DIFFERENTIAL-CRITIC GAN: GENERATING WHAT YOU WANT BY A CUE OF PREFERENCES

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper proposes Differential-Critic Generative Adversarial Network (DiCGAN) to learn the distribution of user-desired data when only partial instead of the entire dataset possesses the desired properties. Existing approaches select the desired samples first and train regular GANs on the selected samples to derive the user-desired data distribution. DiCGAN introduces a differential critic that can learn the preference direction from the pairwise preferences over the entire dataset. The resultant critic would guide the generation of the desired data instead of the whole data. Specifically, apart from the Wasserstein GAN loss, a ranking loss of the pairwise preferences is defined over the critic. It endows the difference of critic values between each pair of samples with the pairwise preference relation. The higher critic value indicates that the sample is preferred by the user. Thus training the generative model for higher critic values would encourage generating the user-preferred samples. Extensive experiments show that our DiCGAN can learn the user-desired data distributions.

1 INTRODUCTION

Learning a good generative model for high-dimensional natural signals, such as images (Zhu et al., 2017), video (Vondrick et al., 2016) and text (Fedus et al., 2018) has long been one of the key milestones of machine learning. Powered by the learning capabilities of deep neural networks, generative adversarial networks (GANs) (Goodfellow et al., 2014) have brought the field closer to attaining this goal. Currently, GANs are applied in a setting where the whole training dataset is of user interest. Therefore, regular GANs no longer meet our requirement when only partial instead of the entire training dataset possesses the desired properties (Killoran et al., 2017). It is more challenging that the given dataset has a small number of desired data.

Current GAN methods require user-defined criteria to select the desired data in order to learn the distribution of the desiderata. Vanilla GAN and Feedback GAN (FBGAN) (Gupta and Zou, 2019) first selects the samples possessing the desired properties and then performs regular GAN training only on the selected samples to derive the desired distribution. However, vanilla GAN would fail when the desired samples are limited. Though FBGAN overcomes the imbalanced class problem by introducing the desired samples generated by the GAN into the training data, it suffers from other two drawbacks. First, bluntly eliminating undesired samples may lead to a biased representation of real desired data distribution. Because the undesired samples can also reveal useful clues about what is not desired. Suppose we want to generate old face images while there are a few old face images but many young face images in the training data, the young face images can be used as negative sampling (Mikolov et al., 2013) to learn subtle aging features (e.g. wrinkles, pigmented skin, etc.), which guides the generation of the desired old face images. Second, a ready-to-use selector is not always accessible for soliciting desired samples in real applications. The conditional variants of GAN, like CGAN (Mirza and Osindero, 2014) and ACGAN (Odena et al., 2017), introduce condition variables to model the conditional desired data distribution. However, the generation performance of condition-based GAN is governed by the respective conditions with sufficient training observations. When the desired data is limited, the conditional modeling would be dominated by major classes, i.e., undesired data, leading to the failure of capturing the desired distribution.

Instead of explicitly eliminating the undesired data, we consider a more general setting where GAN could be towed towards the distribution of user-desired data by the user preference. Particularly,

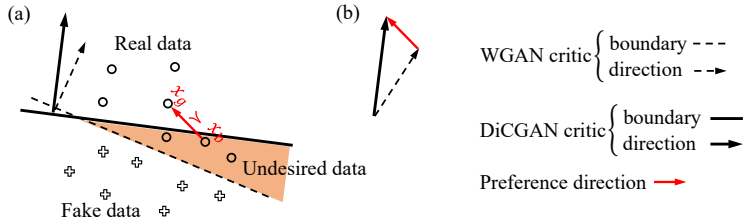

Figure 1: Illustration of why DiCGAN can learn the user-desired data distribution. (a) DiCGAN's critic pushes fake data towards the real desired data while WGAN's critic pushes fake data towards the entire real data. (b) The change of DiCGAN's critic direction is driven by the preference direction. Note that the preference direction is learned from all pairwise preferences.

pairwise preferences are the most popular form of user preference due to its simplicity and easy accessibility (Lu and Boutilier, 2011). Therefore, our target is to incorporate the pairwise preferences into the learning process of GAN, so as to guide the generation of the desired data.

