Turning Waste into Wealth: Leveraging Low-Quality Samples for Enhancing Continuous Conditional Generative Adversarial Networks

Xin Ding[1], **Yongwei Wang**[2*], **Zuheng Xu**[3]

[1] School of Artificial Intelligence, Nanjing University of Information Science & Technology, Nanjing, China
[2] Shanghai Institute for Advanced Study, Zhejiang University, Shanghai, China
[3] Department of Statistics, University of British Columbia, Vancouver, BC, Canada
dingxin@nuist.edu.cn, yongwei.wang@zju.edu.cn, zuheng.xu@stat.ubc.ca

Abstract

Continuous Conditional Generative Adversarial Networks (CcGANs) enable generative modeling conditional on continuous scalar variables (termed regression labels). However, they can produce subpar fake images due to limited training data. Although Negative Data Augmentation (NDA) effectively enhances unconditional and class-conditional GANs by introducing anomalies into real training images, guiding the GANs away from low-quality outputs, its impact on CcGANs is limited, as it fails to replicate negative samples that may occur during the CcGAN sampling. We present a novel NDA approach called Dual-NDA specifically tailored for CcGANs to address this problem. Dual-NDA employs two types of negative samples: visually unrealistic images generated from a pre-trained CcGAN and label-inconsistent images created by manipulating real images' labels. Leveraging these negative samples, we introduce a novel discriminator objective alongside a modified CcGAN training algorithm. Empirical analysis on UTKFace and Steering Angle reveals that Dual-NDA consistently enhances the visual fidelity and label consistency of fake images generated by CcGANs, exhibiting a substantial performance gain over the vanilla NDA. Moreover, by applying Dual-NDA, CcGANs demonstrate a remarkable advancement beyond the capabilities of state-of-the-art conditional GANs and diffusion models, establishing a new pinnacle of performance. Our codes can be found at https://github.com/UBCDingXin/Dual-NDA.

Introduction

The objective of Continuous Conditional Generative Modeling (CCGM), as illustrated in Figure 2, is to estimate the distribution of high-dimensional data, such as images, in terms of continuous scalar variables, referred to as regression labels. However, this task is very challenging due to insufficient or even zero training images for certain regression labels and the absence of a suitable label input mechanism.

In a recent breakthrough, Ding et al. (2021, 2023b) introduced the pioneering model for this purpose, termed Continuous Conditional Generative Adversarial Networks (CcGANs), showcasing their superiority over conventional conditional GANs across various regression datasets. CcGANs have a wide spectrum of practical applications, including engineering inverse design (Heyrani Nobari, Chen, and Ahmed 2021; Fang, Shen, and Wang 2023), remote sensing image analysis (Giry-Fouquet et al. 2022), model compression (Ding et al. 2023a; Shi et al. 2023), point cloud generation (Triess et al. 2022), carbon sequestration (Stepien et al. 2023), data-driven solutions for poroelasticity (Kadeethum et al. 2022), etc. However, it's important to note that while CcGANs have shown success in these tasks, challenges remain when dealing with extremely sparse or imbalanced training data, leaving ample room to improve CcGAN models further.

To alleviate the data sparsity or imbalance issue, traditional data augmentation techniques for GANs (Zhao et al. 2020; Karras et al. 2020; Tran et al. 2021; Jiang et al. 2021; Tseng et al. 2021; Liu et al. 2023) often employ geometric transformations on real images, such as flipping, translation, and rotation, to guide GANs in learning "what to generate". However, Sinha et al. (2021) introduced a distinctive approach, known as Negative Data Augmentation (NDA), for unconditional or class-conditional GANs. The approach, depicted in Figure 3, intentionally crafts negative samples via transformations applied to real images, encompassing techniques like Jigsaw (Noroozi and Favaro 2016), Stitching, Mixup (Zhang et al. 2018), Cutout (DeVries and Taylor 2017), and CutMix (Yun et al. 2019). These negative samples, akin to those produced by the generator, are presented as fake images (in contrast to real images from the training set) and incorporated into the discriminator's training, instructing GANs on "what to avoid". Nevertheless, NDA's application is limited in the context of CcGANs, as it cannot replicate negative samples that may arise from pre-trained CcGANs, illustrated by two types of representative low-quality images in Figure 4.

To address this challenge, we introduce a novel NDA strategy termed Dual-NDA in this work, as depicted in Figure 1. Unlike the synthetic images showcased in Figure 3, Dual-NDA enriches the training set of CcGANs with two categories of negative samples, strategically mirroring the low-quality images seen in Figure 4. Type I negative samples are mismatched image-label pairs formed by manipulating labels of real images in the training set. In contrast, Type II negative samples are generated by evaluating fake images from a pre-trained generator and retaining those exhibiting the poorest visual quality. Based on these two types of negative samples, we devise a new soft vicinal discriminator loss tailored to

*Corresponding author.

enhance the training of CcGANs. Our comprehensive experimental investigation demonstrates the effectiveness of these negative samples in improving CcGAN performance, particularly in terms of visual quality and label consistency enhancement. Notably, CcGANs can generally achieve remarkable superiority through Dual-NDA over state-of-the-art class-conditional GANs and diffusion models.

Our major contributions can be summarized as follows: (1) We propose Dual-NDA, a novel NDA strategy tailored specifically for CcGANs. (2) We present a novel Dual-NDA-based framework for CcGAN training, incorporating a new vicinal discriminator loss. (3) Through extensive experiments, we substantiate the efficacy of Dual-NDA, demonstrating its positive impact on CcGANs' performance. (4) Our work extends beyond existing literature (Ding et al. 2021, 2023b) by providing a comprehensive comparative analysis, benchmarking CcGAN-based approaches against cutting-edge conditional generative models.

Related Work

Continuous Conditional GANs

Conditional Generative Adversarial Networks (cGANs) (cGANs), introduced by Mirza and Osindero (2014), extend the vanilla GAN (Goodfellow et al. 2014) to cater to the Conditional Generative Modeling (CGM) scenario. In this context, a condition denoted as y is provided as input to both the generator and discriminator networks. Mathematically, cGANs are designed to estimate the density function $p_r(\boldsymbol{x}|y)$ characterizing the underlying conditional data distribution, with the generator network tasked with approximating this distribution through an estimated density function denoted as $p_g(\boldsymbol{x}|y)$. While conventional cGANs, as explored in various works (Brock, Donahue, and Simonyan 2019; Zhang et al. 2019; Kang et al. 2021; Hou et al. 2022), mainly focus on scenarios involving discrete conditions like class labels or text descriptions, their applicability faces limitations when dealing with CGM tasks involving regression labels as conditions. Such limitations arise primarily due to the scarcity of available real images for certain regression labels and the challenge of encoding regression labels that may take infinitely many values.

