GyF

This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Style Quantization for Data-Efficient GAN Training

Jian Wang; Xin Lan^{*}, Jizhe Zhou, Yuxin Tian, Jiancheng Lv[†] College of Computer Science, Sichuan University Engineering Research Center of Machine Learning and Industry Intelligence, Ministry of Education

Abstract

Under limited data setting, GANs often struggle to navigate and effectively exploit the input latent space. Consequently, images generated from adjacent variables in a sparse input latent space may exhibit significant discrepancies in realism, leading to suboptimal consistency regularization (CR) outcomes. To address this, we propose SQ-GAN, a novel approach that enhances CR by introducing a style space quantization scheme. This method transforms the sparse, continuous input latent space into a compact, structured discrete proxy space, allowing each element to correspond to a specific real data point, thereby improving CR performance. Instead of direct quantization, we first map the input latent variables into a less entangled "style" space and apply quantization using a learnable codebook. This enables each quantized code to control distinct factors of variation. Additionally, we optimize the optimal transport distance to align the codebook codes with features extracted from the training data by a foundation model, embedding external knowledge into the codebook and establishing a semantically rich vocabulary that properly describes the training dataset. Extensive experiments demonstrate significant improvements in both discriminator robustness and generation quality with our method.

1. Introduction

Generative models have made significant strides in creating human-like natural language [7], high-quality images [18, 29, 40], and diverse human speech [52] and music [17]. A key factor behind these achievements is the extensive amount of training data required to capture the distribution of real data. However, data collection can be challenging due to issues related to subject matter, privacy, or copyright. Consequently, training generative models with limited data has garnered increasing attention from the research community.



[†]Corresponding author.



Figure 1. Illustration of the proposed style quantization GAN (*SQ-GAN*) framework. (I) The input latent variables z are mapped to an intermediate latent space W, which is quantized into a compact and structured proxy space W^q by a learnable codebook C. The quantized codes $w^q \in W^q$ are then fed into the synthesis network to generate images. (II) The codebook initialization aligns the codebook codes with features extracted from the training data, embedding external knowledge into the codebook.

Recently, numerous studies [13, 14, 19, 26, 30, 32, 51, 60] have utilized Generative adversarial networks (GANs) to train models with limited image samples, showcasing a promising future. A core issue they aim to address is the discriminator's overfitting problem when trained with limited data. The most direct solution is augmentation, including traditional [50, 61], differentiable [60], and adaptive [30] methods to augment real or generated data. However, these methods require customized manual design for different datasets, lack scalability, and limit augmentation types to a finite range [30]. Additionally, augmentations may distort the original image semantics (see Fig. 4a), leading to a decline in generation quality. As a complement, model regularization techniques address discriminator degradation by enhancing robustness [14, 32, 38, 58, 62] and diversity [13, 14, 32] of the discriminator's representations, stabilizing learning dynamics during training [19, 51]. Among these, robust representations of the discriminator are paramount.

Consistency regularization (CR) methods [58, 62] offer an alternative perspective to encourage the discriminator to capture more robust features. These methods introduce an auxiliary term that compels the discriminator to produce consistent scores for input images (both generated and real) and their augmented versions. Particularly, Zhao et al. [62] propose generating samples from two adjacent latent variables, z and its perturbed version $z + \epsilon$ in the prior space \mathcal{Z} , using them as implicit augmentations to enforce consistency. This strategy circumvents the limitations of manual augmentation design and achieves impressive results. However, under limited data conditions, the effective exploitation of \mathcal{Z} becomes challenging [14, 32]. This inadequacy may result in significant disparities in authenticity between images generated from z and $z + \epsilon$, which, if forced to be evaluated consistently by the discriminator, could impair its effectiveness.

Our analysis highlights a key problem: how to effectively exploit the prior space \mathcal{Z} with limited data such that two adjacent latent variables, z and $z + \epsilon$, generate diverse yet equally authentic images. The challenge resides in the inherent conflict between limited data and the comprehensive exploration of \mathcal{Z} . Drawing inspiration from vector quantization (VQ) techniques [18, 44, 53, 63], we consider compressing the representation of latent variables into a more compact and structured discrete space making each encoding more semantically rich and explicit, thus facilitating subsequent sampling to generate more diverse and consistent images.

In this paper, we introduce a style space quantization technique (SQ-GAN) aimed at optimizing the consistency regularization of the discriminator, enhancing the robustness of its extracted features, and mitigating GAN training degradation in scenarios with limited samples. As illustrated in Fig. 1, we embed the input latent variables zinto an intermediate latent space W, akin to StyleGAN [29]. The intermediate variable $w \in \mathcal{W}$ is typically interpreted as a less entangled "style," governing various high-level attributes of the generated images [29]. We segment winto s sub-vectors, $\{\hat{w}_i\}_{i=1}^s$, each corresponding to different aspects or subsets of attributes. These segmented latent variables are then quantized to their nearest neighbors in a learnable codebook, enabling each quantized code \hat{w}_i^q to control distinct factors of variation. For image generation, the quantized codes are concatenated (w^q = $[\hat{w}_1^q, \hat{w}_2^q, \ldots] \in \mathcal{W}^q$) and fed into the synthesis network. In a limited data setting, exploring this smaller, more compact proxy space \mathcal{W}^q is more feasible than directly exploring \mathcal{W} . This design effectively compels the intermediate latent variables w to be recombined from a limited set of distinct discrete sub-vectors \hat{w}_a , allowing the generator to assign explicit meanings to each sub-vector.

