

ScoreMix: A Scalable Augmentation Strategy for Training GANs With Limited Data

Jie Cao ^{id}, Mandi Luo ^{id}, Junchi Yu ^{id}, Ming-Hsuan Yang ^{id}, *Fellow, IEEE*, and Ran He ^{id}

Abstract—Generative Adversarial Networks (GANs) typically suffer from overfitting when limited training data is available. To facilitate GAN training, current methods propose to use data-specific augmentation techniques. Despite the effectiveness, it is difficult for these methods to scale to practical applications. In this article, we present ScoreMix, a novel and scalable data augmentation approach for various image synthesis tasks. We first produce augmented samples using the convex combinations of the real samples. Then, we optimize the augmented samples by minimizing the norms of the data scores, i.e., the gradients of the log-density functions. This procedure enforces the augmented samples close to the data manifold. To estimate the scores, we train a deep estimation network with multi-scale score matching. For different image synthesis tasks, we train the score estimation network using different data. We do not require the tuning of the hyperparameters or modifications to the network architecture. The ScoreMix method effectively increases the diversity of data and reduces the overfitting problem. Moreover, it can be easily incorporated into existing GAN models with minor modifications. Experimental results on numerous tasks demonstrate that GAN models equipped with the ScoreMix method achieve significant improvements.

Index Terms—Generative adversarial networks, image synthesis, data augmentation, few-shot image-to-image translation

1 INTRODUCTION

IN recent years, Generative Adversarial Networks (GANs) [1] have shown much progress in numerous image synthesis tasks. However, current GAN-based methods [2], [3], [4] heavily rely on vast quantities of training data. Large-scale training datasets, e.g., the CelebA [5] and FFHQ [6] datasets, are necessary for these methods to achieve state-of-the-art results. When the amount of data is limited, recent findings [7], [8] reveal that GANs easily overfit the training set. This issue leads to drastically degraded results in scenarios where collecting sufficient training data is infeasible. Hence, improving the generalization of GANs to the data-limited regime is practically important to the image synthesis tasks.

To tackle the overfitting problem, some recent approaches develop data augmentation schemes [7], [8], [9], [10] tailored for GANs. These schemes use a set of hand-crafted image processing operations (e.g., random cropping, scaling, and color jittering) to augment the training samples. They combine these operations with adaptive [7] or automatic [8] strategies.

- Jie Cao, Mandi Luo, Junchi Yu, and Ran He are with the CRIPAC & NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing 100045, China, and also with the University of Chinese Academy of Sciences, Beijing 101408, China. E-mail: jie.cao@cripac.ia.ac.cn, {luomandi2019, ran.he}@ia.ac.cn, junchiyu@163.com.
- Ming-Hsuan Yang is with the University of California at Merced, Merced, CA 95343 USA. E-mail: mhyang@ucmerced.edu.

Manuscript received 28 November 2021; revised 19 October 2022; accepted 17 December 2022. Date of publication 23 December 2022; date of current version 5 June 2023.

This work was supported in part by the National Natural Science Foundation of China under Grants 6220073425, U21B2045, and U20A20223. The work of Ming-Hsuan Yang was supported by NSF CAREER under Grant 1149783.

(Corresponding author: Ran He.)

Recommended for acceptance by K. M. Lee (EIC).

Digital Object Identifier no. 10.1109/TPAMI.2022.3231649

For some specific tasks [8], these methods can train GANs effectively with hundreds of samples, allowing decent performances.

However, the current methods have some intrinsic limitations: First, these methods are difficult to generalize across datasets. The reason is that domain knowledge is integrated into the hand-crafted augmentation operations of these methods. Such knowledge is valid only for some specific data. For instance, the color distortion strategies for horse images are not suitable for zebra images. Second, these methods include a large number of hyperparameters that require fine-tuning on specific training data. Note that unsuitable choices of the hyperparameters are even harmful to model training [7], [11]. Fine-tuning these hyperparameters, however, is prohibitively expensive due to the large search space. To reduce the computational burden, current methods assign fixed values to most of the hyperparameters based on their experience. Due to these issues, training GANs with limited data in real-world scenarios remains challenging.

To facilitate training GANs with limited data, we propose a mixing-based augmentation strategy named ScoreMix. As illustrated in Fig. 1, given a data distribution p_{data} , we first take a convex combination of two training samples as the mixed sample x' . Then, we optimize x' by minimizing the L2 norm of its Stein score [12]. The Stein score of a sample, denoted as $\nabla \log p_{\text{data}}(x)$, is the gradients of the logarithmic distribution function with respect to data. During the optimization process, the mixed sample x' gradually moves towards the manifold of data distribution. Under mild assumptions, for the solution x^* , the density probability $p_{\text{data}}(x^*)$ has the maximal value in some local neighborhood around the mixed sample x' . This means the augmented sample x^* has the maximum likelihood that belongs to the data distribution.

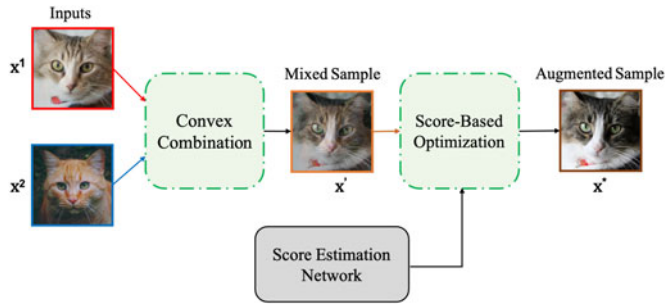


Fig. 1. Illustration of the proposed data augmentation method. Given two data points x_1 and x_2 , we first take the convex combination of the two points as the mixed sample x' . Then, we optimize the mixed sample with a score estimation network. The optimization strategy enforces the mixed sample to move to the high-density regions of the data distribution.

To estimate the Stein scores of samples, we build a deep estimation network using the denoising score matching method [13]. During the training stage, the estimation network learns to predict the Stein scores of the samples perturbed by multi-scale random noises. The minimal noise level is negligible, so that the network can accurately predict the Stein scores of the real samples. We do not introduce any domain knowledge in the training process. Hence, the proposed method is general-purpose in nature. For different image synthesis tasks, we train the score estimation network only with the given training data. Moreover, the proposed method does not require any modifications to the hyperparameters or network architectures.

The ScoreMix method can be incorporated into the GAN models easily. We only need to modify how these methods feed training data. In addition, our approach is fully compatible with existing augmentation methods [8], [14], [15], [16], [17].

We evaluate the proposed method on several image synthesis tasks. Extensive experiments are conducted on the FFHQ dataset [6], the ImageNet dataset [15], MetFace dataset [7], and AFHQ dataset [18]. The results demonstrate that the GAN models equipped with the ScoreMix method achieve significant improvements. Notably, we apply the proposed method to one-shot image-to-image translation tasks, where a single training sample is available for each image domain. We also compare against the state-of-the-art mixing-based augmentation methods [19], [20], [21], and show that the proposed ScoreMix method performs favorably regarding both quantitative evaluations and human evaluations.

The main contributions are summarized as follows:

- We present a score-based data augmentation method scalable to various image synthesis tasks. The augmented samples diversify training data without perturbing the data distribution. Our method reduces the overfitting problem, facilitating the training of GANs with limited data.
- We achieve significant improvements in multiple image synthesis tasks. We produce plausible results with a small amount of training data. In addition, the proposed method works decently in one-shot image synthesis.

This paper is an extension of our previous conference paper [22]. The major improvements over the preliminary

one are in three-folds: (1) The proposed data augmentation method can be applied to general image synthesis tasks. The previous approach is effective only for the image-to-image translation tasks. (2) This work proposes to produce augmented samples with score-based optimization. Moreover, we build a score estimation network based on multi-scale denoising score matching. This approach ensures that the augmented samples are close to the data manifold. In contrast, the previous approach produces augmented samples by linearly interpolating the real training samples. (3) Our approach does not rely on specific designs of the underlying GAN model. Thus, our approach facilitates the training of various GANs. The previous approach, however, requires a predefined content-preserving loss for augmentation.

2 RELATED WORK

2.1 Data Augmentation

Numerous methods have been developed to increase the amount of data for training deep learning models without overfitting. To name a few, Lecun et al. [14] use the affine transformations, including translation, scaling, and shearing, in handwritten character recognition; Bengio et al. [23] introduce the transformations that degrade images (e.g., Gaussian noise, salt and pepper noise, and motion blur) for augmentation; Krizhevsky et al. [15] apply random cropping, horizontal flipping, and color jittering to train deep networks. Using these content-preserving transformations has become a routine data pre-processing step.

The strategies mentioned above can effectively diversify the training data, but the hyperparameters require data-specific tuning, which is non-trivial in practice. Note that an unsuitable augmentation setting even hurts model performance [7], [11]. Some methods [24], [25], [26], [27] propose to automatically search the optimal hyperparameters for generic augmentation operations. However, these methods are prohibitively expensive in many practical settings.