Relativistic GAN (RGAN) (Jolicoeur-Martineau, 2019) considers the critic value as the indicator of sample quality and defines the discriminator with the difference of critic values. Actually, the critic value in RGAN is similar to the ranking score, but used for describing the sample quality. Motivated by this, we consider taking the critic value as the ranking score and define the ranking loss for pairwise preferences based on the critic value directly. Particularly, the difference of critic values for each pair of samples would reflect the user's preference over the samples. This is why we call our critic as **differential critic**, and we propose Differential-Critic GAN (DiCGAN) for learning the user-desired data distribution. As shown in Fig. 1, the differential critic indeed incorporates the user preference direction, which tows the original critic direction towards the real desired data region instead of the entire real data region. The main contributions are summarized as follows:

- We propose DiCGAN to learn the distributions of the desired data from the entire data using pairwise preferences. To our best knowledge, this is the first work to promote the ratio of the desired data by incorporating user preferences directly into the data generation.
- We introduce the differential critic by defining an additional pairwise ranking loss on the WGAN's critic. It endows the difference of critic values between each pair of samples with user preferences.
- Empirical study shows that DiCGAN indeed learns the distribution of user-desired data and the differential critic can derive the preference direction even from the limited preferences.

2 Generative Adversarial Networks

Generative Adversarial Network (GAN) (Goodfellow et al., 2014) performs generative modeling by learning a map from low-dimensional input space $\mathcal{Z}$ to data space $\mathcal{X}$, i.e., $G_\theta : \mathcal{Z} \to \mathcal{X}$, given samples from the training data distribution, namely, $x \sim p_r(x)$. The goal is to find θ that achieves $p_\theta(x) = p_r(x)$, where $p_\theta(x)$ is the fake data distribution $x = G_\theta(z)$. Let $p(z)$ be the input noise distribution and G indicate G_θ. GAN defines a discriminator D that is trained to discriminate real data from fake data to guide the learning of G.

Wasserstein GAN (WGAN) (Arjovsky et al., 2017) proposes to use Wasserstein metric as a critic, which measures the quality of fake data in terms of the distance between the real data distribution and the fake data distribution. The Wasserstein distance (W-distance) is approximated by the difference of average critic values between the real data and the fake data. The empirical experiments show that the W-distance between two distributions well-corresponds to the quality of the generated data. WGAN's objective function is defined as follows:

$$\min_G \max_D \mathbb{E}_{p_r(x)}\left[D(x)\right] - \mathbb{E}_{p_\theta(x)}\left[D\left(x\right)\right], \tag{1}$$

where D is the critic and satisfies 1-Lipschitz.

3 DICGAN FOR USER-DESIRED DISTRIBUTION

No longer learning the distribution of the whole dataset, GAN is applied in a new scenario, where the distribution of the partial dataset is what we desire. The user-desired data may refer to some certain class of data among multiple class datasets, or observations with/without some particular attributes. Such data can be induced from the user preference, which can be represented as an ordering relation between two or more samples in terms of the desired properties. We propose differential-critic GAN to learn the desired data distribution from the user preferences along with the whole dataset.

3.1 LEARNING THE DISTRIBUTION OF USER-DESIRED DATA

Following score-based ranking literature, we suppose that there exists a numeric score associated with each sample, reflecting the user's preference for the sample. A higher score indicates that its corresponding sample is preferred by the user. In detail, let f denote a score function that maps the sample x to the score $f(x)$. Then, if a sample x is desired by the user, its score $f(x)$ exceeds a predefined threshold ϵ, namely, $I(f(x) > \epsilon) = 1$. I is a sign function, which equals 1 if its condition is true and 0 otherwise. For the sake of explanation, we use $p_{\mathrm{r}}(x), p_{\mathrm{d}}(x), p_{\mathrm{u}}(x)$ to denote the distribution of the whole data, the user-desired data and the undesired data, respectively.

FBGAN (Gupta and Zou, 2019) was proposed to learn the distribution of the desired data $p_{\mathrm{d}}(x)$. FBGAN executes alternatively between two steps: (1) construct the desired dataset $\mathrm{X_d} = \{x | I(f(x) > \epsilon) = 1, x \sim p_{\mathrm{r}}(x)\}$; (2) train GAN on $\mathrm{X_d}$ to derive $p_{\mathrm{d}}(x)$. However, the assumption that the score function f is predefined in FBGAN may be too restrictive for real applications, where no universal and explicit criteria exists. Further, the definitions of the desired/undesired samples are highly dependent on the choice of the threshold ϵ. The removal of the so-called undesired samples may result in a biased representation of real desired data distribution.