Furthermore, we introduce a novel codebook initialization (CBI) method, which is crucial for capturing the intrinsic semantic structure of the training data. This method maps the codebook's codes and images from the training dataset to feature space using a foundation model, such as CLIP [42]. Then, it aligns these two sets of features using optimal transport distance, initializing the codebook in a manner that abstractly describes the training dataset. This alignment embeds external knowledge into the codebook, pre-establishing a semantically rich vocabulary of the training dataset. Our contributions are fourfold:

- As far as we know, this is the first work to introduce style space quantization to enhance GAN training with limited data.
- A novel method for quantizing in the intermediate latent space to enhance discriminator consistency regularization.
- An innovative codebook initialization technique that incorporates external knowledge, pre-establishing a semantically rich vocabulary for the codebook.
- Comprehensive experiments demonstrating the effectiveness of our approach.

2. Related work

2.1. Advancements in GANs

Generative adversarial networks (GANs) [20] are a groundbreaking framework designed to minimize the Jensen-Shannon (JS) divergence between real and synthetic data distributions. Since its inception, extensive research has extended this framework by exploring a variety of adversarial loss functions that target different divergences or distances, such as the chi-squared (χ^2) divergence [34], total variation distance [59], and integral probability metrics (IPMs) [3, 47]. These efforts have aimed to enhance the stability and convergence of GAN training, addressing the initial challenges associated with the original GAN framework. The evolution of GANs has led to substantial enhancements in the quality and diversity of generated images. These advancements are attributed not only to the refinement of objective functions but also to the development of sophisticated neural network architectures [36, 37, 57] and robust training methodologies [16, 28, 33, 56]. Notable among these developments are BigGAN [6] and Style-GAN [29, 31], which have set new benchmarks by producing high-resolution images with intricate details and varied styles.

2.2. GAN Training with limited data

Training GANs with limited data poses significant challenges, primarily due to the discriminator's tendency to overfit, which can degrade the quality of generated samples [23, 55]. Data augmentation [30, 49, 58, 60–62] has emerged as a solution by enriching the dataset through transformations that diversify the original data. However, applying augmentation to GANs requires careful consider-

ation to avoid altering the target distribution or introducing artifacts. Recent research has introduced GAN-specific augmentation methods, such as differentiable [60], adaptive [30], and generative [61] augmentations, which address these challenges effectively.

In addition to data augmentation, model regularization offers another strategy to enhance GAN training. Various regularization techniques have been explored, including noise injection [2, 25, 46], gradient penalties [22, 35], spectral normalization [36, 37], and consistency regularization [58, 62]. Recent innovations, such as the LC regularization term [51], which modulates the discriminator's evaluations using two exponential moving average variables and connects to a type of f-divergence known as LeCam divergence [8], have proven beneficial. Additionally, Dig-GAN [19] addresses the gradient discrepancy between real and synthetic images by reducing this gradient gap, thus enhancing GAN performance. Furthermore, studies have explored various methods to introduce external knowledge into the discriminator, such as using pre-trained models as additional discriminators [32], aligning the discriminator's features with those of a pre-trained model through knowledge distillation [14], or fixing the highest resolution layer of the discriminator by pre-training its parameters on a larger dataset [38].

2.3. Codebook learning for image generation

Recent works have focused on image quantization and synthesis through codebook learning. The VQ-VAE model [53] and its successor, VQ-VAE-2 [44], enabled high-resolution image generation by capturing latent representations using discrete symbols, which facilitate efficient indexing and retrieval. Chen et al. [9] improved quality and efficiency within these codebook constraints, while VQ-GAN [18] further enhanced image quality by integrating adversarial and perceptual losses. Building on these foundations, DALL-E [43] showcased the effectiveness of discrete representations in generating images from text. The VQ-Diffusion Model [21] combined vector quantization with diffusion processes for superior results.

In contrast to these approaches, we introduce a novel style space quantization scheme specifically designed for GAN training with limited data. This method leverages the compact and structured properties of codebooks to effectively capture and exploit the intrinsic semantic structures present in small datasets.

3. Preliminary

3.1. Generative adversarial networks

Generative adversarial networks (GANs) [20] consist of two core components: the generator g and the discriminator f_D . The generator synthesizes data samples by transforming a random noise vector drawn from a prior distribution p_z (typically $\mathcal{N}(\mathbf{0}, \mathbf{1})$). The discriminator, on the other hand, evaluates the real samples and these generated samples to assess their authenticity. The performance of the generator and discriminator is optimized through the following loss functions:

$$\mathcal{L}_{adv}(g) = \mathbb{E}[1 - \log f_D(g(\boldsymbol{z}))], \qquad (1)$$

$$\mathcal{L}_{adv}(f_D) = -\mathbb{E}[\log f_D(\boldsymbol{x})] - \mathbb{E}[\log(1 - f_D(g(\boldsymbol{z})))], (2)$$

where $x \in \mathbb{R}^{h \times w \times c}$ represents real data, and h, w, and c denote the height, width, and channel dimensions, respectively.

3.2. Consistency regularization

Recent studies [58] have introduced a latent consistency regularization (CR) term to ensure that the discriminator's predictions for two images generated from a latent variable z and its perturbed version $z + \epsilon$ remain consistent:

$$\mathcal{L}_{\rm cr}(f_D) = \lambda_{f_D} \mathbb{E}\left[\|f_D(g(\boldsymbol{z})) - f_D(g(\boldsymbol{z} + \boldsymbol{\epsilon}))\|^2 \right].$$
(3)

To enhance the generator's sensitivity to variations in the prior distribution, thereby increasing the diversity of generated samples, CR additionally imposes a constraint on the generator:

$$\mathcal{L}_{\rm cr}(g) = -\lambda_g \mathbb{E}\left[\|g(\boldsymbol{z}) - g(\boldsymbol{z} + \boldsymbol{\epsilon})\|^2 \right],\tag{4}$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma \mathbf{I})$, σ represents the perturbation strength, set to $\sigma = 0.1$ in the experiments, and λ_{f_D} and λ_g are hyperparameters that balance the contributions of the respective terms.

