arXiv:1704.04865v1 [cs.CV] 17 Apr 2017

Gang of GANs: Generative Adversarial Networks with Maximum Margin Ranking

Felix Juefei-Xu
Carnegie Mellon University
felixu@cmu.edu

Vishnu Naresh Boddeti
Michigan State University
vishnu@msu.edu

Marios Savvides
Carnegie Mellon University
msavvid@ri.cmu.edu

Abstract

Traditional generative adversarial networks (GAN) and many of its variants are trained by minimizing the KL or JS-divergence loss that measures how close the generated data distribution is from the true data distribution. A recent advance called the WGAN based on Wasserstein distance can improve on the KL and JS-divergence based GANs, and alleviate the gradient vanishing, instability, and mode collapse issues that are common in the GAN training. In this work, we aim at improving on the WGAN by first generalizing its discriminator loss to a margin-based one, which leads to a better discriminator, and in turn a better generator, and then carrying out a progressive training paradigm involving multiple GANs to contribute to the maximum margin ranking loss so that the GAN at later stages will improve upon early stages. We call this method ***Gang of GANs (GoGAN)****. We have shown theoretically that the proposed GoGAN can reduce the gap between the true data distribution and the generated data distribution by at least half in an optimally trained WGAN. We have also proposed a new way of measuring GAN quality which is based on image completion tasks. We have evaluated our method on four visual datasets: CelebA, LSUN Bedroom, CIFAR-10, and 50K-SSFF, and have seen both visual and quantitative improvement over baseline WGAN.*

1. Introduction

Generative approaches can learn from the tremendous amount of data around us and generate new instances that are like the data they have observed, in any domain. This line of research is extremely important because it has the potential to provide meaningful insight into the physical world we human beings can perceive. Take visual perception for instance, the generative models have much smaller number of parameters than the amount of visual data out there in the world, which means that in order for the generative models to come up with new instances that are like the actual true data, they have to search for intrinsic pattern and distill the essence. We can in turn capitalize on that and make machines understand, describe, and model the visual world better. Recently, three classes of algorithms have emerged as successful generative approaches to model the visual data in an unsupervised manner.

Variational autoencoders (VAEs) [17] formalize the generative problem in the framework of probabilistic graphical models where we are to maximize a lower bound on the log likelihood of the training data. The probabilistic graphical models with latent variables allow us to perform both learning and Bayesian inference efficiently. By projecting into a learned latent space, samples can be reconstructed from that space. The VAEs are straightforward to train but at the cost of introducing potentially restrictive assumptions about the approximate posterior distribution. Also, their generated samples tend to be slightly blurry. Autoregressive models such as PixelRNN [32] and PixelCNN [33] get rid of the latent variables and instead directly model the conditional distribution of every individual pixel given previous pixels starting from top-left corner. PixelRNN/CNN have a stable training process via softmax loss and currently give the best log likelihoods on the generated data, which is an indicator of high plausibility. However, they are relatively inefficient during sampling and do not easily provide simple low-dimensional latent codes for images.

Generative adversarial networks (GANs) [10] simultaneously train a generator network for generating realistic images, and a discriminator network for distinguishing between the generated images and the samples from the training data (true distribution). The two players (generator and discriminator) play a two-player minimax game until Nash equilibrium where the generator is able to generate images as genuine as the ones sampled from the true distribution, and the discriminator is no longer able to distinguish between the two sets of images, or equivalently is guessing at random chance. In the traditional GAN formulation, the generator and the discriminator are updated by receiving gradient signals from the loss induced by observing discrepancies between the two distributions by the discriminator.

From our perspective, GANs are able to generate images with the highest visual quality by far. The image details are sharp as well as semantically sound.

Motivation: Although we have observed many successes in applying GANs to various scenarios as well as in many GAN variants that come along, there has not been much work dedicated to improving GAN itself from a very fundamental point of view. Ultimately, we are all interested in the end-product of a GAN, which is the image it can generate. Although we are all focusing on the performance of the GAN generator, we must know that its performance is directly affected by the GAN discriminator. In short, to make the generator stronger, we need a stronger opponent, which is a stronger discriminator in this case. Imagine if we have a weak discriminator which does a poor job telling generated images from the true images, it takes only a little effort for the generator to win the two-player minimax game as described in the original work of GAN [10]. To further improve upon the state-of-the-art GAN method, one possible direction is to enforce a maximum margin ranking loss in the optimization of the discriminator, which will result in a stronger discriminator that attends to the fine details of images, and a stronger discriminator helps obtain a stronger generator in the end.

In this work, we are focusing on how to further improve the GANs by incorporating a maximum margin ranking criterion in the optimization, and with a progressive training paradigm. We call the proposed method Gang of GANs (GoGAN)[1]. Our contributions include (1) generalizing on the Wasserstein GAN discriminator loss with a margin-based discriminator loss; (2) proposing a progressive training paradigm involving multiple GANs to contribute to the maximum margin ranking loss so that the GAN at later GoGAN stages will improve upon early stages; (3) showing theoretical guarantee that the GoGAN will bridge the gap between true data distribution and generated data distribution by at least half; and (4) proposing a new quality measure for the GANs through image completion tasks.

2. Related Work

In this section, we review recent advances in GAN research as well as many of its variants and related work.

Deep convolutional generative adversarial networks (DCGAN) [28] are proposed to replace the multilayer perceptron in the original GAN [10] for more stable training, by utilizing strided convolutions in place of pooling layers, and fractional-strided convolutions in place of image upsampling. Conditional GAN [22] is proposed as a variant of GAN by extending it to a conditional model, where both the generator and discriminator are conditioned on some extra auxiliary information, such as class labels. The conditioning is performed by feeding the auxiliary information into both the generator and the discriminator as additional input layer. Another variant of GAN is called auxiliary classifier GAN (AC-GAN) [25], where every generated sample has a corresponding class label in addition to the noise. The generator needs both for generating images. Meanwhile, the discriminator does two things: giving a probability distribution over image sources, and giving a probability distribution over the class labels. Bidirectional GAN (BiGAN) [7] is proposed to bridge the gap that conventional GAN does not learn the inverse mapping which projects the data back into the latent space, which can be very useful for unsupervised feature learning. The BiGAN not only trains a generator, but also an encoder that induces a distribution for mapping data point into the latent feature space of the generative model. At the same time, the discriminator is also adapted to take input from the latent feature space, and then predict if an image is generated or from the true distribution. There is a pathway from the latent feature $\mathbf{z}$ to the generated data $G(\mathbf{z})$ via the generator G, as well as another pathway from the data $\mathbf{x}$ back to the latent feature representation $E(\mathbf{x})$ via the newly added encoder E. The generated image together with the input latent noise $(G(\mathbf{z}), \mathbf{z})$, and the true data together with its encoded latent representation $(\mathbf{x}, E(\mathbf{x}))$ are fed into the discriminator D for classification. There is a concurrent work proposed in [8] that has the identical model. A sequential variant of the GAN is the Laplacian generative adversarial networks (LAPGAN) [6] model which generates images in a coarse-to-fine manner by generating and upsampling in multiple steps. It is worth mentioning the sequential variant of the VAE is the deep recurrent attentive writer (DRAW) [11] model that generates images by accumulating updates into a canvas using a recurrent network. Built upon the idea of sequential generation of images, the recurrent adversarial networks [15] has been proposed to let the recurrent network to learn the optimal generation procedure by itself, as opposed to imposing a coarse-to-fine structure on the procedure. Introspective adversarial network (IAN) [4] is proposed to hybridize the VAE and the GAN. It leverages the power of the adversarial objective while maintaining the efficient inference mechanism of the VAE. The generative multi-adversarial networks (GMAN) [9] extends the GANs to multiple discriminators. For a fixed generator G, N randomly instantiated copies of the discriminators are utilized to present the maximum value of each value function as the loss for the generator. Requiring the generator to minimize the $\max$ forces G to generate high fidelity samples that must hold up under the scrutiny of all N discriminators. Layered recursive generative adversarial networks (LR-GAN) [35] generates images in a recursive fashion. It first generates a background, and then generates a foreground by conditioning on the back-

[1]Implementation and future updates will be available at http://xujuefei.com/gogan.

ground, along with a mask and an affine transformation that together define how the background and foreground should be composed to obtain a complete image. The foreground-background mask is estimated in a completely unsupervised way without using any object masks for training. Authors of [24] have shown that the generative-adversarial approach in GAN is a special case of an existing more general variational divergence estimation approach, and that any f-divergence can be used for training generative neural samplers. InfoGAN [5] method is a generative adversarial network that also maximizes the mutual information between a small subset of the latent variables and the observation. A lower bound of the mutual information can be derived and optimized efficiently. Rather than a single unstructured noise vector to be input into the generator, InfoGAN decomposes the noise vector into two parts: a source of incompressible noise z and a latent code c that will target the salient structured semantic features of the data distribution, and the generator thus becomes $G(z, c)$. The authors have added an information-theoretic regularization to ensure there is high mutual information between the latent code c and the generator distribution $G(z, c)$. To strive for a more stable GAN training, the energy-based generative adversarial networks (EBGAN) [38] is proposed which views the discriminator as an energy function that assigns low energy to the regions near the data manifold and higher energy to other regions. The authors have shown one instantiation of EBGAN using an autoencoder architecture, with the energy being the reconstruction error. The boundary-seeking GAN (BGAN) [13] aims at generating samples that lie on the decision boundary of a current discriminator in training at each update. The hope is that a generator can be trained in this way to match a target distribution at the limit of a perfect discriminator. Least squares GAN [21] adopts a least squares loss function for the discriminator, which is equivalent to a multi-class GAN with the ℓ_2 loss function. The authors have shown that the objective function yields minimizing the Pearson χ^2 divergence. The stacked GAN (SGAN) [14] consists of a top-down stack of GANs, each trained to generate plausible lower-level representations, conditioned on higher-level representations. Discriminators are attached to each feature hierarchy to provide intermediate supervision. Each GAN of the stack is first trained independently, and then the stack is trained end-to-end.

