

COMPETITIVE-COLLABORATIVE GAN WITH PERFORMANCE GUARANTEE

Anonymous authors

Paper under double-blind review

ABSTRACT

Generative Adversarial Networks (GANs) generate data based on a competition game to minimize the distribution distance between existing and new data. However, such a competition game falls short when insights about data distributions beyond their authenticity are imperative, such as in multi-modal generation and image super resolution. In recognition of the limitations inherent to the pure-competition mechanism, we introduce CCGAN, a Collaborative-Competitive Generative Adversarial Network scheme to enable data generation with additional knowledge beyond the provided dataset distribution. For theoretically preserving the equilibrium point and numerically avoiding training collapse issue, we show the need to convert the regularization term into a divergence metric, so that the modified GAN is well-defined in game theory. By harmonizing the competition and collaboration losses in CCGAN, we effectively reduce the degree complexity of solving the optima, facilitating the establishment of a closed-form equilibrium point. This equilibrium point serves as a guidance for training and hyper-parameter tuning, resulting in consistently high-quality generations. Meanwhile, the regularization breaks the mutual dependency between the generator and discriminator. This newfound independence empowers the CCGAN to explore a broader parameter space, effectively mitigating the training collapse issue. To validate the capabilities of CCGAN, we design comprehensive experiments across four publicly available datasets and systematically compare CCGAN against a range of baseline models. The experiments demonstrate the efficacy of CCGAN on generating satisfactory samples tailored to specific requirements, particularly when applied to the generation of images featuring regularly shaped objects.

1 INTRODUCTION

Generative Adversarial Networks (GANs) Goodfellow et al. (2014; 2020) empower the creation of deceptively real data, showing success across diverse data types, such as images Heusel et al. (2017), text Haidar & Rezagholizadeh (2019), and speech Eskimez et al. (2020). In GAN’s framework, the adversarial dynamic between the generator and discriminator gives rise to a pure-competition game scenario (see Figure 1 (a)): the generator iteratively refines its ability to create samples that are indistinguishable from real data. It does so by learning to deceive the discriminator, who also iteratively refines its ability to distinguish between real (authentic) and fake (synthetic) data.

While GANs gain success in their competitive design to generate synthetic data that follow authentic distribution, this pure-competition nature becomes less effective in certain applications Huang et al. (2022). In such applications, the need extends beyond capturing the authentic data distribution (see Figure 1 (b)). For instance, in multi-modal generation Liu et al. (2021), the focus is on learning one or multiple modes within the data distribution. In image super-resolution Zhu et al. (2020a), one must not only capture the data distribution but also faithfully reproduce the original image. The additional requisites for generation pose challenges for pure-competition GANs, as they lack an inherent incentive for the generator and discriminator to collectively fulfill the additional requisites.

To overcome the challenges inherent in pure-competition GANs, researchers have ventured into the realm of collaboration-competition design of GANs. By re-designing the loss function, one can foster collaboration between the generator and discriminator to collectively fulfill additional requirements regarding generated data Durgadevi et al. (2021). These strategies to reshape the competitive nature of GANs include (1) replacing the standard binary cross-entropy loss with alternative loss functions, such as the least square loss Mao et al. (2017) and Wasserstein loss Arjovsky et al. (2017), (2) introducing regularization terms to promote collaboration between the generator and discriminator (see Figure 1 (c)), such as coordinate GAN Lin et al. (2019) and collaborative GAN (CollaGAN) Lee et al. (2019), and (3) designing multiple generators Zhang et al. (2020) and multiple discriminators Choi & Han (2022) to promote collaboration.

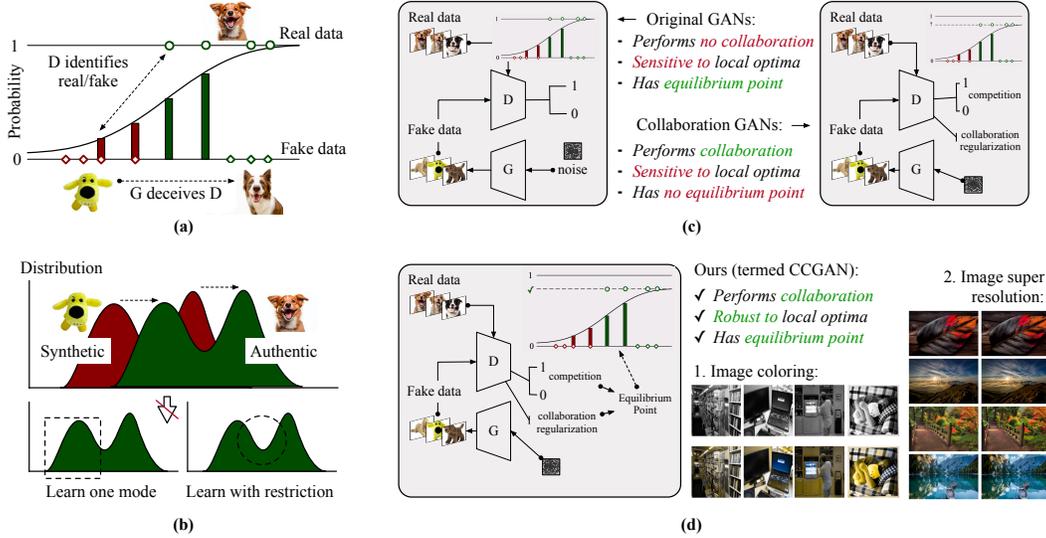


Figure 1: (a) Discriminator D learns a binary classification between fake and real data. Generator G learns to produce deceptively real data to fool D . (b) GAN’s adversarial nature is limited in applications requiring knowledge beyond the “authentic” data distribution. (c) The architecture of original GAN and GAN variants that conduct a collaboration-competition game. (d) The architecture of our proposed model CCGAN and its preliminary experimental results.

However, these collaboration-competition studies have limitations, possibly yielding sub-optimal results. First, they lack a theoretical guarantee. Specifically, the inconsistent integration of the original GAN loss and the collaboration regularization disrupts the original GAN’s theoretical guarantee, resulting in a blind model without proper guidance Lin et al. (2019). Second, they can be numerically stuck in local-optima if training collapses Bang & Shim (2021). Specifically, when generator and discriminator attain a state of local optima, achieving further enhancements becomes increasingly arduous due to their mutual dependent relationship. Overall, the absence of a theoretical guarantee and the training collapse issue can result in misguided collaboration and sub-optimal generated data.