3.3. Challenges in CR

Conceptually, reducing the discriminator's sensitivity to perturbations in the input latent variable to promote invariance to distortions can enable the extraction of more robust features. However, under limited data conditions, the exploration of the generator's input latent space may be insufficient. This insufficiency could result in a discontinuous mapping from latent variables to data points, leaving portions of the latent space underpopulated. Consequently, it is not guaranteed that adjacent latent variables z and $z + \epsilon$ will consistently map to corresponding data points, potentially degrading the effectiveness of CR.

As shown in Fig. 4b, the sparse nature of the input latent space may lead to significant differences in the realism of generated images g(z) and $g(z + \epsilon)$. In such cases, if the discriminator is forced to evaluate these images with consistent criteria, it may encounter difficulties, leading to suboptimal results in CR. To mitigate this issue, we propose a novel approach that incorporates intermediate latent space quantization.



Figure 2. The overall framework of SQ-GAN. (a) **Style quantization.** For a batch of intermediate variables w, we apply an intermediate latent space quantization technique, segmenting and quantizing them using a hyperspherical codebook. The concatenated discrete codes are directly fed into the generator for image synthesis. (b) **Knowledge-enhanced codebook initialization.** We perform an alignment strategy grounded in optimal transport distance to embed semantic knowledge from foundation models into the codebook. The transformation and transmission of latent variables and features occur within the corresponding latent space, as illustrated at the **top** of the figure. Our framework constructs a vocabulary-rich codebook, which ensures that the entries within the codebook adequately represent a diverse and compact set of image features, suitable for image generation in limited data scenarios.

4. Style Quantization

Our goal is to provide GANs with a more compact and structured input latent space, enabling more feasible and effective exploitation of the input latent space under limited data conditions, thereby enhancing the performance of the consistency regularization (CR). We propose an intermediate latent space quantization technique that segments the intermediate latent variables and quantizes them into a learnable codebook. To ensure each codebook entry explicitly controls different factors of variation, we introduce uniformity constraints. Finally, we present a novel quantizationbased CR method.

4.1. Style space

In the generator architecture, the input latent variable $z \in \mathcal{Z}$ is mapped to a less entangled intermediate latent space \mathcal{W} using a non-linear mapping function $f_{\mathcal{W}} : \mathbb{R}^{d_z} \to \mathbb{R}^{d_w}$, akin to StyleGAN [29]. Here, d_z and d_w denote the dimensions of the input latent variable z and the mapped intermediate latent variable $w \in \mathcal{W}$, respectively. The intermediate latent variable w influences each layer of the subsequent synthesis network g via affine transformations, effectively functioning as "style" to control various high-level attributes of the generated images. Therefore, we term the intermediate latent space \mathcal{W} as the *style space*.

4.2. Proxy space

To establish an ideal mapping between the style space W and the data points, we propose a novel style space quanti-

zation method, termed style quantization GAN (*SQ-GAN*), to construct a compact, structured, and discrete *proxy space* W^q that better fits limited datasets.

As illustrated in Fig. 2a, to achieve finer-grained control over the feature space and capture more detailed information, we partition each intermediate latent variable w into s sub-vectors $\{\hat{w}_i\}_{i=1}^s \subset \mathbb{R}^{s \times (d_w/s)}$. We set $d_w/s = 4$ to balance the trade-off between the number of partitions and image quality. Each sub-vector \hat{w}_i is optimized to manage specific details or attributes. Subsequently, we quantize each sub-vector \hat{w}_i to the nearest item in a learnable codebook C, resulting in a quantized sub-vector \hat{w}_i^q :

$$\hat{\boldsymbol{w}}_{i}^{q} = \left(\arg\min_{\boldsymbol{c}_{j}\in\mathcal{C}} \|\hat{\boldsymbol{w}}_{i} - \boldsymbol{c}_{j}\|\right) \in \mathbb{R}^{(d_{\boldsymbol{w}}/s)}, \qquad (5)$$

where c_j is the *j*-th entry in the codebook $C \subset \mathbb{R}^{k \times (d_w/s)}$, and *k* denotes the number of codebook entries. This approach represents each original w_i as a set of *s* quantized discrete representations $\{\hat{w}_i^q\}_{i=1}^s$. Finally, the synthesis network *g* generates a high-quality image $\tilde{x} = g(\hat{w}^q)$ from the latent code $w^q \in W^q$, formed by concatenating the quantized codes $\{\hat{w}_i^q\}_{i=1}^s$:

$$\boldsymbol{w}^{q} = [\hat{\boldsymbol{w}}_{1}^{q}, \hat{\boldsymbol{w}}_{2}^{q}, \dots, \hat{\boldsymbol{w}}_{s}^{q}] \in \mathbb{R}^{d_{\boldsymbol{w}}}.$$
 (6)

Thus, the proxy space W^q is defined as the Cartesian product of the *s* codebooks, $W^q = C^1 \times C^2 \times \ldots \times C^s$.