Perhaps the most seminal GAN-related work since the inception of the original GAN [10] idea is the Wasserstein GAN (WGAN) [3]. Efforts have been made to fully understand the training dynamics of generative adversarial networks through theoretical analysis [2], which leads to the creation of the WGAN. The two major issues with the original GAN and many of its variants are the vanishing gradient issues and the mode collapse issue. By incorporating a smooth Wasserstein distance metric and objective, as opposed to the KL-divergence and JS-divergence, the WGAN is able to overcome the vanishing gradient and mode collapse issues. WGAN also has made training and balancing between the generator and discriminator much easier in the sense that one can now train the discriminator till optimality, and then gradually improve the generator. Moreover, it provides an indicator (based on the Wasserstein distance) for the training progress, which correlates well with the visual image quality of the generated samples.

Other applications include cross-domain image generation [30] through a domain transfer network (DTN) which employs a compound of loss functions including a multi-class GAN loss, an f-constancy component, and a regularization component that encourages the generator to map samples from the target domain to themselves. The image-to-image translation approach [16] is based on conditional GAN, and learns a conditional generative model for generating a corresponding output image at a different domain, conditioned on an input image. The image super-resolution GAN (SRGAN) [19] combines both the image content loss and the adversarial loss for recovering high-resolution counterpart of the low-resolution input image. The plug and play generative networks (PPGN) [23] is able to produce high quality images at higher resolution for all 1000 ImageNet categories. It is composed of a generator that is capable of drawing a wide range of image types, and a replaceable condition network that tells the generator what to draw, hence plug and play.

3. Proposed Method: Gang of GANs

In this section we will review the original GAN [10] and its convolutional variant DCGAN [28]. We will then analyze how to further improve the GAN model with WGAN [3], and introduce our Gang of GANs (GoGAN) method.

3.1. GAN and DCGAN

The GAN [10] framework trains two networks, a generator $\mathcal{G}_\theta(\mathbf{z}) : \mathbf{z} \rightarrow \mathbf{x}$, and a discriminator $\mathcal{D}_\omega(\mathbf{x}) : \mathbf{x} \rightarrow [0, 1]$. $\mathcal{G}$ maps a random vector $\mathbf{z}$, sampled from a prior distribution $p_\mathbf{z}(\mathbf{z})$, to the image space. $\mathcal{D}$ maps an input image to a likelihood. The purpose of $\mathcal{G}$ is to generate realistic images, while $\mathcal{D}$ plays an adversarial role to discriminate between the image generated from $\mathcal{G}$, and the image sampled from data distribution p_{data}. The networks are trained by optimizing the following minimax loss function: $\min_{\mathcal{G}} \max_{\mathcal{D}} V(\mathcal{G}, \mathcal{D}) = \mathbb{E}_{\mathbf{x}\sim p_{\text{data}}(\mathbf{x})}[\log(\mathcal{D}(\mathbf{x}))] + \mathbb{E}_{\mathbf{z}\sim p_\mathbf{z}(\mathbf{z})}[\log(1 - \mathcal{D}(\mathcal{G}(\mathbf{z}))]$ where $\mathbf{x}$ is the sample from the p_{data} distribution; $\mathbf{z}$ is randomly generated and lies in some latent space. There are many ways to structure $\mathcal{G}(\mathbf{z})$. The DCGAN [28] uses fractionally-strided convolutions to upsample images instead of fully-connected neurons as shown in Figure 1. The generator $\mathcal{G}$ is updated to fool the dis-

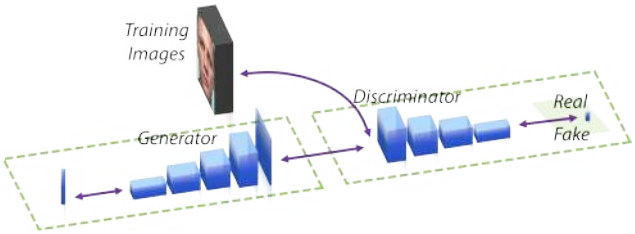

Figure 1: Pipeline of a standard DCGAN with the generator $\mathcal{G}$ mapping a random vector $\mathbf{z}$ to an image and the discriminator $\mathcal{D}$ mapping the image (from true distribution or generated) to a probability value.

criminator $\mathcal{D}$ into wrongly classifying the generated sample, $\mathcal{G}(\mathbf{z})$, while the discriminator $\mathcal{D}$ tries not to be fooled. Here, both $\mathcal{G}$ and $\mathcal{D}$ are deep convolutional neural networks and are trained with an alternating gradient descent algorithm. After convergence, $\mathcal{D}$ is able to reject images that are too fake, and $\mathcal{G}$ can produce high quality images faithful to the training distribution (true distribution p_{data}).

3.2. Wasserstein GAN and Improvement over GAN

In the original GAN, Goodfellow *et al.* [10] have proposed the following two loss functions for the generator: $\mathbb{E}_{\mathbf{z}\sim P_{\mathbf{z}}(\mathbf{z})}[\log(1-\mathcal{D}(\mathcal{G}(\mathbf{z})))]$ and $\mathbb{E}_{\mathbf{z}\sim P_{\mathbf{z}}(\mathbf{z})}[-\log \mathcal{D}(\mathcal{G}(\mathbf{z}))]$. The latter one is referred to as the $-\log \mathcal{D}$ trick [10, 2, 3].

Unfortunately, both forms can lead to potential issues in training the GAN. In short, the former loss function can lead to gradient vanishing problem, especially when the discriminator is trained to be very strong. The real image distribution $\mathbb{P}_r$ and the generated image distribution $\mathbb{P}_g$ have support contained in two closed manifolds that don't perfectly align and don't have full dimension. When the discriminator is near optimal, minimizing the loss of the generator is equivalent to minimizing the JS-divergence between $\mathbb{P}_r$ and $\mathbb{P}_g$, but due to the aforementioned reasons, the JS-divergence will always be a constant $2\log 2$, which allows the existence of an optimal discriminator to (almost) perfectly carve the two distributions, *i.e.*, assigning probability 1 to all the real samples, and 0 to all the generated ones, which renders the gradient of the generator loss to go to 0.

For the latter case, it can be shown that minimizing the loss function is equivalent to minimizing $\text{KL}(\mathbb{P}_g\|\mathbb{P}_r) - 2\text{JS}(\mathbb{P}_r\|\mathbb{P}_g)$, which leads to instability in the gradient because it simultaneous tries to minimize the KL-divergence and maximize the JS-divergence, which is a less ideal loss function design. Even the KL term by itself has some issues. Due to its asymmetry, the penalty for two types of errors is quite different. For example, when $\mathbb{P}_g(\mathbf{x}) \to 0$ and $\mathbb{P}_r(\mathbf{x}) \to 1$, we have $\mathbb{P}_g(\mathbf{x})\log\frac{\mathbb{P}_g(\mathbf{x})}{\mathbb{P}_r(\mathbf{x})} \to 0$, which has almost 0 contribution to $\text{KL}(\mathbb{P}_g\|\mathbb{P}_r)$. On the other hand, when $\mathbb{P}_g(\mathbf{x}) \to 1$ and $\mathbb{P}_r(\mathbf{x}) \to 0$, we have $\mathbb{P}_g(\mathbf{x})\log\frac{\mathbb{P}_g(\mathbf{x})}{\mathbb{P}_r(\mathbf{x})} \to +\infty$, which has gigantic contribution to $\text{KL}(\mathbb{P}_g\|\mathbb{P}_r)$. So the first type of error corresponds to that the generator fails to produce realistic samples, which has tiny penalty, and the second type of error corresponds to that the generator produces unrealistic samples, which has enormous penalty. Under this reality, the generator would rather produce repetitive and 'safe' samples, than samples with high diversity with the risk of triggering the second type of error. This causes the infamous mode collapse.