The mixing-based augmentation methods [19], [20], [21], [28], [29], [30], [31], [32], [33], [34] use the combinations of real samples to generate virtual training samples. These methods can regularize the models from the overfitting problem and improve generalization. The mixup method [19] is one of the most widely used approaches, which uses linear interpolations for augmentation. This method is easy to implement and requires neglectable computation. However, the augmented samples look unnatural and exhibit visible artifacts, which may confuse the training model. Recent approaches have proposed various improved mixing-based methods. Guo et al. [33] present an adaptive mixing policy to reduce the generation of degraded data. The CutMix method [20] augments samples by cutting and pasting image patches within a data batch. The PuzzleMix method [31] proposes to search for an optimal mask, which reveals the most salient objects in the real images, for data mixing. Dabouei et al. [30] present the SuperMix method, which uses the knowledge from a teacher network for data augmentation. The AugMix method [21] proposes a chain of augmentation operations and enforces that the augmented image is similar to the original one. Though the existing mixing-based methods have demonstrated notable

improvements, most methods generate augmented samples that drift off the data manifold. Using such augmented samples might lead to degraded image quality for the image synthesis tasks. To address this issue, we propose to generate the mixed data with the maximum likelihood that belongs to the real data distribution.

2.2 Data-Efficient GANs

GANs [1] are one of the most popular generative models. This method has been successfully applied to many computer vision applications. In particular, the GAN-based methods [35], [36], [37], [38] significantly improve the visual quality and diversity of the synthesized images. The state-of-the-art GAN models [3], [4], [6] yield impressive synthesis results on large-scale datasets. However, a fundamental challenge exists: the generalization ability and training stability of GANs are limited, especially in the data-limited regime.

These issues can be mitigated by collecting a large amount of data. But this can be prohibitively expensive or implausible. Hence, some efforts [7], [8], [9], [10], [39], [40] have recently been made to augment training data for GANs. Karras et al. [7] propose an adaptive data augmentation scheme for the discriminator. During training, this scheme controls the augmentation strength based on the degree of overfitting. The DiffAug method [8] introduces differentiable augmentation operations, which regularize GANs and improve the training stability. Tran et al. [10] propose an augmentation framework with invertible transformations. They also provide theoretical analysis to prove this approach preserves the Jensen-Shannon divergence of original GANs. These methods reduce the overfitting problem and render some plausible results. However, they use data-specific augmentation operations. Recent studies [7], [8], [11] show that these methods are sensitive to the choices of datasets, training methods, and network architectures. In this work, we propose a scalable data augmentation method for GAN training. The proposed augmentation method does not require data-specific tuning of hyperparameters. For different image synthesis tasks, we retrain the score estimation network, which guides the augmentation process.

The adaption-based approaches [41], [42], [43], [44], [45], [46], [47] train GANs with external datasets as an alternative to the target dataset. Liu et al. [41] first learn a semantically related generation model and then adapt it to the target domain. To deal with the content loss problem, Saito et al. [42] propose the content-conditioned style encoder. Moreover, several adaption-based methods focus on generating specific images, including faces [43], [44], scenes [45], human bodies [46], [47], etc. Despite the effectiveness, these approaches require additional image collection. They also assume that the external dataset and target dataset share a great similarity. In this paper, the proposed method trains GANs without any additional training data.

A recent work by Tseng et al. [17] proposes the LeCam-divergence (LC) regularization to improve the generalization of GANs trained with limited data. They impose the LC regularization on the training objective of the discriminator. The strategy can be easily added to the GAN training pipeline. While in this work, we propose to improve the

training of GANs by augmenting the data, which is orthogonal to the LC regularization. We show that applying the LC regularization to the training process of our method brings additional improvements (later in our experiments).

Some recent approaches [48], [49], [50], [50], [51], [52], [53] investigate the one-shot image synthesis task (a.k.a, single image generation). These methods learn a generation model from one sample for each image domain. The one-shot generation methods typically compute the internal statistics of patches within the training sample for model training. Several methods use domain knowledge and solve specific one-shot tasks, e.g., texture synthesis [49] and image super-resolution [48]. Towards generalized one-shot image synthesis, Shocher et al. [50] present InGAN, a unified model for various tasks and data types. Shaham et al. [51] train a pyramid of GANs using a multi-scale pipeline. They produce synthesized results with a coarse-to-fine model inference. Hinz et al. [52] propose a rescaling approach to reduce training time and a refinement method to improve image quality. To encourage content preservation, the SinCUT method [53] introduces a path-based contrastive learning scheme. In this work, the proposed data augmentation method is powerfully complementary to these approaches. Using the proposed method renders better results in the one-shot image synthesis tasks in our experiments.

2.3 Score-Based Generative Model

The score-based generative models [54], [55], [56] represent data distributions through scores, the vector fields which reflect the gradients of the likelihood of data. Similar to GANs, these models can produce high-quality images. They have proven highly effective in multiple image synthesis tasks. Song and Ermon [54] propose to train a neural network to estimate the scores of data. Then, they sample synthesized images using Langevin dynamics. Following this line of research, Song and Ermon [55] later improve the sampling process with a set of techniques that effectively estimate the noise scales from training data. Ho et al. [57] model the image sampling procedure as a reverse diffusion process. Then, Song et al. [56] integrate previous methods into a unified framework with stochastic differential equations. In this work, we train a deep score estimation model for data augmentation. Note that in the sampling process, we use the mixed image as initialization and produce the synthesized image with gradient-based optimization. These differences distinguish our approach from the discussed methods.

3 PROPOSED METHOD

The remainder of this section is organized as follows. Section 3.1 presents the overall data augmentation pipeline. In Section 3.2, we introduce the proposed score-based augmentation function. In Section 3.3, we describe how to train the score estimation network. Section 3.4 provides a comprehensive comparison between our approach and existing mixing-based methods.

3.1 Overall Augmentation Pipeline

GANs [1] aim to model a data distribution p_{data} with a generator and a discriminator. GAN training is a two-player

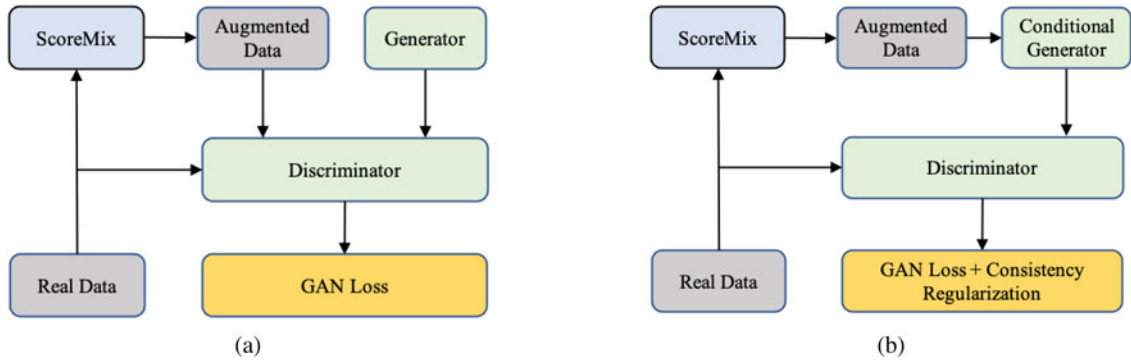


Fig. 2. Overall pipeline of training Generative Adversarial Networks with the proposed ScoreMix method. The noise input to the generator is omitted for simplicity. We synthesize augmented samples with a score-based mixing function (Section 3.2). (a) For unconditional image synthesis, we provide the discriminator with the augmented samples as ground truth. (b) For image-to-image translation, we feed the original samples and augmented samples to the generator.

game: the generator learns to generate samples that resemble the real samples, whereas the discriminator learns to discriminate between the real and fake samples. Formally, the objective function of GANs can be formulated as

$$\min_G \max_D \mathbb{E}_{\mathbf{x}} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z}} [\log (1 - D(G(\mathbf{z})))] \quad (1)$$

where $\mathbf{x} \in \mathbf{R}^n$ is drawn from p_{data} , and z is drawn from a prior distribution. We use G and D to denote the generator and discriminator, respectively. Given a small training dataset, the discriminator D is prone to overfit the real samples. If this happens, the GAN training will collapse quickly since D does not provide useful supervision.

To address this issue, we propose a mixing-based data augmentation strategy. For unconditional GANs, as illustrated in Fig. 2a, we create augmented samples using the combinations of real samples. D regards the augmented samples as the ground truth data during training. When G conditions on the real sample, as shown in Fig. 2b, we augment the input data for G . We use this strategy for practical applications such as image-to-image translation tasks.

The core component of the augmentation strategy is the mixing function. It increases the diversity of training data and prevents overfitting. Recall that one of the most classical mixing-based augmentation strategies is the mixup method [19], which uses the linear interpolation of samples. Concretely, given two real samples $\mathbf{x}_1, \mathbf{x}_2 \sim p_{\text{data}}$, its mixing function can be formulated as

$$\mathbf{x}' = f^{\text{mixup}}(\mathbf{x}_1, \mathbf{x}_2) = \lambda \cdot \mathbf{x}_1 + (1 - \lambda) \cdot \mathbf{x}_2, \quad (2)$$

where $\lambda \in [0, 1]$ is the interpolation weight sampled from the Beta distribution, i.e., $\lambda \sim \text{Beta}(\alpha, \alpha)$ with $\alpha \in (0, +\infty)$ being a hyperparameter. As mentioned in Section 2.1, the mixup method may introduce characteristic image artifacts, leading to manifold intrusion. In the following, we introduce the proposed mixing function to address this issue.