4.3. Loss function

We define the following loss functions to train our model end-to-end:

$$\mathcal{L}(g, f_{\mathcal{W}}, \mathcal{C}, \mathbf{P}) = \mathcal{L}_{adv}(g) + \lambda_{sq}(\mathcal{L}_{sq}(f_{\mathcal{W}}, \mathcal{C}) + \mathcal{L}_{uf}(\mathcal{C}, \mathbf{P})),$$
(7)
$$\mathcal{L}(f_D) = \mathcal{L}_{adv}(f_D) + \lambda_{qcr}\mathcal{L}_{qcr}(f_D),$$
(8)

where λ_{sq} and λ_{qcr} are hyperparameters that balance the contributions of the corresponding terms. In our experiments, we set $\lambda_{sq} = 0.01$ and $\lambda_{qcr} = 0.01$.

Adversarial losses $\mathcal{L}_{adv}(g)$ and $\mathcal{L}_{adv}(f_D)$ The adversarial losses, defined in Eq. (1) and Eq. (2), are designed to optimize the generator g and the discriminator f_D .

Style quantization loss $\mathcal{L}_{sq}(g)$ The style quantization loss [53] optimizes the distance between the split latent variable \hat{w} and the corresponding codebook representation C:

$$\mathcal{L}_{\mathrm{sq}}(f_{\mathcal{W}}, \mathcal{C}) = \|\mathrm{sg}\left(\hat{\boldsymbol{w}}\right) - \boldsymbol{c}\|_{2}^{2} + \beta \|\hat{\boldsymbol{w}} - \mathrm{sg}\left(\boldsymbol{c}\right)\|_{2}^{2}, \quad (9)$$

where $sg(\cdot)$ is the stop-gradient operator, and the second term, scaled by β , is a so-called "commitment loss" that prevents the learned codes in C from growing arbitrarily. To handle the non-differentiable operation in Eq. (9), we employ the straight-through gradient estimator, which allows gradients from the synthesis network g to flow directly to the mapping network f_W during backpropagation.

Uniformity regularization $\mathcal{L}_{uf}(\mathcal{C}, P)$ To prevent codebook collapse during limited data training, we impose a uniformity constraint on the codebook \mathcal{C} . We first project the codebook \mathcal{C} onto a unit hypersphere using a learnable projection matrix P and l_2 normalization, such that $\bar{c}_i = Pc_i/||Pc_i||_2$. The uniformity problem on the unit hypersphere is characterized by minimizing the total pairwise potential based on a specific kernel function [5, 12]. Following [54], we employ the Radial Basis Function (RBF) kernel, which can encapsulate a broad spectrum of kernel functions [11]:

$$\mathcal{L}_{\rm uf}(\mathcal{C}, \boldsymbol{P}) = \log \mathbb{E}\left[\exp\left(-t\|\bar{\boldsymbol{c}}_i - \bar{\boldsymbol{c}}_j\|_2^2\right)\right],\tag{10}$$

where t > 0 is a fixed scale parameter. This regularization promotes uniform distribution of feature codes on the unit hypersphere, ensuring diverse and mean-explicit codes within the codebook throughout training.

Quantization-based consistency regularization \mathcal{L}_{qcr} In quantization-based CR, we enforce consistency in the discriminator's predictions for images generated from quantized intermediate latent variables w^q and their perturbed

versions. The perturbed intermediate latent variable is obtained by adding Gaussian noise ϵ (same as CR) to the original input latent variable z, then mapping it to the intermediate latent variable $w' = f_{\mathcal{W}}(z + \epsilon)$, and quantizing it via the codebook C to get w'^q . The quantization-based CR loss is defined as:

$$\mathcal{L}_{qcr}(f_D) = \mathbb{E}\left[\|f_D(g(\boldsymbol{w}^q)) - f_D(g(\boldsymbol{w}^{\prime q}))\|^2 \right].$$
(11)

This form of regularization ensures that even minor perturbations in the latent space, leading to different quantized representations, result in consistent evaluations by the discriminator.

5. Knowledge-Enhanced Codebook Initialization

Our approach aims to ensure that the codes within the codebook properly represent various image features, subsequently using their combinations to correspond to specific data points. However, deriving such codes from a limited training dataset is a non-trivial task. To address this, we propose an alignment strategy based on optimal transport distance for codebook initialization (CBI). This method leverages the semantic knowledge embedded in foundation models to construct a richer vocabulary for the training dataset.

5.1. Feature extraction

As depicted in Fig. 2b, we begin by utilizing a foundation model to encode both the images in the dataset and the codes (quantized sub-vectors) in the codebook into their respective feature spaces. Given that the initialization of the codebook is essentially a pairing task, we leverage CLIP [42], a widely adopted large-scale vision-language pre-trained model. The extensive pretraining of CLIP enables the exploration of semantic correspondences between visual and linguistic modalities.

Image encoding In the image encoding process, we directly employ CLIP's visual encoder to extract image features. Specifically, given a batch of *n* real images $\{x_i\}_{i=1}^n$ from a limited dataset, we extract the features for each image, $f_i = \phi(x_i) \in \mathbb{R}^{l \times d_e}$, using CLIP's visual encoder ϕ . Here, *l* represents the image token sequence length, and d_e denotes the embedding dimension of the image feature. Consequently, we obtain a set of image features $F = \{f_i\}_{i=1}^n \subset \mathbb{R}^{n \times l \times d_e}$.

Code embedding Consider *m* latent variables $\{z_i\}_{i=1}^m$ sampled from the prior distribution. These are first mapped to the style space \mathcal{W} and then quantized into the discrete proxy space \mathcal{W}^q . To align with the Transformer architecture, the quantized sub-vectors $\{\hat{w}_i^q\}_{i=1}^m \subset \mathbb{R}^{m \times s \times (d_w/s)}$ are passed through a trainable MLP layer, transforming

them into variables with the same embedding dimension as the input tokens. These transformed variables are then processed by a specific Transformer structure. In our implementation, we utilize CLIP's text encoder, ψ , which offers a critical advantage: the features derived from both the visual and text encoders are already closely aligned, thereby reducing feature discrepancy and accelerating training convergence. Thus, we map the *m* quantized discrete codes $\{\hat{w}_i^q\}_{i=1}^m$ to the feature space, resulting in the feature set $T = \{t_i\}_{i=1}^m \subset \mathbb{R}^{m \times s \times d_e}$.