WGAN [2, 3] avoids the gradient vanishing and mode collapse issues in the original GAN and many of its variants by adopting a new distance metric: the Wasserstein-1 distance, or the earth-mover distance as follows:

$$W(\mathbb{P}_r, \mathbb{P}_g) = \inf_{\gamma\in\Gamma(\mathbb{P}_r,\mathbb{P}_g)} \mathbb{E}_{(x,y)\sim\gamma}[\|x-y\|] \qquad (1)$$

where $\Gamma(\mathbb{P}_r, \mathbb{P}_g)$ is the set of all joint distributions $\gamma(x, y)$ whose marginals are $\mathbb{P}_r$ and $\mathbb{P}_g$ respectively. One of the biggest advantages of the Wasserstein distance over KL and JS-divergence is that it is smooth, which is very important in providing meaningful gradient information when the two distributions have support contained in two closed manifolds that don't perfectly aligned don't have full dimension, in which case KL and JS-divergence would fail to provide gradient information successfully. However, the infimum $\inf_{\gamma\in\Gamma(\mathbb{P}_r,\mathbb{P}_g)}$ is highly intractable. Thanks to the Kantorovich-Rubinstein duality [34], the Wasserstein distance becomes: $W(\mathbb{P}_r, \mathbb{P}_g) = \sup_{\|f\|_L\leq 1} \mathbb{E}_{x\sim\mathbb{P}_r}[f(x)] - \mathbb{E}_{x\sim\mathbb{P}_g}[f(x)]$, where the supremum is over all the 1-Lipschitz functions. Therefore, we can have a parameterized family of functions $\{f_w\}_{w\in\mathcal{W}}$ that are K-Lipschitz for some K, and the problem we are solving now becomes: $\max_{w:|f_w|_L\leq K} \mathbb{E}_{x\sim\mathbb{P}_r}[f_w(x)] - \mathbb{E}_{z\sim p(z)}[f_w(g_\theta(x))] \approx K\cdot W(\mathbb{P}_r, \mathbb{P}_g)$. Let the f_w (discriminator) be a neural network with weights w, and maximize $L = \mathbb{E}_{x\sim\mathbb{P}_r}[f(x)] - \mathbb{E}_{x\sim\mathbb{P}_g}[f(x)]$ as much as possible so that it can well approximate the actual Wasserstein distance between real data distribution and generated data distribution, up to a multiplicative constant. On the other hand, the generator will try to minimize L, and since the first term in L does not concern the generator, its loss function is to minimize $-\mathbb{E}_{x\sim\mathbb{P}_g}[f(x)]$, and the loss function for the discriminator is to minimize $\mathbb{E}_{x\sim\mathbb{P}_g}[f(x)] - \mathbb{E}_{x\sim\mathbb{P}_r}[f(x)] = -L$.

3.3. Gang of GANs (GoGAN)

In this section, we will discuss our proposed GoGAN method which is a progressive training paradigm to improve the GAN, by allowing GANs at later stages to contribute to a new ranking loss function that will improve the GAN performance further. Also, at each GoGAN stage, we generalize on the WGAN discriminator loss, and arrive at a margin-based discriminator loss, and we call the network margin GAN (MGAN). The entire GoGAN flowchart is shown in Figure 2, and we will introduce the components involved.

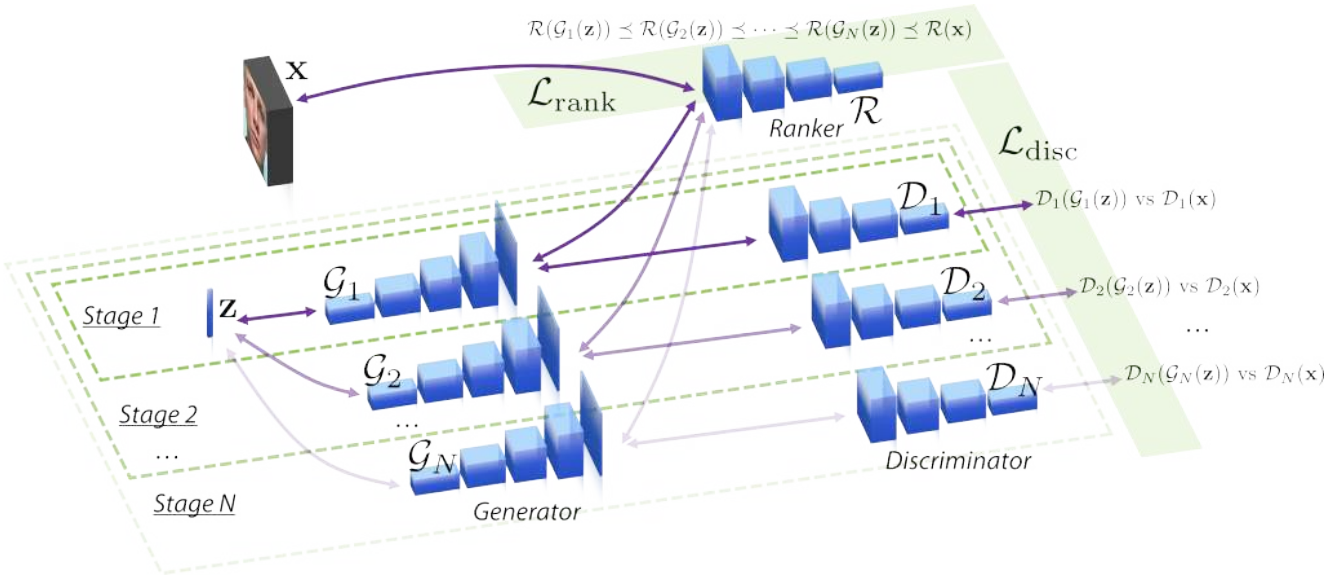

Figure 2: Flowchart of the proposed GoGAN method.

Based on the previous discussion, we have seen that WGAN has several advantages over the traditional GAN. Recall that $\mathcal{D}_{w_i}(\mathbf{x})$ and $\mathcal{D}_{w_i}(\mathcal{G}_{\theta_i}(\mathbf{z}))$ are the discriminator score for the real image $\mathbf{x}$ and generated image $\mathcal{G}_{\theta_i}(\mathbf{z})$ in Stage-$(i+1)$ GoGAN. In order to further improve it, we have proposed a margin-based WGAN discriminator loss:

$$\mathcal{L}_{\text{disc}} = [\mathcal{D}_{w_{i+1}}(\mathcal{G}_{\theta_{i+1}}(\mathbf{z})) + \epsilon - \mathcal{D}_{w_{i+1}}(\mathbf{x})]_+ \qquad (2)$$

where $[x]_+ = \max(0, x)$ is the hinge loss. This MGAN loss function is a generalization of the discriminator loss in WGAN. When the margin $\epsilon \to \infty$, this loss becomes WGAN discriminator loss.

The intuition behind the MGAN loss is as follows. WGAN loss treats a gap of 10 or 1 equally and it tries to increase the gap even further. The MGAN loss will focus on increasing separation of examples with gap 1 and leave the samples with separation 10, which ensures a better discriminator, hence a better generator. We will see next that the MGAN loss can be extended even further by incorporating margin-based ranking when go beyond a single MGAN.

Ranker $\mathcal{R}$: When going from Stage-i GoGAN to Stage-$(i+1)$ GoGAN, we incorporate a margin-based ranking loss in the progressive training of the GoGAN for ensuring that the generated images from later GAN training stage is better than those from previous stages. The idea is fairly straight-forward: the discriminator scores coming from the generated images at *later* stages should be ranked closer to that of the images sampled from the true distribution. The ranking loss is:

$$\mathcal{L}_{\text{rank}} = [\mathcal{D}_{w_i}(\mathcal{G}_{\theta_i}(\mathbf{z})) + 2\epsilon - \mathcal{D}_{w_{i+1}}(\mathbf{x})]_+ \qquad (3)$$

Combing (2) and (3), the $\mathcal{L}_{\text{disc}}$ and $\mathcal{L}_{\text{rank}}$ loss together are equivalent to enforcing the following ranking strategy. Notice that such ranking constraint only happens between adjacent GoGAN pairs, and it can be easily verified that it has intrinsically established an ordering among all the MGANs involved, which will be further discussed in Section 4.

$$\mathcal{D}_{w_{i+1}}(\mathbf{x}) \geq \mathcal{D}_{w_{i+1}}(\mathcal{G}_{\theta_{i+1}}(\mathbf{z})) + \epsilon \qquad (4)$$
$$\mathcal{D}_{w_{i+1}}(\mathbf{x}) \geq \mathcal{D}_{w_i}(\mathcal{G}_{\theta_i}(\mathbf{z})) + 2\epsilon \qquad (5)$$

The weights of the ranker $\mathcal{R}$ and the discriminator $\mathcal{D}$ are tied together. Conceptually, from Stage-2 and onward, the ranker is just the discriminator which takes in extra ranking loss in addition to the discriminator loss already in place for the MGAN. In Figure 2, the ranker is a separate block, but only for illustrative purpose. Different training stages are encircled by green dotted lines with various transparency levels. The purple solid lines show the connectivity within the GoGAN, with various transparency levels in accordance with the progressive training stages. The arrows on both ends of the purple lines indicate forward and backward pass of the information and gradient signal. If the entire GoGAN is trained, the ranker will have achieved the following desired goal: $\mathcal{R}(\mathcal{G}_1(\mathbf{z})) \preceq \mathcal{R}(\mathcal{G}_2(\mathbf{z})) \preceq \mathcal{R}(\mathcal{G}_3(\mathbf{z})) \preceq \cdots \preceq \mathcal{R}(\mathcal{G}_K(\mathbf{z})) \preceq \mathcal{R}(\mathbf{x})$, where $\preceq$ indicates relative ordering. The total loss for GoGAN can be written as: $\mathcal{L}_{\text{GoGAN}} = \lambda_1 \cdot \mathcal{L}_{\text{disc}} + \lambda_2 \cdot \mathcal{L}_{\text{rank}}$, where weighting parameters λ_1 and λ_2 controls the relative strength.

