ResNet-50 and Inception-v3): generators that can fool one MobileNetV2 classifier instance will be able to fool unseen MobileNetV2 classifier instances only half the time.

We note that our ResNet-50, Inception-v3 and MobileNetV2 architectures had 23.5M, 21.8M, and 2.3M parameters, respectively, and hypothesize that this difference in fooling generalization is an effect of the capacity of the classifier the generator is learning to fool. More con-

cretely, ResNet-50 classifiers, a high capacity model (relatively speaking), each learn the bulk of the artifacts available to them, and thus have high “knowledge overlap” between instances. This overlap accounts for the observed generality of fooling ability on behalf of the generators: it’s sufficient to learn to fool one ResNet-50 instance, since all the instances have learned largely the same thing. Conversely, MobileNetV2, a relatively lower capacity model, learns a smaller subset of available artifacts, reducing the probability of overlapping between instances. This reduced overlap means that a generator trained to fool a single MobileNetV2 instance is less likely to fool unseen MobileNetV2 instances when compared to the ResNet-50 case.

Interestingly, the results (Tab. 2, off-diagonal) imply that the sets of artifacts learned by ResNet-50 and MobileNetV2 are different. If this were not the case, we would expect a MobileNetV2 instance to learn a subset of the artifacts that ResNet-50 learns, due to its lower relative capacity, and hence expect GAN generators that fool unseen ResNet-50 classifiers to be able to readily fool unseen MobileNetV2 classifiers. This effect is even more true of Inception-v3: learning to fool Inception-v3 implies fooling unseen Inception-v3 classifiers but not unseen ResNet-50 classifiers and vice versa. Taken together, this suggests that the higher-capacity architectures learn sets of artifacts that are well-conserved within architecture but are largely distinct between architectures.

To quantify the diversity present in MobileNetV2 classifiers, we modify our “fool-all” $\mathcal{L}_{\mathbf{G}^{(i)}}^{\Sigma}$ loss function to accept multiple MobileNetV2 classifiers (each initialized differently and trained independently) from the previous iteration, rather than just one:

$$\begin{aligned} \mathcal{L}_{\mathbf{G}^{(i)}}^{\Sigma\Sigma} &= -[\log(D^{(1)}(G^{(1)}(w))) \\ &\quad + \phi \sum_{j=0}^{i-1} \sum_{k=0}^{t-1} \log(C_k^{(j)}(G^{(1)}(w)))] \\ &= -[\log(D^{(1)}(G^{(1)}(w))) \\ &\quad + \phi \sum_{k=0}^{t-1} \log(C_k^{(0)}(G^{(1)}(w)))] \end{aligned} \tag{5}$$

Because we perform this experiment on iteration $i = 1$, there is only one previous iteration (the first, $i = 0$), so we may drop the summation-over-previous iteration terms. We add a new summation term to include t classifiers from the previous iteration, rather than a single one. When additional MobileNetV2 classifiers are included, we see the fooling ability of the resulting generators begins to generalize, as shown in Fig. 6.

Consistent with our classifier-capacity hypothesis, multiple MobileNetV2 classifiers are required to achieve generalization because each instance may only learn a subset of

Classifier	GAN trained to fool...		
	ResNet-50	Inception-v3	MobileNetV2
ResNet-50	0.05±0.02	0.74±0.12	0.5±0.39
Inception-v3	0.51±0.25	0.16±0.26	0.53±0.34
MobileNetV2	0.31±0.17	0.36±0.16	0.41±0.38

Table 2. **Generalization of SG2 generator classifier-fooling ability varies widely with classifier architecture.** Each entry is the accuracy (mean±std) of classifiers (row) on generated images *only*, sampled from SG2 generators trained to fool other classifiers (column). Each entry is based on the performance of ten (unseen) classifiers measured against each of ten SG2 generators, where each generator is trained to fool a distinct classifier instance. “Fooling” is considered to occur when accuracy ≤ 0.20 (an arbitrary threshold chosen for visualization purposes), and is highlighted in **bold**. Note that, in all cases, classifiers are at least 98% accurate on held-out real images.