To address the aforementioned challenges, which involve three key aspects: (1) pushing the boundary of GAN from pure-competition to collaboration-competition, (2) theoretically deriving the equilibrium point of this collaboration-competition game, and (3) numerically avoiding training collapse, this paper proposes a novel design of regularization for the collaboration need. Specifically, since the existing methods fail to obtain a closed-form equilibrium point due to inconsistent integration between the collaboration and competition loss components, we aim to establish a smooth integration (see Figure 1 (d)), where the collaboration loss function can be transformed into a divergence metric. With this in mind, we present the innovative solution named **Collaboration Competition-based Generative Adversarial Network (CCGAN)**. Our main contributions are:

- Employing collaboration-competition mechanism, we establish a unified optimization procedure that can generate high-fidelity samples catering to specific requisites (see Figure 1 (d)). The numerical experiments also reveal one interesting phenomenon:
 - The generated image samples exhibit superior quality when the objects within the images are of regular shape. This advantage can potentially be attributed to the presence of the closed-form equilibrium point.
- We leveraged Jensen–Shannon divergence to develop a smooth integration between the collaboration and competition loss components. It enables the attainment of a global optimality, i.e., the equilibrium point.
- We prove that the generator and discriminator is not mutually dependent in CCGAN. This non-mutual-dependence empowers CCGAN to explore more extensive parameter space, effectively mitigating the issue of training collapse.

Overall, these findings provide both theoretical and empirical insights into collaboration-competition GANs, paving the way for stable and consistently high-quality data generation with specific requirements that extend beyond the provided data distribution.

2 RELATED WORK

Related work section considers three lines of intellectual inquiry: (1) recent developments of GAN, (2) GANs with regularization, and (3) collaboration-competition design in machine learning.

Recent developments of GAN. Addressing critical challenges identified in GAN Chen (2021), current research trends include the follows. First, there’s a concerted effort to enhance training techniques Zhong et al. (2020) and achieve training stability Neyshabur et al. (2017). Second, researchers are working towards the development of interpretable and controllable GANs while ensuring performance guarantees Tripathy et al. (2020). Scalability is another focus, with endeavors to establish large-scale and high-resolution GANs Brock et al. (2018). Moreover, GANs are expanding their horizons beyond image generation to include data generation in diverse domains Chen et al. (2018), the generation of multi-modal data Liu et al. (2021), and cross-domain data translation Reed et al. (2016). Our paper investigates the interconnected landscape of several of these trends, with the overarching goal of advancing GAN research in training, interpretability, and application.

GANs with regularization. Building upon the original GAN, researchers have introduced various regularization techniques for a multitude of practical objectives, such as fulfilling additional requirements for generated data Lin et al. (2019); Lee et al. (2019), addressing mode collapse issue Nagarajan & Kolter (2017); Bang & Shim (2021), ensuring training stability Brock et al. (2018), and preserving privacy Wu et al. (2019); Chen et al. (2021). However, the introduction of regularization terms may inadvertently alter the optimization landscape of GAN training, compromising the theoretical guarantee of original GAN Nagarajan & Kolter (2017). Hence, there is a demand for the development of a novel regularization term that simultaneously addresses practical objectives while preserving theoretical guarantees.

Collaboration-competition design in machine learning. In machine learning, the notion of collaboration-competition emerges in social network studies Wasserman & Faust (1994); Fu et al. (2008); Lee et al. (2014). The goal was to balance collective advancement and individual progress Lee et al. (2012), promoting cooperation for communal objectives while motivating individuals to excel and outperform their peers Aaldering et al. (2019). Following its success in social networks, the collaboration-competition framework extended its influence to neural network studies Shi et al. (2020); Kopacz et al. (2023). This concept entails coordinating the contributions of various models/networks to collectively boost performance, while preserving a competitive element to drive individual model enhancements. This approach fosters knowledge sharing Gong et al. (2020), information exchange Tan et al. (2019), and collaborative learning among neural networks, resulting in improved generalization, model diversity, and overall network efficacy. In this paper, we leverage the collaboration-competition idea in GAN to combine individual network strengths Poirot et al. (2019) and harness collective intelligence for superior outcomes.

3 PROPOSED MODEL

Pure-competition GANs and challenges. To present the motivation of our method, we start from analysing the mechanism and challenges in original GAN. The original GAN model Goodfellow et al. (2014) conducts an adversarial training between the generator G and discriminator D as

$$\min_G \max_D V_{\text{GAN}}(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

It aims to capture p_{data} , which represents the underlying data distribution of given samples \mathbf{x} . For this purpose, G generates fake samples $G(\mathbf{z})$ from random noises \mathbf{z}^1 , while D acts as a binary classifier to distinguish whether a sample belongs to the given training data or is generated by G (see Figure 1 (a)). From the pure-competition nature of the generator-discriminator relationship, GANs can generate fake data that follow “real” distribution, thus showing success in generating diverse data types. However, there also remains challenges associated with the pure-competition design where learning the “real” distribution is not the only focus Mao et al. (2017); Huang et al. (2022). In particular, certain applications may necessitate the acquisition of knowledge about a single mode within the distribution or a restricted portion of the distribution (see Figure 1 (b)). For example, in the task of multi-modal generation Liu et al. (2021), it requires to learn one or several modes within the data distribution. In image super-resolution Zhu et al. (2020a), it is essential not only to grasp the underlying data distribution but also to replicate the exact original image. For these applications,

¹ \mathbf{z} follows a given distribution (e.g., Gaussian) denoted as $p_{\mathbf{z}}(\mathbf{z})$

the focus is understanding distributions beyond its authenticity, thus necessitating another incentive beyond the pure-competition loss function in GAN.

Collaboration-competition GANs and challenges. To fulfill the specific need beyond the authenticity of the data distribution, recent studies Lin et al. (2019); Zhang et al. (2020); Lee et al. (2019) have introduced regularization terms to GAN. They are designed to align the objectives of both the generator and discriminator, guiding them towards a common need. The general form reads as

$$\max_D V_{\text{GAN}}(D, G) - L_{\text{regularization}}(D, G; \lambda); \quad \min_G V_{\text{GAN}}(D, G) + L_{\text{regularization}}(D, G; \lambda), \quad (2)$$

where $V_{\text{GAN}}(D, G)$ is the original GAN’s loss function in Equation (1), $L_{\text{regularization}}(D, G; \lambda)^2$ the loss tailored to address specific needs, and λ the hyperparameter.

However, the collaboration-competition game in Equation (2) represents a rough combination of regularization $L_{\text{regularization}}(D, G; \lambda)$ to the original GAN loss $V_{\text{GAN}}(D, G)$. As shown in Theorem 1, this fragmented design compromises the equilibrium point of original GAN. Without the equilibrium point, the model will suffer from misguided collaboration and can yield sub-optimal outcomes. For example, the hyper-parameter (λ) tuning can become a challenging task devoid of clear guidance or guarantees Kurach et al. (2018), and the haphazard mixing of loss functions can introduce conflicting objectives, hindering convergence and generating subpar outputs Tran et al. (2018).

