

# Diverse Rare Sample Generation with Pretrained GANs

Subeen Lee<sup>1</sup>, Jiyeon Han<sup>1</sup>, Soyeon Kim<sup>1</sup> and Jaesik Choi<sup>1,2</sup>

<sup>1</sup>Korea Advanced Institute of Science and Technology (KAIST), South Korea

<sup>2</sup>INEEJI, South Korea

{sbrblee7, j.han, soyeon.k, jaesik.choi}@kaist.ac.kr

## Abstract

Deep generative models are proficient in generating realistic data but struggle with producing rare samples in low density regions due to their scarcity of training datasets and the mode collapse problem. While recent methods aim to improve the fidelity of generated samples, they often reduce diversity and coverage by ignoring rare and novel samples. This study proposes a novel approach for generating diverse rare samples from high-resolution image datasets with pretrained GANs. Our method employs gradient-based optimization of latent vectors within a multi-objective framework and utilizes normalizing flows for density estimation on the feature space. This enables the generation of diverse rare images, with controllable parameters for rarity, diversity, and similarity to a reference image. We demonstrate the effectiveness of our approach both qualitatively and quantitatively across various datasets and GANs without retraining or fine-tuning the pretrained GANs.

**Code** — <https://github.com/sbrblee/DivRareGen>

## 1 Introduction

Deep generative models have shown impressive generative capabilities across various domains. The primary focus of current generative model research is on enhancing the fidelity of generated images (DeVries, Drozdal, and Taylor 2020; Karras, Laine, and Aila 2019; Azadi et al. 2018; Turner et al. 2019). However, these approaches often compromise sample diversity and encounter difficulties generating rare samples, primarily due to their limited representation in the training dataset (Sehwag et al. 2022; Lee et al. 2021). In GANs, this issue is worsened by the mode collapse problem (Thanh-Tung and Tran 2020). Investigating rare samples is crucial for several reasons: it enhances the creation of synthetic datasets that embody diversity and creativity (Sehwag et al. 2022; Agarwal, D’souza, and Hooker 2022), ensures fairness in generative processes (Teo, Abdollahzadeh, and Cheung 2023; Hwang et al. 2020), and aligns with the human tendency to favor unique features (Snyder and Lopez 2001; Lynn and Harris 1997). Additionally, exploring edge cases and unusual scenarios is essential in various domains, such as drug discovery or molecular design

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(Sagar et al. 2023; Zeng et al. 2022), and natural hazard analysis (Ma, Mei, and Xu 2024).

Several studies have been conducted to enhance the overall diversity of GAN-generated outputs and to promote the generation of rare samples (Chang et al. 2024; Allahyani et al. 2023; Humayun, Balestrieri, and Baraniuk 2022, 2021; Heyrani Nobari, Rashad, and Ahmed 2021; Ghosh et al. 2018; Tolstikhin et al. 2017; Srivastava et al. 2017; Chen et al. 2016). Due to the high computational cost of training GANs (Karras, Laine, and Aila 2019; Brock, Donahue, and Simonyan 2018), techniques without model retraining are appealing. For example, Humayun, Balestrieri, and Baraniuk (2022) proposed a resampling technique for pretrained GANs with a controllable fidelity-diversity trade-off parameter. However, the proposed method requires extensive sampling to cover the data manifold fully. On the other hand, Chang et al. (2024) proposed a method to obtain diverse samples that satisfy text conditions by optimizing latent vectors with a quality-diversity objective.

Han et al. (2023) proposes a rarity score for samples using relative density measures based on  $k$ -nearest neighbor ( $k$ -NN) manifolds in the feature space of pretrained classifiers. Although  $k$ -NN density estimation is straightforward and reliable (Naeem et al. 2020; Kynkäänniemi et al. 2019), its non-differentiable nature complicates gradient-based optimization. In contrast, normalizing flows (NFs) excel at high-dimensional density estimation in differentiable form (Papamakarios et al. 2021; Kingma and Dhariwal 2018; Dinh, Sohl-Dickstein, and Bengio 2016; Dinh, Krueger, and Bengio 2014). We employ NFs for density estimation in the feature space, which incurs a lower training cost compared to retraining or fine-tuning GANs, and analyze how NF-based density estimates relate to the rarity score.

This study aims to generate diverse rare samples for a given high-resolution image datasets and GANs. Our method does not require any fine-tuning or retraining of the GANs but instead explores the latent space of the given model through gradient-based optimization. Our contributions are as follows:

- Our method can generate diverse versions of rare images utilizing the multi-start method in optimization, without being trapped at the same local optima.
- Rarity and diversity of generated images and similarity to the initial image can be controlled via a multi-objective

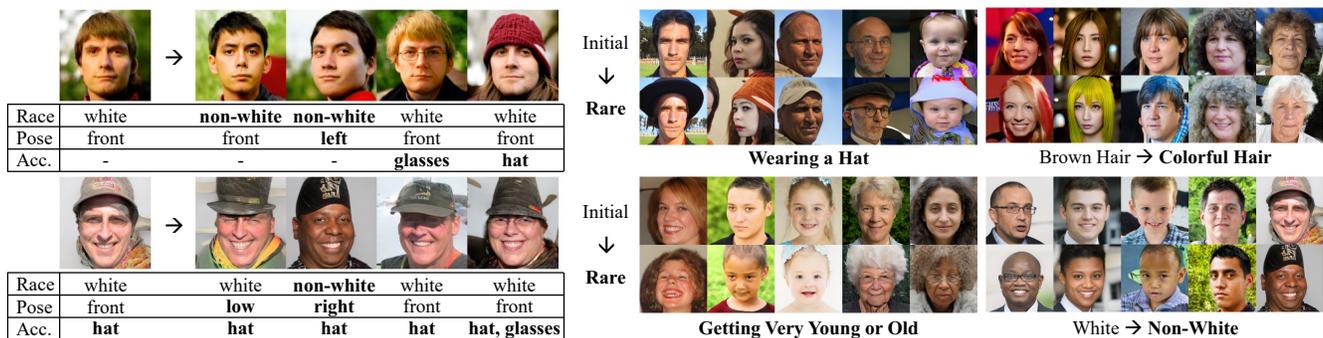


Figure 1: Examples of rare samples generated by our method. **Left:** Our method produces diverse rare images for a single reference, with variations even within the same rare attribute (e.g., hats of different shapes and colors). **Right:** Generated rare attributes include accessories like hats, non-brown hair colors, extreme ages, and non-white races. “Pose” refers to head orientation, and “Acc.” denotes accessories. Rare attributes are highlighted in bold.

optimization framework.

