

AT-DDPM: Restoring Faces degraded by Atmospheric Turbulence using Denoising Diffusion Probabilistic Models

Nithin Gopalakrishnan Nair
Johns Hopkins University
ngopala2@jhu.edu

Kangfu Mei
Johns Hopkins University
kmei1@jhu.edu

Vishal M. Patel
Johns Hopkins University
vpatel136@jhu.edu



Figure 1: Sample qualitative comparisons with single image atmospheric turbulence mitigation methods as well as generative restoration models on the LRFID[27] dataset.

Abstract

Although many long-range imaging systems are designed to support extended vision applications, a natural obstacle to their operation is degradation due to atmospheric turbulence. Atmospheric turbulence causes significant degradation to image quality by introducing blur and geometric distortion. In recent years, various deep learning-based single image atmospheric turbulence mitigation methods, including CNN-based and GAN inversion-based, have been proposed in the literature which attempt to remove the distortion in the image. However, some of these methods are difficult to train and often fail to reconstruct facial features and produce unrealistic results especially in the case of high turbulence. Denoising Diffusion Probabilistic Models (DDPMs) have recently gained some traction because of their stable training process and their ability to generate high quality images. In this paper, we propose the first DDPM-based solution for the problem of atmospheric turbulence mitigation. We also propose a fast sampling technique for reducing the inference times for conditional DDPMs. Extensive experiments are conducted on synthetic and real-world data to show the significance of our model. To facilitate further research,

all codes and pretrained models are publicly available at <http://github.com/Nithin-GK/AT-DDPM>

1. Introduction

Autonomous vehicles and surveillance systems utilize long-range imaging systems to capture faraway scenes. Images captured by such systems are prone to degradation due to variations of the refractive index of air. Atmospheric turbulence is a ubiquitous phenomenon that causes spatially and temporally random variations in the air’s refractive index [34, 14]. This introduces random geometric distortions and blur to the images captured and significantly degrades the performance of computer vision tasks like object detection [32], tracking and recognition [5] applied on these images. Modeling atmospheric turbulence degradation is an extremely complex task. Various methods have been proposed in the literature that attempt to model deformations caused due to turbulence. Under the assumption that the scene and the imaging sensor are both static and degradations caused are due to atmospheric turbulence alone, the degradation due to atmospheric turbulence is modelled as

$$T_k = D_k(H_k(I)) + n_k, \quad k = 1, 2, \dots, N, \quad (1)$$

where T_k is the degraded image at the k^{th} time instant, I is the clean latent image, D_k is the deformation operator which is assumed to deform randomly, n_k is additive noise, and H_k is air turbulence-caused blurring operator [16, 17, 47, 3, 12].

Mitigating the effect of atmospheric turbulence and restoring the latent clean image is extremely challenging due to its ill-posed nature that arises due to both blur and geometric distortions. Restoring images affected by atmospheric turbulence is an extensively researched topic in the literature. To mitigate these effects, there exists two classical solutions – adaptive optics [33, 38, 34] and image processing [10, 47, 1, 39, 2, 26, 19]. In this paper we focus on the latter and propose an image processing-based strategy for atmospheric turbulence mitigation.

In recent years, various deep learning-based methods for mitigating atmospheric turbulence have been proposed in the literature [16, 19, 42]. [42] proposes a deep network utilizing blur and degradation priors to restore the clean image. [16] uses two parallel paths to remove blur and degradation and finally concatenates them to remove the effect of turbulence. These methods tackle atmospheric turbulence degradation on facial images but focus on removing the distortion and does not have enough generative power to generate a realistic face for strong turbulence distortions. Moreover these methods fail to generalize well in case of real world images. Generative prior-based face reconstruction has gained much attention in recent years due to their ability in modelling the distribution of clean facial images [40, 24]. This is usually performed by utilizing Generative Adversarial Networks (GANs) for modelling the distribution of clean faces and then utilizing these GANs as a pretrained prior to generate realistic facial images while performing restoration. But training of GANs is unstable and may also fail due to mode collapse. In this paper, we propose a generative modeling-based solution using Denoising Diffusion Probabilistic Models (DDPMs)[37, 13]. DDPM-based methods have already achieved state-of-the-art results in image generation tasks [9]. When compared to GANs, their training process is also much more stable[13, 9]. Motivated by their ability to perform exact sampling and generating high quality images, we use them as generative priors to restore single facial images degraded by atmospheric turbulence.

We first train a DDPM model to learn the distribution of facial images by training on a large dataset[15]. We then adapt this model to learn the transformation for the simpler restoration problem of image super resolution by using knowledge distillation. With the super-resolution model in hand, we adapt these weights again for the transformation from turbulent images to clean faces by performing knowledge distillation with the super-resolution model. This could be thought of as a continual learning based strategy[45]. Where, the final model has access to information about the

manifold of realistic faces, hence being able to produce realistic face outputs even for strong distortions. Restoration of images degraded by atmospheric turbulence is an ill-posed problem. Hence, multiple possibilities of clean faces exist while reconstructing using a DDPM model that has access to the distribution of clean facial images. Moreover, the inference process of DDPMs is quite slow and time consuming. To this end we enforce a constraint to ensure that the facial features present in the restored image are closer to the ones present in the distorted image during inference by starting from noise turbulence-distorted image rather than pure Gaussian noise.

This paper makes the following contributions:

- We propose the first generative prior-based image restoration method using DDPM for the task of atmospheric turbulence mitigation.
- We propose a progressive training(PT)-based learning framework for better reconstruction and introduce a new sampling procedure for efficient inference.
- We qualitatively and quantitatively evaluate our method on both synthetic and real-world datasets and show that it performs better than the existing state-of-the-art single image-based atmospheric turbulence removal methods.

