

Towards Authentic Face Restoration with Iterative Diffusion Models and Beyond

Yang Zhao Tingbo Hou Yu-Chuan Su Xuhui Jia Yandong Li Matthias Grundmann
Google

{yzhaoeric, tingbo, ycsu, xhjia, yandongli, grundman}@google.com

Abstract

*An authentic face restoration system is becoming increasingly demanding in many computer vision applications, e.g., image enhancement, video communication, and taking portrait. Most of the advanced face restoration models can recover high-quality faces from low-quality ones but usually fail to faithfully generate realistic and high-frequency details that are favored by users. To achieve authentic restoration, we propose **IDM**, an Iteratively learned face restoration system based on denoising Diffusion Models (DDMs). We define the criterion of an authentic face restoration system, and argue that denoising diffusion models are naturally endowed with this property from two aspects: intrinsic iterative refinement and extrinsic iterative enhancement. Intrinsic learning can preserve the content well and gradually refine the high-quality details, while extrinsic enhancement helps clean the data and improve the restoration task one step further. We demonstrate superior performance on blind face restoration tasks. Beyond restoration, we find the authentically cleaned data by the proposed restoration system is also helpful to image generation tasks in terms of training stabilization and sample quality. Without modifying the models, we achieve better quality than state-of-the-art on FFHQ and ImageNet generation using either GANs or diffusion models.*

1. Introduction

Face images convey critical information and are used in every corner of our daily lives, e.g., social networks, conference calls, and taking portrait, to name a few. Humans are very sensitive to delicate facial features. High-quality faces are the most preferred and typically are the key to the success of these applications. However, unknown and complicated image degradations are inevitable in real-world scenarios due to e.g. camera defocus, encoding artifacts, motion blur, and noises. In recent years, rapid progress has been made in developing a *blind face restoration* capability as a post-processing step to restore low-quality photos into high-quality ones faithfully.

Most existing popular works exploit pretrained facial

prior networks [43, 54, 3, 47, 4, 3] in assisting recovering photorealistic facial features even if the input is much deficient. Their success derived from the advanced power of generative models in encapsulating high quality priors, e.g., StyleGANs [9, 19] and VQVAE [42, 31]. Other works use auxiliary information as guidance, such as face landmarks [6], component features [23], and high-quality exemplars [24, 25]. Advanced architectures, including CNNs and Transformers [45, 26] have been explored to alleviate face degradation issues. These works are based on building U-net models that restore low-quality faces into high-quality counterparts. The performance difference is often primarily attributed to the model design.

Although these works have achieved notable progress by generating perceptually realistic faces and reporting good numeric results on the benchmarks like PSNR (Peak Signal-to-Noise Ratio) and FID (Fréchet Inception Distance) [13], they still struggle over several important factors, e.g., fail to preserve delicate identity features, hallucinate uncanny artifacts especially in the presence of severe degradation, lose desired high-frequency details such as around facial hair. A quick test is to let the model perform restoration on slightly degraded faces that can make us distinguish delicate differences before and after restoration. As shown in Figure 1, we evaluate two state-of-the-art works CodeFormer [54] and GFPGAN [43] on four examples from CelebA-HQ data [17, 27]. Both models either lose identity details, e.g., freckles and eyelashes or unexpectedly damage the original high-quality components like mustache and hair. The former issue sometimes raises ethical concerns, and the latter holds users from applying the restoration model broadly to fulfill their demands as the regression is unsatisfactory. We conjecture the aforementioned limitations not only from the design choices of their model, but also the quality of training data, as some ground-truth images are still noisy as seen in Figure 1. In some sense, it contradicts the goal of recovering low-quality images into high-quality images if the ground-truth data is not clean, implying that training a restoration model is an ill-posed problem. Moreover, calculating pairwise metrics like PSNR between the predictions and the ground truth could also be problematic in evaluation. Intuitively, noisy training data can harm the restoration model's performance compared with clean data.



Figure 1: Qualitative comparison on original CelebA-HQ test examples (without degradations). Baselines are biased by losing face details and generating hazy facial features. Our method authentically restores faces and preserves details much better regardless of the input quality. Zoom in to see richer details.

Questions To address above challenges, there are two main questions to answer, which are also the primary factors that should be considered when building a machine learning system. (i) **Model:** When processing above examples, why current model designs are unsatisfactory in balancing degradation removal and detail refinement? (ii) **Data:** Can we automatically clean the data without expensive labor an-

notation and pave the path for authentic restoration and accurate evaluation?