To address these challenges, Ding et al. (2021, 2023b) propose the Continuous Conditional GAN (CcGAN) framework, which incorporates novel empirical cGAN losses and label input mechanisms. To mitigate data insufficiency, CcGANs leverage real images within a hard/soft vicinity of y to estimate $p_r(\boldsymbol{x}|y)$. This leads to the definition of the Hard/Soft Vicinal Discriminator Loss (HVDL/SVDL) and the generator loss for CcGANs as follows:

$$\begin{aligned}\widehat{\mathcal{L}}(D) = & \\ & -\frac{1}{N^r}\sum_{j=1}^{N^r}\sum_{i=1}^{N^r}\mathbb{E}_{\epsilon^r\sim\mathcal{N}(0,\sigma^2)}\left[W_1\log(D(\boldsymbol{x}_i^r, y_j^r+\epsilon^r))\right] \\ & -\frac{1}{N^g}\sum_{j=1}^{N^g}\sum_{i=1}^{N^g}\mathbb{E}_{\epsilon^g\sim\mathcal{N}(0,\sigma^2)}\left[W_2\log(1-D(\boldsymbol{x}_i^g, y_j^g+\epsilon^g))\right],\end{aligned} \tag{1}$$

$$\begin{aligned}\widehat{\mathcal{L}}(G) = & \\ & -\frac{1}{N^g}\sum_{i=1}^{N^g}\mathbb{E}_{\epsilon^g\sim\mathcal{N}(0,\sigma^2)}\log(D(G(\boldsymbol{z}_i, y_i^g+\epsilon^g), y_i^g+\epsilon^g)),\end{aligned} \tag{2}$$

where $D(\boldsymbol{x}, y)$ and $G(\boldsymbol{z}, y)$ represent the discriminator and generator networks, $\boldsymbol{x}_i^r$ and $\boldsymbol{x}_i^g$ are real and fake images, y_i^r and y_i^g are real and fake labels, $\boldsymbol{z}_i$ is Gaussian noise, N^r and N^g are sample sizes, σ is a hyperparameter controlling the variance of Gaussian noise ϵ, and the weights W_1 and W_2 are determined by the types of vicinity. For HVDL with the hard vicinity, W_1 and W_2 are defined as follows:

$$W_1 = \frac{\mathbb{1}_{\{|y_j^r+\epsilon-y_i^r|\le\kappa\}}}{N^r_{y_j^r+\epsilon,\kappa}}, \quad W_2 = \frac{\mathbb{1}_{\{|y_j^g+\epsilon-y_i^g|\le\kappa\}}}{N^g_{y_j^g+\epsilon,\kappa}}, \tag{3}$$

where $\mathbb{1}$ is an indicator function, $N_{y,\kappa}$ is the number of the y_i's satisfying $|y-y_i|\le\kappa$. For SVDL with the soft vicinity, W_1 and W_2 are defined as:

$$W_1 = \frac{w(y_i^r, y_j^r+\epsilon)}{\sum_{i=1}^{N^r} w(y_i^r, y_j^r+\epsilon)}, \quad W_2 = \frac{w(y_i^g, y_j^g+\epsilon)}{\sum_{i=1}^{N^g} w(y_i^g, y_j^g+\epsilon)}, \tag{4}$$

where $w(y, y') = e^{-\nu(y-y')^2}$. The hyperparameters κ, ν and σ in the above formulae can be determined using the rule of thumb in Ding et al. (2021, 2023b). It is worth noting that both Eq. (1) and Eq. (2) are derived by reformulating the vanilla cGAN loss (Mirza and Osindero 2014). Similarly, Ding et al. (2023b) also reformulated the GAN hinge loss (Lim and Ye 2017) for the CcGAN training. Furthermore, to address the challenge of encoding regression labels, Ding et al. (2021, 2023b) developed the Naive Label Input (NLI) and Improved Label Input (ILI) mechanisms. The effectiveness of CcGANs has been substantiated across diverse datasets.

Negative Data Augmentation

To enhance unconditional and class-conditional GANs, Sinha et al. (2021) introduced a novel Negative Data Augmentation (NDA) strategy that intentionally creates out-of-distribution negative samples by applying various transformations, such as Jigsaw, Stitching, and Mixup, to realistic training images. Furthermore, Sinha et al. (2021) introduced a new GAN training objective incorporating these negative samples. This training objective guides the generator network away from producing low-quality outputs resembling the negative samples, thereby encouraging the generator to create more realistic and diverse images. An illustrative workflow of NDA is provided in Figure 1. NDA has demonstrated its effectiveness in enhancing the performance of BigGAN (Brock, Donahue, and Simonyan 2019), a popular cGAN model, across both unconditional and class-conditional generative modeling tasks, highlighting NDA's potential for advancing GAN-driven image synthesis. However, as demonstrated in Figure 3 and Figure 4, NDA-generated negative samples markedly differ from the low-quality samples produced by CcGANs, which explains the limited impact of NDA in our experimental study.