3.2 ScoreMix Augmentation

We first introduce the notion of the Stein score. According to Stein's method [12], given a smooth data distribution p_{data} supported on the set of real numbers, we define the Stein

score of a sample $\mathbf{x} \sim p_{\text{data}}$ as follows:

$$S_{p_{\text{data}}}(\mathbf{x}) := \nabla_{\mathbf{x}} p_{\text{data}}(\mathbf{x}) / p(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x}). \quad (3)$$

Equation (3) shows that the Stein score can be interpreted as the gradients of the logarithmic distribution function with respect to data. Here we assume that a deep neural network S_{θ} can estimate the Stein score (discussed later in Section 3.3). Given two real samples \mathbf{x}_1 and \mathbf{x}_2 , we compute the initial value \mathbf{x}' using the mixup method [19] by Equation (2). Then, we search for \mathbf{x}^* with the following objective:

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in U(\mathbf{x}')} \|S_{\theta}(\mathbf{x})\|_2^2 := f^{\text{scoremix}}(\mathbf{x}'; S_{\theta}), \quad (4)$$

where $\|\cdot\|_2^2$ denotes the euclidean L2 norm. We use the gradient descent method to solve Equation (4). The solution \mathbf{x}^* is a stationary point of the logarithmic distribution function $\log p_{\text{data}}(\mathbf{x})$. This implies \mathbf{x}^* is also a stationary point of the distribution function $p_{\text{data}}(\mathbf{x})$. Since \mathbf{x}' produced from the mixup method is close to the manifold of p_{data} [58], \mathbf{x}^* is a local maximum almost surely. Hence, the augmented sample \mathbf{x}^* has the maximum likelihood that belongs to p_{data} in a local neighborhood.

In Equation (4), the initial value \mathbf{x}' provides a strong image prior, which prevents the score-based model from producing samples that deviate from the real samples. This is particularly useful when the training data is insufficient. Moreover, using \mathbf{x}' as initialization accelerates the augmentation process. This allows the score estimation network to optimize an augmented sample from an intermediate result, which is significantly faster than the initialization with random noise in the existing methods (e.g., [54], [55], [56]).

The proposed ScoreMix augmentation follows the Vicinal Risk Minimization principle [59]. Similar to the mixup method [19], given a training data set \mathcal{D} , we optimize the empirical vicinal risk. Concretely, we define the vicinity distribution ν as follows:

$$\begin{aligned} \nu(\mathbf{x}^* | \mathbf{x}_i) &= \frac{1}{n} \sum_j \mathbb{E}_{\lambda} [\delta(\mathbf{x}^* = f^{\text{scoremix}}(\lambda \cdot \mathbf{x}_i + (1 - \lambda) \cdot \mathbf{x}_j; S_{\theta}))] \\ &= \frac{1}{n} \sum_j \mathbb{E}_{\lambda} [\delta(\mathbf{x}^* = f^{\text{scoremix}}(\mathbf{x}'; S_{\theta}))], \end{aligned} \quad (5)$$

where $\delta(\mathbf{x}^* = \mathbf{x})$ is a Dirac mass centered at \mathbf{x} . Compared with the mixup method, our approach produces augmented samples closer to the training data distribution. This is because f^{scoremix} produces a local maximum of the distribution, which implies that $p_{\text{data}}(\mathbf{x}^*) \geq p_{\text{data}}(\mathbf{x}')$.

Algorithm 1 shows the main steps to augment training data with the proposed method. We use Equation (2) to compute initial values and then use Equation (4) to push the initialized augmented samples towards the real data manifold. The proposed augmentation method can be plugged into existing GAN-based methods without changing the underlying algorithms.

Algorithm 1. The ScoreMix Augmentation Method

Input: Image samples \mathbf{x}_1 and \mathbf{x}_2 ;
Score estimation network S_θ ;
Learning rate η , optimization step N ;
Distribution Beta(α, α).

Output: Augmented sample \mathbf{x}^* .

- 1 Sample interpolation weight $\lambda \sim \text{Beta}(\alpha, \alpha)$;
- 2 Compute the mixed sample $\mathbf{x}' = \lambda \cdot \mathbf{x}_1 + (1 - \lambda) \cdot \mathbf{x}_2$;
- 3 Initialize $\mathbf{x} \leftarrow \mathbf{x}'$, and $i \leftarrow 0$;
- 4 **while** $i < N$ **do**
- 5 Calculate loss $\mathcal{L}(\mathbf{x}^*) = \|S_\theta(\mathbf{x}^*)\|_2^2$;
- 6 Calculate gradient $\nabla_{\mathbf{x}^*} \mathcal{L}(\mathbf{x}^*)$;
- 7 $\mathbf{x}^* \leftarrow \mathbf{x}^* - \eta \cdot \nabla_{\mathbf{x}^*} \mathcal{L}(\mathbf{x}^*)$;
- 8 $i \leftarrow i + 1$;
- 9 **end**
- 10 **return** \mathbf{x}^*

3.3 Score Estimation Network

We build a score estimation network S_θ to estimate the Stein score $S_{p_{\text{data}}}(\mathbf{x})$. Following the idea of score matching [60], we do not estimate $p_{\text{data}}(\mathbf{x})$ first. Instead, we directly optimize the network parameters θ by minimizing

$$\mathbb{E}_{\mathbf{x}} \left[\|S_\theta(\mathbf{x}) - \nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x})\|_2^2 \right], \quad (6)$$

where we have $\mathbf{x} \in \mathbf{R}^n$ drawn from p_{data} . However, in Equation (6), the term $\nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x})$ is computationally intractable for real-world data distributions. To address this problem, we use the denoising score matching (DSM) method [13], which is robust in low-density data regions. The DSM method perturbs training samples with a prior distribution whose supports span the whole image space \mathbf{R}^n . This strategy enforces that the perturbed training data distribution has full support over \mathbf{R}^n . Hence, the score-based model gets stable training signals [60] even in low-density data regions.

Concretely, we perturb \mathbf{x} with Gaussian noise $\mathcal{N}(0, \sigma^2)$ which has a zero mean and variance σ^2 . Then, we use S_θ to estimate the Stein score of the perturbed sample $\tilde{\mathbf{x}}$. Let q_σ denote the perturbed data distribution, and the objective can be formulated

$$\mathbb{E}_{\tilde{\mathbf{x}}} \left[\|S_\theta(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log q_\sigma(\tilde{\mathbf{x}} | \mathbf{x})\|_2^2 \right], \quad (7)$$

where we have $\tilde{\mathbf{x}} \sim q_\sigma(\tilde{\mathbf{x}} | \mathbf{x}) = \mathcal{N}(\mathbf{x}, \sigma^2)$. The objective is tractable since we have $\nabla_{\tilde{\mathbf{x}}} \log q_\sigma(\tilde{\mathbf{x}} | \mathbf{x}) = -\frac{1}{\sigma^2} \cdot (\tilde{\mathbf{x}} - \mathbf{x})$. This immediately implies that we can rewrite the objective as

$$\mathbb{E}_{\tilde{\mathbf{x}}} \left[\left\| S_\theta(\tilde{\mathbf{x}}) + \frac{\tilde{\mathbf{x}} - \mathbf{x}}{\sigma^2} \right\|_2^2 \right]. \quad (8)$$

We could minimize Equation (8) to train the estimation network. The optimal solution $S_{\theta^*}(\mathbf{x})$ converges to $\nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x})$ as long as the noise variance σ^2 is small enough [13]. To train the estimation network S_θ more efficiently, we propose to perturb the data samples with multi-scale Gaussian noises. Specifically, the variances of the multi-scale noises are a geometric sequence $\{\sigma_i\}_{i=1}^N$, and we have

$$\sigma_n = \sigma_1 \cdot \gamma^{n-1}, \quad (9)$$

where γ denotes the common ratio, and $n \in [0, N]$. Without loss of generality, we assume that $0 < \gamma < 1$. Then, σ_1 has the maximal value, and σ_N has the minimal value. On the one hand, the perturbed sample with the minimal noise scale is close enough to the real sample. We can regard $\nabla_{\tilde{\mathbf{x}}} \log q_{\sigma_N}(\tilde{\mathbf{x}} | \mathbf{x})$ as the Stein score of \mathbf{x} with negligible estimation error. On the other hand, the Stein score of the perturbed samples with a larger noise scale is easier to estimate. Hence, these samples help to stabilize the training process. Concretely, we train the estimation network S_θ by minimizing

$$\begin{aligned} \mathcal{L}^{\text{SCORE}} &= \sum_{i=1}^N \omega_i \cdot \mathbb{E}_{\tilde{\mathbf{x}}} \left[\left\| S_\theta(\tilde{\mathbf{x}}, \sigma_i) - \nabla_{\tilde{\mathbf{x}}} \log q_\sigma(\tilde{\mathbf{x}} | \mathbf{x}) \right\|_2^2 \right] \\ &= \sum_{i=1}^N \sigma_i^2 \cdot \mathbb{E}_{\tilde{\mathbf{x}}} \left[\left\| S_\theta(\tilde{\mathbf{x}}, \sigma_i) + \frac{\tilde{\mathbf{x}} - \mathbf{x}}{\sigma_i^2} \right\|_2^2 \right], \end{aligned} \quad (10)$$

where $\tilde{\mathbf{x}} \sim q_{\sigma_i}(\tilde{\mathbf{x}} | \mathbf{x}) = \mathcal{N}(\mathbf{x}, \sigma_i^2)$. The score estimation network S_θ conditions on the noise variation σ_i^2 during training. Note that in Algorithm 1, we only sample the noises with the minimal variation σ_N^2 . The weighting term ω_i keeps the expectation values in Equation (10) in the same order of magnitude, which is beneficial for model training [55].