- We demonstrate the effectiveness of our method with various high-resolution image datasets and GANs, both qualitatively and quantitatively.

## 2 Related Work

**Rare Generation** Han et al. (2023) introduced the rarity score to quantify the uniqueness of individual samples, distinguishing it from conventional metrics that primarily evaluate fidelity or diversity in generated samples (Kynkäänniemi et al. 2019; Zhang et al. 2018; Heusel et al. 2017). The rarity score is defined as the minimum  $k$ -nearest neighbor distance ( $k$ -NND) among real samples that are closer to the target sample, with higher scores indicating lower density within the real data manifold. However, obtaining rare samples has received limited attention. Sehwag et al. (2022) addressed this by leveraging a pretrained classifier to estimate likelihoods and adapting the sampling process of diffusion probabilistic models to target low-density regions while maintaining fidelity. Their method focuses on sampling from regions far from class mean vectors in the feature space and penalizes deviations from the overall mean vector of real data. However, this approach is class-conditional and depends on a Gaussian likelihood function.

On the other hand, Humayun, Balestrieri, and Baraniuk (2022) tackled the mode collapse issue in GANs with Polarity sampling, a fidelity-diversity controllable resampling strategy for pretrained GANs. It approximates the GAN’s output space using continuous piecewise affine splines. By tuning  $\rho$ , sampling can focus on modes ( $\rho < 0$ ) or anti-modes ( $\rho > 0$ ), with higher  $\rho$  increasing diversity by targeting low-density regions. However, it does not guarantee the fidelity of the selected samples and requires extensive sampling and Jacobian matrix computations, leading to high computational costs.

**Quality-Preserved Diverse Generation Using Pretrained GANs** Generating rare samples is important, but maintaining quality is also crucial for their usefulness (Amabile 2018). Achieving diverse, high-fidelity samples is similar

to finding multiple solutions in combinatorial optimization. Chang et al. (2024) addressed this by proposing a quality-diversity algorithm that updates the latent vector to balance quality and diversity, using the CLIP score (Radford et al. 2021) to measure similarity and diversity. In our work, we also optimize the latent vector to generate diverse samples; however, we prioritize rarity as the main objective rather than quality and use Euclidean distance in arbitrary feature spaces, avoiding the additional text constraints required by the CLIP score. To prevent low-fidelity samples, we apply a constraint to keep the sampled data within the real data manifold.

**Reference-based Generation** Finding diverse rare variations of a given initial image relates to reference-based image generation, encompassing tasks like domain adaptation (Yang et al. 2023), editing (Xia et al. 2023), and conditional generation (Casanova et al. 2021). While these approaches involve additional training costs for each attribute (Yang et al. 2023) or reference (Xia et al. 2023), or require a different GAN training scheme (Casanova et al. 2021), our method only requires a single training phase for the density estimator across multiple references with pretrained GANs.

**Density Estimation for Images** The rarity score identifies samples in low-density regions of the real data manifold, making it valuable for detecting rare generations. However, the non-differentiable nature of  $k$ -NN-based manifold estimation complicates gradient-based optimization for directly obtaining rare samples. In contrast, extensive research has been conducted to estimate density in high-dimensional spaces. Normalizing flows (NFs), as likelihood-based probabilistic models, use a sequence of invertible functions to transform a simple density into a complex one, potentially representing multi-modal distributions while preserving data relationships (Papamakarios et al. 2021; Kingma and Dhariwal 2018; Dinh, Sohl-Dickstein, and Bengio 2016; Dinh, Krueger, and Bengio 2014). While NFs may struggle with out-of-distribution data due to their focus on low-level features (Kirichenko, Izmailov, and Wilson 2020), training them on the feature space of a pretrained network—which

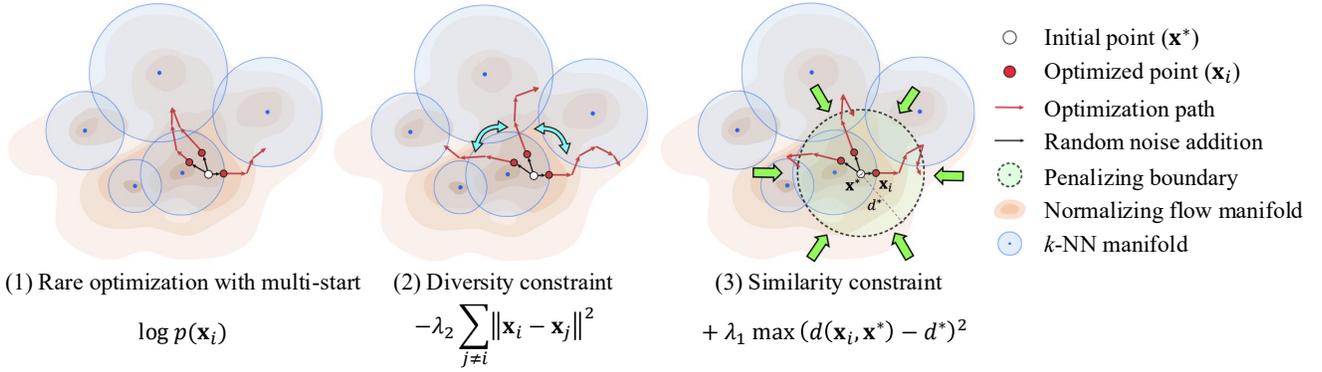


Figure 2: Schematic diagram for the objective function of our method.  $\mathbf{x}^* = f(G(\mathbf{z}^*))$  and  $\mathbf{x}_i = f(G(\mathbf{z}_i))$  for brevity.

includes high-level semantic information—can help mitigate this issue (Esser, Rombach, and Ommer 2020).