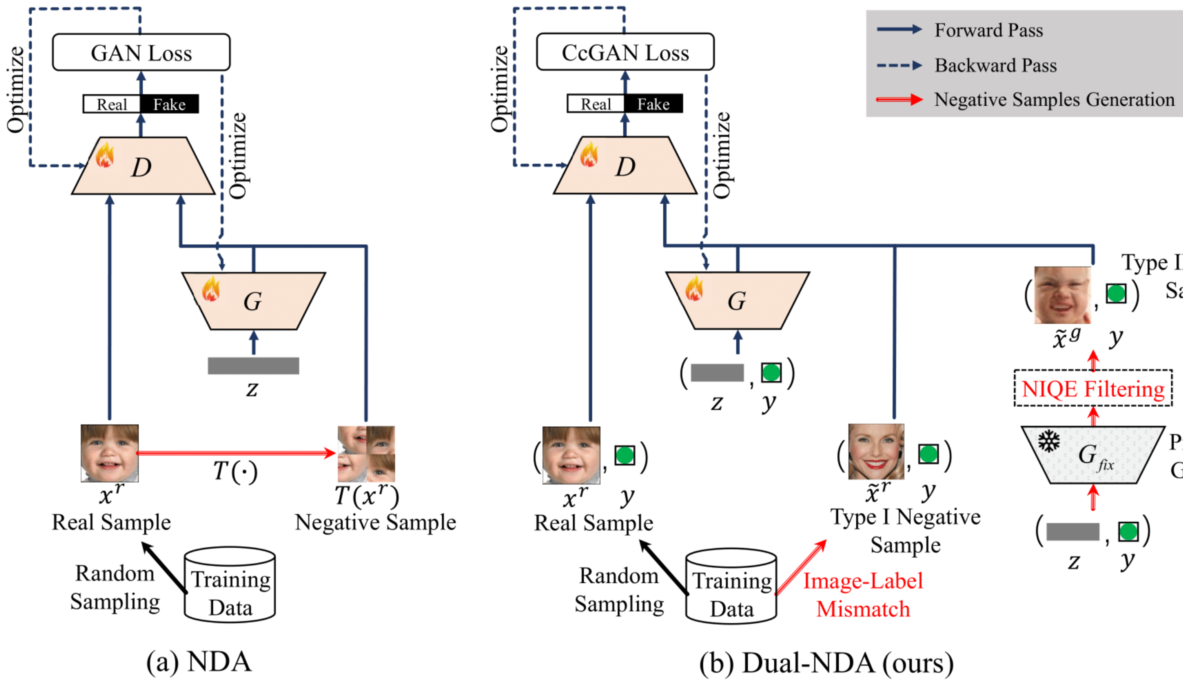

Figure 1: Illustrative workflows for the vanilla NDA (Sinha et al. 2021) and our proposed Dual-NDA.

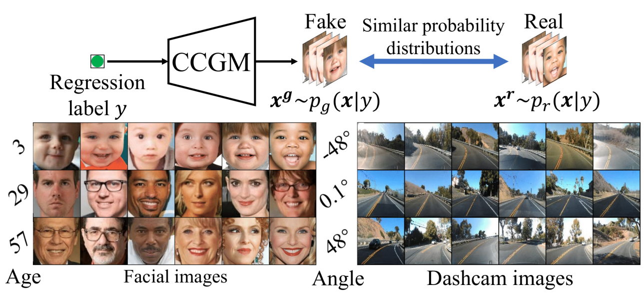

Figure 2: Illustration of the CCGM task and sample images from two regression datasets (UTKFace and Steering Angle).

Figure 3: Example negative samples from NDA by transforming a realistic training image (from Sinha et al. (2021)).

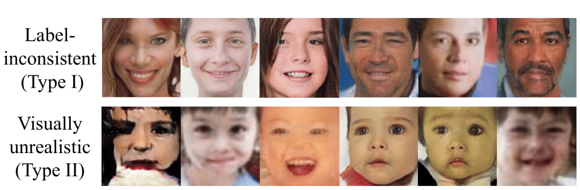

Figure 4: Some actual low-quality fake images generated from a pre-trained CcGAN at "Age=3" on UTKFace. The term "label-inconsistent" indicates that these fake images do not align with the conditioning label.

Consequently, we propose a novel NDA strategy named Dual-NDA, aimed at generating the two distinct types of negative samples illustrated in Figure 4.

Proposed Method: Dual-NDA

Overview

We present Dual-NDA, an innovative NDA strategy tailored to enhance the performance of CcGANs. The comprehensive workflow of Dual-NDA is depicted in Figure 1. This approach introduces a novel training paradigm for CcGANs, comprising two key components: a carefully designed negative sample generation mechanism adept at imitating the low-quality images shown in Figure 4, and a new vicinal discriminator loss that harnesses these negative samples to enhance the visual quality and label consistency of generated images.

Two Types of Negative Samples

Dual-NDA generates two types of negative samples, denoted as Type I and Type II, respectively. Some example negative samples are presented in Appendix C.

Type I: Label-Inconsistent Real Images

Type I negative samples comprise label-inconsistent real images generated through the dynamic mismatching of image-label pairs during the discriminator training. Recall the CcGAN training for $D(\boldsymbol{x}, y)$. When estimating the image distribution conditional on a selected label $y \in [0, 1]$, denoted as $p_r(\boldsymbol{x}|y)$, both real and fake images with labels falling within a hard or soft vicinity of y are selected to train the discriminator $D(\boldsymbol{x}, y)$. However, using a very wide vicinity could potentially lead to label-inconsistency issues. To solve this problem, Dual-NDA generates Type I negative samples labeled by y through the following steps:

- Compute the pairwise absolute distances between y and the labels of all real images in the training set, denoted as $d_i^y = |y - y_i^r|, i = 1, ..., N^r$. These absolute distance values form an array, denoted as $\boldsymbol{d}^y = [d_1^y, d_2^y, ..., d_{N^r}^y]$.
- Subsequently, calculate the q_1-th quantile of $\boldsymbol{d}^y$, denoted as $c_{q_1}^y$, with q_1 being a hyperparameter typically set in the range of 0.3 to 0.9.
- Choose training images with labels greater than $c_{q_1}^y$ to form the Type I negative samples for label y. It's important to note that the actual labels of these selected real images are significantly different from y.

Based on this mechanism, and considering the observed y's in the CcGAN training process, a set of all Type I negative samples can be constructed, represented as follows:

$$\begin{aligned} Q_{\mathrm{I}} =\{(\tilde{\boldsymbol{x}}_i^r, y_i^{\mathrm{I}}) | \tilde{\boldsymbol{x}}_i^r \sim p_r(\boldsymbol{x}|y_i^{\mathrm{I}}), y_i^{\mathrm{I}} \sim p(y), \\ |\tilde{y}_i^r - y_i^{\mathrm{I}}| > c_{q_1}^{y_i^{\mathrm{I}}}, i = 1, ..., N_{\mathrm{I}}^r\}, \end{aligned} \tag{5}$$

where $p_r(\boldsymbol{x}|y_i^{\mathrm{I}})$ denotes the density of real images' distribution conditional on y_i^{I}, $p(y)$ is the density of the labels' distribution, $\tilde{y}_i^r$ is the actual label of $\tilde{\boldsymbol{x}}_i^r$, and N_{I}^r is the sample size. Please note that this process for generating Type I negative samples is seamlessly integrated into the training algorithm of CcGANs, which is outlined in detail in Algorithm 3 provided in Appendix A.