It is important to highlight that we treat these discrete codes as analogous to text encodings composed of words from a vocabulary. Consequently, the initialized codebook can be viewed as a vocabulary that abstractly describes the image features within the dataset.

5.2. Aligning via optimal transport distance

Given that the images and latent variables in each batch are unpaired, and considering the differences in semantic structure and interpretation between the image feature F and the code embedding T, we employ the optimal transport (OT) distance as an effective metric to quantify their discrepancy.

Optimal transport distance Consider two distributions of features generated by a neural network: $\{\mathbf{f}_i^x\}_{i=1}^n$, with corresponding weights (marginal distribution) $\mathbf{p} = \{p_i\}_{i=1}^n$, and $\{\mathbf{f}_j^y\}_{j=1}^m$, with weights $\mathbf{q} = \{q_j\}_{j=1}^m$. The objective is to determine a transport plan $\gamma \in \mathbb{R}^{n \times m}$ that minimizes the total transportation cost of mapping one feature distribution to the other. The OT problem is formally defined as:

$$D(\mathbf{p}, \mathbf{q} | \mathbf{C}) = \min_{\gamma \in \mathbb{R}^{n \times m}} \langle \gamma, \mathbf{C} \rangle$$

$$= \min_{\gamma \in \mathbb{R}^{n \times m}} \sum_{i=1}^{n} \sum_{j=1}^{m} \gamma_{ij} c(\mathbf{f}_{i}^{x}, \mathbf{f}_{j}^{y}),$$

subject to
$$\sum_{j=1}^{m} \gamma_{ij} = p_{i} \quad \forall i \in \{1, \dots, n\},$$

$$\sum_{i=1}^{n} \gamma_{ij} = q_{j} \quad \forall j \in \{1, \dots, m\},$$

$$(12)$$

where $c(\mathbf{f}_i^x, \mathbf{f}_j^y)$ represents the cost of transporting a unit of mass from feature \mathbf{f}_i^x to \mathbf{f}_j^y , and $\langle \gamma, \mathbf{C} \rangle$ denotes the Frobenius inner product of γ and the cost matrix \mathbf{C} . Directly solving this problem is computationally intensive, especially in high-dimensional feature spaces. To mitigate this, we employ the Sinkhorn divergence algorithm [15], which incorporates an entropic regularizer to accelerate computation. Detailed derivations can be found in Appendix A.

Alignment Given the marginal distributions **p** and **q** of discrete code embeddings $T = \{t_i\}_{i=1}^m \subset \mathbb{R}^{m \times s \times d_e}$ and

visual features $\mathbf{F} = {\{\mathbf{f}_i\}_{i=1}^n \subset \mathbb{R}^{n \times l \times d_e}}$, alongside the corresponding cost matrix \mathbf{C} , the optimal transport plan γ^* is obtained by solving equation (12). For simplicity, we initialize \mathbf{p} and \mathbf{q} as uniform distributions, i.e.,

$$p_i = \frac{1}{s}, \quad \forall i \in \{1, \dots, s\}, \quad q_j = \frac{1}{l}, \quad \forall j \in \{1, \dots, l\}.$$
(13)

The cost matrix **C** is computed by:

$$C_{j,k} = c(j,k) := d(t_{j}, f_{k}) \quad 1 \le j \le s, \quad 1 \le k \le l,$$
(14)

where $d(\cdot, \cdot)$ denotes the distance metric (e.g., Euclidean distance or cosine distance) serving as the transport cost function c between corresponding components of feature t_i and f_i .

After computing γ^* , we define the optimal transport loss between the two features, T and F, to achieve the alignment:

$$\mathcal{L}_{\text{ot}}(f_{\mathcal{W}}, \mathcal{C}) = \mathbb{E}\left[d(\boldsymbol{T}, \boldsymbol{F}) \cdot \gamma^*\right].$$
(15)

During the codebook initialization phase, we optimize the weighted sum of the OT loss \mathcal{L}_{ot} along with the two additional loss terms (9) and (10). The overall loss function is defined as:

$$\mathcal{L}(f_{\mathcal{W}}, \mathcal{C}, \mathbf{P}) = \mathcal{L}_{sq}(f_{\mathcal{W}}, \mathcal{C}) + \mathcal{L}_{uf}(\mathcal{C}, \mathbf{P}) + \mathcal{L}_{ot}(f_{\mathcal{W}}, \mathcal{C}).$$
(16)

6. Experiments

Experimental settings To evaluate the performance of our proposed SQ-GAN, we conduct experiments on four widely used datasets in the image generation under limiteddata setting research field: Oxford-Dog (from Oxford-IIIT pet dataset [39], detailed in Appendix B), Flickr-Faces-HQ (FFHQ) [29], MetFaces [30], and BreCa-HAD [1]. The datasets are resized to 256×256 pixels for training. Detailed descriptions of these datasets are provided in Appendix B.

We adopt the evaluation metrics: Inception Score (IS) [45], Fréchet Inception Distance (FID) [24], and Kernel Inception Distance (KID) [4] to evaluate our models. The official implementations of these metrics, as provided by [31], are used for all evaluations. For the formal definition of the metrics, see Appendix B.