**Multi-Start Method for Diverse Solutions** We frame our problem of obtaining diverse rare samples for each given reference as identifying multiple local minima around the reference in the data distribution. A straightforward approach to this problem is the multi-start method, an algorithm that iteratively searches for local optima starting from multiple initial points (Feo and Resende 1995; Rochat and Taillard 1995; Rinnooy Kan and Timmer 1987). This method selects several starting positions and applies a local search algorithm to each, aiming to locate distinct local optima. Although easy to implement, it does not always ensure that different starting points result in different local optima (Tarek and Huang 2022). To address this issue, we add the diversity and similarity constraints to the objective function, ensuring that each initial point converges to a different minimum.

### 3 Methods

#### 3.1 Problem Statement

Given a GAN generator  $G = G(\mathbf{z})$ , an arbitrary initial latent vector  $\mathbf{z}^* \in \mathbb{R}^m$ , a feature extractor  $f = f(\mathbf{I})$ , and a density estimator  $g = g(\mathbf{x})$ , our objective is to find a set of latent vectors  $\{\mathbf{z}_i\}_{i=1}^N$  that generates diverse rare samples which are similar to the image generated from  $\mathbf{z}^*$  (referred to as the reference for the rest of the paper). Here,  $\mathbf{I} \in \mathbb{R}^{w \times h \times 3}$  and  $\mathbf{x} \in \mathbb{R}^n$  denote an image and a feature vector, respectively. For simplicity, we denote  $\mathbf{x}^* = f(G(\mathbf{z}^*))$  and  $\mathbf{x}_i = f(G(\mathbf{z}_i))$ . In this study, we define similarity as the Euclidean distance in the feature space,  $d(\mathbf{x}_1, \mathbf{x}_2) = \|\mathbf{x}_1 - \mathbf{x}_2\|$ , a metric shown to align well with human perception (Zhang et al. 2018).

We propose a multi-objective optimization framework that integrates rarity, diversity, and similarity regularization, as illustrated in Fig. 2, and provide a detailed explanation in the subsequent sections.

#### 3.2 Rare Sample Generation

For rarity, the density estimated from  $g$  is used. NFs are employed due to their remarkable performance in high-dimensional density estimation, though any differentiable

density estimator can be applied. NFs provide the exact log-likelihood of individual samples, allowing us to directly define the objective function as  $\mathcal{L}_{rare}(\mathbf{x}) = g(\mathbf{x}) = \log p(\mathbf{x})$  to minimize. To control the similarity between the generated rare image and the reference image, we incorporate a regularization term inspired by Chang et al. (2024). Specifically, we define the similarity loss as  $\mathcal{L}_{sim}(\mathbf{x}) = (\max(d(\mathbf{x}, \mathbf{x}^*), d^*) - d^*)^2$ , where this term penalizes samples that exceed a predefined boundary, referred to as the *penalizing boundary* throughout this paper, defined by  $d^*$ . Specifically, we use the distance to  $k'$ -nearest neighbor in the fake  $k$ -NN manifold for  $d^*$ . We also strictly accept the sample inside this boundary as well as inside the real  $k$ -NN manifold  $\Phi_{real}$  in optimization. The optimization goal is formulated by combining the two objectives—rarity and similarity regularization—as follows.

$$\begin{aligned} \min_{\mathbf{x}=f(G(\mathbf{z}))} \quad & \mathcal{L}_{rare}(\mathbf{x}) + \lambda_1 \mathcal{L}_{sim}(\mathbf{x}) \\ \text{subject to } \mathbf{x} \in \Phi_{real} \text{ and } & d(\mathbf{x}, \mathbf{x}^*) \leq d^* \end{aligned} \quad (1)$$

#### 3.3 Diverse Rare Sample Generation

We utilize the multi-start method to obtain diverse rare images by adding small random noises to  $\mathbf{z}^*$ , generating multiple starting points for optimization. Specifically,  $\{\mathbf{z}_i\}_{i=1}^N$  are initialized as  $\mathbf{z}_i = \mathbf{z}^* + \epsilon$ , where  $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 I)$ . However, as shown in Fig. 2 (1), this does not guarantee obtaining different rare images, but they might converge to the same local minima. To address this issue, we add a diversity constraint to the objective,  $\mathcal{L}_{div}(\mathbf{x}_i) = -\sum_{j \neq i} d(\mathbf{x}_i, \mathbf{x}_j)^2$ , ensuring that the samples are far from each other, as shown in Fig. 2 (2). This term is inspired by the expected distances in feature space, similar to the concept of Maximum Mean Discrepancy (MMD) (Gretton et al. 2006). Combining all objectives, the multi-objective optimization problem is formulated as follows.

$$\begin{aligned} \min_{\mathbf{x}_i=f(G(\mathbf{z}_i))} \quad & \mathcal{L}_{rare}(\mathbf{x}_i) + \lambda_1 \mathcal{L}_{sim}(\mathbf{x}_i) + \lambda_2 \mathcal{L}_{div}(\mathbf{x}_i) \\ \text{subject to } \mathbf{x}_i \in \Phi_{real} \text{ and } & d(\mathbf{x}_i, \mathbf{x}^*) \leq d^* \end{aligned} \quad (2)$$

## 4 Experimental Results

We validate our proposed method using high-resolution image datasets with a resolution of  $1024 \times 1024$ , including Flickr Faces HQ (FFHQ) (Karras, Laine, and Aila 2019), Animal Faces HQ (AFHQ) (Choi et al. 2020), and Metfaces (Karras et al. 2020a). StyleGAN2 with config-f (Karras et al. 2020b) and StyleGAN2-ADA (Karras et al. 2020a) are utilized. For feature extraction, we employ the VGG16-fc2 architecture (Simonyan and Zisserman 2015). As the density estimator, the Glow architecture (Kingma and Dhariwal 2018) is adapted to accommodate the high dimensionality of the feature space. The optimization is performed using the Adam optimizer (Diederik 2014) with a learning rate of  $2 \times 10^{-2}$ , combined with a StepLR scheduler. The best optimization results are recorded when the lowest loss is achieved according to Equation (2). Additional details including computational cost are provided in Appendix B.