Solution To this end, we propose to use conditional denoising diffusion models (DDMs) [14] with well-designed iterative learning as a unified solution to both questions. Because of the inspiring and unprecedented results from image




Figure 2: Restoring original ImageNet 256×256 image with the proposed iterative DDMs. (Zoom in to see differences before and after.)

synthesis *e.g.*, text-to-image synthesis [32, 35] and image editing [33], DDMs have attracted more and more attention from research community and shown great potentials to land in industry applications. There also have been works exploring DDMs for image restoration [20, 49, 34], however, none of them yields a solution to address the above concerns. In this paper, the proposed iterative DDMs have the following properties:

(i) *Intrinsic iterative refinement*: DDMs comprise of iterative Markov chains in both forward diffusion and reverse denoising. The very dense U-Net architecture design and iterative refinement procedure make it a better choice to naturally preserve high-quality contents and remove degradation patterns.

(ii) *Extrinsic iterative enhancement*: As shown in Figure 1, DDMs are found to balance degradation removal and detail preservation very well. Hence, the DDMs trained for restoration are ready for enhancing the training data which is then used for the next iteration of restoration training and more accurate evaluation.

Beyond restoration Besides, we would rather think about using a model elsewhere than just for a specific goal, here, restoration. If above challenges are well resolved, we are able to get a new copy of higher-quality data in terms of resolution and cleanness. This consequently benefits generative models for training stabilization by reducing the number of outliers in terms of data quality and generating higher-fidelity samples. For example, as is known to all, the resolution of ImageNet can be as small as 75×56 and encoding compression is also noticeable in many images. Figure 2 shows qualitative restored examples using a model trained on ImageNet restoration. Consistently, the sample quality has been increased a lot. More inspiring results from generative models will be detailed in later sections.

Contribution The main contributions of this paper are as follows. First, we point out the issues of existing face restoration systems and formulate the definition of authentic restoration that can quickly examine the effectiveness of any

face restoration model. Second, we build an iterative diffusion model (IDM) to tackle the raised concerns. Finally, we demonstrate strong empirical results on two benchmarks, blind face restoration and image generation, showing that the proposed model consistently outperforms state-of-the-art methods and can be used to benefit the learning of generative models beyond restoration.

2. Related Works

Image restoration Due to the presence of various image degradations, image restoration has attracted considerable attention from multiple aspects, *e.g.*, deblurring [48, 38, 22], denoising [51, 11] and super-resolution [12, 44, 29, 46, 28]. Among these topics, blind face restoration [25, 23, 43, 47] (BFR) is of particular interest to us. Most approaches fall into one stage of learning where a U-Net model directly outputs the prediction given the input. In this category, utilizing the pretrained facial prior based on GANs and VAEs is the most commonly used strategy to handle a wide range of complicated degradations [43, 54, 3, 47]. This line of works assume the model can presumably produce realistic faces by projecting the low-quality faces into a compact low-dimension space of the pretrained generator. However, these models are sub-optimal by hallucinating input objects and damaging high-quality details, as the best GAN inversion method cannot guarantee a nearly same reconstruction even for the high quality inputs. In contrast, our work is trained from scratch and can do an authentic restoration.

There are works exploring multi-stage learning by decomposing the task into several more tractable subtasks that refine the output in a coarse-to-fine manner and show improved performances [5, 50, 30]. Depending on the model design, each subtask may output coarse feature or image predictions to the next stage. Our work is mainly related to this category of work. But different from them, diffusion models’ learning and model design can be viewed as a whole without the need for delicate control over losses and architectures.

Denoising diffusion models Denoising diffusion models (DDMs) though they have been tried for image restoration [20, 49, 34], none of them has shown the potential to surpass current state-of-the-art face restoration works. To our best knowledge, it is the first work that presents a thorough study of using DDMs for authentic face restoration.

3. Methodology

In this section, we start with the problem formulation of authentic face restoration. Next, we introduce the core iterative restoration method from two aspects, (i) *Intrinsic iterative learning* benefits authentic restoration via iterative refinement. (ii) *Extrinsic iterative learning* further improves restoration by automatically enhancing the training data.




Figure 3: Illustration of intrinsic and extrinsic iterative learning. Both intrinsic iterations are conditioned on the synthetically degraded samples from original data. The extrinsic iteration consists of two intrinsic iterations of learning. The resulting model from the first intrinsic iteration is used to enhance the training data only which will then be used as the target data for the second intrinsic iteration.