2. Related Works

Image Processing methods. “Lucky Imaging” techniques [2, 39] are the earliest techniques for mitigating atmospheric turbulence. Lucky imaging works by selecting a few good frames from turbulence degraded videos and fusing them to restore the latent sharp image. Anantrasirichai *et al.* [1] proposed a method to extract regions with fewer degradations from the good frames and fuse them using a dual-tree complex wavelet transform to remove turbulence degradations. Methods based on registration-fusion approaches like [47, 12, 19] were proposed in later years, where a good reference image is initially determined, and the consequent frames are aligned and fused to remove geometric distortions, after which a deconvolution algorithm is applied to remove blur. Zhu *et al.* [47] use a B-spline registration algorithm and temporal kernel regression to compute a degradation-free image. Lou *et al.* [22] proposed a method to reduce deformations between images in a video by applying a temporal smoothing algorithm on video frames sharpened using Sobolev gradient flow. Xie *et al.* [41] used a low-rank image as an intermediate reference image and then obtained the sharp image through a variational model.

Atmospheric turbulence mitigation using deep networks. With the success of CNN-based models in addressing image restoration tasks[25], a few models have been proposed in the literature to remove atmospheric turbulence.

Chak *et al.* [4] utilize an effective data augmentation algorithm to model turbulence, then filter out the corrupted frames using a sub-sampling algorithm, and finally use a Generative adversarial Network (GAN) to restore the clean image. Recently some methods [42, 16, 18, 30] have been proposed for restoring a single face image from a turbulence degraded observations. Lau *et al.* [19] use the commutative relationship between blur and geometric distortions to generate blur and degradation-free images. They further disentangle these images to obtain the clean output. Yasarla & Patel [42] propose a turbulence removal network that estimates blur and degradation priors using Monte Carlo simulations and use them to restore the sharp image.

3. Background

3.1. Denoising Diffusion Probabilistic models

Denoising Diffusion Probabilistic Models (DDPMs) [13, 37] are a class of generative models that perform image generation through variational inference using a Markovian process with a finite number of timesteps ‘ T ’. Each step in this Markovian process is a denoising process. DDPMs consists of two stages – a forward process and a reverse process. In the forward process, a clean image y_0 is sampled, and small Gaussian noises of variance schedules $\{\beta_1, \dots, \beta_T\}$ are added over ‘ T ’ timesteps. The overall forward process and each forward step are defined as

$$q(y_{0,1\dots T}) := q(y_0) \prod_{t=1}^T q(y_t | y_{t-1}) \quad (2)$$

$$q(y_t | y_{t-1}) = \mathcal{N}\left(y_t; \sqrt{1 - \beta_t} y_{t-1}, \sqrt{\beta_t} \mathbf{I}\right) \quad (3)$$

$$= \sqrt{\beta_t} y_{t-1} + \epsilon \sqrt{1 - \beta_t}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (4)$$

where y_t and y_{t-1} are noisy samples generated at timesteps t and $t - 1$ and $\{\beta_i\}$ refers to the variance schedules. Since the sum of t zero mean Gaussians with different variance schedules is again a Gaussian, this Markovian process could be effectively represented in terms of the initial data sample y_0 and expressed as

$$q(y_t | y_0) := \mathcal{N}\left(y_t; \sqrt{\bar{\alpha}_t} y_0, (1 - \bar{\alpha}_t) \mathbf{I}\right) = \sqrt{\bar{\alpha}_t} y_0 + \epsilon \sqrt{1 - \bar{\alpha}_t}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (5)$$

where $\alpha_t = 1 - \beta_t$. In the reverse process we define the joint distribution $p_\theta(y_{0:T})$ with parameters θ . Similar to the forward process, the reverse process is also a Markovian process which is defined as follows

$$p(y_T) := \mathcal{N}(0, \mathbf{I}) \quad (6)$$

$$p(y_{t-1} | y_t) := \mathcal{N}\left(y_{t-1}; \mu_\theta(y_t, t), \sqrt{\beta_t} \mathbf{I}\right). \quad (7)$$

For optimizing the parameters of the network θ , we minimize the variational lower bound of the the negative log likelihood of the distribution of clean images y_0 . In this paper, we use the simplified training objective of DDPMs proposed by Ho *et al.* [13]. For training our model we sample a timestep $t \sim U[1, T]$ and produce the noisy sample

corresponding to this timestep using equation(3), defined by

$$y_t := \sqrt{\bar{\alpha}_t} y_0 + \epsilon \sqrt{1 - \bar{\alpha}_t}, \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (8)$$

The network $p_\theta(\cdot)$ predicts the noise ϵ in this image taking y_t and t as the inputs. The training objective is defined as,

$$L_{simple} := E_{t \sim [1, T], \epsilon \sim \mathcal{N}(0, \mathbf{I})} \left[\|\epsilon - \epsilon_\theta(y_t, t)\|^2 \right]. \quad (9)$$

3.2. Conditional Diffusion Probabilistic Models

The equations mentioned above are developed for the task of image generation. For utilizing DDPMs for low level vision tasks like image restoration, the conditional distribution of the clean image has to be modelled. Saharia et al [35] have proposed a simple technique for modelling the conditional distribution of clean image given the corresponding degraded image. In conditional DDPM, the forward process remains the same as that of the unconditional model. Given a clean image sampled from the dataset, we add random Gaussian noise based on a randomly sampled timestep t . During the reverse process, along with the noisy image and the time t , we also pass the degraded image(x) as input to the neural network. Hence the denoising model is defined by $p_\theta(y_t, x, t)$ and the reverse process is defined by

$$p(y_T) = \mathcal{N}(0, \mathbf{I}) \quad (10)$$

$$p(y_{t-1} | y_t, x) = \mathcal{N}\left(y_{t-1}; \mu_\theta(y_t, x, t), \sqrt{\beta_t} \mathbf{I}\right). \quad (11)$$

The mean $\mu_\theta(y_t, x, t)$ is estimated according to,

$$\mu_\theta(y_t, x, t) = \frac{1}{\sqrt{(1 - \beta_t)}} \left(y_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x, y_t, t) \right). \quad (12)$$

3.3. Proposed method

In this section we detail our proposed method and the training process. For training our model, we perform a multi stage training process as shown in Figure 2. Given a dataset containing clean facial images, we train a diffusion model $p_\theta(\cdot)$ for the task of unconditional generation of facial images. [6, 9] have already released models trained on large amounts of facial data. Hence we make use of this model as the starting point. Once we have a model that has learned the distribution of clean facial images, it is easier to adapt those weights to model the conditional distribution where we aim to restore a slightly degraded image and recover the clean image by conditioning on the distorted image.