Instead of relying on a predefined score function, we propose to learn the desired data distribution straightforward from the user preference. Here, we consider a general auxiliary information, i.e., the pairwise preferences, to represent the user preference, due to its simplicity and easy accessibility. For any two samples $x_1, x_2 \sim p_{\mathrm{r}}(x)$, let $x_1 \succ x_2$ denote x_1 is preferred over x_2 according to the user-defined criteria. Let X be the training samples, i.e., $\mathrm{X} = \{x^i \sim p_{\mathrm{r}}(x)\}$. A collection of pairwise preferences S are obtained by:

$$\mathrm{S} = \big\{ s = (x_1, x_2) | x_1 \succ x_2, x_1, x_2 \in \mathrm{X} \big\}. \tag{2}$$

Definition 1 (Problem Setting)**.** *Given the training samples* X *and the pairwise preferences* S*, the target is to learn a generative model* $p_\theta(x)$ *that is identical to the distribution of the desired data* $p_{\mathrm{d}}(x)$*, i.e.,* $p_\theta(x) = p_{\mathrm{d}}(x)$*.*

3.2 DIFFERENTIAL CRITIC GAN

Instead of WGAN'scritic for quality assessment, we present the differential critic for modelling pairwise preferences. The differential critic can guide the generation of the user-desired data.

3.2.1 PAIRWISE PREFERENCE

In this section, we consider incorporating the pairwise preference into the training of GAN.

Score-based ranking model (Zhou et al., 2008) is used to model the pairwise preferences. It learns the score function f, of which the score value, called ranking score in the model, is the indicator of the user preference. Further, the difference of ranking scores can indicate the pairwise preference relation. That is, for any pair of samples x_1, x_2, if $f(x_1) - f(x_2) > 0$ then $x_1 \succ x_2$ and vice versa. For any pairwise preference $s : x_1 \succ x_2$, the ranking loss we consider is as follows:

$$h(s) = \max\left(0, -\left(f\left(x_1\right) - f\left(x_2\right)\right) + m\right), \tag{3}$$

where m is the ranking margin. Other ranking losses can be referred to (Zhou et al., 2008).

Instead of learning the score function independent of GAN's training, we consider incorporating it into GAN's training, guiding GAN towards the generation of the desired data. The critic in RGAN (Jolicoeur-Martineau, 2019) is similar to the score function, where the critic values are used

to describe the sample quality. We are motivated to take the critic value as the ranking score and define the ranking loss on the critic value directly. Particularly, the difference of critic values for each pair of samples would reflect the user's preference over the samples.

3.2.2 LOSS FUNCTION

We build DiCGAN based on WGAN and the pairwise ranking loss is defined over the WGAN's critic. The loss function for DiCGAN is defined as:

$$\min_{G}\max_{D}\mathbb{E}_{p_{\mathrm{r}}(x)}\left[D(x)\right]-\mathbb{E}_{p_{\theta}(x)}\left[D\left(x\right)\right]-\lambda\frac{1}{|\mathrm{S}|}\sum_{s\in\mathrm{S}}\left[h\left(s\right)\right]. \tag{4}$$

where $h(s)$ is the pairwise ranking loss (equation 3). λ is a balance factor, which will be further discussed in section 3.3. Similar to WGAN, we formulate the objective for the differential critic L_D and the generator L_G as:

$$L_D=\frac{1}{b}\sum_{i=1}^{b}D(x^i)-D(G(z^i))-\lambda\frac{1}{n_{\mathrm{s}}}\sum_{j=1}^{n_{\mathrm{s}}}h(s^j),\qquad L_G=\frac{1}{b}\sum_{i=1}^{b}-D(G(z^i)), \tag{5}$$

where b is the batch size. Same for the fake samples. n_{s} is the number of pairs sampling from S.

The superiority of DiCGAN is twofold: (1) The introduced ranking loss in DiCGAN is defined on the critic value. Apart from WGAN, it can be easily applied to other GAN variants developed based on the critic, e.g., RGAN. (2) DiCGAN leverages the entire dataset. The pairwise preferences are constructed on the whole dataset. Thus the undesired samples are also utilized during the training.