We scale the pixel values of samples to the range $[0, 1]$ and set $\sigma_N = 0.01$. We set the value of σ_1 to the maximum euclidean distance between all pairs of training samples. One may determine the value of the common ratio γ using the heuristic rules proposed by Song et al. [55]. However, we simply use $\gamma = 0.99$ and find this setting works well in the following experiments.

3.4 Comparison With Existing Methods

In this subsection, we show the advantages of the ScoreMix method by making comparisons with previous mixing-based methods:

- The proposed ScoreMix method does not require any data-specific hyperparameter tuning. In contrast, previous methods such as AugMix [21] build data augmentation pipelines consisting of numerous hyperparameters. These hyperparameters need to be tuned for different tasks. Otherwise, the augmentation methods may even hurt model performance.
- We do not rely on any pre-trained model during training. Instead, we first train a score estimation network and then use it to guide the augmentation process. Methods like SaliencyMix [32] and Puzzle Mix [31] generate augmented data with the guidance



Fig. 3. Visual comparison of the augmented images. Mixup [19], CutMix [20], and ScoreMix (ours) synthesize the augmented images using the combinations of inputs x_1 and x_2 . AugMix [21] generates the augmented image with input x_1 . The input images are from the AFHQ v2 dataset [18].

of a pre-trained saliency detection model. But how to obtain an accurate saliency detection model in the data-limited regime remains a challenging problem. The requirement of accurate pre-trained models hinders the application of these methods in real-world scenarios.

- We do not require any human annotations for data augmentation. That means the ScoreMix method is readily applicable to unsupervised image synthesis tasks. In contrast, the SuperMix method [30] and hardmix method [34] improve the data augmentation strategies with label information, which may be expensive or difficult to obtain in practice.

Fig. 3 shows some examples of the augmented results from our approach and several existing methods [19], [20], [21]. The example from the mixup method [19] exhibits characteristic artifacts. Two cats in the augmented image are displaced in relation to each other. For the CutMix method [20], it is easy to notice the edges of the pasted patch, and the augmented image looks unrealistic. We note that the mixup and CutMix methods augment samples using image arithmetic operations (e.g., addition and multiplication), which inevitably produces unrealistic details. The AugMix method [21] aims to generate augmented images without deviating too much from the original ones. However, the augmented result seems to be too similar to the original image. By contrast, the example from our approach is distinguished from the original image. We find that the fur color and pupil shape of the cat are changed. Our approach also subtly changes the shape of the stripes on the forehead of the cat. Moreover, our result expresses natural textures and details.

4 EXPERIMENTS AND ANALYSIS

4.1 Implementation Details

To generate the augmented sample x^* , we use the Adam optimizer with a learning rate of 0.005 and linear decay scheduling. We use 100 gradient descent steps for the

optimization and then finetune x^* by score-based denoising. The score estimation network optimizes x^* and approximates the denoising process at $t_0 = 0.25$. Similar to the mixup method [19], we set the hyperparameter $\alpha = 0.2$ for the beta distribution. We use these settings for all the experiments without hyperparameter tuning.

We build the score estimation network following the architecture design in [54]. Specifically, we adopt the RefineNet [61] architecture and use the conditional instance normalization [62] for every convolutional layer. To train the score estimation network, we use the Adam optimizer with a learning rate of $1e-4$. The training batch size is set to 32.

4.2 Datasets

In this section, we introduce the datasets used in the following experiments.

The Flickr-Faces-HQ (FFHQ) dataset [6] is a benchmark dataset for unconditional image synthesis. This dataset consists of 70 thousand high-quality human face images with considerable variations. The images are numbered from 0 to 69,999, and we use the first one thousand images for model training.

The ImageNet dataset [15] consists of various images organized into 1,000 classes. It is a standard benchmark for conditional image synthesis. In this work, we train the generation models with 10% training images per class.

The MetFace dataset [7] consists of human face images from the Metropolitan Museum of Art. There are 1,336 manually selected images which depict artworks such as paintings, drawings, and statues. In the following, we train the models with the first one thousand images.

The Animal Faces-HQ v2 (AFHQ) dataset [18] provides animal faces of three domains: cat, dog, and wildlife. Each category contains about 5,000 images. Choi et al. [18] split this dataset into 14,336 samples for training and 1,467 samples for testing. We randomly choose 500 training images for each class, which is about 10% of the full training set.

4.3 Unconditional Image Synthesis

We first verify the proposed ScoreMix method by unconditional image synthesis. In this task, we train GANs to map random noise to image samples. We use the StyleGAN2 [3] as the baseline network, and the output resolution is 256×256 . We implement the baseline using the released source codes.¹

We evaluate the quality of synthesized images using Fréchet Inception Distance (FID) [63] and Kernel Inception Distance (KID) [64]. The FID metric measures the Wasserstein distance between two image sets. We use the FID score between the real images and the synthesized images to evaluate the sample quality: a lower score indicates better results. The KID metric measures the squared maximum mean discrepancy between the representations of the real and synthesized image sets. Similarly, a lower KID indicates that the results are better.

In the experiments, we find that the GAN training is unstable when the training data is limited. The FID score

1. <https://github.com/rosinality/stylegan2-pytorch>




Fig. 4. Ablation of the augmentation ratio μ . We train the StyleGAN2 model [3] with the proposed ScoreMix method on the FFHQ dataset [6] under 1 K data setting. We compute Fréchet Inception Distance [63] (FID, lower is better) during training.

starts to deteriorate after a certain number of training iterations. That means the overfitting problem occurs. Then, the generated samples become meaningless noises. Such observations are also reported in previous methods [4], [7], [8]. Hence, we perform the early stopping strategy to train all the models. Specially, we compute the FID score every 100,000 iterations during training. The training is stopped when the FID score increases in five consecutive epochs. Then, we use the training snapshot with the best FID score for evaluation.

Ablation: Augmentation Ratio. We aim to determine the optimal amount of augmented data. Intuitively, this amount depends on the amount of real data. Hence, we study the effect of the augmentation ratio μ , i.e., the ratio of augmented to real data. In this ablation study, we investigate the following two cases:

- The augmentation ratio μ is static. We produce the augmented data before the training starts. The amount of training data increases by a factor of $(\mu + 1)$.
- The augmentation ratio μ grows during training. We augment the training batch with a certain probability at each iteration. The amount of augmented data will increase as the training goes on. Considering that the number of training iterations is huge, we assume the augmentation ratio approaches infinity.

The results are shown in Fig. 4. Overall, GANs benefit from augmented training data. We observe that the FID scores improve, and most of the training processes become more stable. These results show that augmented data can improve model performance. Models with a growing augmentation ratio still collapse quickly. One possible reason is that the amount of augmented data is overwhelming. The models have to neglect the real data to minimize the training loss. This phenomenon indicates that too many augmented samples is not helpful to model training.

The model with a larger static augmentation ratio renders better results, but we observe diminishing returns in terms of sample quality. When we create augmentation ratio μ from 0 to 10, the FID score decreases from 58.10 to 24.21. Whereas, increasing μ to 100 only makes the FID

TABLE 1
Ablation of Baseline Method

	Network	Preprocessing method	FID↓
(a)	StyleGAN2 [3]	none	58.11 / 24.21
(b)	StyleGAN2 _{narrower}	none	55.91 / 25.04
(c)	StyleGAN2 _{wider}	none	64.38 / 23.76
(d)	StyleGAN2 [3]	DiffAug [8]	26.63 / 21.13
(e)	StyleGAN2 _{narrower}	DiffAug [8]	29.62 / 21.26
(f)	StyleGAN2 _{wider}	DiffAug [8]	26.13 / 19.41

We train different baselines on the FFHQ dataset [3] under 1 K data setting. We report FID [63] (lower is better). The former/later number in the table cell denotes the FID score of the baseline method without/with the proposed scheme.

score 1.31 lower. To balance the trade-off between sample quality and computational cost, we set $\mu = 10$ for the following tasks.

Ablation: Baseline Model. In this ablation study, we investigate whether the proposed method consistently improves model performance across different baseline models or not. Specifically, we consider the following two types of variants:

- Baselines with different model capacity. We change the network architecture to increase/reduce model capacity. Concretely, we double/halve the network width (the number of feature channels).
- Baselines with additional augmentation method. We add the DiffAug method [8], which consists of a family of the traditional augmentation operations [14], [15], [23], to the augmentation pipeline.

The evaluation results of the variants are shown in Table 1. We use the subscript wider/narrower to indicate more/fewer feature channels. For all the variants of baselines, the proposed ScoreMix method significantly improves model performance. In addition, combining our approach with the DiffAug method [8] yields further performance gain. StyleGAN2_{wider} performs the best among different network architectures. This shows that increasing model capacity leads to better results if we augment the training data properly. The variant (e), i.e., wider StyleGAN model + DiffAug + ScoreMix, achieves the best FID score. In the following, we consider this variant as the baseline method for unconditional image synthesis unless otherwise specified.