3.1. Problem Formulation

Let X denotes the high-quality image domain and P_X implies the distribution of X . Assume that there exists a non-learnable one-to-many degradation function d that maps X to degraded low-quality image domain X_d : $X \rightarrow X_d$, the goal of face restoration is to learn an inverse function f_θ parameterized by θ : $X_d \rightarrow X$ that satisfies

$$\min_{f_\theta} \mathcal{D}(P_{f_\theta(X_d)} || P_X) + \mathbb{E}_{x \sim X} \mathbb{E}_{x_d \sim f(x)} \kappa(f_\theta(x_d), x), \quad (1)$$

where \mathcal{D} is a distribution distance, *e.g.*, KL (Kullback–Leibler) divergence (or maximum likelihood estimation) [21] or Jensen-Shannon divergence (or adversarial loss) [9]; $\kappa(\cdot)$ is an instance-wise distance between two images, *e.g.*, perceptual loss [16] or L_p distance ($p \sim \{1, 2\}$) [53]. All current face restoration works can be formulated by Eq. 1. For example, GFPGAN [43] and CodeFormer [54] both keep the adversarial loss for high-fidelity restoration and mix different kinds of $\kappa(\cdot)$ for input’s content preservation. In contrast, the proposed approach does not explicitly puts constraint to preserve content. Instead, the content preservation is achieved via injecting conditional signals inside the model plus iterative refinement. Thus, the second terms is omitted.

Authentic restoration In real applications, it is common that the degraded input $x_d \sim X_d$ can be either severely degraded, only has minor degradation, or even as clean as $x \sim X$. Since there is not a reliable metric to evaluate a single image’s quality, *e.g.*, sharpness and realisticness, an authentic restoration model is expected to treat x and x_d the same regardless the actual quality. Thus, we propose the following criterion that an authentic restoration model should obey:

$$f_\theta(x_d) = x, \quad f_\theta(x) = x \quad (2)$$

which implies applying the model f_θ iteratively by starting from x_d should converge to the same high-quality image x ,

$$x = f_\theta \dots f_\theta(f_\theta(x_d)) \quad (3)$$

A mild relaxation is $f_\theta(f_\theta(x_d))$ is not worse than x after two iterations. As is shown in Figure 1, both baselines GFPGAN and CodeFormer fail the authentication test, *e.g.*, hair

detail is lost and freckles are removed. The results demonstrate that the baselines do not meet the criterion in Eq. 2.

In this paper, by assuming the input is always of low quality, we can focus on two iterative restorations to deal with almost all real restoration scenarios. Next, we will describe how iterative diffusion models naturally couple with authentic restoration in Eq. 3 with proper training.

3.2. Intrinsic Iterative Learning

Briefly, DDMs have a Markov chain structure residing in both the forward diffusion process and the reverse denoising chain. To recover the high-quality image, this structure naturally benefits authentic restoration via iterative refinement of the low-quality input, termed as *intrinsic iterative learning*.

Conditional DDMs DDMs are designed to learn a data distribution $q(x_0)$ by defining a chain of latent variable models $p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t)$,

$$x_T \rightarrow x_{T-1} \rightarrow \dots \rightarrow x_1 \rightarrow x_0 \quad (4)$$

where each timestep’s example x_t has the same dimensionality $x_t \in \mathbb{R}^D$. Usually, the chain starts from a standard Gaussian distribution $x_T \sim \mathcal{N}(0, I^D)$ and only the final sample x_0 is stored. For restoration, we are more interested in the conditional DDMs that maps a low-quality source input x_d to a high-quality target image x . The conditional DDMs iteratively edits the noisy intermediate image x_t by learning a conditional distribution $p_\theta(x_{t-1}|x_t, x_d)$ such that $x_0 \sim p(x|x_d)$ where $x_0 \equiv x, x_d = d(x), x \sim P_X$.

Learning Let’s first review the forward diffusion process that can be simplified into a specification of the true posterior distribution

$$q(x_t|x_0) = \mathcal{N}(x_t | \sqrt{\gamma_t}x_0, (1 - \gamma_t)I) \quad (5)$$

where γ_t defines the noise schedule.