In the following, we argue that the differential critic in DiCGAN can guide the generator to learn the user-desired distribution from two aspects. (1) As shown in Fig. 1, the differential critic in DiCGAN provides the direction towards the real desired data. We denote the critic direction as the moving direction of fake data, which is orthogonal to the decision boundary of the critic. Referred to equation 4, DiCGAN's critic loss consists of two terms: the vanilla WGAN loss and the ranking loss. The vanilla WGAN loss imposes the critic direction from the fake data to the real data. Meanwhile, the ranking loss induces a user preference direction, which points from the undesired data to the desired data. Combining these two effects, the critic direction of DiCGAN would target for the region of the real desired data only. (2) DiCGAN assigns high critic values for user-desired data and promotes the generation of samples with high critic values. The vanilla WGAN loss encourages the critic to assign high critic values for real data while low critic values for fake data. Meanwhile, the ranking loss urges to assign high critic values for the real desired data while low critic values for real undesired data. Therefore, the real desired data would achieve high critic values from the critic. Similar to WGAN, the training paradigm of DiCGAN promotes the generation of samples with high critic values, which is equivalent to encourage generating the user-desired data in DiCGAN.

3.3 REFORMULATING DICGAN TO ENSURE DATA QUALITY

Let's revisit the objective of DiCGAN (equation 4). The first two terms of equation 4 can be considered as the WGAN regularisation, which ensures the generated data distribution is close to the whole real data distribution, i.e., $p_{\theta}\approx p_{\mathrm{r}}$. The third term serves as a correction of WGAN, which drives WGAN slightly biased to our target of learning the desired data distribution, i.e., $p_{\theta}=p_{\mathrm{d}}$.

Therefore, the WGAN regularisation serves as the cornerstone of our DiCGAN. Particularly, if the desired data distribution is close to the entire data distribution, the ranking loss would easily correct the WGAN for achieving the desired data distribution. Otherwise, the satisfactory performance of DiCGAN may require online hyperparameter tuning of λ during the training process.

According to the above analysis, we consider reformulating the objective of DiCGAN, i.e., equation 4 into an equivalent objective with a hard WGAN constraint:

$$\min_{G}\max_{D}-\sum_{s\in\mathrm{S}}\left[h\left(s\right)\right],\quad \text{s.t. } d(p_{\mathrm{r}},p_{\theta})=\left|\mathbb{E}_{p_{\mathrm{r}}(x)}\left[D(x)\right]-\mathbb{E}_{p_{\theta}(x)}\left[D\left(x\right)\right]\right|\leq\varepsilon. \tag{6}$$

where $\varepsilon>0$. Note that we impose an explicit non-negative constraint on $d(p_{\mathrm{r}},p_{\theta})$, to highlight it is a distance metric. It is still equivalent to WGAN loss from its definition. Therefore, equation 4 is the

Lagrange function of equation 6. Since equation 6 imposes a hard constraint on the WGAN loss, it is more difficult to optimize compared to equation 4. However, more efficient solutions of DiCGAN can be explored through analyzing equation 6 regarding the hard constraint on $d(p_{\mathrm{r}}, p_{\theta})$.

In terms of a minor correction situation, it means the desired data distribution p_{d} is close to the real data distribution p_{r}. Therefore, the correction would dominate the training goal of DiCGAN. By assigning a proper λ to ensure the constraint satisfied, equation 4 itself can learn the distribution of the user-desired data while ensuring the data quality.

In terms of a major correction situation, it means the desired data distribution p_{d} is quite diverse from the real data distribution p_{r}. Therefore, DiCGAN needs to achieve an equilibrium between the correction, imposed by the ranking loss, and the hard constraint, imposed by the WGAN loss. However, a large correction may not ensure the quality of the generated data, since the WGAN loss, used to guarantee the image quality, is defined between the generated data and the whole real data. To avoid the major correction, we propose to break the major correction into a sequential of minor corrections to ensure data quality. Namely, at each epoch, we first use the generator G to generate n_{g} samples, denoted as X_{g}:

$$\mathrm{X}_{\mathrm{g}}^{e} \leftarrow \{G^{e}(z^{1}), \ldots, G^{e}(z^{n_{\mathrm{g}}})\}, \quad \{z^{i} \sim p(z)\}_{i=1}^{n_{\mathrm{g}}}, \tag{7}$$

where e is the epoch index. Then we replace the old training samples with the generated samples:

$$\mathrm{X}^{e+1} \leftarrow \mathrm{X}^{e} \setminus \mathrm{X}_{\mathrm{o}}^{e} \cup \mathrm{X}_{\mathrm{g}}^{e}, \tag{8}$$

where $\mathrm{X}_{\mathrm{o}}^{e}$ are the old (least-recently added) n_{g} samples in X^{e}.