4. Theoretical Analysis

In WGAN [3], the following loss function involving the weights updating of the discriminator and the generator is

a good indicator of the EM distance during WGAN training: $\max_{w\in\mathcal{W}} \mathbb{E}_{x\sim\mathbb{P}_r}[\mathcal{D}_w(\mathbf{x})] - \mathbb{E}_{\mathbf{z}\sim p_z}[\mathcal{D}_w(\mathcal{G}_\theta(\mathbf{x}))]$. This loss function is essentially the Gap $\mathbf{\Gamma}$ between real data distribution and generated data distribution, and of course the discriminator is trying to push the gap larger. The realization of this loss function for one batch is as follows:

$$\text{Gap} = \mathbf{\Gamma} = \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}_w(\mathbf{x}^{(i)}) - \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}_w(\mathcal{G}_\theta(\mathbf{z}^{(i)})) \quad (6)$$

Theorem 4.1. *GoGAN with ranking loss (3) trained at its equilibrium will reduce the gap between the real data distribution $\mathbb{P}_r$ and the generated data distribution $\mathbb{P}_\theta$* ***at least by half*** *for Wasserstein GAN trained at its optimality.*

Proof. Let $\mathcal{D}^*_{w_1}$ and $\mathcal{G}^*_{\theta_1}$ be the optimally trained discriminator and generator for the original WGAN (Stage-1 GoGAN). Let $\mathcal{D}^*_{w_2}$ and $\mathcal{G}^*_{\theta_2}$ be the optimally trained discriminator and generator for the Stage-2 GoGAN in the proposed progressive training framework.

The gap between real data distribution and the generated data distribution for Stage-1 to Stage-N GoGAN is:

$$\mathbf{\Gamma}_1 = \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_1}(\mathbf{x}^{(i)}) - \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_1}(\mathcal{G}^*_{\theta_1}(\mathbf{z}^{(i)})) \quad (7)$$

$$\mathbf{\Gamma}_N = \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_N}(\mathbf{x}^{(i)}) - \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_N}(\mathcal{G}^*_{\theta_N}(\mathbf{z}^{(i)})) \quad (8)$$

Let us first establish the relationship between gap $\mathbf{\Gamma}_1$ and gap $\mathbf{\Gamma}_2$, and then extends to the $\mathbf{\Gamma}_N$ case.

According to the ranking strategy, we enforce the following ordering:

$$\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) > \mathcal{D}^*_{w_2}(\mathcal{G}^*_{\theta_2}(\mathbf{z}^{(i)})) > \mathcal{D}^*_{w_1}(\mathcal{G}^*_{\theta_1}(\mathbf{z}^{(i)})) \quad (9)$$

which means that

$$\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_2}(\mathcal{G}^*_{\theta_2}(\mathbf{z}^{(i)})) < \mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_1}(\mathcal{G}^*_{\theta_1}(\mathbf{z}^{(i)}))$$

On the left hand side, it is the new gap from Stage-2 GoGAN for one image, and for the the whole batch, this relationship follows:

$$\mathbf{\Gamma}_2 = \frac{1}{m}\sum_{i=1}^{m}\left[\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_2}(\mathcal{G}^*_{\theta_2}(\mathbf{z}^{(i)}))\right] \quad (10)$$

$$< \frac{1}{m}\sum_{i=1}^{m}\left[\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_1}(\mathcal{G}^*_{\theta_1}(\mathbf{z}^{(i)}))\right] \quad (11)$$

$$= \underbrace{\frac{1}{m}\sum_{i=1}^{m}\left[\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_1}(\mathbf{x}^{(i)})\right]}_{\eta_1}$$

$$+ \underbrace{\frac{1}{m}\sum_{i=1}^{m}\left[\mathcal{D}^*_{w_1}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_1}(\mathcal{G}^*_{\theta_1}(\mathbf{z}^{(i)}))\right]}_{\mathbf{\Gamma}_1} \quad (12)$$

Therefore, we have $0 < \mathbf{\Gamma}_2 < \xi_1 + \mathbf{\Gamma}_1$, where the term ξ_1 can be positive, negative, or zero. But only when $\xi_1 \leq 0$, the relation $\mathbf{\Gamma}_2 < \mathbf{\Gamma}_1$ can thus always hold true. In other words, according to the ranking strategy, we have a byproduct relation $\xi_1 \leq 0$ established, which is equivalent to the following expressions:

$$\xi_1 = \frac{1}{m}\sum_{i=1}^{m}\left[\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) - \mathcal{D}^*_{w_2}(\mathbf{x}^{(i)})\right] \leq 0 \quad (13)$$

$$\frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_1}(\mathbf{x}^{(i)}) \geq \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) \quad (14)$$

Combing relations (9) and (14), we can arrive at the new ordering:

$$\frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_1}(\mathbf{x}^{(i)}) \geq \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_2}(\mathbf{x}^{(i)}) > \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_2}(\mathcal{G}^*_{\theta_2}(\mathbf{z}^{(i)})) > \frac{1}{m}\sum_{i=1}^{m}\mathcal{D}^*_{w_1}(\mathcal{G}^*_{\theta_1}(\mathbf{z}^{(i)})) \quad (15)$$

Notice the nested ranking strategy as a result of the derivation. Therefore, when going from Stage-2 to Stage-3 GoGAN, similar relationship can be obtained (for notation simplification, we drop the (i) super script and use bar to represent average over m instances):

$$\overline{\mathcal{D}}^*_{w_2}(\mathbf{x}) \geq \overline{\mathcal{D}}^*_{w_3}(\mathbf{x}) > \overline{\mathcal{D}}^*_{w_3}(\mathcal{G}^*_{\theta_3}(\mathbf{z})) > \overline{\mathcal{D}}^*_{w_2}(\mathcal{G}^*_{\theta_2}(\mathbf{z})) \quad (16)$$

which is equivalent to the following expression when considering the already-existing relationship from Stage-1 to Stage-2 GoGAN:

$$\overline{\mathcal{D}}^*_{w_1}(\mathbf{x}) \geq \overline{\mathcal{D}}^*_{w_2}(\mathbf{x}) \geq \overline{\mathcal{D}}^*_{w_3}(\mathbf{x}) > \overline{\mathcal{D}}^*_{w_3}(\mathcal{G}^*_{\theta_3}(\mathbf{z})) > \overline{\mathcal{D}}^*_{w_2}(\mathcal{G}^*_{\theta_2}(\mathbf{z})) > \overline{\mathcal{D}}^*_{w_1}(\mathcal{G}^*_{\theta_2}(\mathbf{z})) \quad (17)$$

Similar ordering can be established for all the way to Stage-N GoGAN. Let us assume that the distance between the first and last term: $\overline{\mathcal{D}}^*_{w_1}(\mathbf{x})$ and $\overline{\mathcal{D}}^*_{w_1}(\mathcal{G}^*_{\theta_2}(\mathbf{z}))$ is β which is finite, as shown in Figure 3. Let us also assume that the distance between $\overline{\mathcal{D}}^*_{w_i}(\mathbf{x})$ and $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathbf{x})$ is η_i, and the distance between $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathcal{G}^*_{\theta_{i+1}}(\mathbf{z}))$ and $\overline{\mathcal{D}}^*_{w_i}(\mathcal{G}^*_{\theta_i}(\mathbf{z}))$ is φ_i.

Extending the pairwise relationship established by the ranker in (4, 5) to the entire batch, we will have equal margins between the terms $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathbf{x})$, $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathcal{G}^*_{\theta_{i+1}}(\mathbf{z}))$, and $\overline{\mathcal{D}}^*_{w_i}(\mathcal{G}^*_{\theta_i}(\mathbf{z}))$; and the margin between $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathbf{x})$ and $\overline{\mathcal{D}}^*_{w_i}(\mathbf{x})$ remains flexible.

Therefore, we can put the corresponding terms in order as shown in Figure 3, with the distances between the terms η_i and φ_i also showing. The homoscedasticity assumption from the ranker is illustrated by dashed line with the same color. For instance, the distances between adjacent purple dots are the same.