We thus learn the reverse chain with a “denoiser” model f_θ which takes both source image x_d and a intermediate noisy target image x_t by comparing it $p_\theta(x_{t-1}|x_t, x_d)$ with the tractable posterior with $q(x_{t-1}|x_t, x_0)$. Consequently,




Figure 4: Dense U-Net with conditional signals.

following [36], we aim to optimize the following objective that estimates the noise ϵ ,

$$\mathcal{L} = \mathbb{E}_{(x, x_d)} \mathbb{E}_{\epsilon, \gamma} \|f_{\theta}(x_d, \hat{x}, \gamma) - \epsilon\|_p^p, \quad \epsilon \sim \mathcal{N}(0, 1). \quad (6)$$

where $\hat{x} = \sqrt{\gamma}x_0 + \sqrt{1 - \gamma}\epsilon$ as in Eq. 5 and usually we can set $p \in \{1, 2\}$. One can also directly predict the output of f_{θ} to x_0 , a.k.a, regression, which makes (6) becomes:

$$\mathcal{L} = \mathbb{E}_{(x, x_d)} \mathbb{E}_{\gamma} \|f_{\theta}(x_d, \hat{x}, \gamma) - x_0\|_p^p. \quad (7)$$

We find this formulation more efficient in both training and inference since the sequence of network approximates x_0 at each time step starting from different amount of noises. This makes our work differ from previous conditional image-to-image synthesis works SR3 [36] and Palette [34].

Iterative restoration As for the reverse diffusion process during the inference stage, we follow the Langevin dynamics [36],

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \gamma_t}} f_{\theta}(x_t, x_d, \gamma_t) \right) + \sqrt{1 - \alpha_t} \epsilon_t \quad (8)$$

where α_t is hyper-parameter related to $\gamma_t = \prod_{i=1}^t \alpha_i$ [36]. To help understand why DDMs can provide authentic restoration, we justify from two aspects,

- **Iterative refinement:** the above process iteratively refines the input towards the target high-quality image. Different time steps are guided to learn restoration with an annealing noise schedule. That is, it coincides with our authentic restoration motivation in Eq. 3. Theoretically, as α_t is annealed to 1, it makes the sufficiently long denoising chain always converge to the same data point, that is,

$$x_t \approx x_{t-1} \quad (9)$$

- **Dense architecture:** in convention as shown in Figure 4, the design of the predictor f_{θ} is a very dense structure with hierarchical links between encoder and decoder. Conditional signal can also be strong to preserve the high-quality contents without any auxiliary losses. That is, x_t will gradually learn to absorb the clean content from x_d to approximate a clean version x_0 of x_d as in Eq.(7).

3.3. Extrinsic Iterative Learning

Most works choose to use 70K FFHQ (Flickr-Faces-HQ) data [19] as training data. The dataset is currently the largest public high-quality high-resolution face data and is assumed to be clean and sharp. It is collected from internet and filtered by both automatic annotation tools and human workers. However, revisiting the examples in Figure 1 reveals that the training data is not always high-quality. JPEG degradation, blur and Gaussian noise are found to account for around 15% when we examine the first 2000 images. The quality of dataset has proven to be a key factor in learning a restoration model and we will also provide an ablation later. To avoid being trapped in such a dilemma, we herein propose a simple yet effective solution *extrinsic iterative learning*. Thanks to the capability of authentic restoration of DDMs, the solution can automatically restore the training data without damaging the data, especially for those whose quality are already high and facial details are ineligible.

Iterative restoration As shown in Figure 1, DDMs are proven to satisfy Eq. 2. After learning the restoration model f_{θ} , we can apply it to all the training data to produce a new high-quality image domain P_{X^*} ,

$$x^* = f_{\theta}(x), \quad x \sim P_X. \quad (10)$$

We term the method as extrinsic iterative learning because Eq.10 produces a higher quality data for another iteration of restoration model training, which is different from the internal chained process in DDMs.

Note that we still draw low-quality image samples x_d with mapping: $X \rightarrow X_d$ and the target distribution becomes P_{X^*} instead of P_X . Therefore we learn a conditional distribution $p_{\theta^*}(x_{t-1}|x_t, x_d)$ such that $x_0 \sim p(x^*|x_d)$ where $x_0 \equiv x^*$, $x_d = d(x)$, $x \sim P_X$.

Coupling two DDMs Note that the only difference of learning f_{θ} and f_{θ^*} is the target distribution P_X and P_{X^*} . With proper tuning, the pool of P_X can be gradually replaced and filled by new data from P_{X^*} . Given an ideal “stop sign” for learning f_{θ} , we actually can unify the two DDMs f_{θ} and f_{θ^*} into only one learning process. In this paper, we do not adopt this option as we find simply training two separate models works well and saves parameter tuning cost. Nevertheless, we still can view the decoupled DDMs as a single model experimentally since f_{θ^*} is initialized from f_{θ} . Through intrinsic and extrinsic iterative learning, we only keep the latest model f_{θ^*} for inference.