Evaluation Results. We evaluate the proposed ScoreMix method against existing mixing-based approaches, including the mixup [19], CutMix [20], and AugMix [21] methods. Note that for the mixup and CutMix methods, we use them to augment both the real and fake samples for the discriminator. This strategy is necessary for them to avoid the leaking problem [7].

Table 2 reports the FID and KID scores evaluated on the FFHQ [3] and MetFace [7] datasets. The score-based models achieve better FID/KID scores than the GAN model [3] under the 1 K data setting. This provides empirical evidence that the score-based methods are more robust in data-limited regimes. Our approach, combining the score-based data augmentation with GANs, performs best in terms of FID and KID on both datasets.

In addition, ScoreMix_{LC} combines the proposed method with the LC regularization in our prior work [17], and

TABLE 2

Unconditional Image Synthesis Results of Different Methods on the FFHQ [3] and MetFace [7] Datasets Under 1 K Data Setting

Method	FFHQ-1 K		MetFace-1 K	
	FID ↓	KID×10 ³ ↓	FID ↓	KID×10 ³ ↓
GAN: StyleGAN2 [3]	58.11	40.32	64.58	40.60
score-based: NCSNv2 [55]	38.20	21.73	44.56	19.48
score-based: SMLD [56]	31.23	17.19	35.70	13.71
Baseline: [3] + [8]	26.13	14.56	34.60	13.87
Baseline + mixup [19]	25.25	13.87	31.93	11.81
Baseline + CutMix [20]	24.72	13.37	30.11	12.27
Baseline + AugMix [21]	22.95	9.46	25.91	6.53
Baseline + ScoreMix	19.41	5.65	22.13	4.51
Baseline + ScoreMix _{ADA}	18.74	4.38	19.12	2.97
Baseline + ScoreMix _{LC}	18.29	3.91	19.42	3.35

We report FID [63] (lower is better) and KID [64] (lower is better). ScoreMix is our proposed method. ScoreMix_{LC} and ScoreMix_{ADA} combine it with the regularization methods in [17] and [7], respectively.

ScoreMix_{ADA} combines the proposed method with the adaptive discriminator augmentation (ADA) method [7]. The comparison results in Table 2 show that appending the regularization method in [17] or [7] to ScoreMix both leads to further improvements in quantitative metrics.

We visualize some samples synthesized by different methods in Fig. 5. Our synthesized results are diverse and contain fewer artifacts. This indicates the generated distribution of our results is similar to the real distribution.

Table 3 shows the FID [63] scores with different amounts of training data (from 70 K to 1 K) on the FFHQ dataset [3]. The results show that our approaches consistently improve the baseline method under all the data settings. As the number of training samples increases, the baseline method benefits less from our approach. However, under the 70 K setting, ScoreMix still improves the FID score by 5.3% (4.56 → 4.32).



Fig. 5. Unconditional FFHQ [3] (photographs) and MetFace [7] (portrait painting) samples generated by different methods under 1 K data setting. We use the StyleGAN2 model [3] with the DiffAug scheme [8] as the baseline. The results are synthesized by the baseline, baseline + mixup [19], baseline + CutMix [20], baseline + AugMix [21], and baseline + ScoreMix (ours).

TABLE 3
FID [63] Scores (Lower is Better) as a Function of Training Set Size on the FFHQ [3] Dataset

Method	70 k	35 k	10 k	5 k	2 k	1 k
Baseline: [3] + [8]	4.56	5.85	8.85	13.44	17.93	26.13
Baseline + ScoreMix	4.32	5.36	7.57	10.46	13.45	19.41
Baseline + ScoreMix _{ADA}	4.25	5.19	7.08	9.84	12.61	18.74
Baseline + ScoreMix _{LC}	4.29	5.03	6.85	9.54	12.04	18.29

ScoreMix is our proposed method. ScoreMix_{LC} and ScoreMix_{ADA} combine it with the regularization methods in [17] and [7], respectively.

Computational Cost. The proposed ScoreMix method has a reasonable computational cost. In our experiments, training the baseline method (StyleGAN2 [3] + DiffAug [8]) takes 69.3 hours on 4 Tesla V100S GPUs. We follow the setting in [7] that stops the GAN training when 25,000 k real samples are shown to the discriminator. The training that combines the baseline method with our approach takes 76.0 hours, which increases the wall-clock time by around 9.5%. In addition, we take about 120 hours to train the score estimation network. During inference, our approach requires no additional computational cost.

4.4 Class-Conditional Image Synthesis

We also evaluate the proposed method on class-conditional image synthesis using the ImageNet dataset [65], where we train GANs to produce samples with a specified class based on a given label. Since labeled data is usually expensive to collect, solving this conditional task with limited data is important for practical applications. We use BigGAN [4] with the DiffAug method [8] as the baseline. The output resolution is 128×128 . Note that we implement the baseline method on GPU, and we reduce the training batch size to 512, which is 4 times less than the size in the original work [4].

We use FID [63] and Inception Score (IS) [66] as quantitative metrics. We sample generated samples with the truncation trick proposed in [4]. Similar to the unconditional image synthesis task, a lower FID indicates a better performance. IS measures the diversity and quality of synthesized results using entropy. In this task, a higher IS score indicates a better result.

TABLE 4
Class-Conditional Image Synthesis Results of Different Methods on the ImageNet Dataset [65]

Methods	100% data		50% data		25% data		10% data	
	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑	FID ↓	IS ↑
Baseline: [4] + [8]	9.08	91.82	12.88	83.27	17.10	62.15	49.67	21.30
Baseline + mixup [19]	9.01	91.19	12.70	84.05	16.63	64.30	42.19	25.41
Baseline + CutMix [20]	9.03	90.44	12.77	83.98	16.69	64.26	44.62	24.88
Baseline + AugMix [21]	8.86	92.31	12.46	85.12	16.07	66.56	38.76	28.25
Baseline + ScoreMix	8.60	94.37	11.85	87.09	15.12	71.31	24.67	37.48
Baseline + ScoreMix _{ADA}	8.55	98.17	11.78	89.69	14.91	76.23	24.03	40.33
Baseline + ScoreMix _{LC}	8.42	95.55	11.36	87.14	14.43	76.74	22.51	41.12

We report FID [63] (lower is better) and IS [64] (higher is better). We implement BigGAN [4] on GPU with a training batch size of 512. ScoreMix is our proposed method. ScoreMix_{LC} and ScoreMix_{ADA} combine it with the regularization methods in [17] and [7], respectively.

Authorized licensed use limited to: INSTITUTE OF AUTOMATION CAS. Downloaded on August 08, 2023 at 07:40:39 UTC from IEEE Xplore. Restrictions apply.

Evaluation Results. We evaluate our approaches against the mixing-based methods [19], [20], [21]. Table 4 reports the FID and IS scores under 100%, 50%, 25%, and 10% data settings. The comparison results show that our approaches outperform the competing methods in terms of FID and IS. The 10% data setting only has fewer than 130 training samples for each class. Such scarce training data leads to degraded performance of the baseline method. However, the proposed method significantly improves the model performance. Similar to the results in unconditional image synthesis, combining the proposed method with the LC regularization [17] or adaptive discriminator augmentation method [7] brings further performance gain. For instance, ScoreMix improves the FID score by 50.3% ($49.67 \rightarrow 24.67$), and ScoreMix_{LC} further improves it by 8.8% ($24.67 \rightarrow 22.51$) under the 10% data setting.

Fig. 6 shows the synthesized samples from the baseline method and our approach. These results show that our approach improves sample quality. It can be observed that the samples from our approach have more plausible visual details than those from the baseline method (e.g., wings of the butterfly).

4.5 Image-to-Image Translation

Image-to-image translation (I2I) aims at learning the mapping from the source domain to the target domain. In the following, we perform two specific unsupervised I2I tasks in the data-limited regime.

4.5.1 Animal Face Translation

In this subsection, we change the species of the input animal face (the source image). We consider the following two types of translations: (1) reference-guided translation. Given a reference animal face (the reference image), we extract a style representation from it. Then, the generator mixes the style with the content of the source image, producing the translated result. (2) latent-guided translation. Given a one-hot class label, we draw a latent code from a prior distribution conditioned on the label. The generator uses the latent code as the style representation. For this task, we use StarGAN v2 [18] as the baseline method. It can perform both reference-guided translation and latent-guided translation. The training objective is minimizing the combination of the adversarial loss [1], style reconstruction loss [68], and style