The same idea holds for the case of two conditional distributions where the diffusion model learns the conditional distribution of clean facial images. For our experiments we choose the conditional distributions as face super-resolution and turbulence since off the shelf super-resolution model trained on a large datasets are already available[9]. Once we have a model that can perform ($8\times$) super-resolution(SR), we adapt the model for the stronger degradation of atmospheric turbulence. For making the model robust to the turbulence degradation, we take the model trained for super-

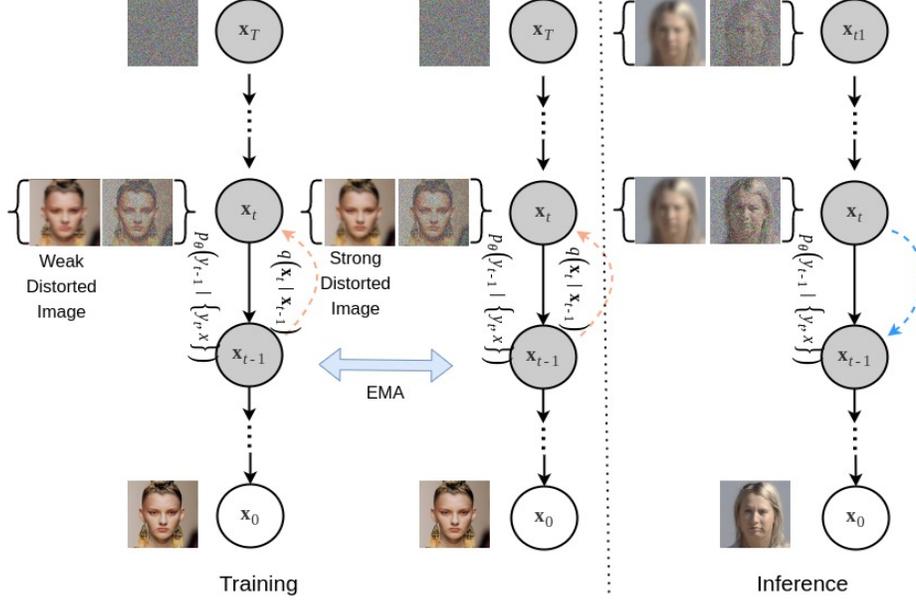


Figure 2: An overview of the proposed approach, during the training process, we perform knowledge distillation to transfer class prior information from a network trained for image super-resolution on a large dataset to the network for removing turbulence degradation. During inference, rather than starting from pure Gaussian noise as in normal DDPM, we begin with noised turbulence degraded images for speed-up in inference times.

resolution and adapt the model to work equally well when the input degraded image is of strong degradation as well as of weak degradation. Let the model trained for the weaker-degradation be represented by $p_\phi(\cdot)$ and the model for restoration of strong turbulence degraded image be denoted by $p_\delta(\cdot)$. For optimizing the parameters δ , we sample an $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ and create the noisy sample y_t of the clean image according to Eq. 5. For each training iteration, the parameters δ are optimized using the loss function L_{final} defined by

$$\begin{aligned} L_T &= \|\epsilon - \epsilon_\delta(y_t, x_{S-turb}, t)\|^2, \epsilon \sim \mathcal{N}(0, \mathbf{I}) \\ L_S &= \|\epsilon_\phi(y_t, x_{W-turb}, t) - \epsilon_\delta(y_t, x_{S-turb}, t)\|^2, \\ L_{final} &= E_{t \sim [1, T]} [L_t + \gamma L_S]. \end{aligned} \quad (13)$$

Here the term L_S ensures that the model focuses on reconstructing the same face regardless of the distortion nature. Once the parameters δ are updated by optimizing the loss function L_{final} , we perform an exponential moving average (EMA) based updating of the weights of the SR model. Please note that the weights ϕ are not updated by optimizing the loss function. The EMA-based weight update for the weights ϕ using the estimated weights δ is according to

$$\phi = \gamma_1 \phi + (1 - \gamma_1) \delta. \quad (14)$$

Hence adapting the model $p_\phi(\cdot)$ for weaker degradation.

Efficient inference for turbulence removal: The usual inference process for diffusion models is very time consuming. However we observed that for the task of conditional image generation, this is not the case and the model creates

good results even for very low number of timesteps ranging $T = 40 - 50$. Furthermore, we noted that the initial steps in DDPMs learn the coarse features. But these features are already present in the turbulence distorted image. Hence rather than starting from pure Gaussian noise, the inference process could be started from the noised turbulence distorted image. Let x denote the turbulence distorted input image, and y_t the sampled inference image after t forward steps in diffusion. As mentioned in Eq. 11, the reverse step starts with Gaussian noise at $t = T$. We start the diffusion process from a time $t = t_1$ than T . We explicitly assign y_{t1} by

$$y_{t1} = q(y_t | y_0), t = t_1. \quad (15)$$

There also exist a hidden advantage of this efficient sampling strategy. Atmospheric turbulence mitigation by itself is a highly ill-posed problem. Since the diffusion process is also stochastic, it tends to increase the ill-posedness by sampling a face that is close to the distorted face during inference. Starting from the turbulence distorted images fixes the coarse features in the image to be sampled hence reducing the ill-posedness.