4. Experiments

We first evaluate our method on blind face restoration. Next, we demonstrate that the proposed method can be used to improve other generation tasks. Finally, we conduct ablation studies on our model.




Figure 5: Qualitative comparison with state-of-the-art restoration models on CelebA-HQ Test.

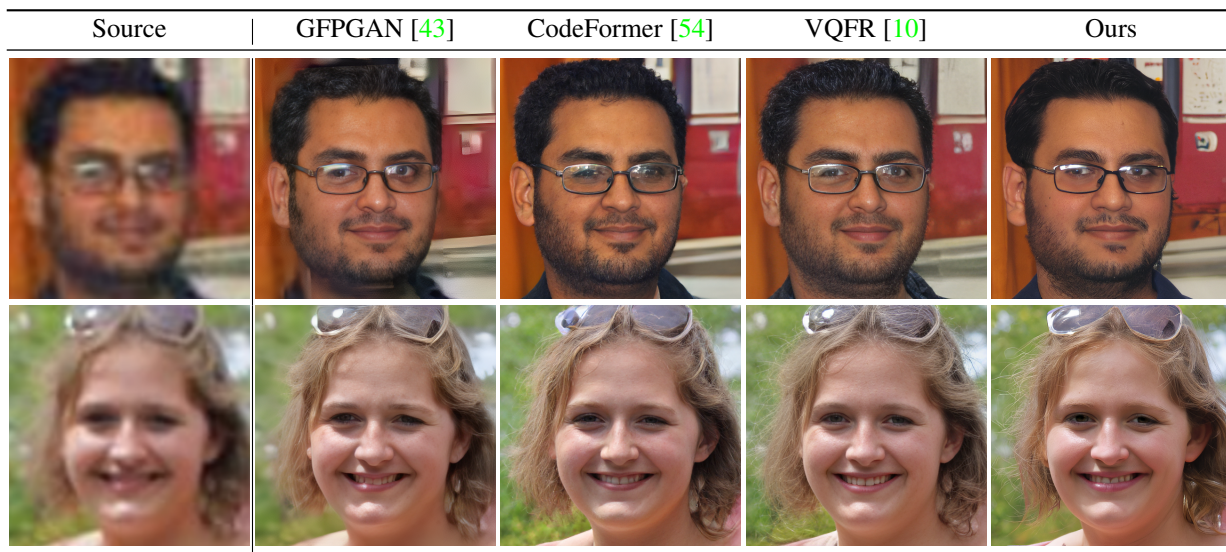


Figure 6: Qualitative comparison with state-of-the-art restoration models on real-world Wider-Test datasets.

4.1. Face restoration

Dataset We train our model on the 70K FFHQ dataset [18] and evaluate on the test split of the CelebA-HQ dataset [17], which contains 30K images following GFP-GAN¹. All images are resized to 512×512 . Following the standard practice in face restoration [25, 43, 23, 47], we synthesize degraded low-quality faces x_d from real high-

quality faces x using the following degradation function:

$$x_d = [(x \otimes \mathbf{k}_\sigma)_{\downarrow r} + \mathbf{n}_\delta]_{\text{JPEG}_q}, \quad (11)$$

i.e. the high-quality image x is first convolved with a Gaussian blur kernel \mathbf{k}_σ with kernel size σ and downsampled by a factor r . Gaussian noise \mathbf{n}_δ with standard deviation δ is then added before applying JPEG compression with quality factor q to obtain the final low-quality image x_d . Although there are possible other types of degradations, we adopt the same degradation model used in prior works [43, 23, 47] for

¹<https://xinntao.github.io/projects/gfpgan>

a fair comparison. We randomly sample σ , r , δ and q from $[0.2, 10]$, $[1, 8]$, $[0, 15]$ and $[50, 100]$ for the degradation function. Degradation pipeline follows GFPGAN [43] implementation². The restoration model is trained with image pairs (x, x_d) following Section 3.

Additionally, as seen in Figure 1, the proposed method has demonstrated the ability to authentically repair the data and CelebA-HQ is known to be pretty noisy, we also prepare a set of clear CelebA-HQ test set as new evaluation set. We will release ground-truth, degraded data and restored results from different models to reproduce. We also evaluate the proposed method on several real-world datasets, LFW, WebPhotos [43] and Wider-Test [54].