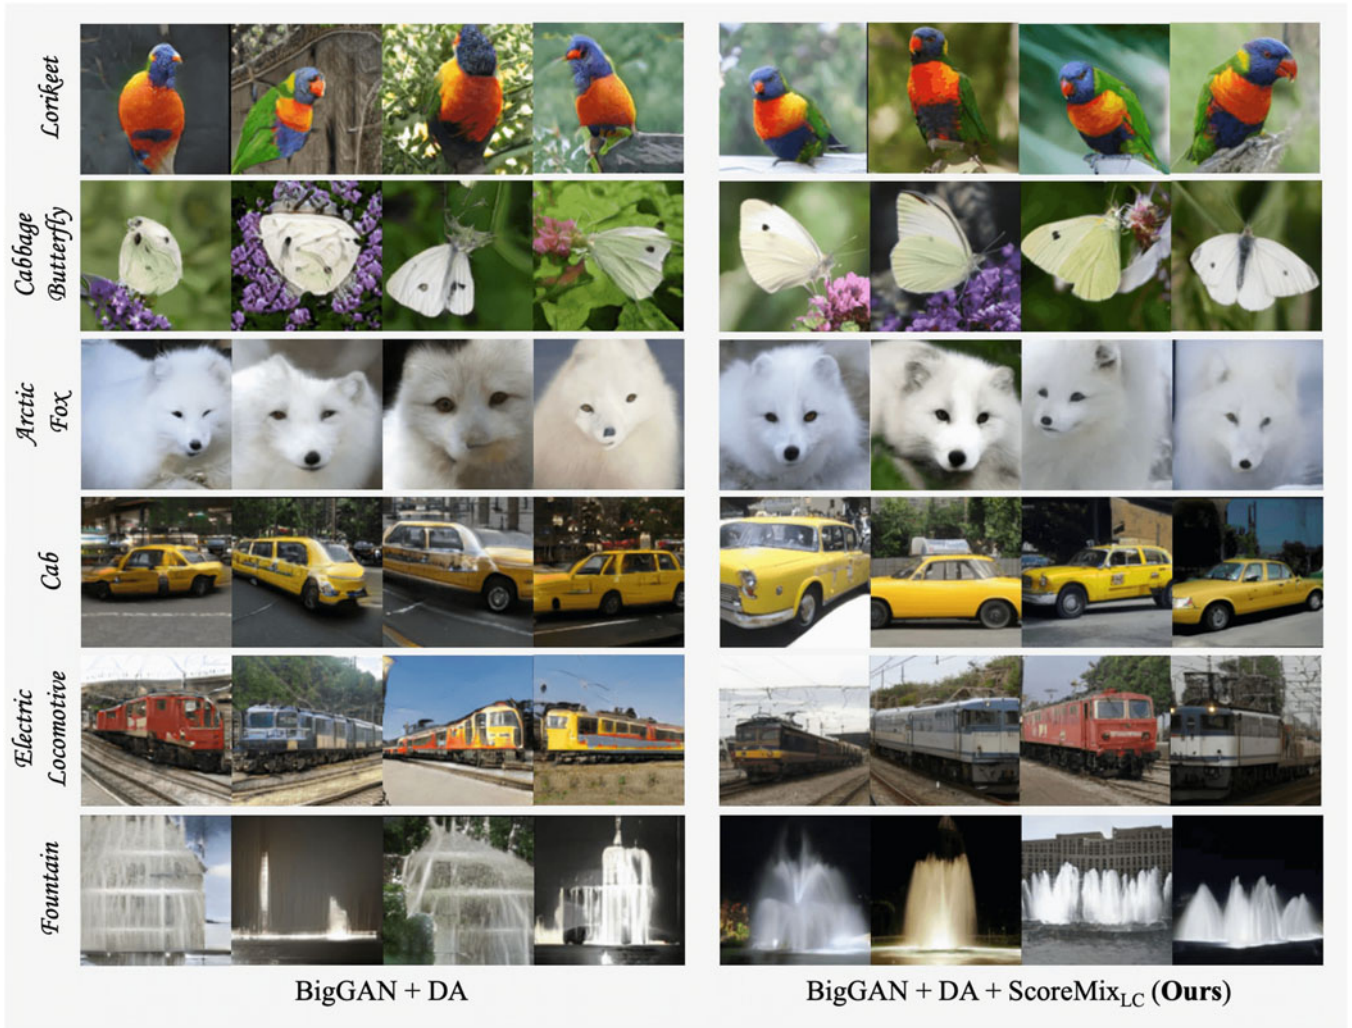


Fig. 6. Class-conditional ImageNet [65] samples generated under 10% data setting. We use the BigGAN model [4] with the DiffAug scheme [8] as the baseline. Applying our approach to the baseline method synthesize more realistic results.

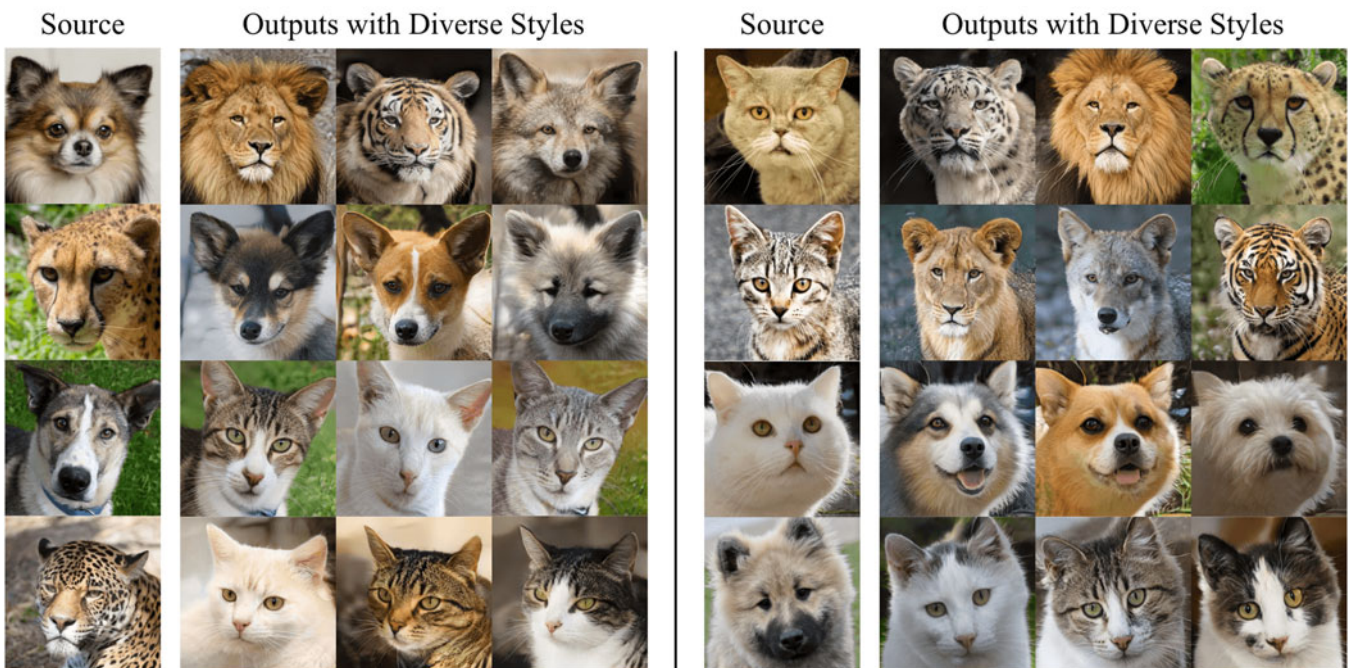


Fig. 7. Diverse translation results on the AFHQ dataset [18]. Our model can learn to generate diverse high-quality results using only 10% data in the training set.



Fig. 8. Visual samples generated by different methods under 10 % data setting of the AFHQ dataset [18]. We use StarGAN v2 [18] as the baseline. From left to right: the reference images, source images, baseline + mixup [19], baseline + CutMix [20], baseline + AugMix [21], baseline + ReMix [22], and baseline + ScoreMix (ours).

diversification loss [69]. The resolutions of the input and output images are 256×256 . To implement the baseline, we use the released official source codes.²

We evaluate the quality of synthesized images using FID [63] and Learned Perceptual Image Patch Similarity (LPIPS) [67]. The LPIPS score measures the diversity of images using the L1 distance in the feature space. In this task, the higher the LPIPS score, the better the synthesized results is. We compute the FID and LPIPS scores for every pair of the image domains (e.g., dog \rightarrow cat, cat \rightarrow dog). Since the AFHQ dataset provides 3 different image domains, we compute these scores on 6 pairs of

domains. Then, we report the average values of the scores as the evaluation metrics.

Evaluation Results. Given a single source image, we generate diverse results by multiple random reference images. Fig. 7 shows diverse translation results by our method. We observe that the synthesized images are obviously different breeds of animals. That means our approach can generate distinctive styles. In addition, the synthesized images preserve the content information such as pose and shape effectively.

We evaluate our method against existing mixing-based approaches, including the mixup [19], CutMix [20], AugMix [21] and ReMix (our previous conference version [22]) schemes. Fig. 8 shows some translation results. The synthesized images are supposed to preserve the low-level semantics (e.g., pose and shape) of the source image and the high-

2. <https://github.com/clovaai/stargan-v2>

TABLE 5
Fréchet Inception Distance (FID, Lower is Better) and Learned Perceptual Image Patch Similarity (LPIPS, Higher is Better) of Different Methods on the AFHQ Dataset [18]

Method	Latent-guided translation				Reference-guided translation			
	100% data		10% data		100% data		10% data	
	FID↓	LPIPS↑	FID↓	LPIPS↑	FID↓	LPIPS↑	FID↓	LPIPS↑
Baseline: StarGAN v2 [18]	16.18	0.450	45.12	0.428	19.78	0.432	39.23	0.407
Baseline + mixup [19]	16.02	0.454	28.15	0.450	18.51	0.453	28.12	0.439
Baseline + CutMix [20]	15.98	0.453	41.36	0.444	17.24	0.449	26.33	0.442
Baseline + AugMix [21]	15.72	0.443	41.36	0.425	16.98	0.434	23.80	0.411
Baseline + DiffAug [8]	15.79	0.468	25.11	0.460	16.89	0.477	23.02	0.459
Baseline + ReMix [22]	15.26	0.482	21.82	0.472	16.03	0.480	22.92	0.458
Baseline + ScoreMix	14.03	0.470	20.49	0.464	15.19	0.484	21.51	0.460
Baseline + ScoreMix _{ADA}	14.18	0.479	19.20	0.469	14.83	0.485	20.14	0.461
Baseline + ScoreMix _{LC}	13.16	0.482	18.83	0.474	14.11	0.485	17.99	0.466

ScoreMix is our proposed method. ScoreMix_{LC} and ScoreMix_{ADA} combine it with the regularization methods in [17] and [7], respectively.