4. Experiments

For quantitative evaluations, we utilize two different type of metrics, namely image quality-based metrics and facial recognition-based metrics. For evaluating the reconstructed image based on the clean target, we use Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index

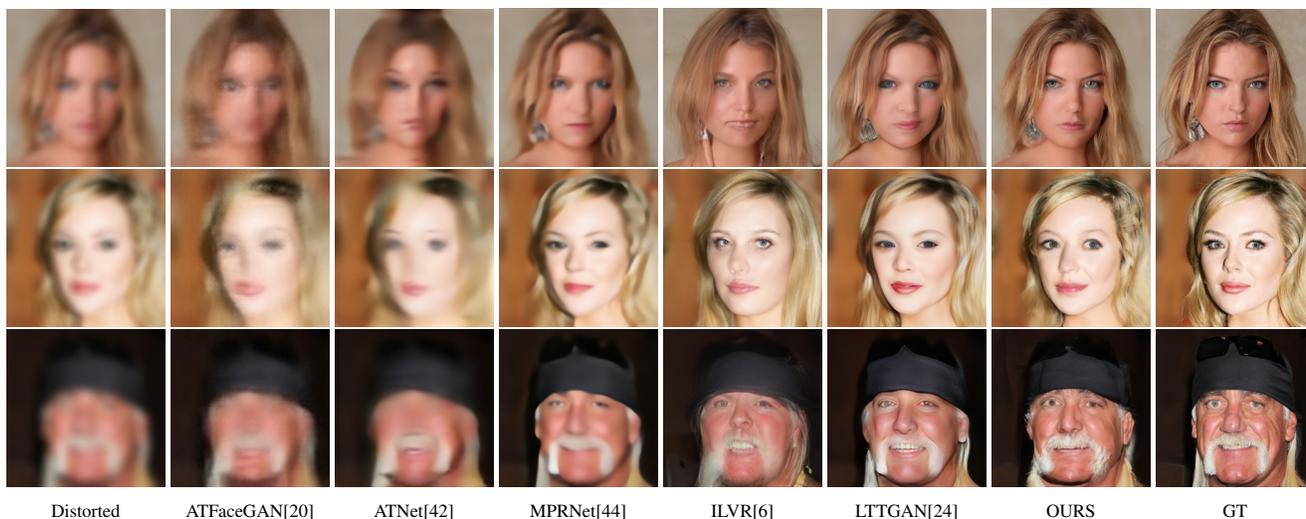


Figure 3: Qualitative comparisons with single image atmospheric turbulence mitigation methods (ATFaceGAN, TDRN) as well as other generative models (ILVR, GFPGAN) on CelebA dataset [21].

(SSIM). PSNR might not always capture the quality of reconstructed faces since blurry images could have higher PSNR than clean images [46]. Hence, for evaluating the quality of the outputs generated by our model, we also utilize other perceptual and no-reference quality metrics. Specifically, we use the Naturalness Image Quality Evaluator (NIQE score) [28] and the LPIPS distance [46] for evaluating the naturalness of the generated image and measuring the perceptual similarity in the feature space. To assess how close the images are to the distribution of clean facial images, we utilize Fréchet Inception Distance (FID score) [11]. As for the facial recognition-based metrics, we use face recognition score calculated using ArcFace [7]. We use three different recognition scores for comparing the performance – Top-1, Top-3 and Top-5. The Top-K score refers to the actual identity being present in the K-nearest faces when the distance between the features of the measured face and the faces in the gallery set is computed.

4.1. Training details

We set the number of training steps as $T = 1000$ and the number of inference steps as $T = 60$. The inference time is reduced by applying the time rescaling strategy mentioned in [31] and $t_1 = 30$ for all our experiments. $\gamma = 0.01$ and $\gamma_1 = 0.9909$ for all our experiments.

4.2. Training dataset

Various methods have been proposed in the literature for simulating atmospheric turbulence degradation [29, 36, 20, 24, 23]. All of these models aim to generate a spatially varying Point Spread Functions (PSFs) and distort each pixel in a clean image. In our work, we follow the model used in *Mei and Patel* [24] which is based on the observation model in Eq. 1. To create a spatially varying PSF and a geometric

distortion at the same time, we use an elastic transformation along with a blur augmentation. Finally, we add an additive white Gaussian noise with a standard deviation $1e^{-4}$. For training we use 35,000 randomly sampled images from the FFHQ dataset [15].

4.3. Testing Datasets

We evaluate our method by performing experiments on one synthetic face dataset, and two real-world datasets (LRFID dataset and BRIAR dataset) consisting of facial images degraded by atmospheric turbulence. We use the CelebA dataset [21] and apply elastic blur [24] to create a synthetic test set. In total, we select 100 identities from the CelebA dataset. For each image, we resize the face to size 256×256 and introduce the degradation. For comparing the performance in real-world scenarios, we use the LRFID dataset [27] and the BRIAR dataset. The LRFID dataset consists of 89 videos in different atmospheric conditions corresponding to 89 different individuals. We extract one image frame randomly from each video and crop it in order to generate the real-world test set. The LRFID dataset also has clean images of the same individuals captured in close range conditions. Since we do not have the clean target pairs of these videos, we utilize clear images of these individuals taken in outdoor conditions as reference and perform facial recognition using ArcFace [7] as well as other perceptual quality metrics to compare the performance of different networks. The BRIAR dataset is another real-world dataset consisting of high-resolution video data of 343 individuals. The videos are shot at different ranges from the individuals varying from 100m to 500m. For our experiments, we utilize a subset of the BRIAR dataset corresponding to ranges 200m and 400m. In total there are 130 identities with data in these ranges. We extract a frame from one video per identity.

Since the resolution of the captured frames are quite high, we utilize retinaface face detection algorithm [8] to extract faces from the center portion of each image. The testset that we use consists of a total of 50 identities.