Baselines Three very recent models, GFPGAN [43], RestoreFormer [45] and CodeFormer [54], are used as primary baselines. GFPGAN is the first successful approach that applies pretrained StyleGAN face prior for improved restoration. RestoreFormer explores transformer architecture as model base. To approximate the high-quality face correspondence, CodeFormer predicts the latent codes with a transformer.

Metrics We mainly apply PSNR, SSIM, LPIPS [52] and Arcface identity score (ID) [7] as quantitative evaluation metrics. In particular, we calculate LPIPS with VGG following official implementation³.

Implementation details The DDM architecture follows that of [14], except that we add attention layers in lower resolution (32×32 to 8×8) blocks. Also, we directly train the model on 512×512 resolution instead of using the common cascade approach [15, 36]. The model comprises of 120M parameters and takes about 2 days (100K iterations) to converge on eight V100 GPUs. We choose $T = 1000$ for training and 10 diffusion steps for efficient inference.

Results We first conduct an analysis on how significant extrinsic iterative learning can improve training dataset quality. As is shown in Figure 7, the trained restoration model can make the data more clear while preserving all contents. We also compute the NIMA score on FFHQ [41], commonly used to evaluate image quality without using reference. After restoration, the image quality of FFHQ dataset improves from 4.51 to 4.85.

Table 1 and Table 2 respectively show the results on the original and clean CelebA-HQ test set with randomly synthetic degradation. Our method consistently outperforms baselines across different metrics. This verifies the effectiveness of the proposed method on face restoration. In Figure 5 and Figure 6, comparing our method and baselines on



Figure 7: An example from FFHQ. Before and after applying the first iteration restoration model.

Models	PSNR \uparrow	SSIM \uparrow	ID \uparrow	LPIPS $_{VGG}\downarrow$
DFDNet [23]	25.81	0.659	0.40	0.418
GFPGAN [43]	26.95	0.715	0.59	0.359
ReFormer [45]	25.83	0.671	0.68	0.360
VQFR [10]	24.60	0.631	0.70	0.341
CodeFormer [54]	25.70	0.664	0.75	0.337
IDM (Ours)	27.42	0.742	0.87	0.307

Table 1: Quantitative comparison on blind face restoration (BFR) using *original* CelebA-HQ.

Models	PSNR \uparrow	SSIM \uparrow	ID \uparrow	LPIPS $_{VGG}\downarrow$
DFDNet [23]	24.10	0.621	0.39	0.404
GFPGAN [43]	24.65	0.646	0.55	0.351
ReFormer [45]	24.06	0.618	0.60	0.359
VQFR [10]	23.31	0.598	0.64	0.336
CodeFormer [54]	23.89	0.613	0.79	0.314
IDM (Ours)	28.25	0.803	0.88	0.291

Table 2: Quantitative comparison on blind face restoration (BFR) using *clean* CelebA-HQ.

both synthetic data and real-world data, it is observed that both baselines tend to generate hazy and unnatural faces. In contrast, our approach can give much more natural and realistic results with rich details.

We also conduct a user study that collects users’ opinions. We ask 10 raters to do a blind model evaluation on restoration results in clear CelebA-HQ test with synthetic degradations. Each rater is asked to select which model shows the best perceptual quality while preserving the facial details given the ground-truth images. As in Table 3, our approach outperforms baselines by a significant margin.

Models	GFPGAN [43]	ReFormer [45]	VQFR [10]	CodeFormer [54]	IDM (Ours)
Preferences \uparrow	4.1%	4.7%	8.4%	11.2%	71.6%

Table 3: Subjective study on blind face restoration (BFR).

4.2. Improved generation

We also analyze the impact of restored dataset on generative models’ performance. We adopt two most popu-

²<https://github.com/TencentARC/GFPGAN>

³<https://github.com/richzhang/PerceptualSimilarity>




Figure 8: Qualitative generated samples of FFHQ 512×512 from StyleGAN2 and ImageNet 128×128 from BigGAN.

lar image generation benchmarks, FFHQ and ImageNet, and evaluate the method on two different kinds of models, GANs and DDMs. For FFHQ generation, we use official StyleGAN2-PyTorch implementation and resume training from the released checkpoint. For ImageNet with BigGAN, we use the official PyTorch BigGAN implementation⁴ and resume the training from the suggested checkpoint for fine-tuning. Numbers from BigGAN are reported as it is without truncation. For DDMs, we train the conditional DDMs on resolution 384×384 by conditioning on ImageNet-1k class labels. Architecture resembles the one being used in restoration model training. Quantitative results of FFHQ and ImageNet are respectively shown in Table 4 and Table 5. The improved FID, precision and recall consistently suggest that refined data by the restoration model can help generative models to learn the data better and obtain higher-quality samples. Besides, in Figure 9, it is shown that the training stability of GANs is also relieved a bit whereas the baseline BigGAN diverges after 140K iterations. We present qualitative generated samples in Figure 8. When comparing uncurated FFHQ samples for StyleGAN, we can see artifacts have been reduced.