### 4.1 Generation of Rare Facial Attributes with StyleGAN2 (FFHQ-StyleGAN2)

**Quantitative Results** As a baseline, 10,000 synthetic samples are generated using latent vectors from StyleGAN2 with a truncation parameter of  $\psi = 1.0$ . With our method, we generate ten rare samples for each of 1,000 initial latent vectors from the baselines, using parameters  $\lambda_1 = 30.0$ ,  $\lambda_2 = 0.002$ ,  $\sigma = 0.1$ , and  $k' = 100^1$ . The choice of the parameters is explained in Appendix C. For Polarity sampling, 250,000 latent vectors and their corresponding Jacobian matrices are obtained from the authors (Humayun, Balestrieri, and Baraniuk 2022), and 10,000 samples are re-sampled using  $\rho = [1.0, 5.0]$  (anti-mode sampling).

Results are evaluated with metrics including the Rarity Score (RS) (Han et al. 2023), precision (Prec.) and recall (Rec.) for fidelity and diversity (Kynkäänniemi et al. 2019), LPIPS score for diversity (Zhang et al. 2018), and FID score (Heusel et al. 2017), as shown in Table 1. Each metric is computed using 10,000 generated samples, with the LPIPS score averaged over 10,000 random sample pairs. Significant differences in LPIPS scores between sampling methods are confirmed using an unpaired t-test. For the real  $k$ -NN manifold,  $k = 3$  is used.

Our method improves both rarity and diversity compared to the baseline, even when the optimization uses only 10% of the baseline samples as references. The FID score decreases because more samples are generated in low-density regions, reducing samples near the data distribution’s modes. Polarity sampling also enhances rarity and diversity but sacrifices precision, as it primarily targets low-density regions in the GAN’s output space rather than the real manifold, often generating out-of-distribution samples (structural zeros; (Kim and Bansal 2023)). Furthermore, in contrast to our objective of generating rare samples similar to a given reference, Polarity sampling is not designed for it. Finally, since Polarity sampling operates by resampling from an initial set, its diversity is heavily dependent on the size of that initial

<sup>1</sup>For the fake  $k$ -NN manifold estimation, 10,000 generated samples are used.

Model	RS $\uparrow$	Prec. $\uparrow$	Rec. $\uparrow$	LPIPS $\uparrow$	FID $\downarrow$
<b>Baseline</b>	18.88	0.69	0.56	0.73	<b>4.17</b>
<b>Polarity (<math>\rho = 1.0</math>)</b>	24.71	0.39	0.70	0.75	33.28
<b>Polarity (<math>\rho = 5.0</math>)</b>	<b>24.83</b>	0.38	<b>0.71</b>	0.75	34.11
<b>Ours</b>	23.50	<b>0.92</b>	0.65	<b>0.76</b>	7.38

Table 1: Quantitative evaluation for Section 4.1.

Model-Based Attribute	FFHQ	Reference	Ours(%)
<b>Age</b>	<i>0-9</i>	7.51	<b>8.56</b>
	<i>10-69</i>	92.23	<b>92.20</b>
	<i>Over70</i>	0.25	<b>0.20</b>
<b>Gender</b>	<i>Male</i>	45.46	<b>48.50</b>
	<i>Female</i>	54.53	<b>52.55</b>
<b>Race</b>	<i>White</i>	62.97	<b>64.66</b>
	<i>NonWhite</i>	37.03	<b>35.33</b>
<b>HeadPose</b>	<i>Front</i>	60.98	<b>55.25</b>
	<i>NotFront</i>	39.01	<b>44.75</b>

Table 2: Percentage of age, gender, race, and head pose attributes predicted by FaceXformer for Section 4.1.

LFWA Attribute	FFHQ	Reference	Ours(%)
<10%	<i>PaleSkin</i>	0.12	<b>0.20</b>
	<i>Mustache</i>	0.39	<b>0.30</b>
	<i>Bald</i>	0.88	<b>1.20</b>
	<i>WearingNecktie</i>	0.96	<b>0.80</b>
	<i>PointyNose</i>	1.42	<b>1.80</b>
	<i>GrayHair</i>	2.43	<b>1.90</b>
	<i>RecedingHairline</i>	3.32	<b>2.40</b>
	<i>BigLips</i>	4.66	<b>5.90</b>
	<i>BlondHair</i>	5.04	<b>5.20</b>
	<i>Eyeglasses</i>	6.10	<b>5.70</b>
	<i>WearingHat</i>	8.07	<b>7.80</b>
>10%	<i>BrownHair</i>	19.76	<b>22.92</b>

Table 3: Percentage of LFWA attributes predicted by FaceXformer for Section 4.1. Sorted in descending order of FFHQ(%). The entire table is in Table 11.

set. An additional comparison with Polarity sampling is in Appendix G.

**Generated Rare Facial Attributes** Our method successfully increases the percentages of rare attributes as shown in Tables 2 and 3. To identify rare facial attributes in FFHQ, we employ FaceXFormer (Narayan et al. 2024), which provides multiple face-related features including age, gender, race, head pose, and the attributes from Deep Learning Face Attributes in the Wild (LFWA) dataset (Liu et al. 2015). Real data attributes with lower percentages include extreme ages (very young or old), male gender, non-white races, non-frontal head poses<sup>2</sup>, non-natural skin colors, hairless or no hair, hair colors other than brown, and accessories.

<sup>2</sup>The head pose is predicted in the form of  $(\theta_1, \theta_2, \theta_3) = (\text{pitch}, \text{yaw}, \text{roll})$ . We define *Front* as the head pose with  $-15^\circ < \theta_1, \theta_2, \theta_3 < 15^\circ$ .

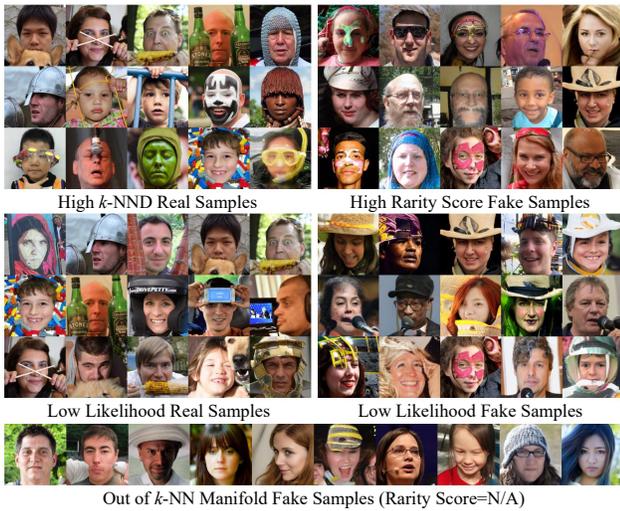


Figure 3: Examples of high- and low-density real and fake samples for Section 4.1.