4.4. Results

We qualitatively and quantitatively compare our method with various deep networks for atmospheric turbulence removal and face restoration. We utilize CNN-based models, GAN-based models and diffusion models. For CNN-based atmospheric turbulence mitigation methods, we use MPRNet [44] and ATNet [43]. We also compare our method with face enhancement networks built upon generative models. Here, we choose two different varieties of generative models, specifically GAN-based method for atmospheric turbulence removal ATFaceGAN [20] as well as two GAN inversion-based models [24, 40] trained for removing atmospheric turbulence. Finally, we compare our method with one diffusion model for image restoration [6]. All the above mentioned methods are retrained with our dataset for fair comparisons. The quantitative analysis can be found in Table 2 and the qualitative results on the CelebA dataset [21], the LRFID dataset [27] and the BRIAR dataset can be found in Figure 3, 1 and 4, respectively.

4.5. Qualitative Analysis

Synthetic dataset evaluation: The quantitative results for celebA dataset [21] are presented in Table 2. The values marked in dark green colour for each metric denotes the best value among all comparison methods and the ones in light green colour denote the second best value. As can be seen from Table 2, the highest PSNR and SSIM are achieved by [44]. This is due to the fact that CNNs are generally trained to optimize the Mean square error(MSE) loss between the network prediction and the clean target. Reducing the MSE loss is proportional to increasing the PSNR. Hence an image with a higher PSNR compared to the clean target could be in fact more distorted than an image with lower PSNR. This could be clearly seen from Figure 3. Hence we make use of the no reference metric NIQE that estimates the naturalness of the predicted output. Lower the NIQE score, more natural is the image. We can see that GFPGAN [40] generates the outputs with the lowest NIQE score. Our method generates outputs with the next lowest NIQE score values. The next metric we utilize is the LPIPS [46] score which compares the perceptual similarity of the predictions from the networks and the set of clean images is the LPIPS score. As we can see, the GAN-inversion based method LTTGAN [24] has the best LPIPS distance and our method produces outputs with the next best LPIPS distance.

To estimate how the predicted distribution of the networks vary from the actual distribution of clean facial images, we compare the predictions from the networks in

| Dataset | CelebA dataset[21] | | | | |
|------------------|--------------------|---------|---------|----------|--------|
| | PSNR(↑) | SSIM(↑) | NIQE(↓) | LPIPS(↓) | FID(↓) |
| GT | inf | 1.000 | 5.750 | 0 | 0 |
| degraded | 22.76 | 0.6517 | 21.35 | 0.4642 | 146.72 |
| CNN based models | | | | | |
| MPRNET[43] | 24.28 | 0.6951 | 10.29 | 0.3885 | 121.37 |
| ATNet[43] | 21.88 | 0.6271 | 13.20 | 0.5065 | 169.52 |
| GAN based models | | | | | |
| ATFaceGAN[20] | 22.11 | 0.5821 | 11.49 | 0.4887 | 182.18 |
| GFPGAN[40] | 21.74 | 0.5925 | 6.083 | 0.3341 | 75.338 |
| LTTGAN[24] | 22.77 | 0.6458 | 7.884 | 0.3192 | 82.182 |
| Diffusion models | | | | | |
| ILVR[6] | 21.40 | 0.5916 | 7.762 | 0.3689 | 76.702 |
| OURS | 22.41 | 0.6262 | 7.186 | 0.3223 | 73.127 |

Table 1: Quantitative results on synthetic turbulence degraded images prepared using CelebA dataset[21]

terms of the FID score. The generative models create outputs with much better FID score when compared to non-generative models as can be seen from Table 2. Our method generates outputs with the best FID score and GFPGAN generates outputs with the second best FID score. The qualitative results on the CelebA dataset can be seen in Figure 3. As we can see the CNN-based methods ATNet [43] and MPRNet [44] and GAN restoration method ATFaceGAN [20] remove the turbulence distortion to some extent but fail to reconstruct relevant facial features like eyes, nose and hair and overly smoothen the facial features. Compared to these techniques, the generative modelling based facial restoration techniques LTTGAN [24] and ILVR [6] reconstruct images with much better facial features. ILVR [6] reconstructs faces with good features but creates faces that are much different from the original faces. In case of LTTGAN [24], although much better faces are reconstructed, the reconstructed image differs in multiple factors like eye color, shape of face which can be clearly seen from the degraded images. In contrast our method is able to reconstruct more accurate faces.

Results on the LRFID dataset [27]: Since the ground truth targets for the LRFID dataset [27] are not available, we compare the performance of the networks qualitatively using no reference metrics, perceptual quality metrics as well as facial recognition accuracy. By comparing the naturalness of the generated outputs using the NIQE score, we can see that the generative methods work much better than the CNN-based models. Our method creates the most natural images and we are able to obtain a NIQE score of 1.094 above the previous best method. Based on the LPIPS metric, the GAN inversion method LTTGAN [24] works the best and our model performs the second best. To evaluate how the generated images are close to real-world facial images, we use the FID score. Regarding the FID score, it can be seen from Table 2 that our model works much better than the other models with an FID score 13 below the next best method. Also we can see that the generative models generate images with much better facial features. To further