Due to the ranking loss, the generated data distribution p_{θ}^{e} is closer to the desired data distribution p_{d}, compared to the constructed p_{r}^{e} at each epoch. Therefore, the iterative replacement (equation 8) can gradually shift the real data distribution p_{r} towards the desired data distribution p_{d}. Namely, $d(p_{\mathrm{r}}, p_{\mathrm{d}}) > \cdots > d(p_{\mathrm{r}}^{e}, p_{\mathrm{d}}) > d(p_{\mathrm{r}}^{e+1}, p_{\mathrm{d}}) > \cdots$. So that only a minor correction needs to be imposed on p_{θ}^{e} through optimizing equation 4 at each epoch. Iteratively, the generated distribution p_{θ} would shift towards p_{d}. The training algorithm is summarized in Algorithm 1 in the appendix. For the sake of easy optimization, we pretrain the differential critic D and the generator G using WGAN.

4 Case Study on Synthetic Data

To gain an intuitive understanding of the difference between our DiCGAN and WGAN regarding the critic, generator, we conduct a case study on the synthetic dataset. Note that all methods are trained with the whole dataset.

The synthetic dataset consists of two concentric circles by adding Gaussian noise with a standard deviation of 0.05 (See Fig. 2a). The samples located on the inner circle are considered as the desired data, while the samples on the outer circle are defined as the undesired data. By labeling the desired data as $y = 1$ and undesired data as $y = 0$, we can construct the pairwise preference for two samples x_1 and x_2 based on their labels. Namely $x_1 \succ x_2$ if $y_1 = 1 \wedge y_2 = 0$, and vice versa. The pairs are constructed within each mini-batch. Our target is to learn the distribution of the desired data (i.e., samples on the inner circle), using the whole data along with the constructed pairwise preferences.

4.1 WGAN vs DiCGAN on Critic

Experiment setting: we fix the generator and simulate fake data as the 2D Gaussian blob with a standard deviation of 0.05 (green pluses). We first train the critic to converge. Then, we project the output on the second last layer of the critic into 1D space using kernel principal components analysis (Schölkopf et al., 1997), to derive the projected features. To explore the difference between the critics of WGAN and DiCGAN, we draw the curve of critic values versus the projected features for WGAN and DiCGAN, respectively (Fig. 2b, 2c).

From Fig. 2b, 2c, we can see: (1) In terms of the real data and the fake data, the critic of both WGAN and DiCGAN can achieve perfect discrimination. Meanwhile, the projected features of real data and that of fake data are also completely separated. (2) In terms of the real desired data and the real undesired data, the critic of DiCGAN assigns higher values to the desired samples, compared to undesired samples. It is because our ranking loss expects a higher ranking score (i.e., critic values)

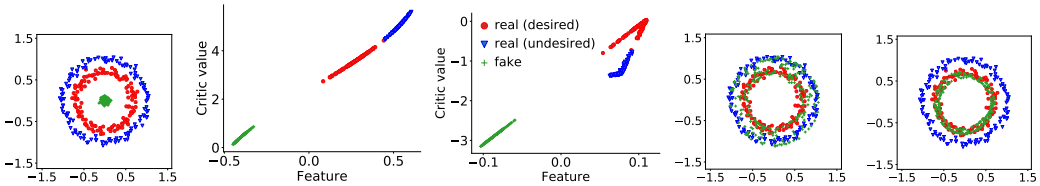

(a) Data (b) WGAN (c) DiCGAN (d) WGAN (e) DiCGAN

Figure 2: (a-c) Comparison of the critic in DiCGAN and WGAN. DiCGAN's critic can assign higher critic values for real desired data than real undesired data while WGAN's critic cannot. (d-e) The visualization of the generated samples from DiCGAN and WGAN. Learning the desired distribution means that the fake data should overlap with the real desired data only.

for the corresponding desired data. (3) In contrast, the critic of WGAN assigns lower values to the desired data, since desired data is closer to the fake data comparing to undesired data.