Theorem 1. [No equilibrium point]. Equation (2) lacks a closed-form equilibrium point when $L_{\text{regularization}}(D, G; \lambda)$ is squared loss, and the true data distribution is Gaussian.

The proof of Theorem 1 can be found in Appendix A. It demonstrates that because of the inconsistent combination of $L_{\text{regularization}}(D, G; \lambda)$ and $V_{\text{GAN}}(D, G)$, the equilibrium equation is a polynomial equation with a degree exceeding five, thus lacking an analytical solution Ramond (2022). To resolve this, it’s crucial to “harmonize” $L_{\text{regularization}}(D, G; \lambda)$ and $V_{\text{GAN}}(D, G)$ in Equation (2), ensuring that $L_{\text{regularization}}(D, G; \lambda)$ doesn’t introduce extra complexity in finding an equilibrium point.

3.1 CCGAN FRAMEWORK

This section presents CCGAN, a GAN model that jointly considers the collaboration requirements and the equilibrium point, aiming to enhance the collaboration, interpretability, and overall optimality of the GAN model. The key design, as depicted in Equation (3)-(5), is chosen to enable the conversion from the loss function to a divergence metric. This transformation can prevent the introduction of extra degrees of complexity when seeking an equilibrium point.

Motivated by the work Mao et al. (2017) and Nowozin et al. (2016), we consider the regularization in Equation (3), which mimics the form of loss function in Equation (1). The corresponding optimization problem between discriminator D and generator G is shown in Equation (4)-(5):

$$L_{\text{regularization}}(D, G; \lambda) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} \log(D(\mathbf{x}) - \lambda) + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \log(D(G(\mathbf{z})) - \lambda), \quad (3)$$

$$\max_D C_D(D, G; \lambda) = V_{\text{GAN}}(D, G) - L_{\text{regularization}}(D, G; \lambda), \quad (4)$$

$$\min_G C_G(D, G; \lambda) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log(D(\mathbf{x}))] + \frac{\lambda - 1}{\lambda} \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (5)$$

The discriminator loss $C_D(D, G; \lambda)$ is the original loss $V_{\text{GAN}}(D, G)$ combined with the regularization term $L_{\text{regularization}}(D, G; \lambda)$, and the generator loss function $C_G(D, G; \lambda)$ is the re-scaled version of $V_{\text{GAN}}(D, G)$. In Equation (4)-(5), the goal of D and G is not completely opposite (unlike Equation (1)), making it a mixed-motive game which contains both the competition and the collaboration mechanism. The collaboration is enforced through the regularization term $L_{\text{regularization}}(D, G; \lambda)$, in which the two networks jointly explore distributions spanning from “authentic” to “synthetic”, with the balance controlled by the hyperparameter λ . We denote the optimization problem in Equation (4)-(5) as **Collaboration Competition-based Generative Adversarial Network (CCGAN)**.

Global optimality. The regularization term in Equation (3) has two advantages. First, it encourages the discriminator to map true data to probability λ . In contrast to the original discriminator, which always attempts to assign a probability 1 to real data, this additional regularization empowers the model to explore the distributions spanning the spectrum from “true” to “half true”. The “half true”

²Since $L_{\text{regularization}}(D, G; \lambda)$ is a collaborative pursuit involving both the discriminator and generator, its impact is manifested with an opposite sign in Equation (2).

distribution reflects the “true” distribution with specific requisites, and is mathematically mapped to probability λ . Hence, the hyperparameter λ is configured to implement the desired level of specific requisites for the generated data. Second, this regularization term facilitates the transformation of the loss function into a divergence metric. The transformation prevents the introduction of additional complexities when searching for an equilibrium point. This outcome leads to the establishment of global optimality, as demonstrated in Theorem 2.

Theorem 2 (Global optimality of CCGAN). *The optimization problem in Equation (4)-(5) is equivalent to minimizing the Jensen–Shannon divergence $JSD(p_{data} \parallel \frac{\lambda-1}{\lambda} \cdot p_{gen})$ at optimality. It leads to $p_{gen} = \frac{\lambda}{\lambda-1} \cdot p_{data}$, where p_{data} is true data distribution, p_{gen} the distribution of generated data.*

Proof. The training criterion for the discriminator D , given a fixed generator G , is to maximize the quantity $V_{GAN}(D, G) - L_{\text{regularization}}(D, G; \lambda)$ as

$$\begin{aligned} & \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log(D(\mathbf{x})) - \log(D(\mathbf{x}) - \lambda)] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z}))) - \log(\lambda - D(G(\mathbf{z})))] \\ &= \int_{\mathbf{x}} p_{data}(\mathbf{x}) [\log(D(\mathbf{x})) - \log(D(\mathbf{x}) - \lambda)] d\mathbf{x} + \int_{\mathbf{x}} p_{gen}(\mathbf{x}) [\log(1 - D(\mathbf{x})) - \log(\lambda - D(\mathbf{x}))]. \end{aligned}$$

Then, the optimal discriminator D_G^* is obtained when

$$\left(\frac{p_{data}(\mathbf{x})}{D_G^*(\mathbf{x})} - \frac{p_{data}(\mathbf{x})}{D_G^*(\mathbf{x}) - \lambda} \right) + \left(\frac{p_{gen}(\mathbf{x})}{D_G^*(\mathbf{x}) - 1} - \frac{p_{gen}(\mathbf{x})}{D_G^*(\mathbf{x}) - \lambda} \right) = 0 \implies D_G^*(\mathbf{x}) = \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + \frac{\lambda-1}{\lambda} \cdot p_{gen}(\mathbf{x})}.$$

Given the optimal D_G^* , the generator is to minimize the quantity in Equation (5) as

$$\begin{aligned} & \min_G \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log(D_G^*(\mathbf{x}))] + \frac{\lambda-1}{\lambda} \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D_G^*(G(\mathbf{z})))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} \left[\log \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + \frac{\lambda-1}{\lambda} \cdot p_{gen}(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim \frac{\lambda-1}{\lambda} \cdot p_{gen}(\mathbf{x})} \left[\log \frac{\frac{\lambda-1}{\lambda} \cdot p_{gen}(\mathbf{x})}{p_{data}(\mathbf{x}) + \frac{\lambda-1}{\lambda} \cdot p_{gen}(\mathbf{x})} \right] \\ &= -\log 4 + 2 \cdot JSD(p_{data} \parallel \frac{\lambda-1}{\lambda} \cdot p_{gen}). \quad \square \end{aligned}$$

Avoidance of training collapse. Besides inheriting the theoretical advantage of original GAN, CCGAN can also numerically alleviate the training collapse issue, which occurs when the generator consistently produces limited and repetitive samples, resulting in a lack of diversity and low-quality generated data. The root cause of this problem in traditional GANs lies in the pure-competition relationship between the generator and discriminator. Their interdependence can lead to a scenario where stagnation in one network hinders improvements in the other. For addressing this challenge, we argue that the regularization term introduced in Equation (3) breaks the mutual dependency between the generator and discriminator, as demonstrated in Theorem 3. This newfound independence provides the networks with additional incentive to improve even if another network is stagnant. The proof of Theorem 3 is in Appendix B.