Resolution	Methods	FID↓	Precision↑	Recall↑
×512	StyleGAN2 [19]	3.08	0.676	0.478
	IDM (Ours)	2.78	0.683	0.479
×256	U-Net GAN [37]	7.48	-	-
	VQGAN [8]	11.40	-	-
	INR-GAN [39]	9.57	-	-
	CIPS [1]	4.38	-	-
	StyleGAN2 [19]	3.83	0.711	0.422
	IDM (Ours)	3.54	0.720	0.425

Table 4: Comparison with the state-of-the-art methods on FFHQ in two resolutions 512×512 and 256×256 .

⁴<https://github.com/ajbrock/BigGAN-PyTorch>

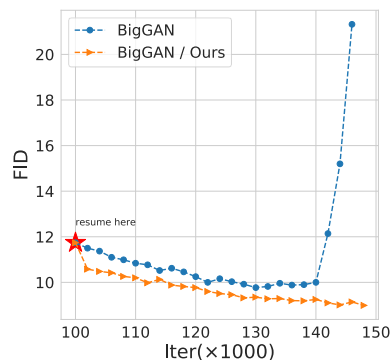


Figure 9: BigGAN training is stabilized on ImageNet 128×128 .

Resolution	Methods	FID↓
×128	BigGAN [2]	9.77
	IDM (Ours)	8.97
	DDM	14.43
	IDM (Ours)	11.10

Table 5: Comparison with the BigGAN methods on conditional ImageNet 128×128 generation. For DDM-based experiment, we report the half-way results under the same number of iterations (100K).

4.3. Ablations

As for ablations, we are mainly interested in how the setting of intrinsic and extrinsic iterations boost the restoration model’s performance.

Intrinsic iteration We study the effect of the number of sampling steps and sampling schedule, *e.g.*, DDPM [14] and DDIM [40]. For the sampling schedule, DDIM yields inferior results such that we use DDPM everywhere in this

paper. Table 6 compares the performances by varying the number of diffusion steps. Increasing number of diffusion steps > 10 doesn't give benefit so we choose 10 steps in this paper. It is equivalent to 2s to restore per image on V100 GPU. We also note that most single-forward pass methods take $< 0.5s$. So the efficiency might be one limitation of the proposed approach.

Steps	5	10	15	20
PSNR \uparrow	13.11	27.42	27.50	27.48
LPIPS _{VGG} \downarrow	0.475	0.307	0.311	0.308

Table 6: Varying the number of sampling steps.

Extrinsic iteration We have claimed that extrinsic iteration enhances training data and improves restoration with another iteration of training. An example of the advantage of using extrinsic iterative learning is given below. Obviously, extrinsic learning helps increase clearness and reduce artifacts. An user study of evaluating this strategy also suggests that results from extrinsic learning are preferred.



5. Conclusion

Face restoration has been an ill-posed problem. The importance of training data, testing data and evaluation criterion have been neglected. In this paper, we advocate an authentic restoration system with iteratively learned denoising diffusion models. The intrinsic iterative refinement mechanism can automatically maintain the high-quality details while erasing degradations. With the refined training data by such a learned model, we find it not only benefits restoration with another extrinsic iteration but also helps stabilize generation. Extensive experiment results have validated the effectiveness and authenticity of the system.