4.2 WGAN vs DiCGAN on Generator

Experiment setting: we train the critic and the generator following the regular GANs' training procedure. The generation results of WGAN and DiCGAN is shown in Fig. 2d, 2e.

DiCGAN (shown in Fig. 2e) only generates the user-desired data, i.e., generated data covering the inner circles, while WGAN (shown in Fig. 2d) generates all data, i.e., generated data covering the inner and outer circles. As the critic in DiCGAN can guide the fake data towards the real data region while away from the undesired data region, the generator thus produces data alike real desired data. Because the critic in WGAN pushes the fake data to the real data region only, the generator finally produces whole real-alike data.

5 Experimental Study

DiCGAN is applied to learn the distribution of the desired data on MNIST (Lecun et al., 1998) and CelebA-HQ (Karras et al., 2018) datasets. Due to the space limit, we present more experimental results in the appendix.

Networks & **Hyperparameters** The balance factor λ is set to 1. The ranking margin m is set to 1 for MNIST and 5 for CelebA-HQ. We adopt the same network structures as in (Gulrajani et al., 2017). See the appendix for other settings.

Evaluation Metric For evaluating the performance of learning the desired data distribution, we calculate the percentage of user-desired data in GAN's generation, i.e., the ratio of the desired generated data to the whole generated data (D/W). We use inception score (IS) (Salimans et al., 2016) and multiscale structural similarity (MS-SSIM) (Odena et al., 2017) to evaluate the quality and intra-class diversity for GANs' generation, respectively.

Baselines We compare DiCGAN with WGAN (Gulrajani et al., 2017), CWGAN (Mirza and Osindero, 2014) and FBGAN (Gupta and Zou, 2019). WGAN is only trained with the desired data to derive the desired distribution. CWGAN is the extension of GAN with a conditional label c. To train CWGAN, we split the training data into the desired class ($c = 1$) and the undesired class based on a predefined user criteria ($c = 0$). Then $p(x|c = 1)$ is the desired data distribution. FBGAN adopts an iterative training paradigm to derive the desired data distribution. At each training epoch, FBGAN resorts to an extra selector to select the desired samples from the generated samples and performs regular GAN training using the selected desired samples. ACGAN (Odena et al., 2017) also shows poor results like CWGAN when the desired data is limited, so we do not report here.

5.1 Learning the Distribution of Desired Digits

We design the experiment to learn the distribution of small digits in MNIST. We use $50K$ 28×28 images as the training data. As MNIST contains digits zero to nine, zero is the most desirable data.

As for WGAN and CWGAN, zero digits in the training data are regarded as the desired samples ($c = 1$), whose size is $4,950$. The other digits are labeled as the undesired samples, whose size is $45,050$ ($c = 0$). WGAN is only trained with the constructed desired data. CWGAN conditions on c to model a conditional data distribution $p(x|c)$ for MNIST dataset. FBGAN and our DiCGAN resort to a classifier, pre-trained for digits classification, to obtain the labels for the generated samples. At every training epoch, FBGAN generates 50000 samples and requests the classifier to label them.

(a) Epoch 0 (b) Epoch 1 (c) Epoch 2 (d) Epoch 5 (e) Epoch 10

Figure 3: Generated digits of DiCGAN on MNIST. DiCGAN learns the distribution of small digits.

Then the images are ranked using the predicted label, with the smaller digit ranked higher. The generated images with digits ranked top 50%, i.e., small digits, are selected as the desired data to replace old training data. As for DiCGAN, the pairwise comparison can be obtained for two images x_1 and x_2 according to their predicated label y_1 and y_2, namely $x_1 \succ x_2$ if $y_1 < y_2$, and vice versa. At each iteration, we construct 32 pairwise preferences for each mini-batch with 64 training samples.