We can establish the following iterative relationship:

$$\varphi_1 = \frac{\beta - \eta_1}{2}, \varphi_2 = \frac{\varphi_1 - \eta_2}{2}, \varphi_N = \frac{\varphi_{N-1} - \eta_N}{2} \quad (18)$$

The total gap reduction $\mathrm{TGR}(N+1)$ all the way to Stage-$(N+1)$ GoGAN is: $\mathrm{TGR}(N+1) = \sum_{i=1}^{N}(\eta_i + \varphi_i)$. $\mathrm{TGR}(\cdot)$ is an increasing function $\mathrm{TGR}(N+1) > \mathrm{TGR}(N)$, and we have:

$$\begin{aligned}\mathrm{TGR}(N+1) > \mathrm{TGR}(2) &= \eta_1 + \varphi_1 = \eta_1 + \frac{1}{2}(\beta - \eta_1) \\ &= \frac{\beta}{2} + \frac{\eta_1}{2} > \frac{\beta}{2} \end{aligned} \quad (19)$$

Therefore, GoGAN with ranking loss (3) trained at its equilibrium will reduce the gap between the real data distribution and the generated data distribution at least by half for Wasserstein GAN trained at its optimality. □

Corollary 4.2. *The total gap reduction up to Stage-$(N+1)$ GoGAN is equal to $\beta - \varphi_N$.*

Proof. Recall the iterative relation from (18):

$$\varphi_N = \frac{1}{2}\varphi_{N-1} - \frac{1}{2}\eta_N \quad (20)$$

$$\Rightarrow 2\varphi_N + \eta_N = \varphi_{N-1} \quad (21)$$

Combining (20) and (21), we can have the following:

$$\varphi_N + \eta_N = \frac{1}{2}\varphi_{N-1} + \frac{1}{2}\eta_N \quad (22)$$

$$\varphi_{N-1} + \eta_{N-1} = \frac{1}{2}\varphi_{N-2} + \frac{1}{2}\eta_{N-1} \quad (23)$$

$$\cdots$$

$$\varphi_2 + \eta_2 = \frac{1}{2}\varphi_1 + \frac{1}{2}\eta_2 \quad (24)$$

Summing up all the LHS and RHS gives (notice the changes in lower and upper bound of summation):

$$\sum_{i=2}^{N}(\varphi_i + \eta_i) = \frac{1}{2}\sum_{i=1}^{N-1}\varphi_i + \frac{1}{2}\sum_{i=2}^{N}\eta_i \quad (25)$$

$$\sum_{i=1}^{N}(\varphi_i + \eta_i) = \frac{1}{2}\sum_{i=1}^{N-1}\varphi_i + \frac{1}{2}\sum_{i=2}^{N}\eta_i + (\varphi_1 + \eta_1) \quad (26)$$

$$\begin{aligned}\sum_{i=1}^{N}(\varphi_i + \eta_i) = \frac{1}{2}\sum_{i=1}^{N}\varphi_i + \frac{1}{2}\sum_{i=1}^{N}\eta_i + (\varphi_1 + \eta_1) \\ - \frac{1}{2}\varphi_N - \frac{1}{2}\eta_1 \end{aligned} \quad (27)$$

$$\mathrm{TGR}(N+1) = \frac{1}{2}\mathrm{TGR}(N+1) - \frac{\varphi_N}{2} + \frac{\eta_1}{2} + \varphi_1 \quad (28)$$

$$\mathrm{TGR}(N+1) = (2\varphi_1 + \eta_1) - \varphi_N = \beta - \varphi_N \quad (29)$$

Therefore, the total gap reduction up to Stage-$(N+1)$ GoGAN is equal to $\beta - \varphi_N$. □

5. Experiments

5.1. Evaluating GANs via Image Completion Tasks

There hasn't been a universal metric to quantitatively evaluate the GAN performance, and often times, we rely on visual examination. This is largely because of the lack of an objective function: what are the generated images gonna be compared against, since there is no corresponding ground-truth images for the generated ones? These are the questions needed to be addressed.

During the WGAN training, we have seen a successful gap indicator that correlates well with image quality. However, it is highly dependent on the particular WGAN model it is based on, and it will be hard to fairly evaluate generated image quality across different WGAN models. We need a metric that is standalone and do not depend on the models.

Perhaps, the Inception score [29] is by far the best solution we have. The score is based on pretrained Inception model. Generated images are pass through the model and those containing *meaningful objects* will have a conditional label distribution $p(y|\boldsymbol{x})$ with low entropy. At the same time, the marginal $\int p(y|\boldsymbol{x} = G(z))dz$ should have high entropy because we expect the GAN to generate varied images. However, we argue that the Inception score will be biased towards the seen objects during the Inception model training, and it measures more of the *"objectness"* in the generated images, rather than the *"realisticity"* the GAN is intended to strive towards.

In this work, we propose a new way to evaluate GAN performance. It is simple and intuitive. We ask the GANs to carry out image completion tasks [36], and the GAN performance is measured by the fidelity (PSNR, SSIM) of the completed image against its ground-truth. There are several advantages: (1) this quality measure works on image level, rather than on the image distribution; (2) the optimization in the image completion procedure utilizes both the generator and the discriminator of the trained GAN, which is a direct indicator of how good the GAN model is; (3) having 1-vs-1 comparison between the ground-truth and the completed image allows very straightforward visual examination of the GAN quality, and also allows head-to-head comparison between various GANs; (4) this is a direct measure of the "realisticity" of the generated image, and also the diversity. Imagine a mode collapse situation happens, the generated images would be very different from the ground-truth images since the latter ones are diverse.

5.2. Details on the Image Completion Tasks

As discussed above, we propose to use the image completion tasks as a quality measure for various GAN models. In short, the quality of the GAN models can be quantitatively measured by the image completion fidelity, in terms of PSNR and SSIM. The motivation is that the image com-

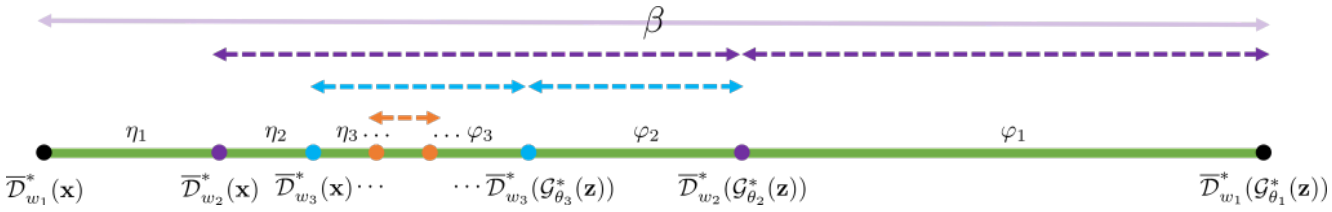

Figure 3: Discriminator scores ordering. η_i is the distance b/t $\overline{\mathcal{D}}^*_{w_i}(\mathbf{x})$ and $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathbf{x})$ and φ_i is the distance b/t $\overline{\mathcal{D}}^*_{w_i}(\mathcal{G}^*_{\theta_i}(\mathbf{z}))$ and $\overline{\mathcal{D}}^*_{w_{i+1}}(\mathcal{G}^*_{\theta_{i+1}}(\mathbf{z}))$.

pletion tasks require both the discriminator $\mathcal{D}$ and the generator $\mathcal{G}$ to work well in order to reach high quality image completion results, as we will see next.

To take on the missing data challenge such as the image completion tasks, we need to utilize both the $\mathcal{G}$ and $\mathcal{D}$ networks from the GoGAN (and its benchmark WGAN), pre-trained with uncorrupted data. After training, $\mathcal{G}$ is able to embed the images from p_{data} onto some non-linear manifold of $\mathbf{z}$. An image that is not from p_{data} (*e.g.* images with missing pixels) should not lie on the learned manifold. Therefore, we seek to recover the "closest" image on the manifold to the corrupted image as the proper image completion. Let us denote the corrupted image as $\mathbf{y}$. To quantify the "closest" mapping from $\mathbf{y}$ to the reconstruction, we define a function consisting of contextual loss and perceptual loss, following the work of Yeh *et al*. [36].

The **contextual loss** is used to measure the fidelity between the reconstructed image portion and the uncorrupted image portion, which is defined as:

$$\mathcal{L}_{\text{contextual}}(\mathbf{z}) = \|\mathbf{M} \odot \mathcal{G}(\mathbf{z}) - \mathbf{M} \odot \mathbf{y}\|_1 \tag{30}$$

where $\mathbf{M}$ is the binary mask of the uncorrupted region and $\odot$ denotes the Hadamard product operation.

The **perceptual loss** encourages the reconstructed image to be similar to the samples drawn from the training set (true distribution p_{data}). This is achieved by updating $\mathbf{z}$ to fool $\mathcal{D}$, or equivalently to have a small gap between $\mathcal{D}(\mathbf{x})$ and $\mathcal{D}(\mathcal{G}(\mathbf{z}))$, where $\mathbf{x}$ is sampled from the real data distribution. As a result, $\mathcal{D}$ will predict $\mathcal{G}(\mathbf{z})$ to be from the real data with a high probability. The same loss for fooling $\mathcal{D}$ as in WGAN and the proposed GoGAN is used:

$$\mathcal{L}_{\text{perceptual}}(\mathbf{z}) = \mathcal{D}(\mathbf{x}) - \mathcal{D}(\mathcal{G}(\mathbf{z})) \tag{31}$$

The corrupted image with missing pixels can now be mapped to the closest $\mathbf{z}$ in the latent representation space with the defined perceptual and contextual losses. $\mathbf{z}$ is updated using back-propagation with the total loss:

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} (\mathcal{L}_{\text{contextual}}(\mathbf{z}) + \lambda \mathcal{L}_{\text{perceptual}}(\mathbf{z})) \tag{32}$$

where λ (set to $\lambda = 0.1$ in our experiments) is a weighting parameter. After finding the optimal solution $\hat{\mathbf{z}}$, the image completion $\mathbf{y}_{\text{completed}}$ can be obtained by:

$$\mathbf{y}_{\text{completed}} = \mathbf{M} \odot \mathbf{y} + (1 - \mathbf{M}) \odot \mathcal{G}(\hat{\mathbf{z}}) \tag{33}$$

5.3. Methods to be Evaluated and Dataset

The **WGAN** baseline uses the Wasserstein discriminator loss [3]. The **MGAN** uses margin-based discriminator loss function discussed in Section 3.3. It is exactly the **Stage-1 GoGAN**, which is a baseline for subsequent GoGAN stages. **Stage-2 GoGAN** incorporates margin-based ranking loss discussed in Section 3.3. These 3 methods will be evaluated on three large-scale visual datasets.