Figure 6. **Reliably fooling unseen MobileNetV2 classifiers requires learning to fool multiple MobileNetV2 classifiers.** Multiple trained instances of the low-capacity (relatively speaking) MobileNetV2 classifiers are required during generator training to achieve generalizable fooling ability. SG2 generators are trained using the loss function in Eq. (5). We report the mean accuracy of 10 held-out MobileNetV2 classifier instances against generated images from a single trained generator, and we start with fooling a classifier where it completely fails to fool held-out classifiers.

the artifacts available to it, and so do not overlap as often as higher-capacity architectures. To measure the between-instance overlap of MobileNetV2, we conducted a pairwise comparisons over 10 independent MobileNetV2 classifiers. To this end, we trained 10 generators to fool one of 10 trained MobileNetV2 classifiers, then we tested that generator’s fooling ability against the other MobileNetV2 instances (as well as the one they were trained to fool). The results (Tab. 3) are striking: rather than randomly sampling artifacts to learn, which would result in fairly uniform off-diagonal values in the table, we see clear clusters emerging that are “mutually-fooling”: a generator trained to fool one will fool the rest, which we take as implying the classifiers within a cluster learned a shared subset of generator artifacts. For instance, classifier instances 6 and 7 are (almost) mutually fooling, as are classifiers 0 through 5; however, a generator trained

MobileNetV2	GAN trained to fool MobileNetV2 classifier #...									
Classifier	0	1	2	3	4	5	6	7	8	9
0	0.04	0.06	0.05	0.08	0.07	0.06	0.98	0.98	0.47	0.95
1	0.16	0.05	0.11	0.10	0.08	0.10	0.99	0.99	0.64	0.97
2	0.09	0.08	0.03	0.10	0.08	0.08	0.99	0.99	0.44	0.94
3	0.09	0.05	0.05	0.03	0.05	0.05	0.99	0.99	0.53	0.96
4	0.09	0.05	0.09	0.06	0.04	0.10	0.98	0.98	0.60	0.94
5	0.22	0.15	0.14	0.20	0.20	0.05	1.00	1.00	0.70	0.98
6	0.53	0.45	0.59	0.37	0.52	0.35	0.01	0.21	0.37	0.03
7	0.28	0.29	0.45	0.17	0.17	0.21	0.02	0.00	0.12	0.03
8	0.13	0.14	0.06	0.11	0.17	0.09	0.92	0.91	0.03	0.74
9	0.44	0.42	0.37	0.34	0.40	0.33	0.60	0.70	0.28	0.02

Table 3. **MobileNetV2 classifiers appear to form clusters based on the subset of artifacts learned.** Each entry (i, j) is the accuracy of the classifier i (row) against images sampled from SG2 generators trained to fool the classifier j (column). A value of 0.0 means the classifier was completely fooled while a value of 1.0 means the classifier was never fooled. An accuracy ≤ 0.20 (an arbitrary threshold) is considered “fooled” and is highlighted in **bold**. Note that, in all cases, classifiers are at least 98% accurate on held-out real images.

to fool 6 or 7 is totally unable to fool classifiers 0 through 5. In the table, the two clusters mentioned are highlighted using boxes. Perhaps more surprising, the table lacks diagonal symmetry: suggesting that some classifiers learn partial subsets of the artifacts learned by others.

4.4. SG2 generators of subsequent iterations do not change in image quality or learning behavior

Secs. 4.1 to 4.3 are essentially concerned with the first iteration and the first stage of the second iteration, as detailed in Figs. 1 and 2. If we continue conducting iterations in the SG2 setting, new dynamics emerge.

Regardless of the loss function used (either Eq. 2 or 3), the training process does not appreciably change for the generators, nor do they require increased training time. Further, this does not result in a drop in the visual quality of the sampled images either, whether measured qualitatively by visual inspection or quantitatively by FID [11], as included in Tab. 4. We do not show the images from the FFHQ dataset or image outputs from models trained on the dataset in this work.