Additionally, other rare attributes can be identified qualitatively. We select and visualize the top- and bottom-ranked samples based on  $k$ -NN-based and likelihood-based density estimates in Fig. 3. For real samples,  $k$ -NND (Loftsgaarden and Quesenberry 1965) is employed, while the rarity score (Han et al. 2023) is used for fake samples. Likelihoods are estimated using the NF model, excluding samples outside the real  $k$ -NN manifold. Rare samples exhibit characteristics such as objects obscuring faces, face painting, various hats, colorful eyeglasses, and artifacts. Notably, samples with undefined (N/A) rarity scores may include high-fidelity, artifact-free images, which arise from underestimated regions in the  $k$ -NN manifold.

**Qualitative Results** As shown in Fig. 1 (Right) and 4, our method generates samples with rare attributes, including hats, hair colors other than natural brown, very young or old age, non-white races such as Black, Indian, and Asian, non-frontal head poses, eyeglasses, bald or receding hairline, colorful backgrounds or T-shirts, hair accessories, and unique skin colors. Moreover, the rare samples generated by our method show diversity, as shown in Fig. 1 (Left) and 5. Starting from initial vectors with small noise variations, we generate diverse rare images that retain perceptual similarity to the reference. Additional examples are provided in Appendix D.

## 4.2 Animal Face and Artwork Generation with StyleGAN2-ADA

**Quantitative & Qualitative Results** As a baseline, 5,000 synthetic samples are generated using latent vectors from StyleGAN2-ADA with a truncation parameter of  $\psi = 1.0$ . With our method, we generate five rare samples for each of 1,000 initial latent vectors from the baseline, using parameters  $\lambda_1 = 200.0$ ,  $\lambda_2 = 0.02$ ,  $\sigma = 0.01$  and  $k' = 100$ .

We also evaluate the results using various metrics, as presented in Table 4. Given that the AFHQ and Met-



Figure 4: Examples of rare samples generated by our method for Section 4.1.

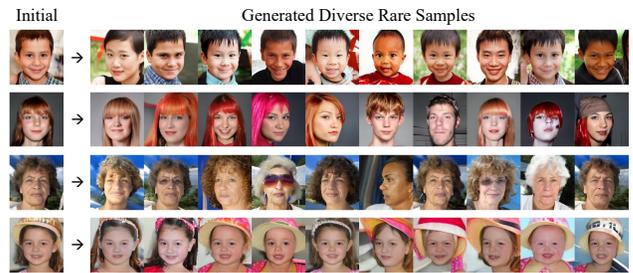


Figure 5: Examples of diverse rare samples generated by our method for Section 4.1.

Faces datasets are relatively small, we use the KID score (Bińkowski et al. 2018) instead of the FID score, as the KID score is inherently unbiased (Karras et al. 2020a). Our method effectively enhances rarity and diversity across all three datasets compared to the baseline. In Fig. 7, we visualize the examples generated by our method. More examples are in Fig.14 and 15.

### Generated Rare Attributes of Animal Face and Artwork

To identify the rare cat and dog breeds in AFHQ datasets, we employed Model Soups (Wortsman et al. 2022) for zero-shot classification of cat and dog classes in the ImageNet dataset (Deng et al. 2009), which includes five cat classes and 120 dog classes. In the AFHQ Dog dataset, each dog class represents less than 5% of the total, so we grouped the dogs into broader categories based on appearance: Toy, Hound, Scent Hound, Terrier, Sporting, Non-Sporting, Herding, and Working. Further details are provided in Appendix E. The classified result is shown in Table 5, and our method successfully increases the percentages of minor classes.

We also apply the FaceXFormer and observe similar rare attributes in FFHQ, as shown in Table 6.

To further identify rare attributes within the datasets, we visualize the high- and low-likelihood samples in Fig. 6. For the AFHQ-cat dataset, high-likelihood samples predominantly consist of brown-colored Tabby cats, whereas low-

Data	Model	RS $\uparrow$	Prec. $\uparrow$	Rec. $\uparrow$	LPIPS $\uparrow$	KID $\downarrow$ $\times 10^3$
AFHQ Cat	Ref.	17.47	<b>0.76</b>	0.49	0.73	<b>0.46</b>
	Ours	<b>20.71</b>	<b>0.76</b>	<b>0.68</b>	<b>0.78</b>	2.03
AFHQ Dog	Ref.	21.19	0.77	0.56	0.75	<b>1.10</b>
	Ours	<b>27.32</b>	<b>0.85</b>	<b>0.76</b>	<b>0.80</b>	2.73
MetFaces	Ref.	16.90	0.80	0.44	0.74	<b>0.97</b>
	Ours	<b>21.25</b>	<b>0.86</b>	<b>0.62</b>	<b>0.78</b>	2.15

Table 4: Quantitative evaluation for Section 4.2.