The **CelebA** dataset [20] is a large-scale face attributes dataset with more than 200K celebrity images. The images in this dataset cover large pose variations and background clutter. The dataset includes 10,177 number of subjects, and 202,599 number of face images. We pre-process and align the face images using dLib as provided by the OpenFace [1]. The **LSUN Bedroom** dataset [37] is meant for large-scale scene understanding. We use the bedroom portion of the dataset, with 3,033,042 images. The **CIFAR-10** [18] is an image classification dataset containing a total of 60K 32×32 color images, which are across the following 10 classes: airplanes, automobiles, birds, cats, deers, dogs, frogs, horses, ships, and trucks. The processed image size is 64×64, and the training-testing split is 90-10.

5.4. Training Details of GoGAN

For all the experiments presented in this paper we use the same generator architecture and parameters. We use the DCGAN [28] architecture for both the generator and the discriminator at all stages of the training. Both the generator and the discriminator are learned using optimizers (RMSprop [31]) that are not based on momentum as recommended in [3] with a learning rate of 5e-5. For learning the model at Stage-2 we initialize it with the model learned from Stage-1. In the second stage the model is updated with the ranking loss while the model from stage one held fixed. Lastly, no data augmentation was used for any of our experiments. Different GoGAN stages are trained with the same number total epochs for fair comparison. We will make our implementation publicly available, so readers can refer to it for more detailed hyper-parameters, scheduling, *etc*.

5.5. Results and Discussion

The GoGAN framework is designed to sequentially train generative models and reduce the gap between the true data

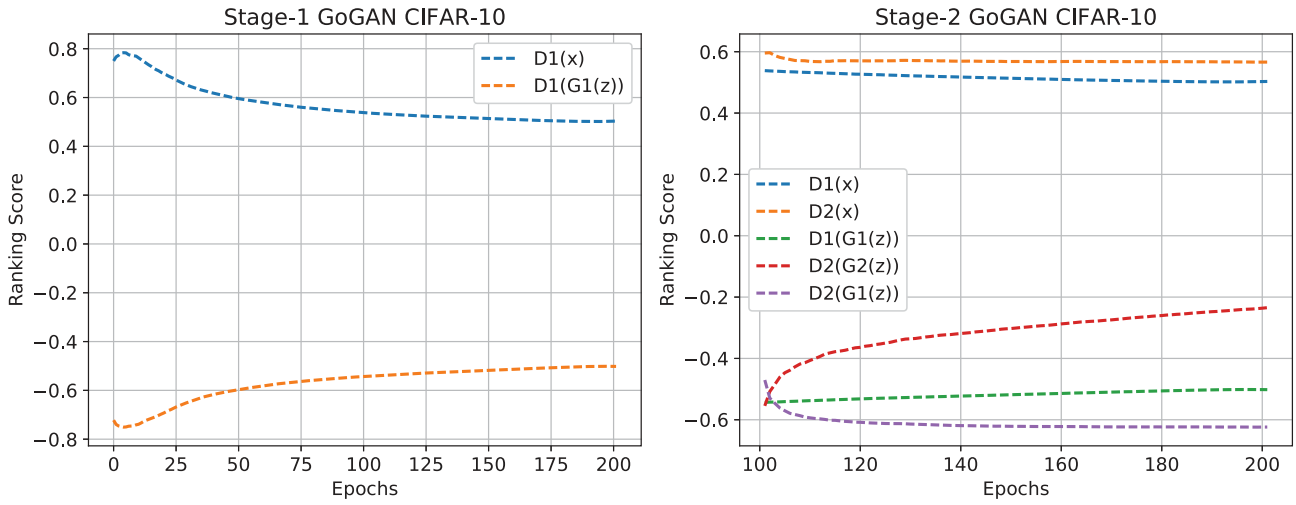

Figure 4: Ranking Scores for Stage-1 and Stage-2 of GoGAN. In the second stage the ranking loss helps ensure that the Stage-2 generator is guaranteed to be stronger than the generator at Stage-1. This is clearly noticeable in the gap between the stage-1 and stage-2 generators.

Table 1: Qualitative results for image completion.

distribution the learned generative model. Figure 4 demonstrates this effect of our proposed approach where the gap between the discriminator scores between the true distribution and the generated distribution reduces from Stage-1 to Stage-2. To quantitatively evaluate the efficacy of our approach we consider the task of image completion *i.e.*, missing data imputation through the generative model. This task is evaluated on three different visual dataset by varying the amount of missing data. We consider five different level of occlusions, occluding the center square region (9%, 25%, 49%, 64%, and 81%) of the image. The image completion task is evaluated by measuring the fidelity between the generated images and the ground-truth images through two metrics: PSNR and SSIM. The results are consolidated in Table 2, 3, and 4 for the 3 datasets respectively. GoGAN consistently outperforms WGAN with the Stage-2 model

	9%	25%	49%	64%	81%
Occluded	23.80	19.18	15.96	14.63	13.41
WGAN	27.26	22.18	18.57	16.93	14.65
Stage-1 GoGAN	27.11	22.27	18.46	16.73	14.63
Stage-2 GoGAN	27.66	22.84	18.77	16.96	14.94
Occluded	0.8679	0.6498	0.3403	0.1578	0.0510
WGAN	0.8985	0.7302	0.4991	0.3488	0.1820
Stage-1 GoGAN	0.8985	0.7385	0.4887	0.3370	0.1847
Stage-2 GoGAN	0.9026	0.7453	0.5017	0.3480	0.1963

Table 2: PSNR (top 4 rows) and SSIM for Celeb-A

	9%	25%	49%	64%	81%
Occluded	22.89	18.15	15.15	14.07	13.08
WGAN	24.22	18.80	15.44	14.10	12.23
State-1 GoGAN	24.28	18.92	15.49	14.07	12.87
State-2 GoGAN	24.34	18.71	15.55	14.32	13.31
Occluded	0.8681	0.6504	0.3330	0.1560	0.0505
WGAN	0.8825	0.6840	0.4025	0.2449	0.1384
State-1 GoGAN	0.8832	0.6909	0.4100	0.2459	0.1295
State-2 GoGAN	0.8835	0.6836	0.4071	0.2475	0.1431

Table 3: PSNR (top 4 rows) and SSIM for LSUN-Bedroom

	9%	25%	49%	64%	81%
Occluded	23.19	18.38	15.05	13.81	12.70
WGAN	22.45	16.76	13.73	12.90	12.17
Stage-1 GoGAN	23.42	17.68	14.13	13.78	11.83
Stage-2 GoGAN	23.68	18.09	14.31	12.90	12.26
Occluded	0.8655	0.6484	0.3207	0.1354	0.0405
WGAN	0.8702	0.6690	0.3901	0.2327	0.1317
Stage-1 GoGAN	0.8777	0.6723	0.3904	0.1980	0.1152
Stage-2 GoGAN	0.8781	0.6806	0.3877	0.2313	0.1034

Table 4: PSNR (top 4 rows) and SSIM for CIFAR-10

also demonstrating improvements over the Stage-1 generator. Our results demonstrate that by enforcing a margin based ranking loss, we can learn sequentially better generative models. We also show qualitative image completion results of Stage-1 GoGAN and Stage-2 GoGAN in Table 1.

5.6. Ablation Studies

In this section, we provide additional experiments and ablation studies on the proposed GoGAN method, and show its improvement over WGAN. For this set of experiments we collect a single-sample dataset containing 50K frontal face images from 50K individuals, which we call the 50K-SSFF dataset. They are sourced from several frontal face datasets including the FRGC v2.0 dataset [27], the MPIE dataset [12], the ND-Twin dataset [26], and mugshot dataset from Pinellas County Sheriff's Office (PCSO). Training and testing split is 9-1, which means we train on 45K images, and test on the remaining 5K. This dataset is single-sample, which means there is only image of a particular subject throughout the entire dataset. Images are aligned using two anchor points on the eyes, and cropped to 64×64.