Tab. 4 details five iterations, where all classifiers are ResNet-50 architecture and generators are trained with the “fool-all” loss $\mathcal{L}_{\mathbf{G}(i)}^{\Sigma}$ (Eq. (2)). In each iteration, generators must fool a classifier from all previous iterations, and accordingly, high-capacity classifiers from iteration i are wholly unable to detect generated images from generators of subsequent iterations $k > i$ (upper-right in Tab. 4). The converse is not true: higher-iteration classifiers sometimes, but not always, detect lower-iteration generators (lower-left in Tab. 4). This phenomenon lacks a readily identifiable pattern but is highly replicable. For instance, across 10 independent trials, iteration 2 classifiers could detect iteration 1 generators

about 73% of the time, while iteration 4 classifiers could detect iteration 3 generators only 7% of the time. Furthermore, classifiers of subsequent iterations completely fail to detect GANs of iteration 0 (the first column in Tab. 4). The fact that this effect is iteration specific suggests underlying complexity.

Along the diagonal, we see high performance: iteration i classifiers are able to learn to identify images synthesized by iteration i generators, unseen by the classifier during training. We do note a potential trend of diminishing classification test accuracy in later iterations, which suggests the generators do begin to de-correlate after a number of iterations. It is not surprising that the generators still produce artifacts even after training against the classifiers. It is surprising, however, that the generalization ability of the classifiers remains high over the course of several iterations. This means that not only are the generators producing new artifacts but these artifacts are shared with the other randomly-initialized and randomly-trained generator instances of the same iteration. In other words, the generators shift consistently and largely in unison across iterations. Taken together, this suggests the existence of an “artifact preference” on the part of the generators. If the generation of some set of artifacts is precluded (for instance, by the need to fool classifiers that have learned them), the generators will not only begin to generate new artifacts but *largely the same* set of new artifacts.

The notion that generators produce artifacts according to an orderly preference, where precluding one set of artifacts leads predictably to the generation of a new set of artifacts, is reinforced by our experiments with iterations trained with the “memoryless” loss function (Eq. (3)). In this regime, generators must fool only a classifier from the previous iteration. Unsurprisingly, SG2 generators trained in this way

Classifier	GAN Instances				
	Iteration 0	Iteration 1	Iteration 2	Iteration 3	Iteration 4
Iteration 0	0.993±0.002	0.028±0.004	0.037±0.006	0.047±0.011	0.053±0.007
Iteration 1	0.000±0.000	0.996±0.005	0.013±0.003	0.015±0.002	0.008±0.002
Iteration 2	0.001±0.000	0.734±0.138	0.957±0.053	0.014±0.003	0.029±0.007
Iteration 3	0.004±0.001	0.868±0.115	0.184±0.133	0.931±0.069	0.018±0.004
Iteration 4	0.007±0.001	0.713±0.140	0.510±0.141	0.068±0.027	0.858±0.112
Mean FID	36.98	36.62	36.67	36.39	36.83

Table 4. **Classifier performance in sequential iterations.** We train 5 iterations of SG2 GANs and their classifier counterparts. The SG2 generators of an iteration are trained with a classifier model of each of the previous iterations (using Eq. (2)). i th iteration classifier models are trained on 15 trained instances of i th iteration SG2 generators. The entry at (i, j) is the accuracy mean±std of a iteration i classifier evaluated on images generated with a balanced source of 10 held-out iteration j SG2 generator instances. A value of 0.0 means the classifier is always fooled; a value of 1.0 means it is never fooled. Note that, in all cases, classifiers are at least 95% accurate on held-out real images.

fool the previous iteration’s classifiers ($G^{(i)}$ fools $C^{(i-1)}$), however, they are detectable by the classifier from *two* iterations ago: $G^{(i)}$ is readily detected by classifier $C^{(i-2)}$. This suggests that, fittingly, in the “memoryless” training regime generators oscillate between one of two clusters of artifacts depending on the parity of the iteration.