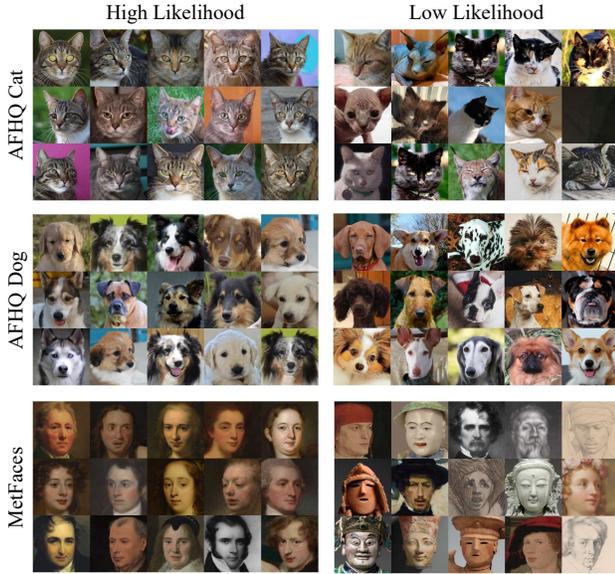


Figure 6: Examples of high- and low-likelihood real samples from the AFHQ Cat, Dog, and MetFaces dataset.

likelihood samples encompass a broader range of classes. In the AFHQ-dog dataset, high-likelihood samples are primarily drawn from the Herding and Sporting groups, including breeds such as Shetland Sheepdogs, Collies, and Retrievers. In contrast, low-likelihood samples span a variety of groups and exhibit greater diversity in head poses, backgrounds, and facial expressions. In the Metfaces dataset, high-likelihood samples predominantly include European-style oil paintings, while low-likelihood samples include statues, drawings, and other styles of paintings. As shown in Fig. 7, such rare attributes can be also observed in the results of our method.

### 4.3 Ablation Study on the Objective Function

Our objective function includes three components: rarity, similarity, and diversity terms. To evaluate their effectiveness, we conduct an ablation study using the FFHQ dataset and StyleGAN2, keeping the density estimator and parameters consistent. We optimize ten samples for each of the 100 initial latent vectors across different objective combinations.

First, we assess results using only the  $\mathcal{L}_{rare}$ . Adding the  $\mathcal{L}_{sim}$  ensures that samples stay within a similarity boundary to the reference, potentially finding rarer and more diverse

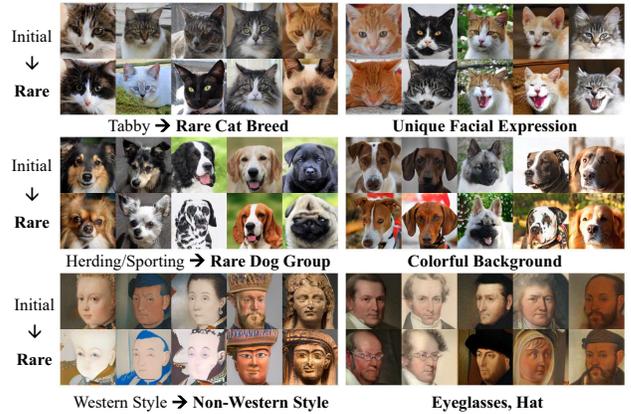


Figure 7: Examples of rare samples generated by our method for Section 4.2.

ImageNet Attribute	Real	Reference	Ours(%)
Cat	Tabby	56.76	<b>58.00</b>
	Others	43.23	<b>44.90</b>
Dog	Herding	21.48	<b>21.70</b>
	Sporting	18.50	<b>17.30</b>
	Working	15.44	14.60
	Toy	14.93	17.30
	Terrier	9.11	<b>9.60</b>
	Non-Sporting	7.02	5.90
	Scent Hound	6.92	8.10
Hound	3.05	2.60	<b>3.76</b>

Table 5: Percentage of the cat-related breeds and dog-related groups in ImageNet classes predicted by Model Soups. Sorted in descending order of Real(%). The entire table is in Table 12.

Model-Based Attribute	MetF.	Reference	Ours(%)
Age	<i>0-9</i>	5.19	3.23
	<i>10-69</i>	94.20	<b>95.96</b>
	<i>Over70</i>	0.60	0.80
Gender	<i>Male</i>	42.84	41.71
	<i>Female</i>	57.15	<b>58.28</b>
Race	<i>White</i>	73.64	<b>77.97</b>
	<i>Indian</i>	9.86	<b>9.19</b>
	<i>Black</i>	3.38	2.02
	<i>Asian</i>	2.86	4.04
	<i>Others</i>	10.24	<b>6.76</b>
HeadPose	<i>Front</i>	42.54	<b>46.06</b>
	<i>NotFront</i>	57.45	53.93
LFWA	<i>Eyeglasses</i>	0.07	0.10
	<i>WearingHat</i>	12.34	9.89

Table 6: Percentage of age, gender, race, head pose, *Eyeglasses*, and *WearingHat* attributes predicted by FaceX-Former. MetF. refers to the MetFaces dataset. The entire table is in Table 17.

Objective	RS	LPIPS
$\mathcal{L}_{rare}$	18.99	0.752
$\mathcal{L}_{rare} + \lambda_1 \mathcal{L}_{sim}$	21.11	0.766
$\mathcal{L}_{rare} + \lambda_1 \mathcal{L}_{sim} + \lambda_2 \mathcal{L}_{div}$	<b>21.28</b>	<b>0.768</b>

Table 7: Rarity scores (RS) and LPIPS scores for the ablation study on the objective function.

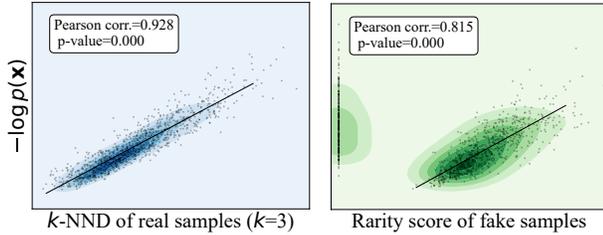


Figure 8: Correlation plot for the  $k$ -NND / rarity score and negative log-likelihood estimated by the normalizing flow.

samples inside the boundary. Finally, incorporating the  $\mathcal{L}_{div}$  completes the full objective. The results in Table 7 demonstrate the effectiveness of each term.

#### 4.4 Relationship with Rarity Score

We use the rarity score (Han et al. 2023) to measure sample rarity and demonstrate that our method improves rarity compared to other sampling methods. Although directly optimizing the rarity score is challenging, our likelihood-based objective effectively guides samples to locally low-density regions. To compare  $k$ -NN-based density measures ( $k$ -NND for real samples and rarity score for fake samples) with NF-estimated density measures, we visualize the scatter plot and compute the Pearson correlation coefficient as represented in Fig. 8. We observe a high Pearson correlation coefficient of 0.928 for real samples and 0.815 for fake samples, with  $p$ -values  $< 10^{-8}$ , excluding samples with undefined rarity scores.