| Dataset | BRIAR dataset | | | | | | LRFID dataset[27] | | | | | | |
|------------------|---------------|---------|----------|--------|----------|----------|-------------------|---------|----------|--------|----------|----------|----------|
| | Metric | NIQE(↓) | LPIPS(↓) | FID(↓) | Top-1(↑) | Top-3(↑) | Top-5(↑) | NIQE(↓) | LPIPS(↓) | FID(↓) | Top-1(↑) | Top-3(↑) | Top-5(↑) |
| GT | 5.867 | 0 | 0 | 100.0 | 100.00 | 100.0 | 5.729 | 0 | 0 | 100.0 | 100.0 | 100.0 | 100.0 |
| degraded | 16.68 | 0.7754 | 317.82 | 26.0 | 38.0 | 46.0 | 12.90 | 0.6293 | 195.71 | 35.3 | 62.2 | 71.2 | |
| CNN based models | | | | | | | | | | | | | |
| MPRNET[43] | 11.76 | 0.7051 | 220.73 | 24.0 | 46.0 | 64.0 | 11.13 | 0.5755 | 176.41 | 34.1 | 64.6 | 74.4 | |
| ATNet[43] | 8.506 | 0.7265 | 314.82 | 14.0 | 28.0 | 38.0 | 12.24 | 0.6128 | 202.45 | 36.5 | 64.6 | 74.4 | |
| GAN based models | | | | | | | | | | | | | |
| ATFaceGAN[20] | 10.66 | 0.7459 | 260.25 | 22.0 | 38.0 | 50.0 | 9.434 | 0.6300 | 169.60 | 47.5 | 65.8 | 82.3 | |
| GFPGAN[40] | 6.854 | 0.6414 | 204.76 | 26.0 | 58.0 | 60.0 | 7.918 | 0.5587 | 124.55 | 57.3 | 79.2 | 85.3 | |
| LTTGAN[24] | 5.500 | 0.5969 | 150.24 | 20.0 | 54.0 | 62.0 | 7.970 | 0.4803 | 119.23 | 58.5 | 81.7 | 85.3 | |
| Diffusion models | | | | | | | | | | | | | |
| ILVR[6] | 8.145 | 0.6523 | 167.78 | 22.0 | 44.0 | 56.0 | 12.068 | 0.5661 | 161.38 | 31.7 | 59.7 | 67.0 | |
| OURS | 6.283 | 0.6368 | 152.57 | 32.0 | 56.0 | 66.0 | 7.576 | 0.5234 | 112.64 | 62.2 | 81.7 | 87.8 | |

Table 2: Quantitative results on on real world turbulence degraded datasets: BRIAR dataset and LRFID dataset [27]

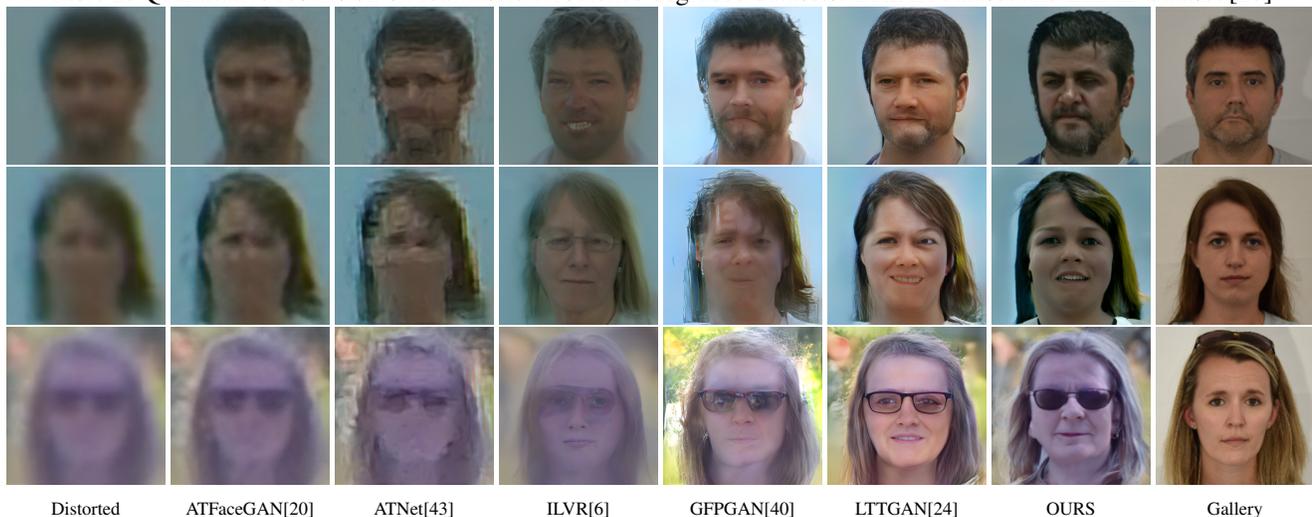


Figure 4: Qualitative comparisons with single image atmospheric turbulence mitigation methods on the BRIAR dataset.

analyze the quality of the restored images, we compare the facial recognition accuracy of the restored images using Arcface facial recognition framework [7] with resnet34 backbone. Based on the facial recognition scores, we can see that the Top-1 facial recognition accuracy is obtained from the images restored by LTTGAN[24]. We obtain the second best facial recognition accuracy, which is just 1.2 below LTTGAN. Our method is able to obtain much better facial recognition accuracies compared to the other methods. The results on the real world LRFID dataset [27] can be seen in Figure 1. As we can see from this figure, ATNet [43] is able to remove small amount of distortions, but is not able to reconstruct facial features well. Compared to this ATNet works a bit better and is able to reconstruct faces with much better features. The diffusion-based model [6] is primarily defined for super-resolution and fails on real-world turbulence distorted images. As we can see from Figure 1, GFPGAN [40] and LTTGAN [24] slightly change the color of the reconstructed images. LTTGAN [24] also creates vertical line artifacts on top of the image. In contrast, our model doesn't cause a colour shift on the reconstructed image and is able to remove turbulence degradation.

5. Results on the BRIAR dataset

Figure 4 shows the qualitative results on the BRIAR dataset. As we can see the images in the BRIAR dataset are very much distorted hence making it very difficult to restore these images. Because of the high level of distortion, most methods fail on the BRIAR dataset. The existing generative modelling-based methods fail to reconstruct the relevant facial features in the image. Our method although cannot reconstruct the exact identity properly, is able to reconstruct a structured face that could be retrieved from the given input image.