TABLE 6
FID [63] (Lower is Better) and LPIPS [67] (Higher is Better) as a Function of Training Set Size on the AFHQ [3] Dataset

Method	Latent-guided translation				Reference-guided translation			
	100% data	50% data	25% data	10% data	100% data	50% data	25% data	10% data
Baseline: [3] + [8]	16.18	19.54	28.34	45.12	19.78	22.76	28.11	39.23
Baseline + ScoreMix	14.03	15.42	18.76	20.49	15.19	17.09	18.30	21.51
Baseline + ScoreMix _{ADA}	14.18	15.79	18.51	19.20	14.83	16.69	17.70	20.14
Baseline + ScoreMix _{LC}	13.16	14.77	17.11	18.83	14.11	15.88	16.57	17.99

ScoreMix is the proposed method. ScoreMix_{LC} and ScoreMix_{ADA} combine it with the regularization methods in [17] and [7], respectively.

level semantics (e.g., color and fur) of the reference image. The baseline model suffers from the overfitting problem and generates some unrealistic texture details. Overall, the proposed method synthesizes images with higher visual quality than other schemes.

Table 5 shows the FID and LPIPS scores of the evaluated methods. Our approach performs favorably against existing augmentation methods in terms of the quantitative metrics. The FID scores indicate that our results are more similar to the real data. The LPIPS score of ScoreMix with 10% data is even higher than that of the baseline with 100% data. These results demonstrate that the proposed method is effective for diverse and realistic image translation. Moreover, combining our approach with additional schemes [7], [17] further improves model performance. ScoreMix_{LC} performs the best among all the competing methods.

Table 6 presents the FID and LPIPS scores with different amounts of training data. Under all the data settings, our approaches provide notable performance improvements. Even for the 100% data setting (about 14 K samples), our approach is still particularly beneficial: ScoreMix improves the FID score by 13.3% (16.18 → 14.03) on latent-guided translation and 23.2% (19.78 → 15.19) on reference-guided translation.

4.5.2 One-Shot Image-to-Image Translation

In this subsection, we conduct I2I in the single image setting. Some recent methods [50], [51], [53] have explored this

challenging setting: only one source image and one target image are available for model training. Note that we do not use the pre-trained model on external data like the adaptation-based I2I approaches [41], [42].

For this task, we use the official implementation³ of the SinCUT method [53] as the baseline. This method trains the I2I model in the patch-wise manner. Concretely, the input image is randomly cropped into 16 patches. Then, these patches are resized to 128 × 128 for model training. The translation model takes these patches as a batch of training data. The training is with the standard augmentation operations such as random scaling and horizontal flipping. In addition, we use the proposed ScoreMix method to increase the diversity of training image patches. During the test stage, the translation model takes the full resolution image as the input. This approach can translate a single image of arbitrary size.

We use Single Image Fréchet Inception Distance (SIFID) [51] to evaluate the synthesized results. This metric calculates the FID score between the internal distributions of two images. We extract the features using the Inception-v3 network [70]. Similar to the FID score, a lower SIFID score between the synthesized image and target image indicates a better result.

Evaluation Results. Fig. 9 shows some translation results and the corresponding SIFID scores. We use the collected

3. <https://github.com/taesungp/contrastive-unpaired-translation>




Fig. 9. Image-to-image translation results from two unpaired training images. We train the translation model with *one source image* and *one target image*. The results are synthesized by the SinCUT method [53] and our approach. We report Single Image Fréchet Inception Distance [51] (SIFID, lower is better).

source and target images from Luan et al. [71]. The translated images are supposed to preserve the general structure of the source image and match the textures of the target image. Although this task is challenging due to the limited data, the proposed method achieves reasonable results. The SinCUT method [53] suffers from the overfitting problem and synthesizes distorted images. In contrast, the proposed method transfers the style of the target image more faithfully and shows fewer artifacts. Moreover, the proposed method obtains significantly lower SIFID scores. These results show that our synthesized images are more similar to the real data.

We conduct a user study to validate the proposed method. First, we use the SinCUT method [53] and our approach to perform a total of 60 one-shot image-to-image translation tasks. The source-target image pairs for training are publicly available.⁴ Each method synthesizes 60 translated results. Then, we ask evaluators to vote for the synthesized images from the SinCUT method and our approach based on image quality. We keep the translation methods anonymous in the evaluation. The source and target images are provided for the evaluators as references. We distribute questionnaires to online users and finally get valid feedbacks from 71 human evaluators. We did not collect any private information.

Table 7 shows the competition results. We compute the SIFID scores and record user preferences. Then, we count

the number of times each method performs better than its competitor and report the percentage in the table. For each task, any method with the majority of the user votes is better in the human evaluations. As for the SIFID score, lower is better. The competition results show that the proposed method generally performs better regarding both quantitative evaluations and human evaluations.

5 DISCUSSION

In this work, we have verified that combining the proposed method with additional augmentation operations, including [7], [8], [14], [15], [16], [17], further improves model performance. Incorporating recent methods that are orthogonal to this work, e.g., the RL-based augmentations [24], [25], may perform well. In addition, the proposed method could

TABLE 7
Our Approach versus the SinCut Method [53]

Metrics	Human Evaluation		
	Ours > SinCut	Ours < SinCut	
Ours > SinCut	83.3%	6.6%	
SIFID	Ours < SinCut	1.7%	8.4%

We evaluate the two methods on 60 different one-shot image-to-image translation tasks. The metrics include Single Image Fréchet Inception Distance [51] (SIFID) and human evaluation.

4. <https://github.com/luanfujun/deep-photo-styletransfer>

be easily used in the semi-supervised learning framework [72], [73], [74]. We leave the study to future work.

The focus of this paper is how to generate images plausible to human visual perception. We have shown that the proposed data augmentation method improves the sample quality using automatic metrics [51], [63], [64], [67] and subjective evaluation. However, whether or not the proposed method is beneficial still depends on the applications. It might be less suited for some particular applications, such as medical imaging [75] and biology [76].

6 CONCLUSION

This article introduces a mixing-based data augmentation method to tackle the overfitting problem of GANs. We propose to optimize the augmented samples by minimizing the norms of the data scores. This approach prevents manifold intrusion and produces high-quality augmented samples. Hence, it increases the diversity of image samples, facilitating GAN training in image synthesis. In addition, the proposed approach is scalable to various image synthesis tasks and requires no data-specific fine-tuning. We demonstrate that our method vastly improves image quality and quantitative metrics in numerous tasks, especially when the training data is limited.