Theorem 3 (Avoid training collapse). *Given the generator, the optimal discriminator is obtained as*

$$D^*(\mathbf{x}) = \frac{[\lambda\gamma(p_{data}(\mathbf{x}) + p_{gen}(\mathbf{x})) + \lambda(p_{data}(\mathbf{x}) - p_{gen}(\mathbf{x}))]^2 + \sqrt{4\lambda\gamma p_{data}(\mathbf{x})((1-\gamma)p_{gen}(\mathbf{x}) - \lambda p_{data}(\mathbf{x}))}}{2(1-\gamma)p_{gen}(\mathbf{x}) - 2\lambda p_{data}(\mathbf{x})}, \quad (6)$$

if the regularization adopts $\mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} \log(D(\mathbf{x}) - \lambda) + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \log(D(G(\mathbf{z})) - \gamma)$.

Overall, CCGAN offers three advantages as follows, and is summarized into Algorithm 1.

- It allows the generator and discriminator to work together to fulfill specific requirements of generated samples that lies beyond the authenticity of the “true” data distribution.
- It preserves the equilibrium point as in original GAN. The presence of an equilibrium point can guide the selection of hyperparameter λ and bolster convergence, facilitating the generation of consistently high-quality samples.
- It empowers the generator and discriminator to explore a more extensive parameter space, effectively mitigating the issue of training collapse.

4 EXPERIMENTS

In this section we evaluate the advancements of CCGAN, exploring in particular its ability of generating consistently satisfactory samples and avoiding training collapse issue. The experiments were conducted on a single NVIDIA GPU.

Algorithm 1 CCGAN: Collaboration Competition-based Generative Adversarial Network**Input:** Dataset $\mathcal{D} = \{\mathbf{x}^n\}_{n=1}^N$.**Initialize:** Hyperparameter λ , generator training gap k , number of iterations T , batch size n_0 .**for** T training iterations **do** **for** k steps **do** Sample minibatch of n_0 images from $\mathcal{D} : \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n_0)}\}$. Sample minibatch of n_0 noises from Gaussian distribution : $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(n_0)}\}$. Update the discriminator $D(\cdot)$ by ascending its stochastic gradient:

$$\nabla_D \frac{1}{n_0} \sum_{i=1}^{n_0} \left[\log D(\mathbf{x}^{(i)}) - \log [D(\mathbf{x}^{(i)}) - \lambda] + \log (1 - D(G(\mathbf{z}^{(i)}))) - \log (\lambda - D(G(\mathbf{z}^{(i)}))) \right].$$

 Update the generator $G(\cdot)$ by descending its stochastic gradient:

$$\nabla_G \frac{1}{n_0} \sum_{i=1}^{n_0} \left[\log (1 - D(G(\mathbf{z}^{(i)}))) \right]$$

Output: CCGAN with the ability to generate customized images.

4.1 SETTINGS

Datasets and Applications. To assess CCGAN’s ability in handling diverse data types, we use four image datasets and one power dataset, each with specific collaboration requirements. The image datasets cover various subjects, including humans, animals, buildings, and natural landscapes. (1) **Coloring grayscale shapes.** We start from simple objects like rectangles and circles. For this, we synthesize a dataset **Syn** containing grayscale images of standardized geometric shapes, where the task is applying suitable colors for different shapes. (2) **Coloring grayscale images.** Then, we proceed to colorize grayscale images of common objects. For this, we choose the Common Objects in Context (**COCO**) dataset Lin et al. (2014). (3) **Image super resolution.** Besides coloring, we consider the task of producing a high-resolution image from its low-resolution counterpart, using the dataset **SuperRes** Bevilacqua et al. (2012); Zeyde et al. (2012) containing life-scenes and natural landscapes. (4) **Multi-modal generation.** We explore multi-modal generation, aiming to create coherent and realistic outputs based on the input information. We use the **Facade** dataset Tyleček & Šára (2013); Isola et al. (2017) and **deepFashion** dataset Liu et al. (2016); Zhu et al. (2020b), which provides paired inputs/outputs for building facades and human clothes, respectively. (5) **False data injection attack.** Besides image data, we study power data from power systems, specifically for an adversarial attack task. It conducts imperceptible (for power utility) modifications to power measurements using the dataset **FDIA** Costilla-Enriquez & Weng (2021).

Benchmark methods. We consider baselines in GAN that implement a regularization to promote collaboration between the generator and discriminator. The following methods are utilized: (1) the original GAN (**GAN**) Goodfellow et al. (2014; 2020), (2) conditional coordinating GAN (**COCO-GAN**) Lin et al. (2019), (3) collaborative GAN (**CollaGAN**) Lee et al. (2019), and (4) cooperation GAN (**Co-GAN**) Zhang et al. (2020).

Implementing details of CCGAN. For the generator and discriminator architecture, we employ five convolutional layers and two fully-connected layers that incorporates a skip connection He et al. (2016). Each layer comprises approximately ten neurons, and the neurons are activated through pooling layers and Rectified Linear Units (ReLU) Agarap (2018). We set the number of maximal iterations to $T = 300$ for sufficient training. Additionally, for every $k = 5$ iterations will we train the generator so that we prioritize training the discriminator to allow for better convergence. For each iteration, we sample $n_0 = 50$ mini-batches to compute gradients for advanced searching for parameters. We update these parameters using the Adam optimizer with a learning rate of 2×10^{-4} .

4.2 CHALLENGES: NO PERFORMANCE GUARANTEE AND TRAINING COLLAPSE ISSUE

Prior to evaluating CCGAN’s improvements, we show two challenges that may arise in existing GAN models. First, the lack of performance guarantee may hinder a successful joint optimization of competition and collaboration in GAN. Specifically, an inconsistent regularization term $L_{\text{regularization}}(D, G; \lambda)$ can result in poor collaboration, leading to the generation of subpar outcomes.