6.1. Comparison with SOTA

We conducted an extensive evaluation of our proposed SQ-GAN model against several state-of-the-art models under diverse experimental conditions. To ensure the validity of our comparisons, we used StyleGAN2 [31] as the foundational architecture, retraining all models from scratch across all datasets. All metrics were computed using the same tool, ensuring consistency in measurement. The results of this comprehensive comparison are presented in Table 1

Mathad	Oxfore	d-Dog	FFHQ-2.5k					
Method	FID ↓	IS ↑	$FID \Downarrow$	IS ↑				
StyleGAN2 [31]	64.26	48.11 3.50						
GAN loss functions								
+ Wasserstein [22]	82.18	9.94	38.64	4.11				
+ LS [34]	216.42	2.69	213.93	2.17				
+ RaHinge [27]	38.68	11.14	32.34	3.96				
Regularization act on discriminator								
+ LeCam [51]	102.87	8.02	68.85	3.53				
+ DigGAN [19]	61.84	10.07	48.45	3.68				
+ KD-DLGAN [14]	54.06	9.73	49.31	3.67				
+ CR [58]	48.73	10.47	41.43	4.06				
+ SQ-GAN (ours)	36.30	11.52	25.38	4.17				
+ SO-GAN + CBI (ours)	35.01	12.44	22.04	4.20				

Table 1. Comprehensive comparison with SOTA methods across various settings on the Oxford-Dog and FFHQ-2.5K datasets, implemented within the StyleGAN framework, as they are all architecture-agnostic. CBI denotes our knowledge-enhanced codebook initialization process. ↑ denotes that higher values are preferable, while U indicates that lower values are better.



Figure 3. Evolutions of FID (\Downarrow) scores during training on the Oxford-Dog and FFHQ-2.5K datasets.

and Fig. 3. Our baseline includes a variety of GAN loss functions, architectures, and regularization techniques. The findings demonstrate that our approach consistently outperforms the benchmarks across all metrics and datasets, showcasing its superior capability in limited data environments. Further generated images can be found in Appendix C.

In addition, we conducted further experiments on two particularly small datasets, MetFaces and BreCaHAD, consisting of only 1,203 and 1,750 training images, respectively. The outcomes, as shown in Table 2, reveal that our model effectively addresses the challenges of extremely limited data scenarios, surpassing the performance of other state-of-the-art methods. To ensure a comprehensive evaluation, we also integrated the ADA strategy [30] into our experiments. As illustrated in Table 2, the combination of our method with ADA outperforms the configuration of CR with ADA, as well as all other tested configurations. This further emphasizes the effectiveness of our approach when paired with advanced augmentation techniques.

In Table 3, we compare our approach with various GAN architectures and techniques, including DC-GAN [41] and SNGAN [37], as well as different improve-

Method	N	letFac	es	BreCaHAD			
Methou	FID ↓	IS ↑	$KID \Downarrow$	FID ↓	IS ↑	$KID \Downarrow$	
StyleGAN2 [31]	53.21	3.16	0.035	97.06	3.10	0.095	
+ LeCam [51]	56.67	2.43	0.102	83.74	2.51	0.046	
+ CR [58]	48.89	3.24	0.029	80.72	2.92	0.058	
+ DigGAN [19]	53.97	3.09	0.031	105.45	3.05	0.095	
+ KD-DLGAN [14]	54.05	3.11	0.029	86.25	3.26	0.083	
+ ADA [30]	28.10	4.29	0.006	23.90	3.04	0.019	
+ CR + ADA	29.91	4.20	0.008	22.69	2.81	0.011	
+ SQ-GAN (ours)	41.12	3.31	0.024	48.44	2.98	0.041	
+ SQ-GAN + CBI (ours)	35.44	4.02	0.022	42.42	2.93	0.037	
+ SQ-GAN + ADA (ours)	24.77	4.93	0.008	22.61	3.15	0.010	

Table 2. Quantitative comparison on small (extremely limited) datasets, MetFaces and BreCaHAD.

Method (64	FID ↓	IS ↑	
DCGAN	Vanilla [41]	107.33	3.23
	+ Auxiliary rotations [10]	103.43	2.85
	+ Dropout [48]	77.35	4.11
	+ Noise [46]	86.13	5.40
SNGAN	Vanilla [37]	57.08	3.42
	+ Auxiliary rotations [10]	54.85	3.63
	+ Dropout [48]	48.54	4.02
	+ Noise [46]	55.99	3.33
StyleGAN2	Vanilla [31]	24.16	10.37
	+ SQ-GAN (ours)	18.96	11.41

Table 3. Comparison of various GAN architectures and techniques on the Oxford-Dog dataset (64×64) .

ment strategies such as auxiliary rotation prediction [10], dropout [48], and noise injection [46]. The results demonstrate that our method significantly outperforms the traditional convolution-based DCGAN and the ResNet-based SNGAN, highlighting the effectiveness of our approach.

6.2. Semantic similarity analysis

To further investigate the effectiveness of our method, we conducted a semantic similarity analysis using the CLIP model [42]. As shown in Figs. 4c and 4d, the semantic similarity between the images and the corresponding category text (e.g., *a photo of a* $\{dog\}$) is significantly improved when using our method. This suggests that our model can effectively disentangle and represent the semantic information within the latent space, leading to more semantically meaningful and interpretable features. The results of this analysis are consistent with the quantitative evaluation, further validating the effectiveness of our method.

6.3. Ablation study

In this section, we begin with ablation studies to examine the effect of code dimension and the role of uniformity regularization within SQ-GAN. We then delve into additional factors that could influence overall performance. Specifically, we document the usage of the codebook during the inference phase (i.e., the proportion of codebook entries actually utilized by the quantized vectors) to indicate how efficiently the codebook is being utilized and the extent of ef-



Figure 4. Distribution of semantic similarity. We compute the cosine similarity between the features extracted from the image and semantic information from category text using the CLIP model, denoted as $\Phi(\cdot)$.