Type II: Visually Unrealistic Fake Images

Type II negative samples are visually unrealistic fake images generated using a combination of a frozen CcGAN generator and a NIQE filtering mechanism. The generation process involves the following procedures:

- Dual-NDA initiates the process by sampling a large number of fake image-label pairs from the frozen CcGAN generator, denoted as $Q^g = \{(\boldsymbol{x}_i^g, y_i^g) | i = 1, ..., M\}$.
- Then, the NIQE filtering mechanism assesses the visual quality of fake images using the Naturalness Image Quality Evaluator (NIQE) (Mittal, Soundararajan, and Bovik 2012), where NIQE is a popular metric used to gauge the quality of an image based on its realism and natural appearance. The corresponding NIQE model has been pre-trained on the training set for CcGANs.
- In accordance with the NIQE metric, the filtering mechanism selects fake images that possess NIQE scores surpassing a predetermined threshold, effectively targeting images with the poorest visual quality.

The fake image-label pairs that are chosen through this process constitute the Type II negative samples, denoted as

$$\begin{aligned} Q_{\mathrm{II}} = \{(\tilde{\boldsymbol{x}}_i^g, \tilde{y}_i^g) | \mathrm{NIQE}(\tilde{\boldsymbol{x}}_i^g) > c_{q_2}, (\tilde{\boldsymbol{x}}_i^g, \tilde{y}_i^g) \in Q^g, \\ i = 1, ..., N_{\mathrm{II}}^g, N_{\mathrm{II}}^g < M\}, \end{aligned} \tag{6}$$

where c_{q_2} is a threshold determined by a hyperparameter q_2 and may have a connection with $\tilde{y}_i^g$. Dual-NDA typically works well if q_2 locates in $[0.5, 0.9]$. It's important to note that we adopt two distinct strategies to determine c_{q_2}:

- For labels with integer-valued observations (e.g., ages or counting numbers), a separate c_{q_2} value is computed for each distinct integer value of the regression label. Specifically, assuming the regression label of interest has K distinct integer values, denoted as $\{1, ..., k..., K\}$, we denote the corresponding fake images labeled by k as $\boldsymbol{X}_k^g = \{\boldsymbol{x}_1^{g,k}, ..., \boldsymbol{x}_{M_k}^{g,k}\}$. Then, we compute the NIQE scores of images in $\boldsymbol{X}_k$, and store the scores in an array, denoted as $\boldsymbol{s}^k = [s_1^k, s_2^k, ..., s_{M_k}^k]$. Subsequently, we calculate the q_2-th quantile of $\boldsymbol{s}^k$, denoted as $c_{q_2}^k$. Finally, we choose fake images in $\boldsymbol{X}_k$ with NIQE scores larger than $c_{q_2}^k$. The above procedures are executed for each of the K distinct integer values and summarized in Algorithm 1 in Appendix A.
- For strictly continuous labels (e.g., angles), a single global threshold c_{q_2} is chosen instead of using K separate values. In this case, c_{q_2} is the q_2-th quantile of the NIQE scores of all fake images in Q^g, and fake images with NIQE scores larger than c_{q_2} are kept as Type II negative samples. We summarize this process in Algorithm 2 in Appendix A.

A Modified Training Mechanism for CcGANs

With the Type I and Type II negative samples separated respectively into Q_{I} and Q_{II}, we introduce a modified training mechanism for CcGANs, illustrated in-depth in Algorithm 3 in Appendix A. The core of this training mechanism is an innovative vicinal discriminator loss, delineated as follows:

$$\begin{aligned} &\widetilde{\mathcal{L}}(D) = \\ &-\frac{1}{N^r}\sum_{j=1}^{N^r}\sum_{i=1}^{N^r}\mathbb{E}_{\epsilon\sim\mathcal{N}(0,\sigma^2)}\left[W_1\log(D(\boldsymbol{x}_i^r, y_j^r+\epsilon))\right] \\ &-\frac{1-\bar{\lambda}}{N^g}\sum_{j=1}^{N^g}\sum_{i=1}^{N^g}\mathbb{E}_{\epsilon\sim\mathcal{N}(0,\sigma^2)}\left[W_2\log(1-D(\boldsymbol{x}_i^g, y_j^g+\epsilon))\right] \\ &-\frac{\lambda_1}{N_{\mathrm{I}}}\sum_{i=1}^{N_{\mathrm{I}}}\log(1-D(\tilde{\boldsymbol{x}}_i^r, y_i^{\mathrm{I}}))) \\ &-\frac{\lambda_2}{N_{\mathrm{II}}}\sum_{j=1}^{N_{\mathrm{II}}}\sum_{i=1}^{N_{\mathrm{II}}}\mathbb{E}_{\epsilon\sim\mathcal{N}(0,\sigma^2)}\left[W_3\log\left(1-D\left(\tilde{\boldsymbol{x}}_i^g, \tilde{y}_j^g+\epsilon\right)\right)\right], \end{aligned} \tag{7}$$

where λ_1 and λ_2 are two hyperparameters with values in $[0, 1]$, $\bar{\lambda} = \lambda_1 + \lambda_2$, $(\tilde{\boldsymbol{x}}_i^r, y_i^{\mathrm{I}}) \in Q_{\mathrm{I}}$, $(\tilde{\boldsymbol{x}}_i^g, \tilde{y}_i^g) \in Q_{\mathrm{II}}$. While the weights W_1 and W_2 have been defined in Eq. (4), W_3 aligns with Eq. (3), taking the form:

$$W_3 = \frac{\mathbb{1}_{\{|\tilde{y}_j^g+\epsilon-\tilde{y}_i^g|\leq\kappa\}}}{N^g_{\tilde{y}_j^g+\epsilon,\kappa}}. \tag{8}$$

Within Eq. (7), the first two terms stem from the SVDL presented in Eq. (1). The third and fourth terms are grounded in the Type I and Type II negative samples, respectively. The hyperparameters λ_1 and λ_2 control the influence of negative samples on the CcGAN training. Such an effect has been carefully studied in our experiment. Furthermore, the hyperparameters κ, ν, and σ in Eq. (7) are determined using the empirically validated guidelines in literature (Ding et al. 2023b). It is important to note that, in the new training mechanism, the generator loss is consistent with Eq. (2). Furthermore, an alternative version of Eq. (7) based on hinge loss is also provided in Appendix A.