For instance, in image super-resolution (see Figure 2 (a)) and image colorization (see Figure 2 (b)), an improper hyperparameter (λ) selection to balance competition and collaboration is shown in the third row and indicated by blue arrows. It yields unexpected and unsatisfactory results in both tasks.

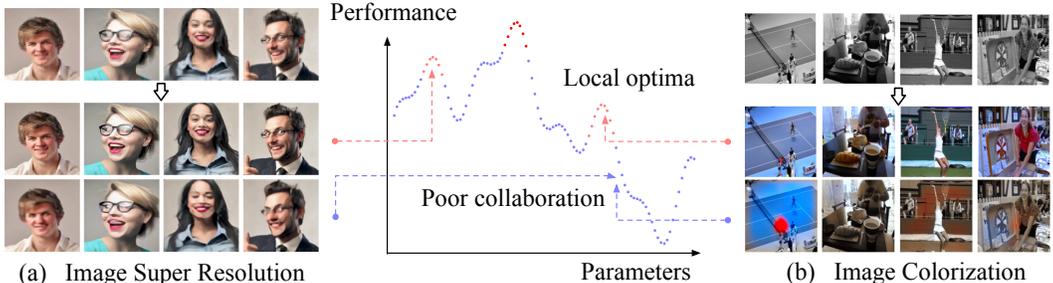


Figure 2: Two challenges in existing GANs: poor collaboration (blue arrows) and local optima (pink arrows), are simulated in the task of (a) image super resolution and (b) image colorization.

Second, in pure-competition GANs, the mutual dependence between the generator and discriminator often leads to a susceptibility to getting trapped in local optima, a phenomenon known as training collapse. This occurs when the networks reach a state of local optima and lose the incentive to further improve due to their competitive mechanism. For instance, in Figure 2, the seemingly good results in the second row might represent local optima resulting from insufficient training of the two networks, as indicated by the pink arrows.

4.3 EVALUATION OF CCGAN

This section evaluates the ability of **CCGAN** to resolve the above two challenges. For this purpose, we provide qualitative and quantitative results as follows.

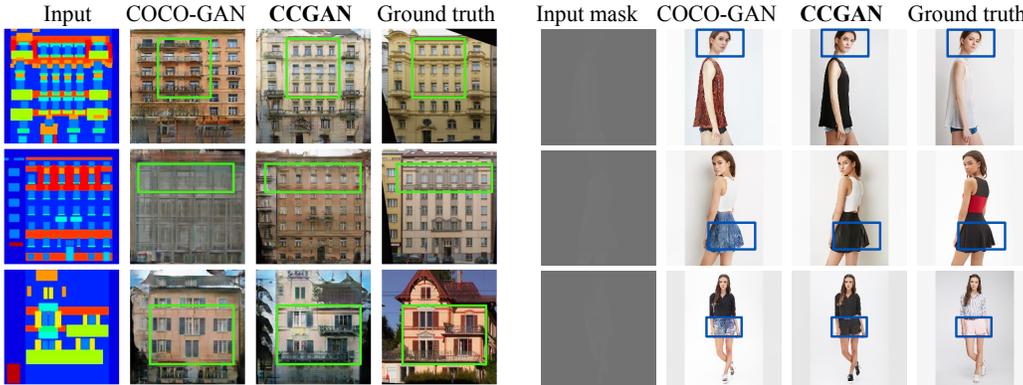
The benefits of performance guarantee. In the colorization task using the **COCO** dataset Lin et al. (2014), we verify the benefit of CCGAN’s equilibrium point result as outlined in Theorem 2. To do this, we colorize images to a predefined color spectrum, specifically the yellow and blue channels. This prior knowledge allows us to optimize the hyperparameter λ in the regularization term $L_{\text{regularization}}(D, G; \lambda)$, avoiding the need for random selection and post-training observation as in other baselines. As depicted in Figure 3, the baseline model **CollaGAN** may achieve good results after exploring the hyperparameter space, but it remains susceptible to suboptimal hyperparameter choices that lead to unexpected outcomes. In contrast, CCGAN is equipped with optimized hyperparameter due to the existence of equilibrium point. Thus, it can accurately infer and apply pre-defined colors to grayscale inputs. This eliminates the troublesome process of hyperparameter tuning and the associated risk of selecting suboptimal values.



Figure 3: Image colorization: CCGAN leads to optimal with equilibrium point in Theorem 2.

The avoidance of training collapse. To evaluate CCGAN’s efficacy in mitigating training collapse, we consider the multi-modal generation task. This task involves generating a new image that is realistic and coherent with an input image, necessitating the model to capture multiple modes within the underlying distribution. To achieve this, GAN models often incorporate regularization terms, striking a balance between fidelity and diversity. However, regularization can also lead to training collapse when applied inconsistently, resulting in unrealistic or incoherent generated images. The

results of multi-modal generation on both the **Facade** dataset Tyleček & Šára (2013); Isola et al. (2017) and **deepFashion** dataset Liu et al. (2016); Zhu et al. (2020b) are shown in Figure 4.



(a) : Image-to-image generation: buildings (b) : Image-to-image generation: human clothes
 Figure 4: The multi-modal reconstruction results of CCGAN and baseline method.

As depicted in Figure 4, the images produced by CCGAN exhibit superior clarity and coherence when contrasted with those generated by COCO-GAN. This distinction may stem from the inherent nature of the regularization loss within our design, which allows the model to explore more modes in the data distribution, as discussed in Theorem 3. Furthermore, our model benefits from the performance guarantee established in Theorem 2, allowing it to achieve more precise reconstructions with smoothly delineate image edges.

Enhanced performance of CCGAN with regular-shaped objects. After evaluating CCGAN’s capabilities, we want to identify the scenarios where CCGAN exhibits enhanced efficacy. We hypothesize that CCGAN, benefiting from its equilibrium point, thrives when processing images containing “convex” objects. These objects can be represented by convex functions within neural networks, allowing models with global optimality guarantee to excel. To test our hypothesis, we created a synthetic dataset **Syn** comprising various geometric shapes, from regular rectangles to irregular circles. We argue that rectangles, describable by multiple line segments, exhibit “convexity” more than irregular circles. In Figure 5, we compare CCGAN and Co-GAN in applying colors to these synthetic geometric shapes. The results reveal that Co-GAN struggles more with rectangles than circles. This may be attributed to inconsistent regularization, leading to unclear and jagged edges in generated images, thereby diminishing their quality and realism. In contrast, our model, benefiting from the equilibrium point, excels at discerning objects with well-defined shapes and edges.