6. Ablation Studies

6.1. Effect of Progressive Training:-

To compare the performance improvements obtained by the newly proposed way of training, we retrain the network in a simple manner by directly conditioning the diffusion model with the turbulence distorted image. We train the model for twice the number of iterations required for training our base network. We then test on the real-world LRFID dataset[27] since it contains images much different than

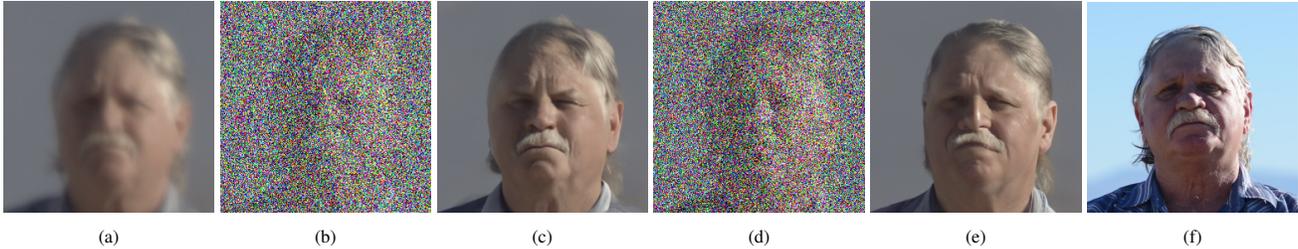


Figure 5: (a) Input turbulence-degraded image. (b) Generated image corresponding to $T = 20$ starting from Gaussian noise. (c) Generated image corresponding to $T = 0$ starting from Gaussian noise. (d) Turbulence image added with Gaussian noise corresponding to $T = 20$. (e) Final reconstructed Image starting from turbulence degraded image. (f) Clean Image.

| Metric | Top-1(\uparrow) | Top-3(\uparrow) | Top-5(\uparrow) | NIQE(\downarrow) | LPIPS(\downarrow) |
|------------|---------------------|---------------------|---------------------|----------------------|-----------------------|
| Without PT | 48.7 | 74.2 | 79.3 | 6.985 | 0.5358 |
| OURS | 54.8 | 84.1 | 86.5 | 6.824 | 0.5255 |

Table 3: An ablation study with and without the proposed progressive training on the LRFID dataset [27].

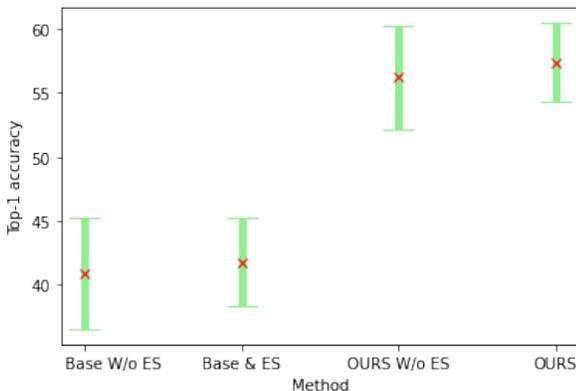


Figure 6: Error bar plot illustrating the effect of efficient sampling.

those used for training. A visualization of improvement brought by the progressive training procedure is shown in Fig. 7. The qualitative comparisons for both the models can be seen from Table 3. We can see that there is a significant boost in performance with our method of training.

6.2. Effect of efficient sampling

Figure 5 shows an illustration of our efficient sampling method. As can be seen from Fig 5, at $T = 20$ the generated image looks like a course face with noise added to it. As we can see from Fig 5 (d) the noised turbulence degraded image looks very similar to this. Figures 5 (c) and 5 (e) show a visualization of the final reconstructed outputs using these starting initializations. As we can see, qualitatively our method is able to preserve the identity pretty well. Moreover, we create 10 samples per inference image and plot the error bar for Top-1 facial verification accuracies in Fig.6. The facial verification accuracies are lower when the efficient sampling technique is used. Hence validating our claims.

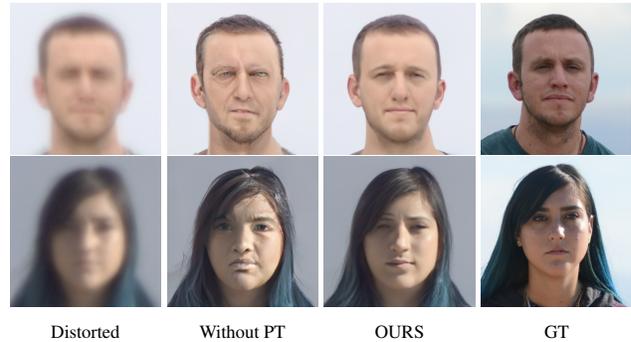


Figure 7: Illustration of improvement obtained by the proposed progressive training method.

7. Conclusion

In this paper, we proposed an effective generative modelling based solution for reconstructing facial images degraded by atmospheric turbulence. A new strategy is proposed for training diffusion models for inverse problems. We make use of a model trained for super resolution on a large facial dataset and adapt it to learn the conditional distribution of clean images given turbulence degraded images. We also introduce a new efficient sampling strategy to reduce the inference time and stabilize the DDPM outputs during inference. We achieve the state-of-the-art results on one synthetic and two real-world datasets. Furthermore, we perform extensive ablation analysis to show the improvement obtained by different parts of our approach.