Experiments

Experimental Setup

Datasets. Following Ding et al. (2023b), we utilize two regression datasets, namely UTKFace (Zhang, Song, and Qi 2017) and Steering Angle (Chen 2018), for our experimental study. The UTKFace dataset includes human facial images annotated with ages for regression analysis. We use the preprocessed version (Ding et al. 2023b), which consists of 14,723 RGB images. Age labels span from 1 to 60, showcasing a broad age distribution with varying sample sizes per age group (ranging from 50 to 1051 images). The Steering Angle dataset, derived from a subset of an autonomous driving dataset (Chen 2018), comprises 12,271 RGB images. These images are annotated with 1,774 distinct steering angles, spanning from $-80.0°$ to $80.0°$. All Steering Angle images are captured using a dashboard camera on a car, with simultaneous recording of the corresponding steering wheel rotation angles. Both datasets offer two versions, encompassing resolutions of 64×64 and 128×128, respectively, resulting in a total of four datasets. Illustrative training images are visualized in Figure 2, while data distributions are presented in Appendix B.

Baseline Methods. To substantiate the effectiveness of Dual-NDA, we choose the following state-of-the-art conditional generative models as baselines for comparison: (1) Two class-conditional GANs are chosen in the comparison, including ReACGAN (Kang et al. 2021) and ADCGAN (Hou et al. 2022). (2) Two cutting-edge class-conditional diffusion models, including classifier guidance (ADM-G) (Dhariwal and Nichol 2021) and classifier-free guidance (CFG) (Ho and Salimans 2021) models, are also included in the comparison. (3) Three CcGAN-based methods are compared, consisting of CcGAN (SVDL+ILI) w/o NDA (Ding et al. 2023b), CcGAN w/ NDA (Sinha et al. 2021), and the proposed Dual-NDA.

Training Setup. In the implementation of class-conditional GANs and diffusion models, we bin the regression labels of UTKFace into 60 classes and those of Steering Angle into 221 classes. For low-resolution experiments, we re-implement the baseline CcGANs, while for high-resolution experiments, we utilize the checkpoints provided by Ding et al. (2023b). To implement NDA, we set the hyperparameter λ to 0.25, adhering to the setup in Sinha et al. (2021). *For Dual-NDA, we utilize the pre-trained generator of "CcGAN w/o NDA" in conjunction with the NIQE filtering process to generate Type II negative samples.* We successfully generate 60,000 and 17,740 Type II negative samples, respectively, in the UTKFace and Steering Angle experiments. Regarding Dual-NDA's hyperparameters, in UTKFace experiments, λ_1 is set to 0.05 and λ_2 to 0.15. In Steering Angle experiments at 64×64 resolution, λ_1 is set to 0.1 and λ_2 to 0.2, while at 128×128 resolution, λ_1 is set to 0.2 and λ_2 to 0.3. Additionally, we set q_1 to 0.9 for UTKFace and 0.5 for Steering Angle, with q_2 consistently set to 0.9. More detailed training setups can be found in Appendix B.

Evaluation Setup. Following Ding et al. (2023b), we generate 60,000 fake images for the UTKFace experiments and 100,000 fake images for the Steering Angle experiments from each compared method. These generated images are subject to evaluation using both an overall metric and three separate metrics. Specifically, the Sliding Fréchet Inception Distance (SFID) (Ding et al. 2023b) is taken as the overall metric. Notably, the radius utilized for SFID computations is set to 0 for the UTKFace experiments and 2 for the Steering Angle experiments. Furthermore, NIQE (Mittal, Soundararajan, and Bovik 2012), Diversity (Ding et al. 2023b), and Label Score (Ding et al. 2023b) metrics are employed to respectively gauge visual fidelity, diversity, and label consistency of the generated images.

Experimental Results

We present a performance comparison of various methods on the four datasets in Table 1, complemented by illustrative figures presented in Figure 5 and Figure 6. Analysis of the table and figures reveal the following findings:

- Among the assessed methods, Dual-NDA demonstrates superior performance across all four datasets in terms of SFID and NIQE metrics. Its distinct advantage becomes more pronounced on the Steering Angle datasets, thereby underscoring its effectiveness in tackling more intricate CCGM scenarios.
- Compared with the baseline "CcGAN w/o NDA", Dual-NDA exhibits reduced values in terms of NIQE and Label Score metrics. This trend is further supported by Figure 5 and Figure 6, which illustrate a substantial reduction in NIQE and Label Score values across the entire range of regression labels. These outcomes suggest that incorporating Type I and II negative samples effectively enhances the visual quality and label consistency of the generated samples.
- NDA consistently fails to yield desirable effects, and in all cases, it leads to a decline in the performance of CcGANs.
- Class-conditional GANs and diffusion models often exhibit limited effectiveness. Particularly when applied to the Steering Angle datasets, most of them suffer from the mode collapse problem. This outcome once more underscores the prevailing superiority of CcGANs over class-conditional generative models.
- A crucial observation to highlight is that certain methods, like ADCGAN, might exhibit higher Diversity scores than Dual-NDA. However, it's important to note that many of their diverse generated images might be label-inconsistent, consequently leading to high Label Score values.

Ablation Study

We also conduct comprehensive ablation studies on the Steering Angle (64×64) dataset. These studies are designed to systematically investigate the impact of individual components and key hyperparameters within the framework of Dual-NDA, as outlined below. The standard deviations for the evaluation scores are provided in Appendix C.