5. Conclusion

GAN-generated images exhibit “artifacts” that distinguish them from real images, even if such artifacts are not apparent to the human eye. Many of these artifacts are “knowledge gaps:” rather than being sample-specific, they are present in most or all of the samples from a given generator. We study two settings: using DCGAN with the MNIST dataset, and StyleGAN2 with the FFHQ dataset. Our results suggest that, far from being random or just instance-specific, some artifacts are produced in a regular, repeatable way *across* independent generators of the same GAN architecture (and trained on the same dataset), comprising an “artifact space.” This is evident from the fact that all classifiers we tested needed samples from a relatively small pool of generators in order to reliably generalize to unseen generator instances.

Between the two settings, and in stark contrast, DCGAN generators trained on MNIST were unable to fool unseen classifiers without compromising output quality. But StyleGAN2 generators trained on FFHQ learned to fool unseen classifiers of the same architecture with high reliability and often from exposure to a single classifier instance. However, this fooling ability does not generalize to classifier instances of other architectures. This suggests that the subset of the available artifacts learned by a classifier is similar within architecture but different between architecture. Further, StyleGAN2 generators’ fooling ability generalizes more reliably when the classifier in question is high capacity (ResNet-50, Inception-v3). Hence, high-capacity architectures learn a large proportion of the artifacts available to

them, resulting in correlated behavior. However, a classifier like MobileNetV2 only learns a portion of the artifact space available to them, and so multiple trained classifier instances are required when training a generator in order to reliably fool held-out MobileNetV2 instances. This does not imply that MobileNetV2-based classifiers learn an arbitrary subset of the artifact space: pairwise comparisons indicate that they instead tend to fall into clusters that are “mutually-fooling.”

When iterating the process in the StyleGAN2 setting, we find that the StyleGAN2 generators continue to quickly learn to fool the classifiers. Similarly, the classifiers require samples from only a few of the newly trained generators to learn how to detect unseen generators reliably. This persistence of detector generalization suggesting that the constraint of needing to fool a classifier induces a consistent transformation on the artifact space across generators (rather than, say, inducing a random transformation specific to each generator). Thus, the StyleGAN2 generators of each new iteration produce artifact spaces that are (mostly) distinct compared to previous iteration but are largely the same *within* the iteration. This suggests an induced preference or ordering over artifact spaces, and merits further study. Our results also hint that this process doesn’t continue indefinitely: after a sufficient number of iterations, the StyleGAN2 generators may begin to “decorrelate” in terms of their artifact space.

Lastly, we discuss the limitations and societal impact of our work. While we expect our findings to be a property of GAN generators broadly, verifying this is left for future work. In line with previous work [9, 31], it would be instructive to investigate the overlap in the “artifact spaces” across different GAN architectures, among multiple iterations. From a misinformation mitigation perspective, we hope our findings will motivate further research into detection of GAN-generated images. However, as with all published work on detection, we cannot conclusively say that there is no potential this could benefit bad actors.