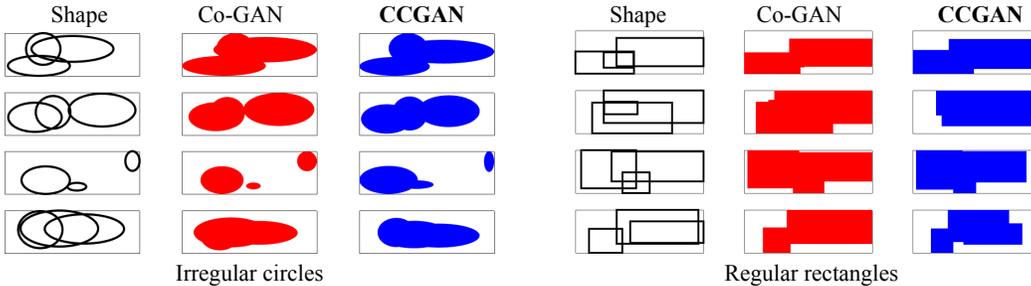


Figure 5: Image colorization on synthetic shapes.

In addition to testing this phenomenon with synthetic data, we extend our validation to real-world datasets, specifically in the context of image super-resolution task. This task often incorporate regularization into GAN models to faithfully reproduce the low-resolution input. Figure 6 presents visual results comparing CCGAN and baselines. The comparison highlights CCGAN’s ability to produce clearer, higher-resolution reconstructions. Notably, CCGAN exhibits superior performance when processing objects with regular shapes and well-defined edges. This heightened performance can be attributed to the inherent performance guarantee of the CCGAN design, which promotes greater accuracy when the depicted objects in the image exhibit convex characteristics.

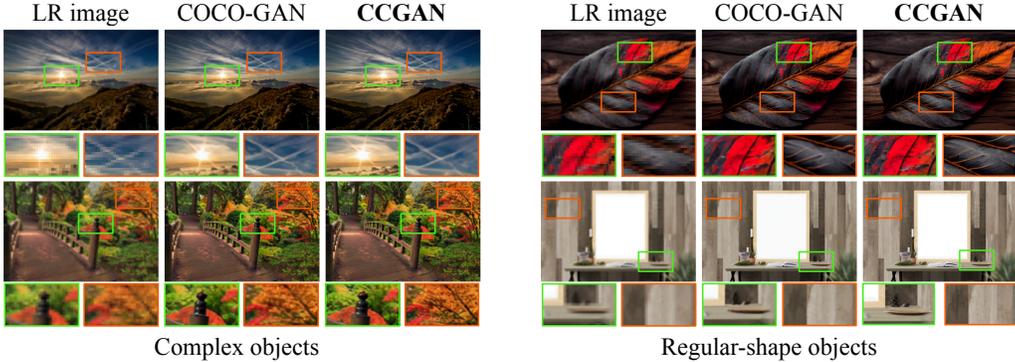


Figure 6: CCGAN and COCO-GAN reconstruction results and reference low-resolution (LR) image.

For image super-resolution task, we also evaluate the result using commonly used metrics, including peak signal-to-noise ratio (PSNR) Yang et al. (2007), structural similarity (SSIM) Wang et al. (2004) and mean opinion score (MOS) . The results in Table 1 validate our performance is better than the baselines, verifying the capability of CCGAN to generate satisfactory data catering to specific need.

Table 1: Comparison of CCGAN and baselines and the original High Resolution (SR) image.

SuperRes	COCO-GAN	CollaGAN	Co-GAN	CCGAN	HR
PSNR \uparrow	29.53	28.91	29.40	30.43	∞
SSIM \uparrow	0.8621	0.8835	0.8472	0.9011	1
MOS \uparrow	3.19	2.85	3.64	3.69	4.25

Performance of CCGAN on generating power data. Beyond image datasets, we explore power data from power systems. We consider a scenario known as the false data injection attack. This task assesses CCGAN’s ability to create deceptive power data based on real power data, capable of passing the chi-square test for power utility evaluation. Notably, power data differs from image data, representing continuous time-series data, which enriches our validation on CCGAN. Figure 7 displays CCGAN’s superior performance in generating deceptive data based on the **FDIA** dataset, as it yields a higher probability of passing the chi-square test compared to baseline models.

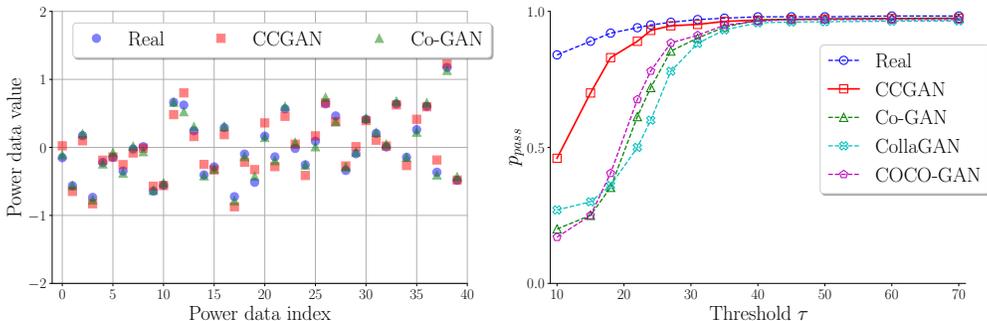


Figure 7: False data injection results from CCGAN and baselines

5 CONCLUSION

We present a technique that enables collaboration in Generative Adversarial Networks. Unlike prior work, our proposed framework preserves the performance guarantee of original GAN, and mitigates the training collapse issue. In specific, we introduced an elegant transformation from the collaboration regularization term to a distribution divergence metric, avoiding addition complexity when seeking an equilibrium point. Our work not only enhances collaboration, interpretability, and the overall optimality of GAN models, but also paves the way for exploring generative models capable of meeting additional requirements with a performance guarantee.