Dataset	Type	Method	SFID ↓	NIQE ↓	Diversity ↑	Label Score ↓
UTKFace (64×64)	Class-conditional (60 classes)	ReACGAN (Neurips'21)	0.548 (0.136)	1.679 (0.170)	1.206 (0.240)	6.846 (5.954)
		ADCGAN (ICML'22)	0.573 (0.218)	1.680 (0.140)	**1.451 (0.019)**	17.574 (12.388)
		ADM-G (Neurips'21)	0.744 (0.195)	2.856 (0.225)	0.917 (0.318)	7.583 (6.066)
		CFG (Neurips'21)	2.155 (0.638)	1.681 (0.303)	0.858 (0.413)	8.477 (7.820)
	CcGAN (SVDL+ILI)	w/o NDA (T-PAMI'23)	0.413 (0.155)	1.733 (0.189)	1.329 (0.161)	8.240 (6.271)
		w/ NDA (ICLR'21)	0.491 (0.157)	1.757 (0.183)	1.399 (0.130)	8.229 (6.713)
		Dual-NDA (ours)	**0.396 (0.153)**	**1.678 (0.183)**	1.298 (0.187)	**6.765 (5.600)**
UTKFace (128×128)	Class-conditional (60 classes)	ReACGAN (Neurips'21)	0.445 (0.098)	1.426 (0.064)	1.152 (0.304)	**6.005** (5.182)
		ADCGAN (ICML'22)	0.468 (0.143)	1.231 (0.048)	**1.365 (0.035)**	15.777 (11.572)
		ADM-G (Neurips'21)	0.997 (0.208)	3.705 (0.409)	0.831 (0.271)	11.618 (8.754)
		CFG (Neurips'21)	1.521 (0.333)	1.888 (0.263)	1.170 (0.174)	11.430 (9.917)
	CcGAN (SVDL+ILI)	w/o NDA (T-PAMI'23)	0.367 (0.123)	1.113 (0.033)	1.199 (0.232)	7.747 (6.580)
		w/ NDA (ICLR'21)	1.136 (0.244)	1.125 (0.049)	0.986 (0.471)	6.384 (5.324)
		Dual-NDA (ours)	**0.361 (0.127)**	**1.081 (0.042)**	1.257 (0.238)	6.310 (5.194)
Steering Angle (64×64)	Class-conditional (221 classes)	ReACGAN (Neurips'21)	3.635 (0.491)	2.099 (0.072)	0.543 (0.366)	27.277 (21.508)
		ADCGAN (ICML'22)	2.960 (1.083)	2.015 (0.003)	0.930 (0.018)	40.535 (24.031)
		ADM-G (Neurips'21)	2.890 (0.547)	2.164 (0.200)	0.205 (0.160)	24.186 (20.685)
		CFG (Neurips'21)	4.703 (0.894)	2.070 (0.022)	0.923 (0.119)	56.663 (39.914)
	CcGAN (SVDL+ILI)	w/o NDA (T-PAMI'23)	1.334 (0.531)	1.784 (0.065)	1.234 (0.209)	14.807 (14.297)
		w/ NDA (ICLR'21)	1.381 (0.527)	1.994 (0.081)	1.231 (0.167)	**10.717 (10.371)**
		Dual-NDA (ours)	**1.114 (0.503)**	**1.738 (0.055)**	**1.251 (0.172)**	11.809 (11.694)
Steering Angle (128×128)	Class-conditional (221 classes)	ReACGAN (Neurips'21)	3.979 (0.919)	4.060 (0.643)	0.250 (0.269)	36.631 (38.592)
		ADCGAN (ICML'22)	3.110 (0.799)	5.181 (0.010)	0.001 (0.001)	44.242 (29.223)
		ADM-G (Neurips'21)	1.593 (0.449)	3.476 (0.153)	1.120 (0.121)	32.040 (27.836)
		CFG (Neurips'21)	5.425 (1.573)	2.742 (0.109)	0.762 (0.121)	50.015 (34.640)
	CcGAN (SVDL+ILI)	w/o NDA (T-PAMI'23)	1.689 (0.443)	2.411 (0.100)	1.088 (0.243)	18.438 (16.072)
		w/ NDA (ICLR'21)	1.736 (0.562)	2.435 (0.160)	1.022 (0.247)	**12.438 (11.612)**
		Dual-NDA (ours)	**1.390 (0.421)**	**2.135 (0.065)**	**1.133 (0.217)**	14.099 (12.097)

Table 1: Average quality of fake images from compared methods with standard deviations in the parentheses. "↓" ("↑") indicates lower (higher) values are preferred. The best and second-best results are marked respectively in boldface and underlined. We re-implement all compared methods except for "CcGAN w/o NDA" in the 128 × 128 experiments.

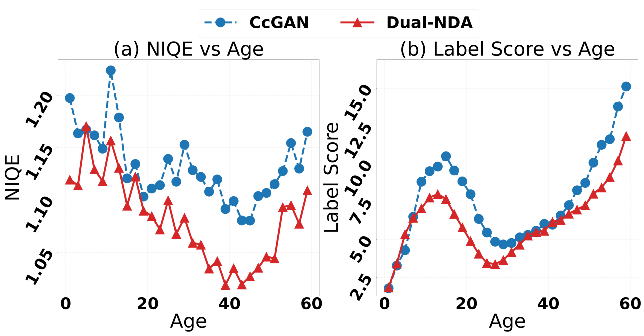

Figure 5: Line graphs of NIQE/Label Score versus Age for the UTKFace (128×128) experiment.

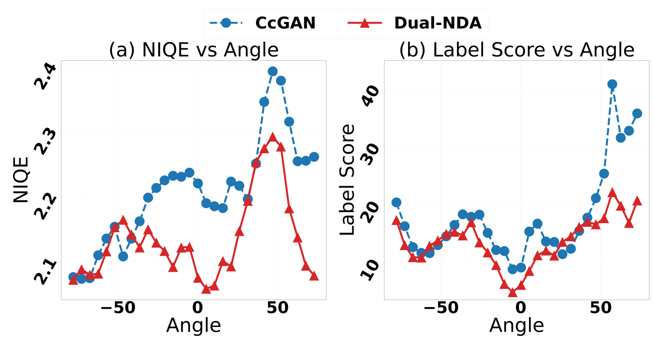

Figure 6: Line graphs of NIQE/Label Score versus Angle for the Steering Angle (128×128) experiment.