Although the NF estimates likelihood across the feature space, the rarity score is undefined outside the real  $k$ -NN manifold. This allows for out-of-manifold samples with sufficient quality, as shown at the bottom of Fig. 3. In Fig. 9, we visualize an optimization example that starts with an undefined rarity score but eventually gains and increases the rarity score. The sample, initially in an underestimated  $k$ -NN region, becomes rare by moving to a low-likelihood region, altering the reference image to achieve curlier blonde hair and a non-frontal head pose. We plot the NF-estimated density using RBF kernel interpolation and UMAP (McInnes, Healy, and Melville 2018) dimensionality reduction on the real feature space and its inverse transformation function. Further details are provided in Appendix F.

## 5 Conclusion

We proposed a novel algorithm that generates diverse rare samples using multi-start gradient-based optimization, avoiding low-quality samples. Users can control rarity, di-

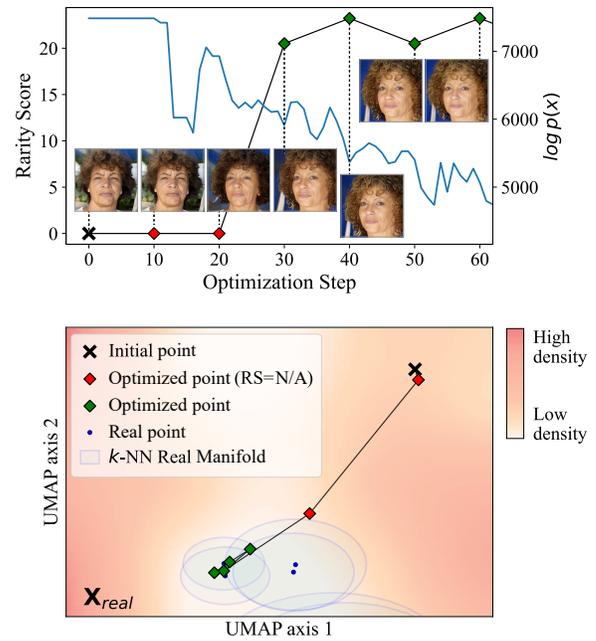


Figure 9: Example of the optimization path with a real  $k$ -NN manifold and a heatmap of likelihoods estimated by the normalizing flow. Notably, the local  $k$ -NN manifold includes only the three nearest real data points ( $k = 3$ ) for each point, rather than the entire manifold.

versity, and similarity to the reference through a multi-objective approach. Our method successfully increased the prevalence of rare attributes in various image generation domains. We also provide an experimental comparison between  $k$ -NN-based and normalizing flow-based density estimation methods. We hope this work contributes to advancing creativity in deep generative models. However, there are some limitations that could be improved. The results rely on the GAN’s capabilities and require an additional density estimator. Exploring other generative models might improve outcomes and eliminate the need for extra training. Additionally, our method alters multiple attributes simultaneously; integrating it with other image manipulation techniques could allow for more controlled manipulation.

## Acknowledgments

This work was partly supported by KAIST-NAVER Hypercreative AI Center, and from the Korean Institute of Information & Communications Technology Planning & Evaluation and the Korean Ministry of Science and ICT under grant agreement No. RS-2019-II190075 (Artificial Intelligence Graduate School Program(KAIST)), No. RS-2022-II220984 (Development of Artificial Intelligence Technology for Personalized Plug-and-Play Explanation and Verification of Explanation), and No.RS-2022-II220184 (Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics).

## References

- Agarwal, C.; D'souza, D.; and Hooker, S. 2022. Estimating example difficulty using variance of gradients. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10368–10378.
- Allahyani, M.; Alsulami, R.; Alwafi, T.; Alafif, T.; Ammar, H.; Sabban, S.; and Chen, X. 2023. DivGAN: A diversity enforcing generative adversarial network for mode collapse reduction. *Artificial Intelligence*, 317: 103863.
- Amabile, T. M. 2018. *Creativity in context: Update to the social psychology of creativity*. Routledge.
- Azadi, S.; Olsson, C.; Darrell, T.; Goodfellow, I.; and Odena, A. 2018. Discriminator rejection sampling. *arXiv preprint arXiv:1810.06758*.
- Bińkowski, M.; Sutherland, D. J.; Arbel, M.; and Gretton, A. 2018. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*.
- Brock, A.; Donahue, J.; and Simonyan, K. 2018. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.
- Casanova, A.; Careil, M.; Verbeek, J.; Drozdal, M.; and Romero Soriano, A. 2021. Instance-conditioned gan. *Advances in Neural Information Processing Systems*, 34: 27517–27529.
- Chang, A.; Fontaine, M. C.; Booth, S.; Matarić, M. J.; and Nikolaidis, S. 2024. Quality-Diversity Generative Sampling for Learning with Synthetic Data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(18): 19805–19812.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in neural information processing systems*, 29.
- Choi, Y.; Uh, Y.; Yoo, J.; and Ha, J.-W. 2020. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8188–8197.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- DeVries, T.; Drozdal, M.; and Taylor, G. W. 2020. Instance selection for gans. *Advances in Neural Information Processing Systems*, 33: 13285–13296.
- Diederik, P. K. 2014. Adam: A method for stochastic optimization. (*No Title*).
- Dinh, L.; Krueger, D.; and Bengio, Y. 2014. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*.
- Esser, P.; Rombach, R.; and Ommer, B. 2020. A disentangling invertible interpretation network for explaining latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9223–9232.
- Feo, T. A.; and Resende, M. G. 1995. Greedy randomized adaptive search procedures. *Journal of global optimization*, 6: 109–133.
- Ghosh, A.; Kulharia, V.; Nambodiri, V. P.; Torr, P. H.; and Dokania, P. K. 2018. Multi-agent diverse generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8513–8521.
- Gretton, A.; Borgwardt, K.; Rasch, M.; Schölkopf, B.; and Smola, A. 2006. A kernel method for the two-sample-problem. *Advances in neural information processing systems*, 19.
- Han, J.; Choi, H.; Choi, Y.; Kim, J.; Ha, J.-W.; and Choi, J. 2023. Rarity Score: A New Metric to Evaluate the Uncommonness of Synthesized Images. In *International Conference on Learning Representations (ICLR)*. International Conference on Learning Representations.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans Trained by a Two Time-scale Update Rule Converge to a Local Nash Equilibrium. *Advances in neural information processing systems*, 30.
- Heyrani Nobari, A.; Rashad, M. F.; and Ahmed, F. 2021. Creative-gan: Editing generative adversarial networks for creative design synthesis. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 85383, V03AT03A002. American Society of Mechanical Engineers.
- Humayun, A. I.; Balestrieri, R.; and Baraniuk, R. 2021. MaG-NET: Uniform sampling from deep generative network manifolds without retraining. *arXiv preprint arXiv:2110.08009*.
- Humayun, A. I.; Balestrieri, R.; and Baraniuk, R. 2022. Polarity sampling: Quality and diversity control of pre-trained generative networks via singular values. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10641–10650.
- Hwang, S.; Park, S.; Kim, D.; Do, M.; and Byun, H. 2020. Fairfacegan: Fairness-aware facial image-to-image translation. *arXiv preprint arXiv:2012.00282*.
- Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020a. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33: 12104–12114.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020b. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.
- Kim, E.-J.; and Bansal, P. 2023. A deep generative model for feasible and diverse population synthesis. *Transportation Research Part C: Emerging Technologies*, 148: 104053.
- Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31.
- Kirichenko, P.; Izmailov, P.; and Wilson, A. G. 2020. Why normalizing flows fail to detect out-of-distribution data. *Advances in neural information processing systems*, 33: 20578–20589.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. In *NeurIPS*.
- Lee, J.; Kim, H.; Hong, Y.; and Chung, H. W. 2021. Self-diagnosing gan: Diagnosing underrepresented samples in generative adversarial networks. *Advances in Neural Information Processing Systems*, 34: 1925–1938.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Loftsgaarden, D. O.; and Quesenberry, C. P. 1965. A nonparametric estimate of a multivariate density function. *The Annals of Mathematical Statistics*, 36(3): 1049–1051.