Fig. 3 presents the generated MNIST images randomly sampling from the generator of DiCGAN. It shows that the generated MNIST digits gradually shift to smaller digits during the training, and converges to digit zero. We sample $50K$ samples from the generators of various GANs and respectively calculate the percentage of digit zero and digits zero to four among the generated digits for quantitative evaluation. In Table 1, only small digits are generated by DiCGAN and FBGAN; WGAN and CWGAN can also learn the distribution of the desired digit since the dataset is simple and has relatively sufficient data for the desired digit. However, WGAN and CWGAN do not exhibit a smooth convergence to digit zero like FBGAN and DiCGAN (See Fig. 3 in the appendix). In addition, when the dataset is complex and the desired data is insufficient, WGAN and CWGAN would fail, which will describe in the next section.

5.1.1 Superiority of DiCGAN Over FBGAN

Though FBGAN achieves a good performance on learning the desired data distribution, it requires a lot of supervision information from the selector. We calculate the number of effective pairs (#EP) used in DiCGAN and FBGAN, respectively. #EP in DiCGAN denotes the total number of explicitly constructed pairs during the training, i.e., #EP $= \sum_{i=1}^{n_\mathrm{e}} \sum_{j=1}^{n_\mathrm{i}} n_\mathrm{s}$. FBGAN selects the desired samples from the generated samples. #EP can be induced by the implicit pairs implied by the desired samples versus the undesired samples, i.e., #EP $= \sum_{i=1}^{n_\mathrm{e}} n_\mathrm{gd} \times n_\mathrm{gu}$, where n_e is the number of training epoch. where n_gd and n_gu denote the number of desired samples and undesired samples in the generation, respectively.

Fig. 4a shows that (1) the #EP used in DiCGAN is much smaller than that in FBGAN at each training epoch. (2) The total #EP used in DiCGAN is significantly less than that in FBGAN, which can be reflected from the shadow area. In total, DiCGAN used $9.53e4$ effective pairs while FBGAN used $2.02e8$ effective pairs. Our DiCGAN is scalable to the large training dataset, e.g. MNIST. #EP in DiCGAN is linearly correlated to the training size. In contrast, #EP in FBGAN is determined by n_gd and n_gu, which are both linearly correlated to the training size. Thus #EP in FBGAN thus is quadratically correlated to the training size.

5.1.2 Comparing DiCGAN and FBGAN Given the limited Supervision

We compare DiCGAN and FBGAN on MNIST dataset given the limited supervision. Specifically, the query amount of resorting to the pre-trained classifier for obtaining the prediction of the generated samples are restricted to $5,000$ both for FBGAN and DiCGAN.

Fig. 4d shows that DiCGAN can learn the desired distribution while FBGAN fails, only generating 10.3% digit zero, which are consistent with the visual results in Fig. 4b and Fig. 4c. This verifies that

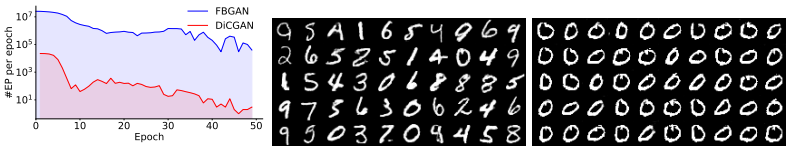

(a) (b) (c)

Method	MNIST top 1	MNIST top 5
FBGAN	10.3	52.6
DiCGAN	**99.7**	**99.9**

(d)

Figure 4: (a) Comparison between DiCGAN and FBGAN w.r.t. #EP per epoch on MNIST in Section 5.1. (b) The generated results of FBGAN in Section 5.1.2. (c) The generated results of FBGAN in Section 5.1.2. (d) % of desired digits (D/W) of FBGAN and DiCGAN in Section 5.1.2.

(a) WGAN (b) CWGAN (c) FBGAN (d) DiCGAN

Figure 5: Generated images by WGAN, CWGAN, FBGAN and DiCGAN. The green boxes refer to the images which are classified as old images by a trained classifier.

the negative samples are beneficial to learn the user-desired distribution. Particularly, the preference direction can be captured by our differential critic even when the supervision is limited, which guides the generation towards the desired data.

5.2 Learning the Distribution of Face with the Desired Attribute

We consider old face images as the desired data and design the experiment to learn the distribution of old face images on CelebA-HQ. CelebA-HQ dataset is the high quality version of a subset from Celeb Faces Attributes (CelebA) dataset. It consists of 30K images face images of celebrities, annotated with 40 binary attributes such as age. We resize the images to 64×64. The training setting for each GAN is similar as the mentioned above. See the appendix for details.