- **The Effects of Type I and II Negative Samples**: We individually delve into the distinct impacts of Type I and Type II negative samples. The experimental results are summarized in Figure 7. From this visual result, we can deduce that the integration of Type I negative samples during training effectively enhances the label consistency of generated images (evidenced by decreased Label Score values), yet it might potentially lead to a compromise in visual quality (reflected by higher NIQE values relative to the blue dashed line). Interestingly, the influence of Type II negative samples exhibits a unique pattern. These samples, designed to enhance visual quality, surprisingly contribute to an improvement in label consistency (as reflected by a Label Score value below the red dashed line). This phenomenon can be rationalized by the insights provided by Ding et al. (2023a), which suggest that certain label-inconsistent fake images often exhibit poor visual quality. Consequently, the removal of visually unrealistic images may lead to an enhancement in label consistency. Combining both Type I and Type II negative samples ultimately leads to the most favorable overall performance.

- **The Combined Impact of λ_1 and λ_2**: Instead of examining the individual effects of λ_1 and λ_2, we delve into their combined impact on the SFID scores of CcGANs. We achieve this by varying $\lambda_1 + \lambda_2$ within the range of 0.2 to 0.7. The experimental outcomes, shown in Figure 8, highlight a consistent trend: Dual-NDA consistently outperforms CcGAN without NDA across a wide spectrum of $\lambda_1 + \lambda_2$ values in terms of SFID. Notably, adjustments to $\lambda_1 + \lambda_2$ within the interval of $[0.2, 0.5]$ often yield optimal performance for CcGANs.
- **The Effect of q_1**: We also experiment to analyze the impact of q_1 on Dual-NDA's performance. We set a grid of values for q_1, and the corresponding quantitative result is presented in Table 2. These findings show that Dual-NDA's performance is not significantly affected by variations in q_1. Empirically, q_1 with values such as 0.5 or 0.9 often produces favorable outcomes.
- **The Effect of q_2**: In the last ablation study, we examine the influence of q_2 on Dual-NDA's performance and report the quantitative results in Table 3. These results reveal that Dual-NDA is insensitive to the values of q_2, so we let $q_2 = 0.9$ throughout our experiments.
- **Creating Type II negative samples from other generators**: As indicated in Table 4, we incorporate generators from ADCGAN and ReACGAN to generate Type II negative samples. Nonetheless, it is noteworthy that both ADCGAN and ReACGAN exhibit an adverse impact on the NIQE values, indicating a decrease in visual quality.

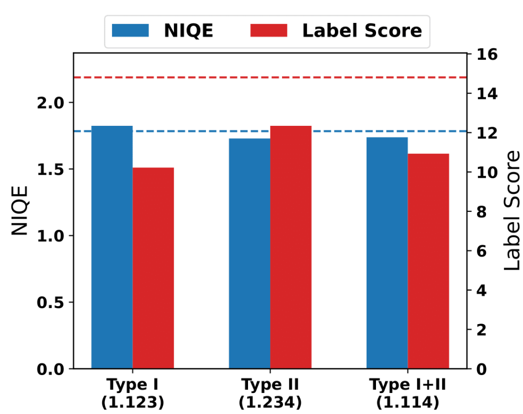

Figure 7: The individual impacts of Type I and Type II negative samples on Dual-NDA in the Steering Angle (64×64) experiment. The SFID scores for compared settings are shown in parentheses. The blue and red dashed lines represent the NIQE and Label Score of "CcGAN w/o NDA", respectively.

Conclusion

In this paper, we present an innovative NDA approach called Dual-NDA, aimed at enhancing the performance of CcGANs in the realm of continuous conditional generative modeling. Our approach involves two strategies for generating negative samples, simulating two categories of low-quality images that may arise during CcGAN sampling. Furthermore, we introduce a modified CcGAN training mechanism that relies on these negative samples to steer the model away from generating undesirable outputs. Through comprehensive experimentation, we demonstrate that Dual-NDA effectively improves the performance of CcGANs, surpassing widely-used class-conditional GANs and diffusion models.

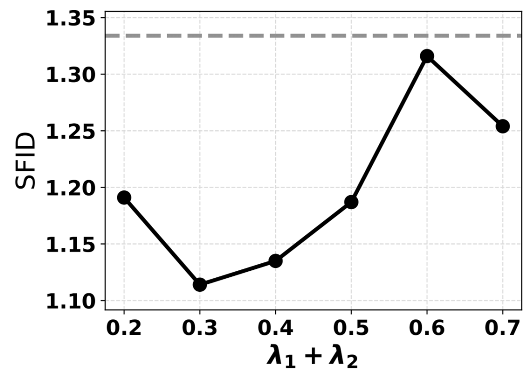

Figure 8: The effect of $\lambda_1 + \lambda_2$ in the Steering Angle (64×64) experiment. The gray dashed line represents "w/o NDA".

q_1	SFID	NIQE	Diversity	Label Score
w/o NDA	1.334	1.784	1.234	14.807
0.3	1.125	1.759	1.271	11.365
0.5	1.114	1.738	1.251	11.809
0.7	1.161	1.763	1.227	11.162
0.9	1.187	1.763	1.270	11.550

Table 2: The impact of parameter q_1 on the Dual-NDA performance in the Steering Angle (64x64) experiment. "CcGAN w/o NDA" is included in the table as a baseline for reference.

q_2	SFID	NIQE	Diversity	Label Score
w/o NDA	1.334	1.784	1.234	14.807
0.5	1.107	1.717	1.301	11.881
0.6	1.193	1.733	1.278	11.459
0.7	1.150	1.763	1.278	13.161
0.8	1.174	1.756	1.254	11.207
0.9	1.114	1.738	1.251	11.809

Table 3: The impact of parameter q_2 on the Dual-NDA performance in the Steering Angle (64x64) experiment. "CcGAN w/o NDA" is included in the table as a baseline for reference.

Generator	SFID	NIQE	Diversity	Label Score
None	1.334	1.784	1.234	14.807
ADCGAN	1.212	**1.832**	1.222	10.884
ReACGAN	1.203	**1.800**	1.273	11.486
CcGAN	1.114	1.738	1.251	11.809

Table 4: Type II negative samples from different generators.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62306147) and Pujiang Talent Program (Grant No. 23PJ1412100).

References

Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*.