Figure 5. Evolutions of FID (\Downarrow) scores during training on the Oxford-Dog and FFHQ-2.5K datasets with different code dimensions.

Config.	Code dimension						Uniformity	mity Initialization	Licogo(%)	Oxford-Dog	
	1	2	4	8	16	32	Uniformity	Initialization	Usage(%)	FID ↓	IS ↑
	\checkmark						 Image: A start of the start of	X	58.98	37.93	11.23
		\checkmark					~	×	75.50	37.71	11.06
			\checkmark				~	×	76.63	36.30	11.02
				\checkmark			~	×	69.18	38.60	11.02
					\checkmark		~	X	72.11	42.45	11.26
						\checkmark	~	×	71.01	45.33	11.01
			\checkmark				×	X	69.90	40.17	10.98
SQ-GAN			\checkmark				 ✓ 	×	76.63	36.30	11.52
SQ-GAN + CBI			~				√	 ✓ 	78.55	35.01	12.44

Table 4. Ablation Studies of variant code dimension, uniformity regularization and codebook initialization on the Oxford-Dog dataset.

fective feature compression. Unless otherwise specified, all the ablation experiments are conducted on the Oxford-Dog dataset [39].

Effects of Code Dimension In Fig. 5 and Table 4, we incrementally increased the code dimension of the codebook (d_w/s) in our SQ-GAN from 1 to 32. The results show that when the code dimension is set to 4, the model's performance significantly improves. However, further increasing the code dimension does not yield additional improvement. This suggests that a code dimension of 4 is sufficient to capture the essential features of the data distribution, while higher dimensions may introduce unnecessary complexity and stability issues. This finding is consistent with the results of our comprehensive comparison.

Effects of codebook uniformity As demonstrated in Table 4, the improvement in the usage of codebook entries

is markedly greater compared to settings without uniformity regularization, indicating that the discrete features learned in the codebook are more compact and representative. Additionally, we observe enhanced performance metrics when uniformity regularization is employed, highlighting how uniformity contributes to improving the representational quality of the codebook.

Effects of codebook initialization In Table 4, we conducted an ablation study by replacing this knowledge-based initialization (i.e., "+ CBI") with random initialization. The results show a significant reduction in codebook usage (compared to the "+ CBI" setup with a code dimension of 4), along with a decreased performance on FID. However, even with random initialization, the performance of the codebook remains competitive with other methods (as shown in Table 1). This indicates that our approach generates a more robust and expressive codebook when the knowledge-enhanced initialization method is employed. A more detailed discussion and theoretical analysis are provided in Appendix A.

7. Conclusion

In this study, we introduced a style space quantization method, leveraging codebook learning to enhance consistency regularization within the discriminator, particularly in limited-data scenarios. Furthermore, we advocate a pioneering approach to codebook initialization that incorporates external semantic knowledge from foundation models, thereby crafting a vocabulary-rich and perceptually diverse codebook. Extensive experiments across various tasks and datasets demonstrate that our approach results in highquality image generation under limited data conditions. The discrete proxy space \mathcal{W}^q we constructed serves as an advanced abstraction of the original W space, offering a more compact representation for image generation. Consequently, this paves the way for future research on stylized image generation, editing, and studies on attribute disentanglement and interpretability within the W space.

Acknowledgements

This work is supported by the Fundamental Research Funds for the Central Universities under Grant 1082204112364, the National Major Scientific Instruments and Equipments Development Project of National Natural Science Foundation of China under Grant 62427820 and the Science Fund for Creative Research Groups of Sichuan Province Natural Science Foundation (No. 2024NSFTD0035).

References

- Alper Aksac, Douglas J Demetrick, Tansel Ozyer, and Reda Alhajj. Brecahad: a dataset for breast cancer histopathological annotation and diagnosis. *BMC research notes*, 12:1–3, 2019. 6
- Martín Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In 5th International Conference on Learning Representations, 2017.
 3
- [3] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings* of the 34th International Conference on Machine Learning, pages 214–223, 2017. 2
- [4] Mikolaj Binkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD gans. In 6th International Conference on Learning Representations, 2018.
 6
- [5] Sergiy V Borodachov, Douglas P Hardin, and Edward B Saff. Discrete energy on rectifiable sets. Springer, 2019. 5
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In 7th International Conference on Learning Representations, 2019. 2
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, 2020. 1
- [8] Lucien Le Cam. Asymptotic Methods in Statistical Decision Theory. Springer, 1986. 3
- [9] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1691–1703, 2020. 3
- [10] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised gans via auxiliary rotation loss. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12154–12163, 2019. 7

- [11] Henry Cohn and Abhinav Kumar. Universally optimal distribution of points on spheres. *Journal of the American Mathematical Society*, pages 99–148, 2006. 5
- [12] A H Cook. Foundations of modern potential theory. *Physics Bulletin*, 24(11):681–681, 1973. 5
- [13] Kaiwen Cui, Jiaxing Huang, Zhipeng Luo, Gongjie Zhang, Fangneng Zhan, and Shijian Lu. Genco: Generative cotraining for generative adversarial networks with limited data. In *Thirty-Sixth AAAI Conference on Artificial Intelligence*, pages 499–507, 2022. 1
- [14] Kaiwen Cui, Yingchen Yu, Fangneng Zhan, Shengcai Liao, Shijian Lu, and Eric P. Xing. KD-DLGAN: data limited image generation via knowledge distillation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3872–3882, 2023. 1, 2, 3, 7
- [15] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States, pages 2292–2300, 2013. 6
- [16] Emily L. Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, pages 1486– 1494, 2015. 2
- [17] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *CoRR*, abs/2005.00341, 2020. 1
- [18] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021. 1, 2, 3
- [19] Tiantian Fang, Ruoyu Sun, and Alexander G. Schwing. Diggan: Discriminator gradient gap regularization for GAN training with limited data. In Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, pages 31782–31795, 2022. 1, 3, 7
- [20] 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 27: Annual Conference on Neural Information Processing Systems 2014, pages 2672–2680, 2014. 2, 3
- [21] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10686–10696, 2022. 3
- [22] Ishaan Gulrajani, Faruk Ahmed, Martín Arjovsky, Vincent Dumoulin, and Aaron C. Courville. Improved training of wasserstein gans. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, pages 5767–5777, 2017. 3, 7