REFERENCES

- Lukas Jan Aaldering, Jens Leker, and Chie Hoon Song. Competition or collaboration?—analysis of technological knowledge ecosystem within the field of alternative powertrain systems: A patent-based approach. *Journal of cleaner production*, 212:362–371, 2019.
- Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.
- Duhyeon Bang and Hyunjung Shim. Mggan: Solving mode collapse using manifold-guided training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2347–2356, 2021.
- Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. 2012.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Haiyang Chen. Challenges and corresponding solutions of generative adversarial networks (gans): a survey study. In *Journal of Physics: Conference Series*, pp. 012066. IOP Publishing, 2021.
- Liqun Chen, Shuyang Dai, Chenyang Tao, Haichao Zhang, Zhe Gan, Dinghan Shen, Yizhe Zhang, Guoyin Wang, Ruiyi Zhang, and Lawrence Carin. Adversarial text generation via feature-mover’s distance. *Advances in Neural Information Processing Systems*, 31, 2018.
- Zhenfei Chen, Tianqing Zhu, Ping Xiong, Chenguang Wang, and Wei Ren. Privacy preservation for image data: a gan-based method. *International Journal of Intelligent Systems*, 36(4):1668–1685, 2021.
- Jinyoung Choi and Bohyung Han. Mcl-gan: Generative adversarial networks with multiple specialized discriminators. *Advances in Neural Information Processing Systems*, 35:29597–29609, 2022.
- Napoleon Costilla-Enriquez and Yang Weng. Exposing cyber-physical system weaknesses by implicitly learning their underlying models. In *Asian Conference on Machine Learning*, pp. 1333–1348. PMLR, 2021.
- M Durgadevi et al. Generative adversarial network (gan): a general review on different variants of gan and applications. In *2021 6th International Conference on Communication and Electronics Systems (ICCES)*, pp. 1–8. IEEE, 2021.
- Sefik Emre Eskimez, Dimitrios Dimitriadis, Robert Gmyr, and Kenichi Kumanati. Gan-based data generation for speech emotion recognition. In *INTERSPEECH*, pp. 3446–3450, 2020.
- Chun-Hua Fu, Zeng-Ping Zhang, Hui Chang, Jian-Ru Tao, Zhuo-Hui Chen, Yun-Long Dai, Wei Zhang, and Da-Ren He. A kind of collaboration–competition networks. *Physica A: Statistical Mechanics and its Applications*, 387(5-6):1411–1420, 2008.
- Chao Gong, Fuhong Lin, Xiaowen Gong, and Yueming Lu. Intelligent cooperative edge computing in internet of things. *IEEE Internet of Things Journal*, 7(10):9372–9382, 2020.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. URL <https://arxiv.org/abs/1406.2661>.

- Md Akmal Haidar and Mehdi Rezagholizadeh. Textkd-gan: Text generation using knowledge distillation and generative adversarial networks. In *Advances in Artificial Intelligence: 32nd Canadian Conference on Artificial Intelligence, Canadian AI 2019, Kingston, ON, Canada, May 28–31, 2019, Proceedings 32*, pp. 107–118. Springer, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Xun Huang, Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Multimodal conditional image synthesis with product-of-experts gans. In *European Conference on Computer Vision*, pp. 91–109. Springer, 2022.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.
- Anikó Kopacz, Lehel Csató, and Camelia Chira. Evaluating cooperative-competitive dynamics with deep q-learning. *Neurocomputing*, pp. 126507, 2023.
- Karol Kurach, Mario Lucic, Xiaohua Zhai, Marcin Michalski, and Sylvain Gelly. The gan landscape: Losses, architectures, regularization, and normalization. 2018.
- Dongwook Lee, Junyoung Kim, Won-Jin Moon, and Jong Chul Ye. Collagan: Collaborative gan for missing image data imputation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2487–2496, 2019.
- In-Won Lee, Richard C Feiock, and Youngmi Lee. Competitors and cooperators: A micro-level analysis of regional economic development collaboration networks. *Public administration review*, 72(2):253–262, 2012.
- Sunwha Lee, Sungho Lee, Yoojin Lee, Sanghoo Park, and Jinwoo Kim. Effect of competition and collaboration in social network game on intimacy among players. In *Proceedings of HCI Korea*, pp. 425–433. 2014.
- Chieh Hubert Lin, Chia-Che Chang, Yu-Sheng Chen, Da-Cheng Juan, Wei Wei, and Hwann-Tzong Chen. Coco-gan: Generation by parts via conditional coordinating. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4512–4521, 2019.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pp. 740–755. Springer, 2014.
- Linlin Liu, Haijun Zhang, and Dongliang Zhou. Clothing generation by multi-modal embedding: A compatibility matrix-regularized gan model. *Image and Vision Computing*, 107:104097, 2021.
- Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1096–1104, 2016.
- Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2794–2802, 2017.
- Vaishnavh Nagarajan and J Zico Kolter. Gradient descent gan optimization is locally stable. *Advances in neural information processing systems*, 30, 2017.
- Behnam Neyshabur, Srinadh Bhojanapalli, and Ayan Chakrabarti. Stabilizing gan training with multiple random projections. *arXiv preprint arXiv:1705.07831*, 2017.

- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, 29, 2016.
- Maarten G Poirot, Praneeth Vepakomma, Ken Chang, Jayashree Kalpathy-Cramer, Rajiv Gupta, and Ramesh Raskar. Split learning for collaborative deep learning in healthcare. *arXiv preprint arXiv:1912.12115*, 2019.
- Paul Ramond. The abel–ruffini theorem: Complex but not complicated. *The American Mathematical Monthly*, 129(3):231–245, 2022.
- Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International conference on machine learning*, pp. 1060–1069. PMLR, 2016.
- Zhenjiang Shi, Wei Gao, Shangwei Zhang, Jiajia Liu, and Nei Kato. Machine learning-enabled cooperative spectrum sensing for non-orthogonal multiple access. *IEEE Transactions on Wireless Communications*, 19(9):5692–5702, 2020.
- Tian Tan, Feng Bao, Yue Deng, Alex Jin, Qionghai Dai, and Jie Wang. Cooperative deep reinforcement learning for large-scale traffic grid signal control. *IEEE transactions on cybernetics*, 50(6): 2687–2700, 2019.
- Ngoc-Trung Tran, Tuan-Anh Bui, and Ngai-Man Cheung. Dist-gan: An improved gan using distance constraints. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 370–385, 2018.
- Soumya Tripathy, Juho Kannala, and Esa Rahtu. Icfac: Interpretable and controllable face reenactment using gans. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 3385–3394, 2020.
- Radim Tyleček and Radim Šára. Spatial pattern templates for recognition of objects with regular structure. In *Pattern Recognition: 35th German Conference, GCPR 2013, Saarbrücken, Germany, September 3-6, 2013. Proceedings 35*, pp. 364–374. Springer, 2013.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*. 1994.
- Yifan Wu, Fan Yang, Yong Xu, and Haibin Ling. Privacy-protective-gan for privacy preserving face de-identification. *Journal of Computer Science and Technology*, 34:47–60, 2019.
- Qingxiong Yang, Ruigang Yang, James Davis, and David Nistér. Spatial-depth super resolution for range images. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8. IEEE, 2007.
- Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Curves and Surfaces: 7th International Conference, Avignon, France, June 24-30, 2010, Revised Selected Papers 7*, pp. 711–730. Springer, 2012.
- Long Zhang, Jieyu Zhao, Xulun Ye, and Yu Chen. Cooperation: a new force for boosting generative adversarial nets with dual-networkstructure. *IET Image Processing*, 14(6):1073–1080, 2020.
- Jiachen Zhong, Xuanqing Liu, and Cho-Jui Hsieh. Improving the speed and quality of gan by adversarial training. *arXiv preprint arXiv:2008.03364*, 2020.
- Xining Zhu, Lin Zhang, Lijun Zhang, Xiao Liu, Ying Shen, and Shengjie Zhao. Gan-based image super-resolution with a novel quality loss. *Mathematical Problems in Engineering*, 2020:1–12, 2020a.
- Zhen Zhu, Zhiliang Xu, Ansheng You, and Xiang Bai. Semantically multi-modal image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5467–5476, 2020b.