- Lynn, M.; and Harris, J. 1997. Individual differences in the pursuit of self-uniqueness through consumption. *Journal of Applied Social Psychology*, 27(21): 1861–1883.
- Ma, Z.; Mei, G.; and Xu, N. 2024. Generative deep learning for data generation in natural hazard analysis: motivations, advances, challenges, and opportunities. *Artificial Intelligence Review*, 57(6): 160.
- McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Naeem, M. F.; Oh, S. J.; Uh, Y.; Choi, Y.; and Yoo, J. 2020. Reliable fidelity and diversity metrics for generative models. In *ICML*.
- Narayan, K.; VS, V.; Chellappa, R.; and Patel, V. M. 2024. FaceXFormer: A Unified Transformer for Facial Analysis. *arXiv preprint arXiv:2403.12960*.
- Papamakarios, G.; Nalisnick, E.; Rezende, D. J.; Mohamed, S.; and Lakshminarayanan, B. 2021. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57): 1–64.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rinnooy Kan, A.; and Timmer, G. T. 1987. Stochastic global optimization methods part I: Clustering methods. *Mathematical programming*, 39: 27–56.
- Rochat, Y.; and Taillard, É. D. 1995. Probabilistic diversification and intensification in local search for vehicle routing. *Journal of heuristics*, 1: 147–167.
- Sagar, D.; Risheh, A.; Sheikh, N.; and Forouzes, N. 2023. Physics-Guided Deep Generative Model For New Ligand Discovery. In *Proceedings of the 14th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, 1–9.
- Sehwag, V.; Hazirbas, C.; Gordo, A.; Ozgenel, F.; and Canton, C. 2022. Generating high fidelity data from low-density regions using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11492–11501.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*.
- Snyder, C. R.; and Lopez, S. J. 2001. *Handbook of positive psychology*. Oxford university press.
- Srivastava, A.; Valkov, L.; Russell, C.; Gutmann, M. U.; and Sutton, C. 2017. Veegan: Reducing mode collapse in gans using implicit variational learning. *Advances in neural information processing systems*, 30.
- Tarek, M.; and Huang, Y. 2022. Simplifying deflation for non-convex optimization with applications in Bayesian inference and topology optimization. *arXiv preprint arXiv:2201.11926*.
- Teo, C. T.; Abdollahzadeh, M.; and Cheung, N.-M. 2023. Fair generative models via transfer learning. *Proceedings of the AAAI conference on artificial intelligence*, 37(2): 2429–2437.
- Thanh-Tung, H.; and Tran, T. 2020. Catastrophic forgetting and mode collapse in GANs. In *2020 international joint conference on neural networks (ijcnn)*, 1–10. IEEE.
- Tolstikhin, I. O.; Gelly, S.; Bousquet, O.; Simon-Gabriel, C.-J.; and Schölkopf, B. 2017. Adagan: Boosting generative models. *Advances in neural information processing systems*, 30.
- Turner, R.; Hung, J.; Frank, E.; Saatchi, Y.; and Yosinski, J. 2019. Metropolis-hastings generative adversarial networks. In *International Conference on Machine Learning*, 6345–6353. PMLR.
- Wortsman, M.; Ilharco, G.; Gadre, S. Y.; Roelofs, R.; Gontijo-Lopes, R.; Morcos, A. S.; Namkoong, H.; Farhadi, A.; Carmon, Y.; Kornblith, S.; et al. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, 23965–23998. PMLR.
- Xia, M.; Shu, Y.; Wang, Y.; Lai, Y.-K.; Li, Q.; Wan, P.; Wang, Z.; and Liu, Y.-J. 2023. FEditNet: few-shot editing of latent semantics in GAN spaces. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2919–2927.
- Yang, C.; Shen, Y.; Zhang, Z.; Xu, Y.; Zhu, J.; Wu, Z.; and Zhou, B. 2023. One-shot generative domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7733–7742.
- Zeng, X.; Wang, F.; Luo, Y.; Kang, S.-g.; Tang, J.; Lightstone, F. C.; Fang, E. F.; Cornell, W.; Nussinov, R.; and Cheng, F. 2022. Deep generative molecular design reshapes drug discovery. *Cell Reports Medicine*, 3(12).
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.