- [23] Ishaan Gulrajani, Colin Raffel, and Luke Metz. Towards GAN benchmarks which require generalization. In 7th International Conference on Learning Representations, 2019.
- [24] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, pages 6626–6637, 2017. 6
- [25] Simon Jenni and Paolo Favaro. On stabilizing generative adversarial training with noise. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12145–12153, 2019. 3
- [26] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Deceive D: adaptive pseudo augmentation for GAN training with limited data. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, pages 21655–21667, 2021. 1
- [27] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard GAN. In 7th International Conference on Learning Representations, 2019. 7
- [28] 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, 2018. 2
- [29] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 2, 4, 6
- [30] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, 2020. 1, 2, 3, 6, 7
- [31] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8107–8116, 2020. 2, 6, 7
- [32] Nupur Kumari, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Ensembling off-the-shelf models for GAN training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10641–10652, 2022. 1, 2, 3
- [33] Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, and Antonio Torralba. Diverse image generation via selfconditioned gans. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14274–14283, 2020. 2
- [34] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *IEEE International Conference on Computer Vision*, pages 2813–2821, 2017. 2, 7
- [35] Lars M. Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *Proceedings of the 35th International Conference* on Machine Learning, pages 3478–3487, 2018. 3

- [36] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. In 6th International Conference on Learning Representations, 2018. 2, 3
- [37] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In 6th International Conference on Learning Representations, 2018. 2, 3, 7
- [38] Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Freeze discriminator: A simple baseline for fine-tuning gans. *CoRR*, abs/2002.10964, 2020. 1, 3
- [39] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, pages 3498– 3505, 2012. 6, 8
- [40] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *IEEE/CVF International Conference* on Computer Vision, pages 4172–4182, 2023. 1
- [41] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In 4th International Conference on Learning Representations, 2016. 7
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings* of the 38th International Conference on Machine Learning, pages 8748–8763, 2021. 2, 5, 7
- [43] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the* 38th International Conference on Machine Learning, pages 8821–8831, 2021. 3
- [44] Ali Razavi, Aäron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. In Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, pages 14837–14847, 2019. 2, 3
- [45] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, pages 2226–2234, 2016. 6
- [46] Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, and Ferenc Huszár. Amortised MAP inference for image super-resolution. In 5th International Conference on Learning Representations, 2017. 3, 7
- [47] Jiaming Song and Stefano Ermon. Bridging the gap between f-gans and wasserstein gans. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9078– 9087, 2020. 2
- [48] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res., 15(1):1929–1958, 2014. 7
- [49] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung. Towards good

practices for data augmentation in GAN training. *CoRR*, abs/2006.05338, 2020. 2

- [50] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung. On data augmentation for GAN training. *IEEE Trans. Image Process.*, 30:1882–1897, 2021. 1
- [51] Hung-Yu Tseng, Lu Jiang, Ce Liu, Ming-Hsuan Yang, and Weilong Yang. Regularizing generative adversarial networks under limited data. In *IEEE Conference on Computer Vision* and Pattern Recognition, pages 7921–7931, 2021. 1, 3, 7
- [52] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *The 9th ISCA Speech Synthesis Workshop*, page 125, 2016. 1
- [53] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, pages 6306–6315, 2017. 2, 3, 5
- [54] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9929–9939, 2020.
- [55] Ryan Webster, Julien Rabin, Loïc Simon, and Frédéric Jurie. Detecting overfitting of deep generative networks via latent recovery. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 11273–11282, 2019. 2
- [56] Han Zhang, Tao Xu, and Hongsheng Li. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *IEEE International Conference on Computer Vision*, pages 5908–5916, 2017. 2
- [57] Han Zhang, Ian J. Goodfellow, Dimitris N. Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *Proceedings of the 36th International Conference* on Machine Learning, pages 7354–7363, 2019. 2
- [58] Han Zhang, Zizhao Zhang, Augustus Odena, and Honglak Lee. Consistency regularization for generative adversarial networks. In 8th International Conference on Learning Representations. OpenReview.net, 2020. 1, 2, 3, 7
- [59] Junbo Jake Zhao, Michaël Mathieu, and Yann LeCun. Energy-based generative adversarial networks. In 5th International Conference on Learning Representations, 2017. 2
- [60] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient GAN training. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, 2020. 1, 2, 3
- [61] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for GAN training. *CoRR*, abs/2006.02595, 2020. 1, 3
- [62] Zhengli Zhao, Sameer Singh, Honglak Lee, Zizhao Zhang, Augustus Odena, and Han Zhang. Improved consistency regularization for gans. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*, pages 11033–11041, 2021. 1, 2, 3
- [63] Shangchen Zhou, Kelvin C. K. Chan, Chongyi Li, and Chen Change Loy. Towards robust blind face restoration

with codebook lookup transformer. In Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, 2022. 2