Chen, S. 2018. The Steering Angle Dataset @ONLINE. https://github.com/SullyChen/driving-datasets. Accessed: 2020-12-01.

DeVries, T.; and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.

Dhariwal, P.; and Nichol, A. 2021. Diffusion Models Beat GANs on Image Synthesis. *Advances in Neural Information Processing Systems*, 34: 8780–8794.

Ding, X.; Wang, Y.; Xu, Z.; Wang, Z. J.; and Welch, W. J. 2023a. Distilling and Transferring Knowledge via cGAN-Generated Samples for Image Classification and Regression. *Expert Systems with Applications*, 213: 119060.

Ding, X.; Wang, Y.; Xu, Z.; Welch, W. J.; and Wang, Z. J. 2021. CcGAN: Continuous Conditional Generative Adversarial Networks for Image Generation. In *International Conference on Learning Representations*.

Ding, X.; Wang, Y.; Xu, Z.; Welch, W. J.; and Wang, Z. J. 2023b. Continuous Conditional Generative Adversarial Networks: Novel Empirical Losses and Label Input Mechanisms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7): 8143–8158.

Dufumier, B.; Gori, P.; Victor, J.; Grigis, A.; Wessa, M.; Brambilla, P.; Favre, P.; Polosan, M.; Mcdonald, C.; Piguet, C. M.; et al. 2021. Contrastive Learning with Continuous Proxy Meta-Data for 3D MRI Classification. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, 58–68.

Fang, X.; Shen, H.-S.; and Wang, H. 2023. Diverse 3D Auxetic Unit Cell Inverse Design with Deep Learning. *Applied Physics Reviews*, 10(3).

Giry-Fouquet, Y.; Baussard, A.; Enderli, C.; and Porges, T. 2022. SAR Image Synthesis with GAN and Continuous Aspect Angle and Class Constraints. In *European Conference on Synthetic Aperture Radar*, 1–6.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27*, 2672–2680.

Heyrani Nobari, A.; Chen, W.; and Ahmed, F. 2021. PcDGAN: A Continuous Conditional Diverse Generative Adversarial Network for Inverse Design. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 606–616.

Ho, J.; and Salimans, T. 2021. Classifier-Free Diffusion Guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.

Hou, L.; Cao, Q.; Shen, H.; Pan, S.; Li, X.; and Cheng, X. 2022. Conditional GANS with Auxiliary Discriminative Classifier. In *International Conference on Machine Learning*, 8888–8902.

Jiang, L.; Dai, B.; Wu, W.; and Loy, C. C. 2021. Deceive D: Adaptive pseudo augmentation for gan training with limited data. *Advances in Neural Information Processing Systems*, 34: 21655–21667.

Kadeethum, T.; O'Malley, D.; Choi, Y.; Viswanathan, H. S.; Bouklas, N.; and Yoon, H. 2022. Continuous Conditional Generative Adversarial Networks for Data-Driven Solutions of Poroelasticity with Heterogeneous Material Properties. *Computers & Geosciences*, 167: 105212.

Kang, M.; Shim, W.; Cho, M.; and Park, J. 2021. Rebooting ACGAN: Auxiliary Classifier GANs with Stable Training. *Advances in neural information processing systems*, 34: 23505–23518.

Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33: 12104–12114.

Lim, J. H.; and Ye, J. C. 2017. Geometric GAN. *arXiv preprint arXiv:1705.02894*.

Liu, H.; Li, B.; Wu, H.; Liang, H.; Huang, Y.; Li, Y.; Ghanem, B.; and Zheng, Y. 2023. Combating Mode Collapse via Offline Manifold Entropy Estimation. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'23/IAAI'23/EAAI'23. AAAI Press. ISBN 978-1-57735-880-0.

Mirza, M.; and Osindero, S. 2014. Conditional Generative Adversarial Nets. *arXiv preprint arXiv:1411.1784*.

Mittal, A.; Soundararajan, R.; and Bovik, A. C. 2012. Making a "Completely Blind" Image Quality Analyzer. *IEEE Signal Processing Letters*, 20(3): 209–212.

Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, 69–84.

Shi, Z.; Ding, X.; Ding, P.; Yang, C.; Huang, R.; and Song, X. 2023. Regression-Oriented Knowledge Distillation for Lightweight Ship Orientation Angle Prediction with Optical Remote Sensing Images. *arXiv preprint arXiv:2307.06566*.

Sinha, A.; Ayush, K.; Song, J.; Uzkent, B.; Jin, H.; and Ermon, S. 2021. Negative Data Augmentation. In *International Conference on Learning Representations*.

Stepien, M.; Ferreira, C. A.; Hosseinzadehsadati, S.; Kadeethum, T.; and Nick, H. M. 2023. Continuous Conditional Generative Adversarial Networks for Data-Driven Modelling of Geologic CO2 Storage and Plume Evolution. *Gas Science and Engineering*, 115: 204982.

Tran, N.-T.; Tran, V.-H.; Nguyen, N.-B.; Nguyen, T.-K.; and Cheung, N.-M. 2021. On Data Augmentation for GAN Training. *IEEE Transactions on Image Processing*, 30: 1882–1897.

Triess, L. T.; Bühler, A.; Peter, D.; Flohr, F. B.; and Zöllner, M. 2022. Point Cloud Generation with Continuous Conditioning. In *International Conference on Artificial Intelligence and Statistics*, 4462–4481.

Tseng, H.-Y.; Jiang, L.; Liu, C.; Yang, M.-H.; and Yang, W. 2021. Regularizing generative adversarial networks under limited data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7921–7931.

Yun, S.; Han, D.; Oh, S. J.; Chun, S.; Choe, J.; and Yoo, Y. 2019. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6023–6032.

Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*.

Zhang, H.; Goodfellow, I.; Metaxas, D.; and Odena, A. 2019. Self-Attention Generative Adversarial Networks. In *International Conference on Machine Learning*, 7354–7363.

Zhang, Z.; Song, Y.; and Qi, H. 2017. Age Progression/Regression by Conditional Adversarial Autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5810–5818.

Zhao, S.; Liu, Z.; Lin, J.; Zhu, J.-Y.; and Han, S. 2020. Differentiable Augmentation for Data-Efficient GAN Training. *Advances in Neural Information Processing Systems*, 33: 7559–7570.