One-shot Learning: Different from commonly used

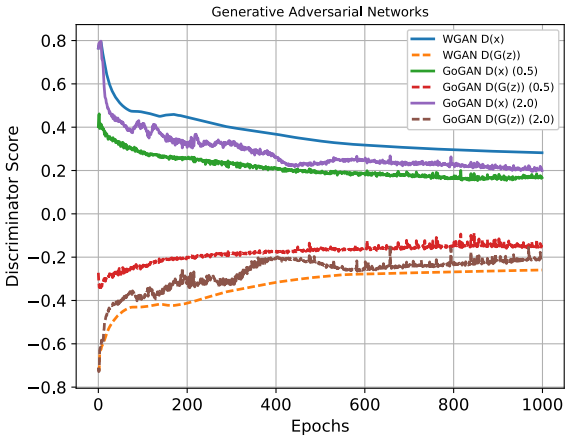

Figure 5: Training Progression Comparison: Here we compare the margin of separation between WGAN and Stage-1 GoGAN with different margin (shown in brackets) in the hinge loss.

celebrity face dataset such as **CelebA** [20], our collected 50K-SSFF dataset is dedicated for one-shot learning in the GAN context due to its single-sample nature. We will explore how the proposed GoGAN method performs under the one-shot learning setting. The majority of the single-sample face images in this dataset are PCSO mugshots, and therefore, we draw a black bar on the original and generated images (see Figures 9, 10, 11) for the sake of privacy protection and is not an artifact of the GAN methods studied.

Training: The GAN models were trained for 1000 epochs each which corresponds to about 135,000 iterations of generator update for a batch size of 64 images. We used the same DCGAN architecture as in the rest of the experiments in the ablation studies.

Margin of Separation: Here we study the impact of the choice of margin in the hinge loss. Figure 5 compares the margin of separation as WGAN and Stage-1 GoGAN are trained to optimality. Figure 6 compares the generators through the image completion task with 49% occlusion. Figure 7 compares the generators through the image completion task with 25% occlusion.

Image Completion with Iterations: Here we show the quality of the image generator as the training proceeds by evaluating the generated models on the image completion task. Figure 8 compares the generators through the image completion task with 25% occlusion.

Qualitative Results: We first show some example real and generated images (64×64) in Figure 9. The real images shown in this picture are used for the image completion task. Figure 10 shows qualitative image completion results with 25% occlusion. Figure 11 shows qualitative image completion results with 49% occlusion.

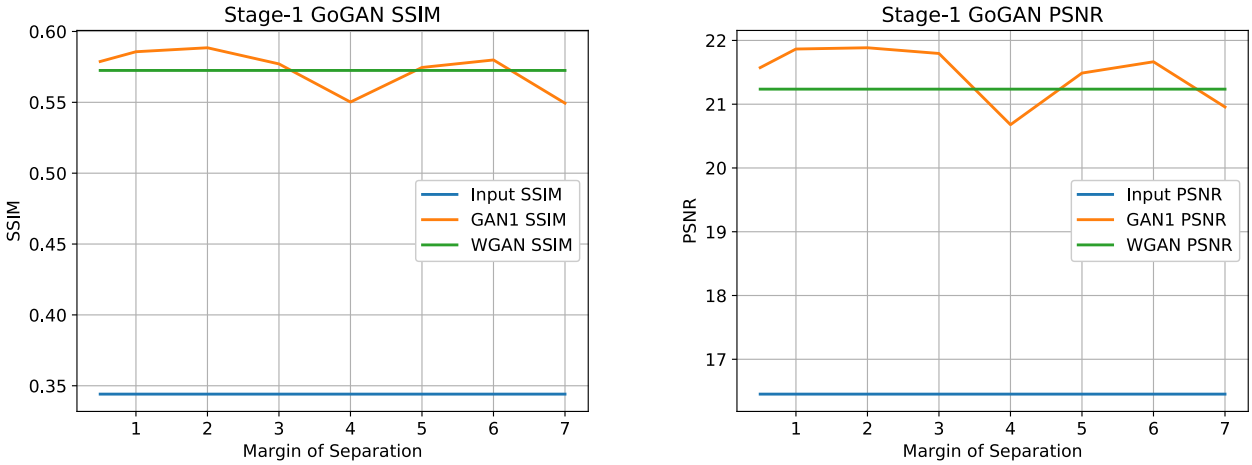

(a) SSIM for Image Completion with 49% Occlusion.

(b) PSNR for Image Completion with 49% Occlusion.

Figure 6: SSIM and PSNR on image completion with 49% occlusion using various margin of separation in GoGAN, benchmarked against WGAN.

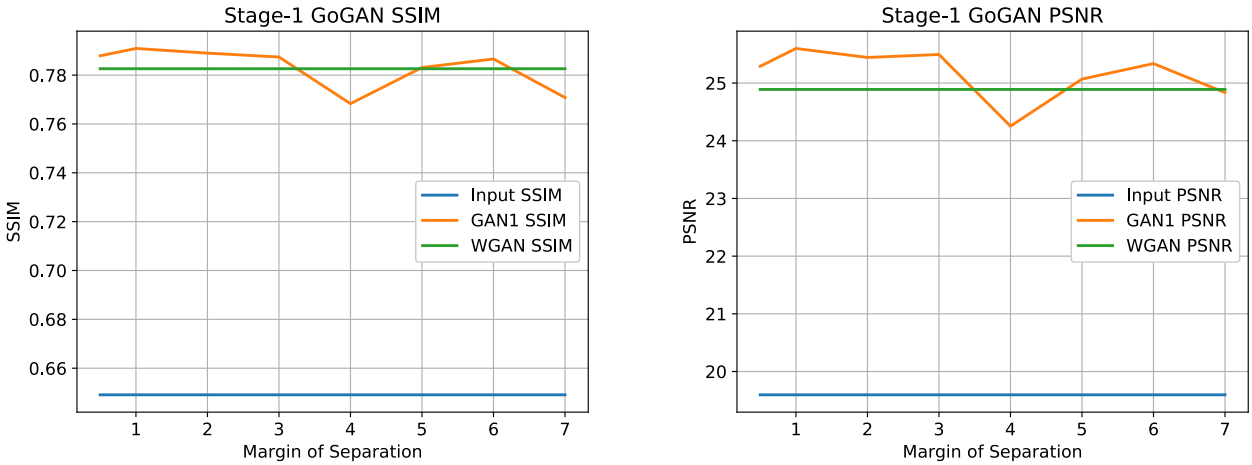

(a) SSIM for Image Completion with 25% Occlusion.

(b) PSNR for Image Completion with 25% Occlusion.

Figure 7: SSIM and PSNR on image completion with 25% occlusion using various margin of separation in GoGAN, benchmarked against WGAN.

Quantitative Results: We compare the quality of the image generators of WGAN, Stage-1 GoGAN and Stage-2 GoGAN through the image completion task. We measure the fidelity of the image completions via PSNR and SSIM. Table 5 shows results for our test set consisting of 5000 test faces, averaged over 10 runs, with 25% and 49% occlusions respectively.

Metric	SSIM		PSNR	
Occlusion	25%	49%	25%	49%
Occluded	0.6491	0.3441	19.5962	16.4536
WGAN	0.7826	0.5725	24.8892	21.2361
Stage-1 GoGAN	0.7908	0.5857	25.5998	21.8653
Stage-2 GoGAN	**0.7966**	**0.5963**	**25.7065**	**22.0040**

Table 5: Quantitative comparison of GAN models through image completion task.

6. Conclusions

In order to improve on the WGAN, we first generalize its discriminator loss to a margin-based one, which leads to a better discriminator, and in turn a better generator, and then carry out a progressive training paradigm involving multiple GANs to contribute to the maximum margin rank-

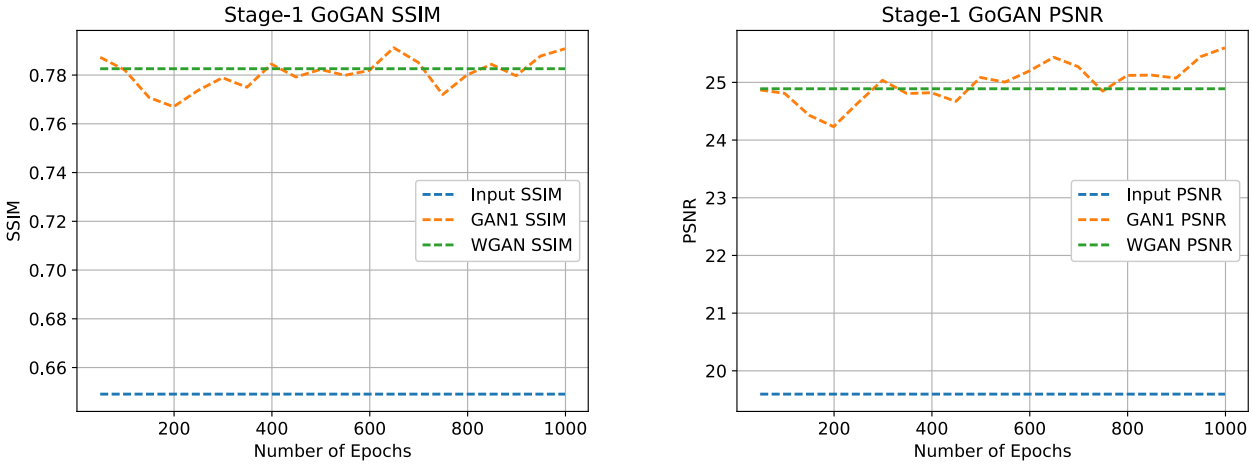

(a) SSIM for Image Completion with 25% Occlusion.

(b) PSNR for Image Completion with 25% Occlusion.