REFERENCES

- [1] I. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [2] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [3] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of StyleGAN," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8107–8116.
- [4] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [5] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 3730–3738.
- [6] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4401–4410.
- [7] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila, "Training generative adversarial networks with limited data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 1015.
- [8] S. Zhao, Z. Liu, J. Lin, J.-Y. Zhu, and S. Han, "Differentiable augmentation for data-efficient GAN training," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 634.
- [9] H. Zhang, Z. Zhang, A. Odena, and H. Lee, "Consistency regularization for generative adversarial networks," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [10] N.-T. Tran, V.-H. Tran, N.-B. Nguyen, T.-K. Nguyen, and N.-M. Cheung, "On data augmentation for GAN training," *IEEE Trans. Image Process.*, vol. 30, pp. 1882–1897, Jan. 2021.
- [11] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 1597–1607.
- [12] C. Stein, "A bound for the error in the normal approximation to the distribution of a sum of dependent random variables," in *Proc. Berkeley Symp. Math. Statist. Probability*, 1972, pp. 583–602.
- [13] P. Vincent, "A connection between score matching and denoising autoencoders," *Neural Comput.*, vol. 23, no. 7, pp. 1661–1674, 2011.
- [14] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [16] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," 2017, *arXiv:1708.04552*.
- [17] H.-Y. Tseng, L. Jiang, C. Liu, M.-H. Yang, and W. Yang, "Regularizing generative adversarial networks under limited data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 7917–7927.
- [18] Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, "StarGAN V2: Diverse image synthesis for multiple domains," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8188–8197.
- [19] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [20] S. Yun, D. Han, S. Chun, S. J. Oh, Y. Yoo, and J. Choe, "CutMix: Regularization strategy to train strong classifiers with localizable features," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 6022–6031.
- [21] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan, "AugMix: A simple data processing method to improve robustness and uncertainty," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [22] J. Cao, L. Hou, M.-H. Yang, R. He, and Z. Sun, "ReMix: Towards image-to-image translation with limited data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 15018–15027.
- [23] Y. Bengio et al., "Deep learners benefit more from out-of-distribution examples," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2011, pp. 164–172.
- [24] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le, "AutoAugment: Learning augmentation strategies from data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 113–123.
- [25] S. Lim, I. Kim, T. Kim, C. Kim, and S. Kim, "Fast autoaugment," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 6665–6675.
- [26] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "RandAugment: Practical automated data augmentation with a reduced search space," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshop*, 2020, pp. 3008–3017.
- [27] X. Zhang, Q. Wang, J. Zhang, and Z. Zhong, "Adversarial autoaugment," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [28] T. DeVries and G. W. Taylor, "Dataset augmentation in feature space," in *Proc. Int. Conf. Learn. Representations Workshop*, 2017.
- [29] C. Beckham et al., "Adversarial mixup resynthesizers," in *Proc. Int. Conf. Learn. Representations Workshop*, 2019.
- [30] A. Dabouei, S. Soleymani, F. Taherkhani, and N. M. Nasrabadi, "SuperMix: Supervising the mixing data augmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 13794–13803.
- [31] J.-H. Kim, W. Choo, and H. O. Song, "Puzzle mix: Exploiting saliency and local statistics for optimal mixup," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 5275–5285.
- [32] A. F. M. S. Uddin, M. S. Monira, W. Shin, T. Chung, and S.-H. Bae, "SaliencyMix: A saliency guided data augmentation strategy for better regularization," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [33] H. Guo, Y. Mao, and R. Zhang, "Mixup as locally linear out-of-manifold regularization," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 3714–3722.
- [34] Y. Kalantidis, M. B. Sariyildiz, N. Pion, P. Weinzaepfel, and D. Larlus, "Hard negative mixing for contrastive learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 1829.
- [35] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of wasserstein GANs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5767–5777.
- [36] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 214–223.
- [37] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet, "Are GANs created equal? A large-scale study," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 698–707.
- [38] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for GANs do actually converge?," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 3481–3490.
- [39] N.-T. Tran, V.-H. Tran, N.-B. Nguyen, T.-K. Nguyen, and N.-M. Cheung, "Towards good practices for data augmentation in GAN training," 2020, *arXiv:2006.05338*.
- [40] Z. Zhao, Z. Zhang, T. Chen, S. Singh, and H. Zhang, "Image augmentations for GAN training," 2020, *arXiv:2006.02595*.
- [41] M.-Y. Liu et al., "Few-shot unsupervised image-to-image translation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 10551–10560.


- [42] K. Saito, K. Saenko, and M.-Y. Liu, "COCO-FUNIT: Few-shot unsupervised image translation with a content conditioned style encoder," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 382–398.
- [43] Q. Gu, G. Wang, M. T. Chiu, Y.-W. Tai, and C.-K. Tang, "LADN: Local adversarial disentangling network for facial makeup and de-makeup," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 10480–10489.
- [44] E. Zakhharov, A. Shysheya, E. Burkov, and V. Lempitsky, "Few-shot adversarial learning of realistic neural talking head models," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9458–9467.
- [45] M. Wang et al., "Example-guided style-consistent image synthesis from semantic labeling," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1495–1504.
- [46] T.-C. Wang, M.-Y. Liu, A. Tao, G. Liu, J. Kautz, and B. Catanzaro, "Few-shot video-to-video synthesis," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 5014–5025.
- [47] X. Han, Z. Wu, Z. Wu, R. Yu, and L. S. Davis, "VITON: An image-based virtual try-on network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7543–7552.
- [48] A. Shocher, N. Cohen, and M. Irani, "Zero-shot super-resolution using deep internal learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 3118–3126.
- [49] Y. Zhou, Z. Zhu, X. Bai, D. Lischinski, D. Cohen-Or, and H. Huang, "Non-stationary texture synthesis by adversarial expansion," in *Proc. Int. Conf. Comput. Graph. Interactive Techn.*, 2018, Art. no. 49.
- [50] A. Shocher, S. Bagon, P. Isola, and M. Irani, "InGAN: Capturing and remapping the "DNA," of a natural image," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 4492–4501.
- [51] T. R. Shaham, T. Dekel, and T. Michaeli, "SinGAN: Learning a generative model from a single natural image," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 4569–4579.
- [52] T. Hinz, M. Fisher, O. Wang, and S. Wermter, "Improved techniques for training single-image GANs," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2021, pp. 1300–1309.
- [53] T. Park, A. A. Efros, R. Zhang, and J.-Y. Zhu, "Contrastive learning for unpaired image-to-image translation," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 319–345.
- [54] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 11895–11907.
- [55] Y. Song and S. Ermon, "Improved techniques for training score-based generative models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 1043.
- [56] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [57] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 574.
- [58] L. Zhang, Z. Deng, K. Kawaguchi, A. Ghorbani, and J. Zou, "How does mixup help with robustness and generalization?," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [59] O. Chapelle, J. Weston, L. Bottou, and V. Vapnik, "Vicinal risk minimization," in *Proc. Adv. Neural Inf. Process. Syst.*, 2000, pp. 395–401.
- [60] A. Hyvärinen and P. Dayan, "Estimation of non-normalized statistical models by score matching," *J. Mach. Learn. Res.*, vol. 6, no. 4, pp. 695–709, 2005.
- [61] G. Lin, A. Milan, C. Shen, and I. Reid, "RefineNet: Multi-path refinement networks for high-resolution semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1925–1934.
- [62] V. Dumoulin, J. Shlens, and M. Kudlur, "A learned representation for artistic style," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [63] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local nash equilibrium," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6626–6637.
- [64] M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton, "MMD demystifying GANs," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [65] O. Russakovsky et al., "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.
- [66] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training GANs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2234–2242.
- [67] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 586–595.
- [68] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2223–2232.
- [69] Q. Mao, H.-Y. Lee, H.-Y. Tseng, S. Ma, and M.-H. Yang, "Mode seeking generative adversarial networks for diverse image synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1429–1437.
- [70] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2818–2826.
- [71] F. Luan, S. Paris, E. Shechtman, and K. Bala, "Deep photo style transfer," 2017, *arXiv:1703.07511*.
- [72] J. Li, R. Socher, and S. C. H. Hoi, "DivideMix: Learning with noisy labels as semi-supervised learning," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [73] K. Sohn et al., "FixMatch: Simplifying semi-supervised learning with consistency and confidence," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 51.
- [74] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "MixMatch: A holistic approach to semi-supervised learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 5049–5059.
- [75] J. T. Guibas, T. S. Virdi, and P. S. Li, "Synthetic medical images from dual generative adversarial networks," 2017, *arXiv:1709.01872*.
- [76] P. Costa et al., "End-to-end adversarial retinal image synthesis," *IEEE Trans. Med. Imag.*, vol. 37, no. 3, pp. 781–791, Mar. 2018.



Jie Cao received the BE degree in automation from North China Electric Power University, Beijing, China, in 2016. He is currently working toward the PhD degree with the Center for Research on Intelligent Perception and Computing, National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing. His research interests include biometrics, pattern recognition, computer vision, and machine learning.



Mandi Luo received the BE degree in automation engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2017, and the BSc and MSc degrees in electronic engineering from Katholieke Universiteit te Leuven, Leuven, Belgium, in 2017 and 2018, respectively. She is currently working toward the PhD degree in computer application technology with the University of Chinese Academy of Sciences, Beijing, China. Her research interests include biometrics, pattern recognition, and computer vision.



Junchi Yu received the BE degree from the School of Electrical Engineering, Wuhan University, Wuhan, China. He is currently working toward the PhD degree with the Institute of Automation, Chinese Academy of Sciences, Beijing, China. His main research interests include interpretable geometric deep learning and AI+X. He has published several papers on machine learning top conferences/journals including ICLR, CVPR, IJCAI, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *Pattern Recognition*, etc. He has served as a reviewer for NeurIPS, CVPR and *Pattern Recognition* etc.



Ming-Hsuan Yang (Fellow, IEEE) is a professor in electrical engineering and computer science with the University of California, Merced, Merced, California. He served as a program co-chair of IEEE International Conference on Computer Vision (ICCV), in 2019, program co-chair of Asian Conference on Computer Vision (ACCV), in 2014, and general co-chair of ACCV 2016. He served as an associate editor of *IEEE Transactions on Pattern Analysis and Machine Intelligence* from 2007 to 2011 and is an associate

editor of *International Journal of Computer Vision, Image and Vision Computing*, and *Journal of Artificial Intelligence Research*. He received an NSF CAREER award, in 2012 and Google Faculty Award, in 2009. He is a fellow of the ACM.



Ran He received BE degree in computer science from the Dalian University of Technology, the MS degree in computer science from the Dalian University of Technology, and the PhD degree in pattern recognition and intelligent systems from CASIA, in 2001, 2004, and 2009, respectively. Since September 2010, he has joined NLPR, where he is currently a full professor. His research interests focus on information theoretic learning, pattern recognition, and computer vision. He served as an associate editor of *IEEE Transactions on Image Processing* and *IEEE Transactions on Biometrics, Behavior, and Identity Science* and as an area chair of CVPR/ECCV/NeurIPS. He is a fellow of the IAPR.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.**