References

- [1] Nicholas Carlini and Hany Farid. Evading deepfake-image detectors with white-and black-box attacks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 658–659, 2020. **2**
- [2] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. In *European Conference on Computer Vision*, pages 103–120. Springer, 2020. **2**
- [3] Davide Cozzolino, Diego Gragnaniello, Giovanni Poggi, and Luisa Verdoliva. Towards Universal GAN Image Detection. In *2021 International Conference on Visual Communications and Image Processing (VCIP)*, pages 1–5. IEEE, 2021. **2**
- [4] Davide Cozzolino, Justus Thies, Andreas Rössler, Christian Riess, Matthias Nießner, and Luisa Verdoliva. ForensicTransfer: Weakly-supervised Domain Adaptation for Forgery Detection. *arXiv preprint arXiv:1812.02510*, 2018. **2**
- [5] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *International Conference on Machine Learning*, pages 3247–3258. PMLR, 2020. **2**
- [6] Joel Frank and Thorsten Holz. [RE] CNN-generated images are surprisingly easy to spot... for now. *arXiv preprint arXiv:2104.02984*, 2021. **2**
- [7] Michael Goebel and BS Manjunath. Adversarial Attacks on Co-Occurrence Features for GAN Detection. *arXiv preprint arXiv:2009.07456*, 2020. **2**
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. **1**
- [9] Diego Gragnaniello, Davide Cozzolino, Francesco Marra, Giovanni Poggi, and Luisa Verdoliva. Are GAN generated images easy to detect? A critical analysis of the state-of-the-art. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. **1, 2, 8**
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **1**
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6629–6640, 2017. **7**
- [12] Kashmir Hill and Jeremy White. Designed to Deceive: Do These People Look Real to You?, 2020. **1**
- [13] Chih-Chung Hsu, Chia-Yen Lee, and Yi-Xiu Zhuang. Learning to detect fake face images in the wild. In *2018 International Symposium on Computer, Consumer and Control (IS3C)*, pages 388–391. IEEE, 2018. **2**
- [14] Nils Hulzebosch, Sarah Ibrahim, and Marcel Worring. Detecting cnn-generated facial images in real-world scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 642–643, 2020. **1**
- [15] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations*, 2018. **2**
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. **1, 2, 4**
- [17] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. **1**
- [18] Federica Lago, Cecilia Pasquini, Rainer Böhme, Hélène Dumont, Valérie Goffaux, and Giulia Boato. More Real Than Real: A Study on Human Visual Perception of Synthetic Faces [Applications Corner]. *IEEE Signal Processing Magazine*, 39(1):109–116, 2021. **1**
- [19] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010. **1**
- [20] Chi Liu, Huajie Chen, Tianqing Zhu, Jun Zhang, and Wanlei Zhou. Making DeepFakes more spurious: evading deep face forgery detection via trace removal attack. *arXiv preprint arXiv:2203.11433*, 2022. **2**
- [21] Francesco Marra, Diego Gragnaniello, Davide Cozzolino, and Luisa Verdoliva. Detection of gan-generated fake images over social networks. In *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 384–389. IEEE, 2018. **2**
- [22] Francesco Marra, Diego Gragnaniello, Luisa Verdoliva, and Giovanni Poggi. Do gans leave artificial fingerprints? In *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 506–511. IEEE, 2019. **2**
- [23] Huaxiao Mo, Bolin Chen, and Weiqi Luo. Fake faces identification via convolutional neural network. In *Proceedings of the 6th ACM workshop on information hiding and multimedia security*, pages 43–47, 2018. **2**
- [24] Sophie Nightingale and Hany Farid. AI-synthesized faces are indistinguishable from real faces and more trustworthy. *Proceedings of the National Academy of Sciences of the United States of America*, 119, 2022. **1**
- [25] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. **1**
- [26] 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. *International journal of computer vision*, 115(3):211–252, 2015. **2**
- [27] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520. IEEE, 2018. **1**
- [28] C Szegedy, V Vanhoucke, S Ioffe, J Shlens, and ZB Wojna. Rethinking the Inception Architecture for Computer Vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, volume 2016, pages 2818–2826. IEEE, 2016. **1**
- [29] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan

- Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [2](#)
- [30] Shahroz Tariq, Sangyup Lee, Hoyoung Kim, Youjin Shin, and Simon S Woo. Detecting both machine and human created fake face images in the wild. In *Proceedings of the 2nd international workshop on multimedia privacy and security*, pages 81–87, 2018. [2](#)
- [31] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A. Efros. CNN-Generated Images Are Surprisingly Easy to Spot... for Now. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#), [2](#), [8](#)
- [32] Ning Yu, Larry S Davis, and Mario Fritz. Attributing fake images to gans: Learning and analyzing gan fingerprints. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7556–7566, 2019. [2](#)
- [33] Xinwei Zhao and Matthew C Stamm. Making GAN-Generated Images Difficult To Spot: A New Attack Against Synthetic Image Detectors. *arXiv preprint arXiv:2104.12069*, 2021. [2](#)