8. Acknowledgement

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via [2022-21102100005]. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The US. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Nantheera Anantrasirichai, Alin Achim, Nick G Kingsbury, and David R Bull. Atmospheric turbulence mitigation using complex wavelet-based fusion. *IEEE Transactions on Image Processing*, 22(6):2398–2408, 2013.
- [2] Mathieu Aubailly, Mikhail A Vorontsov, Gary W Carhart, and Michael T Valley. Automated video enhancement from a stream of atmospherically-distorted images: the lucky-region fusion approach. In *Atmospheric Optics: Models, Measurements, and Target-in-the-Loop Propagation III*, volume 7463, page 74630C. International Society for Optics and Photonics, 2009.
- [3] Mikhail S Belen’kii, John M Stewart, and Patti Gillespie. Turbulence-induced edge image waviness: theory and experiment. *Applied optics*, 40(9):1321–1328, 2001.
- [4] Wai Ho Chak, Chun Pong Lau, and Lok Ming Lui. Subsampled turbulence removal network. *arXiv preprint arXiv:1807.04418*, 2018.
- [5] Eli Chen, Oren Haik, and Yitzhak Yitzhaky. Detecting and tracking moving objects in long-distance imaging through turbulent medium. *Applied optics*, 53(6):1181–1190, 2014.
- [6] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021.
- [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] Jiankang Deng, Jia Guo, Zhou Yuxiang, Jinke Yu, Irene Kotzia, and Stefanos Zafeiriou. Retinaface: Single-stage dense face localisation in the wild. In *arxiv*, 2019.
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [10] Md Hasan Furhad, Murat Tahtali, and Andrew Lambert. Restoring atmospheric-turbulence-degraded images. *Applied optics*, 55(19):5082–5090, 2016.
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [12] Michael Hirsch, Suvrit Sra, Bernhard Schölkopf, and Stefan Harmeling. Efficient filter flow for space-variant multiframe blind deconvolution. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 607–614. IEEE, 2010.
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [14] RE Hufnagel and NR Stanley. Modulation transfer function associated with image transmission through turbulent media. *JOSA*, 54(1):52–61, 1964.
- [15] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [16] Chun Pong Lau, Carlos D Castillo, and Rama Chellappa. Atfacegan: Single face semantic aware image restoration and recognition from atmospheric turbulence. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2021.
- [17] C. P. Lau, A. Kumar, and R. Chellappa. Semi-supervised landmark-guided restoration of atmospheric turbulent images. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):204–215, 2021.
- [18] Chun Pong Lau, Amit Kumar, and Rama Chellappa. Semi-supervised landmark-guided restoration of atmospheric turbulent images. *IEEE Journal of Selected Topics in Signal Processing*, 2021.
- [19] Chun Pong Lau, Yu Hin Lai, and Lok Ming Lui. Restoration of atmospheric turbulence-distorted images via rpca and quasinormal maps. *Inverse Problems*, 35(7):074002, 2019.
- [20] Chun Pong Lau, Hossein Souiri, and Rama Chellappa. Atfacegan: Single face image restoration and recognition from atmospheric turbulence. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 32–39. IEEE, 2020.
- [21] 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.
- [22] Yifei Lou, Sung Ha Kang, Stefano Soatto, and Andrea L Bertozzi. Video stabilization of atmospheric turbulence distortion. *Inverse Problems & Imaging*, 7(3):839, 2013.
- [23] Zhiyuan Mao, Nicholas Chimitt, and Stanley H Chan. Accelerating atmospheric turbulence simulation via learned phase-to-space transform. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14759–14768, 2021.
- [24] Kangfu Mei and Vishal M Patel. Ltt-gan: Looking through turbulence by inverting gans. *arXiv preprint arXiv:2112.02379*, 2021.
- [25] Kangfu Mei, Vishal M Patel, and Rui Huang. Deep semantic statistics matching (d2sm) denoising network. *arXiv preprint arXiv:2207.09302*, 2022.
- [26] Enric Meinhardt-Llopis and Mario Micheli. Implementation of the centroid method for the correction of turbulence. *Image Processing On Line*, 4:187–195, 2014.
- [27] Kevin J Miller, Bradley Preece, Todd W Du Bosq, and Kevin R Leonard. A data-constrained algorithm for the emulation of long-range turbulence-degraded video. In *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXX*, volume 11001, page 110010J. International Society for Optics and Photonics, 2019.
- [28] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012.
- [29] Nithin Gopalakrishnan Nair, Kangfu Mei, and Vishal M Patel. A comparison of different atmospheric turbulence simulation methods for image restoration. *arXiv preprint arXiv:2204.08974*, 2022.

- [30] Nithin Gopalakrishnan Nair and Vishal M Patel. Confidence guided network for atmospheric turbulence mitigation. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1359–1363. IEEE, 2021.
- [31] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- [32] Omar Oreifej, Xin Li, and Mubarak Shah. Simultaneous video stabilization and moving object detection in turbulence. *IEEE transactions on pattern analysis and machine intelligence*, 35(2):450–462, 2012.
- [33] James E Pearson. Atmospheric turbulence compensation using coherent optical adaptive techniques. *Applied optics*, 15(3):622–631, 1976.
- [34] Michael C Roggemann and Byron M Welsh. *Imaging through turbulence*. CRC press, 2018.
- [35] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021.
- [36] Armin Schwartzman, Marina Alterman, Rotem Zamir, and Yoav Y Schechner. Turbulence-induced 2d correlated image distortion. In *2017 IEEE International Conference on Computational Photography (ICCP)*, pages 1–13. IEEE, 2017.
- [37] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [38] Robert K Tyson. *Principles of adaptive optics*. CRC press, 2015.
- [39] Mikhail A Vorontsov and Gary W Carhart. Anisoplanatic imaging through turbulent media: image recovery by local information fusion from a set of short-exposure images. *JOSA A*, 18(6):1312–1324, 2001.
- [40] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9168–9178, 2021.
- [41] Yuan Xie, Wensheng Zhang, Dacheng Tao, Wenrui Hu, Yanyun Qu, and Hanzi Wang. Removing turbulence effect via hybrid total variation and deformation-guided kernel regression. *IEEE Transactions on Image Processing*, 25(10):4943–4958, 2016.
- [42] Rajeev Yasarla and Vishal M Patel. Learning to restore a single face image degraded by atmospheric turbulence using cnns. *arXiv preprint arXiv:2007.08404*, 2020.
- [43] Rajeev Yasarla and Vishal M Patel. Learning to restore images degraded by atmospheric turbulence using uncertainty. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1694–1698. IEEE, 2021.
- [44] 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.
- [45] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017.
- [46] 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.
- [47] Xiang Zhu and Peyman Milanfar. Removing atmospheric turbulence via space-invariant deconvolution. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):157–170, 2012.