A PROOF OF THEOREM 1

Proof. The formulation in Equation (2) can be seen as the following optimization over the generator and discriminator:

$$\min_{p_{\text{gen}}} \underbrace{\text{JS}(p_{\text{data}}, p_{\text{gen}})}_{\text{learn "true" distribution}} + \lambda \cdot \underbrace{L_{\text{regularization}}(p_{\text{need}}, p_{\text{gen}})}_{\text{fulfill "other" requirements}}, \quad (7)$$

where the original loss $V_{\text{GAN}}(D, G)$ is equivalent Goodfellow et al. (2014) to the Jensen–Shannon (JS) divergence $\text{JS}(p_{\text{data}}||p_{\text{gen}})$ between true data distribution p_{data} and generated data distribution p_{gen} . Minimizing the JS divergence $\text{JS}(p_{\text{data}}||p_{\text{gen}})$ can make learned distribution p_{gen} match the true distribution p_{gen} . Minimizing the $L_{\text{regularization}}$ fulfills specific need p_{need} beyond the authenticity of the data distribution. Under the assumption of Gaussianity, we denote $p_{\text{data}} = \mathcal{N}(\mu_1, \sigma_1^2)$, $p_{\text{need}} = \mathcal{N}(\mu_2, \sigma_2^2)$ and $p_{\text{gen}} = \mathcal{N}(\mu, \sigma^2)$. Then, Equation (7) reads as

$$\begin{aligned} & \min_{\mu, \sigma} \text{JS}(\mathcal{N}(\mu_1, \sigma_1^2), \mathcal{N}(\mu, \sigma^2)) + \lambda \cdot L_{\text{regularization}}(\mathcal{N}(\mu_2, \sigma_2^2), \mathcal{N}(\mu, \sigma^2)) \\ & = \min_{\mu, \sigma} \log \frac{\sigma_1^2 + \sigma^2}{2\sigma_1\sigma} + \frac{(\mu_1 - \mu)^2}{\sigma_1^2 + \sigma^2} + \lambda \cdot (\mu_2 - \mu)^2 + \lambda \cdot (\sigma_2^2 - \sigma^2)^2. \end{aligned} \quad (8)$$

For deriving the equilibrium point of (μ, σ) , we assign corresponding derivative for μ to zero as

$$\frac{2(\mu_1 - \mu)}{\sigma_1^2 + \sigma^2} + 2\lambda \cdot (\mu_2 - \mu) = 0 \Rightarrow \mu = \frac{\mu_1 + \lambda(\sigma_1^2 + \sigma^2)\mu_2}{1 + \lambda(\sigma_1^2 + \sigma^2)}. \quad (9)$$

Taking the above result into Equation (8) and assign derivative for σ to zero, we have

$$\frac{\sigma^2 - \sigma_1^2}{(\sigma^2 + \sigma_1^2)\sigma} - 4\lambda\sigma(\sigma_2^2 - \sigma^2) + \lambda(\mu_2 - \mu)^2 \frac{-2\lambda\sigma}{[1 + (\sigma^2 + \sigma_1^2)]^2} = 0, \quad (10)$$

which is a polynomial equation of degree higher than five, where an analytic (closed-form) solution is absent Ramond (2022). \square

B PROOF OF THEOREM 3

Proof. The training criterion for the discriminator D , given a fixed generator G , is to maximize the quantity $V_{\text{GAN}}(D, G) - L_{\text{regularization}}(D, G; \lambda)$ as

$$\begin{aligned} & \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log(D(\mathbf{x})) - \log(D(\mathbf{x}) - \lambda)] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z}))) - \log(\gamma - D(G(\mathbf{z})))] \\ & = \int_{\mathbf{x}} p_{\text{data}}(\mathbf{x}) [\log(D(\mathbf{x})) - \log(D(\mathbf{x}) - \lambda)] d\mathbf{x} + \int_{\mathbf{z}} p_{\mathbf{z}}(\mathbf{z}) [\log(1 - D(G(\mathbf{z}))) - \log(\gamma - D(G(\mathbf{z})))] \\ & = \int_{\mathbf{x}} p_{\text{data}}(\mathbf{x}) [\log(D(\mathbf{x})) - \log(D(\mathbf{x}) - \lambda)] d\mathbf{x} + \int_{\mathbf{x}} p_{\text{gen}}(\mathbf{x}) [\log(1 - D(\mathbf{x})) - \log(\gamma - D(\mathbf{x}))]. \end{aligned}$$

Then, the optimal discriminator D_G^* is obtained when

$$\begin{aligned} & \left(\frac{p_{\text{data}}(\mathbf{x})}{D_G^*(\mathbf{x})} - \frac{p_{\text{data}}(\mathbf{x})}{D_G^*(\mathbf{x}) - \lambda} \right) + \left(\frac{p_{\text{gen}}(\mathbf{x})}{D_G^*(\mathbf{x}) - 1} - \frac{p_{\text{gen}}(\mathbf{x})}{D_G^*(\mathbf{x}) - \gamma} \right) = 0 \\ & \Rightarrow -\lambda p_{\text{data}}(\mathbf{x})(D_G^*(\mathbf{x}) - 1)(D_G^*(\mathbf{x}) - \gamma) + (1 - \gamma)p_{\text{gen}}(\mathbf{x})D_G^*(\mathbf{x})(D_G^*(\mathbf{x}) - \lambda) \\ & \Rightarrow D_G^*(\mathbf{x}) = \frac{1}{2} \frac{[\lambda\gamma(p_{\text{data}}(\mathbf{x}) + p_{\text{gen}}(\mathbf{x})) + \lambda(p_{\text{data}}(\mathbf{x}) - p_{\text{gen}}(\mathbf{x}))]^2 + \sqrt{4\lambda\gamma p_{\text{data}}(\mathbf{x})((1 - \gamma)p_{\text{gen}}(\mathbf{x}) - \lambda p_{\text{data}}(\mathbf{x}))}}{(1 - \gamma)p_{\text{gen}}(\mathbf{x}) - \lambda p_{\text{data}}(\mathbf{x})}. \end{aligned}$$

\square