In Fig. 5, we visualize the generated face images randomly sampled from the generator of each model. (1) WGAN has poor generation since its training data, i.e., the desired subset is insufficient. (2) CWGAN presents good quality of generation but fails to capture the desired distribution. There are only two old face image out of 9 randomly sampled images. (3) All sampled images from FBGAN and our DiCGAN are the desired old faces. We sample $10K$ samples from the generator and calculate the percentage of old face images among the generated samples for quantitative evaluation. In Table 1, (1) shown by D/W, almost all images generated by DiCGAN and FBGAN are old face images. WGAN and CWGAN both fail to capture the distribution of old face images. (2) As for the quality and the diversity, we first calculate the IS and MS-SSIM of the old faces in the training data (denoted as "original" in Table 1). The IS and MS-SSIM of the generations do not exhibit a big difference from the "original", which means that the quality and the diversity of the generations are relatively good. WGAN achieves the poorest quality due to the insufficient training data.

Method	MNIST		CelebA-HQ		
	D/W		D/W	IS	MS-SSIM
	top 1	top 5			
Original	9.9	51.1	22.1	**2.81**	**0.52**
WGAN (subset)	97.3	98.2	73.8	1.77	0.60
CWGAN	95.0	96.4	16.4	2.12	0.65
FBGAN	**100.0**	**100.0**	98.2	2.05	0.60
DiCGAN	**100.0**	**100.0**	**99.5**	1.93	0.60

Table 1: Results on the MNIST and CelebA-HQ dataset. Top 1 refers to digit zero. Top 5 refers to digits zero to four.

6 Discussions

Current GAN (Goodfellow et al., 2014; Mirza and Osindero, 2014; Odena et al., 2017; Gupta and Zou, 2019) based methods require user-defined criteria to select the desired data in order to derive the distribution of the desiderata. There are two limitations for these methods: (1) the criteria are not always accessible in real applications; (2) eliminating the handcrafted undesired samples loses useful information about what is not desired. Other works (Gómez-Bombarelli et al., 2018; Engel et al., 2018) proposed to find the subset of the latent space corresponding to the desired data and generates data from the latent subset. However, they also rely on user-defined criteria for labeling the desired data or a ready-to-use score function for finding the subset of the latent space with high scores.

This paper proposes DiCGAN to learn the distributions of the user-desired data from the entire data using the pairwise preferences. We empirically demonstrate the efficacy of DiCGAN in terms of promoting samples with user-desired properties on MNIST and CelebA-HQ datasets, respectively. One promising future direction for DiCGAN could be the minority population promotion for imbalanced data tasks, such as imbalance classification problems, few-shot learning, one-shot learning or even open-set problem. Another interesting direction of DiCGAN could be the desired policy generation in imitation learning if the pairwise comparison between the policies could be properly designed.

REFERENCES

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.

Jesse Engel, Matthew Hoffman, and Adam Roberts. Latent constraints: Learning to generate conditionally from unconditional generative models. In *International Conference on Learning Representations*, 2018.

William Fedus, Ian J Goodfellow, and Andrew M Dai. MaskGAN: Better Text Generation via Filling in the _____. In *International Conference on Learning Representations*, 2018.

Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2):268–276, 2018. doi: 10.1021/acscentsci.7b00572. URL https://doi.org/10.1021/acscentsci.7b00572.

Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.

Anvita Gupta and James Zou. Feedback gan for dna optimizes protein functions. *Nature Machine Intelligence*, 1(2):105, 2019.

Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard GAN. In *International Conference on Learning Representations*, 2019.

Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

Nathan Killoran, Leo J Lee, Andrew Delong, David Duvenaud, and Brendan J Frey. Generating and designing dna with deep generative models. *arXiv preprint arXiv:1712.06148*, 2017.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998.

Tyler Lu and Craig Boutilier. Learning mallows models with pairwise preferences. In *International Conference on Machine Learning*, pages 145–152, 2011.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, pages 2642–2651, 2017.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.

Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel principal component analysis. In *International conference on artificial neural networks*, pages 583–588. Springer, 1997.

Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *Advances in neural information processing systems*, pages 613–621, 2016.

Ke Zhou, Gui-Rong Xue, Hongyuan Zha, and Yong Yu. Learning to rank with ties. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '08, page 275–282, New York, NY, USA, 2008. Association for Computing Machinery.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.