Figure 8: SSIM and PSNR on image completion with 25% occlusion using GoGAN as the training progresses, benchmarked against WGAN.

ing loss so that the GAN at later stages will improve upon early stages. We have shown theoretically that the proposed GoGAN can reduce the gap between the true data distribution and the generated data distribution by at least half in an optimally trained WGAN. We have also proposed a new way of measuring GAN quality which is based on image completion tasks. We have evaluated our method on four visual datasets: CelebA, LSUN Bedroom, CIFAR-10, and 50K-SSFF, and have seen both visual and quantitative improvement over baseline WGAN. Future work may include extending the GoGAN for other GAN variants and study how other divergence-based loss functions can benefit from the ranking loss and progressive training.

References

[1] B. Amos, L. Bartosz, and M. Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. Technical report, CMU-CS-16-118, CMU School of Computer Science, 2016. 8

[2] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. *ICLR (under review)*, 2017. 3, 4

[3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017. 3, 4, 5, 8

[4] A. Brock, T. Lim, J. Ritchie, and N. Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016. 2

[5] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, and P. Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *arXiv preprint arXiv:1606.03657*, 2016. 3

[6] E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015. 2

[7] J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016. 2

[8] V. Dumoulin, I. Belghazi, B. Poole, A. Lamb, M. Arjovsky, O. Mastropietro, and A. Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016. 2

[9] I. Durugkar, I. Gemp, and S. Mahadevan. Generative multi-adversarial networks. *arXiv preprint arXiv:1611.01673*, 2016. 2

[10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. 1, 2, 3, 4

[11] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015. 2

[12] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010. 10

[13] R. D. Hjelm, A. P. Jacob, T. Che, K. Cho, and Y. Bengio. Boundary-seeking generative adversarial networks. *arXiv preprint arXiv:1702.08431*, 2017. 3

[14] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie. Stacked generative adversarial networks. *arXiv preprint arXiv:1612.04357*, 2016. 3

[15] D. J. Im, C. D. Kim, H. Jiang, and R. Memisevic. Generating images with recurrent adversarial networks. *arXiv preprint arXiv:1602.05110*, 2016. 2

[16] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004*, 2016. 3

[17] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1

[18] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *CIFAR-10 Database*, 2009. 8

[19] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016. 3

[20] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 8, 10

[21] X. Mao, Q. Li, H. Xie, R. Y. Lau, and Z. Wang. Least squares generative adversarial networks. *arXiv preprint arXiv:1611.04076*, 2017. 3

[22] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2

[23] A. Nguyen, J. Yosinski, Y. Bengio, A. Dosovitskiy, and J. Clune. Plug & play generative networks: Conditional iterative generation of images in latent space. *arXiv preprint arXiv:1612.00005*, 2016. 3

[24] S. Nowozin, B. Cseke, and R. Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. *arXiv preprint arXiv:1606.00709*, 2016. 3

[25] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier gans. *arXiv preprint arXiv:1610.09585*, 2016. 2

[26] P. J. Phillips, P. J. Flynn, K. W. Bowyer, R. W. V. Bruegge, P. J. Grother, G. W. Quinn, and M. Pruitt. Distinguishing identical twins by face recognition. In *Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, pages 185–192. IEEE, 2011. 10

[27] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *Computer vision and pattern recognition, 2005. CVPR 2005. IEEE computer society conference on*, volume 1, pages 947–954. IEEE, 2005. 10

[28] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 2, 3, 8

[29] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2226–2234, 2016. 7

[30] Y. Taigman, A. Polyak, and L. Wolf. Unsupervised cross-domain image generation. *arXiv preprint arXiv:1611.02200*, 2016. 3

[31] T. Tieleman and G. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2), 2012. 8

[32] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016. 1

[33] A. van den Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu. Conditional image generation with pixelcnn decoders. *arXiv preprint arXiv:1606.05328*, 2016. 1

[34] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008. 4

[35] J. Yang, A. Kannan, B. Batra, and D. Parikh. Lr-gan - layered recursive generative adversarial networks for image generation. *ICLR (under review)*, 2017. 2

[36] R. Yeh, C. Chen, T. Y. Lim, M. Hasegawa-Johnson, and M. N. Do. Semantic image inpainting with perceptual and contextual losses. *arXiv preprint arXiv:1607.07539*, 2016. 7, 8

[37] F. Yu, Y. Zhang, S. Song, A. Seff, and J. Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 8

[38] J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016. 3

(a) Real Images.

(b) WGAN Generated Images.

(c) Stage-1 GoGAN Generated Images.

(d) Stage-2 GoGAN Generated Images.

Figure 9: Generated Images. Black bars are drawn on the real face images for privacy protection.

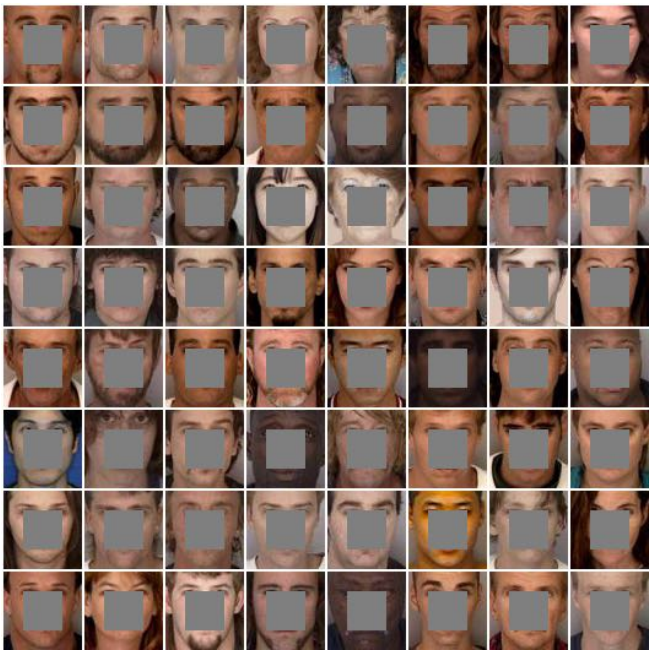

(a) Images with 25% Occlusion.

(b) WGAN Completion.

(c) Stage-1 GoGAN Completion.

(d) Stage-2 GoGAN Completion.

Figure 10: Image Completion with 25% Occlusion. Black bars are drawn on the completed face images for privacy protection.

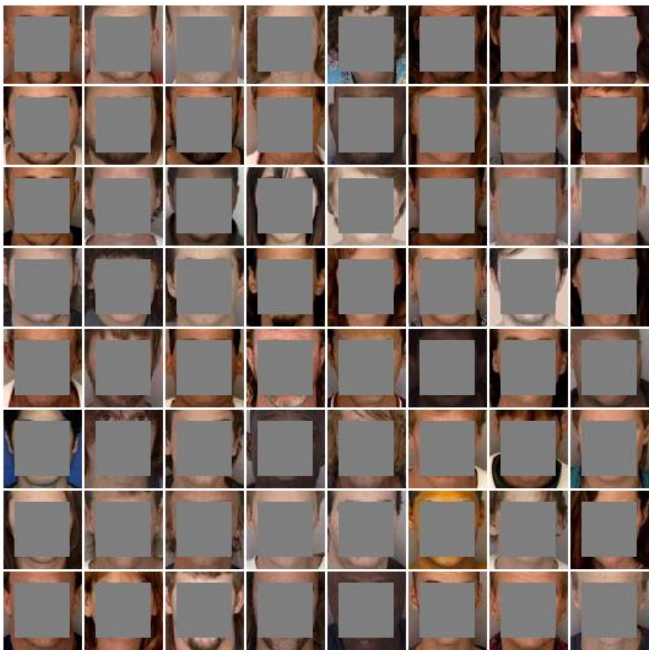

(a) Images with 49% Occlusion.

(b) WGAN Completion.

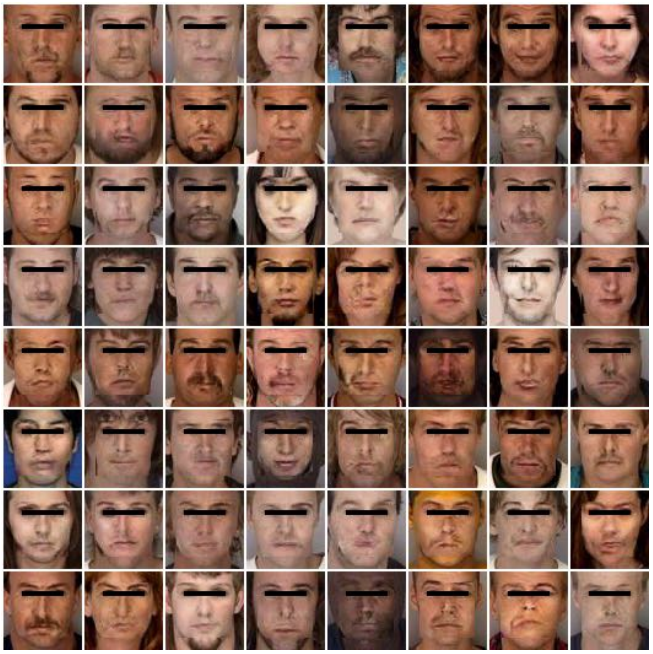

(c) Stage-1 GoGAN Completion.

(d) Stage-2 GoGAN Completion.

Figure 11: Image Completion with 49% Occlusion. Black bars are drawn on the completed face images for privacy protection.