References

- [1] Ivan Anokhin, Kirill Demochkin, Taras Khakhulin, Gleb Sterkin, Victor Lempitsky, and Denis Korzhenkov. Image generators with conditionally-independent pixel synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14278–14287, 2021.
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018.
- [3] Kelvin CK Chan, Xintao Wang, Xiangyu Xu, Jinwei Gu, and Chen Change Loy. Glean: Generative latent bank for large-factor image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14245–14254, 2021.
- [4] Kelvin C.K. Chan, Xiangyu Xu, Xintao Wang, Jinwei Gu, and Chen Change Loy. Glean: Generative latent bank for image super-resolution and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3154–3168, 2023.
- [5] Chaofeng Chen, Xiaoming Li, Lingbo Yang, Xianhui Lin, Lei Zhang, and Kwan-Yee K Wong. Progressive semantic-aware style transformation for blind face restoration. In *CVPR*, 2021.
- [6] Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, and Jian Yang. Fsrnet: End-to-end learning face super-resolution with facial priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2492–2501, 2018.
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019.
- [8] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- [10] Yuchao Gu, Xintao Wang, Liangbin Xie, Chao Dong, Gen Li, Ying Shan, and Ming-Ming Cheng. Vqfr: Blind face restoration with vector-quantized dictionary and parallel decoder. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVIII*, pages 126–143. Springer, 2022.
- [11] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. In *CVPR*, 2019.
- [12] Tiantong Guo, Hojjat Seyed Mousavi, Tiep Huu Vu, and Vishal Monga. Deep wavelet prediction for image super-resolution. In *CVPR Workshops*, 2017.
- [13] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [15] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23:47–1, 2022.
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
- [17] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018.

- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- [20] Bahjat Kwar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *arXiv preprint arXiv:2201.11793*, 2022.
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [22] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *ICCV*, 2019.
- [23] Xiaoming Li, Chaofeng Chen, Shangchen Zhou, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. Blind face restoration via deep multi-scale component dictionaries. In *ECCV*, 2020.
- [24] Xiaoming Li, Wenyu Li, Dongwei Ren, Hongzhi Zhang, Meng Wang, and Wangmeng Zuo. Enhanced blind face restoration with multi-exemplar images and adaptive spatial feature fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2706–2715, 2020.
- [25] Xiaoming Li, Ming Liu, Yuting Ye, Wangmeng Zuo, Liang Lin, and Ruigang Yang. Learning warped guidance for blind face restoration. In *ECCV*, 2018.
- [26] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1833–1844, 2021.
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [28] Cheng Ma, Zhenyu Jiang, Yongming Rao, Jiwen Lu, and Jie Zhou. Deep face super-resolution with iterative collaboration between attentive recovery and landmark estimation. In *CVPR*, 2020.
- [29] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *CVPR*, 2020.
- [30] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3883–3891, 2017.
- [31] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *NeurIPS*, 2019.
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [33] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022.
- [34] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022.
- [35] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- [36] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [37] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva. A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8207–8216, 2020.
- [38] Ziyi Shen, Wei-Sheng Lai, Tingfa Xu, Jan Kautz, and Ming-Hsuan Yang. Deep semantic face deblurring. In *CVPR*, 2018.
- [39] Ivan Skorokhodov, Savva Ignatyev, and Mohamed Elhoseiny. Adversarial generation of continuous images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10753–10764, 2021.
- [40] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [41] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE transactions on image processing*, 27(8):3998–4011, 2018.
- [42] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [43] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *CVPR*, 2021.
- [44] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *ECCV Workshops*, 2018.
- [45] Zhouxia Wang, Jiawei Zhang, Runjian Chen, Wenping Wang, and Ping Luo. Restoreformer: High-quality blind face restoration from undegraded key-value pairs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17512–17521, 2022.
- [46] Lingbo Yang, Shanshe Wang, Siwei Ma, Wen Gao, Chang Liu, Pan Wang, and Peiran Ren. Hifacegan: Face renovation via collaborative suppression and replenishment. In *ACM Multimedia*, 2020.
- [47] Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. Gan prior embedded network for blind face restoration in the wild. In *CVPR*, 2021.
- [48] Rajeev Yasarla, Federico Perazzi, and Vishal M Patel. Deblurring face images using uncertainty guided multi-stream

- semantic networks. *IEEE Transactions on Image Processing*, 29:6251–6263, 2020.
- [49] Yueqin Yin, Lianghua Huang, Yu Liu, and Kaiqi Huang. Diffgar: Model-agnostic restoration from generative artifacts using image-to-image diffusion models. *arXiv preprint arXiv:2210.08573*, 2022.
- [50] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14821–14831, 2021.
- [51] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018.
- [52] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [53] Yang Zhao, Yu-Chuan Su, Chun-Te Chu, Yandong Li, Marius Renn, Yukun Zhu, Changyou Chen, and Xuhui Jia. Rethinking deep face restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7652–7661, 2022.
- [54] Shangchen Zhou, Kelvin CK Chan, Chongyi Li, and Chen Change Loy. Towards robust blind face restoration with codebook lookup transformer. *arXiv preprint arXiv:2206.11253*, 2022.