

This ACCV 2024 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2024 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

# DeTurb: Atmospheric Turbulence Mitigation with Deformable 3D Convolutions and 3D Swin Transformers

Zhicheng Zou [0009-0004-6619-7333] and Nantheera Anantrasirichai [0000-0002-2122-5781]

Visual Information Laboratory, University of Bristol, Bristol, UK frederick.zou@outlook.com, N.Anantrasirichai@bristol.ac.uk

Abstract. Atmospheric turbulence in long-range imaging significantly degrades the quality and fidelity of captured scenes due to random variations in both spatial and temporal dimensions. These distortions present a formidable challenge across various applications, from surveillance to astronomy, necessitating robust mitigation strategies. While model-based approaches achieve good results, they are very slow. Deep learning approaches show promise in image and video restoration but have struggled to address these spatiotemporal variant distortions effectively. This paper proposes a new framework that combines geometric restoration with an enhancement module. Random perturbations and geometric distortion are removed using a pyramid architecture with deformable 3D convolutions, resulting in aligned frames. These frames are then used to reconstruct a sharp, clear image via a multi-scale architecture of 3D Swin Transformers. The proposed framework demonstrates superior performance over the state of the art for both synthetic and real atmospheric turbulence effects, with reasonable speed and model size.

Keywords: Atmospheric turbulence · Video restoration · Transformer.

## 1 Introduction

Light propagation through the various layers of the atmosphere, which differ in temperature, pressure, humidity, and wind speed, introduces diffraction-related blurring and random refractions. These variations cause fluctuation in intensity and random phase distortions in the wavefront of light waves, significantly degrading the performance of imaging systems. Most image and video enhancement and restoration techniques have been proposed, addressing specific problems like denoising and deblurring. However, these methods may not be directly used to solve the problem of atmospheric turbulence, as it involves multiple types of distortion, making it challenging to model its degradation accurately. Successful approaches will be invaluable in many applications, including air-to-ground imaging, long-range terrestrial video surveillance, creative industries such as natural history filmmaking, and other computer vision applications, including object recognition and tracking.

Traditional methods for turbulence video restoration typically involve: (i) removing pixel offset caused by tilt, often using the optical flow method; (ii) employing lucky image fusion to select and combine the clearest pixel blocks within a specific interval; and/or (iii) applying blind deconvolution algorithms to remove residual fuzziness [7,8,47,52,53]. However, these methods have notable limitations. They rely heavily on a large quantity of measured data for fusion, use general and non-optimized point spread function (PSF) priors in blind deconvolution, and struggle to effectively address the complex statistical behavior of atmospheric turbulence [8, 12, 19]. Additionally, they are prone to artifacts from inaccurate flow estimates and are very slow.

In recent years, deep learning approaches have emerged to tackle the challenge of atmospheric turbulence mitigation. These methods include deep-stacked autoencoder neural network models and convolutional UNet-like architectures [2, 8,13]. Although these models have shown promising results in simulations, they often rely on simplified assumptions about atmospheric turbulence and lack sufficient real-world datasets, limiting their generalization to diverse scene reconstructions. The effectiveness of these models on actual measured data remains an area of active research.

In this work, we propose a novel framework called DeTurb that integrates two modules: (i) a non-rigid registration module to reduce wavy effects and temporal distortion caused by atmospheric turbulence, and (ii) a feature fusion module to select and fuse useful features for enhanced visualization. Several existing video restoration methods include alignment and feature fusion modules [21,28,29,41], which we found to be more effective when separated rather than combined into a single process. The two-step approach has also proven effective in dealing with atmospheric turbulence issues. Model-based methods, such as space-invariant deconvolution (SID) [52] and complex wavelet-based fusion (CLEAR) [1], employ non-rigid image registration to reduce random perturbations. SID then applies deblurring to the near-diffraction-limited image, while CLEAR produces a sharp image through wavelet-based image fusion. With the advent of deep learning, similar strategies have continued (sometimes referred to as tilt-blur models), such as the CNN-based method AT-Net [43] and the transformer-based method TMT [49].

In our DeTurb framework, both modules employ UNet-like architectures. The first module learns different levels of turbulence distortion, while the second module extracts features from different scales, reconstructing local details related to semantic meanings. A multi-scale approach (like UNet) is essential because the levels of turbulence distortion are unpredictable (atmospheric turbulence is quasi-periodic [25]). Additionally, distortion levels can vary spatially due to the varying distances between objects and the camera—the further an object is to the camera, the more distortion it shows.

To address displacement among frames caused by turbulence and moving objects, deformable convolutions were applied in [17]. However, our architecture differs by processing data in a 3D manner, handling features in both spatial and temporal dimensions simultaneously. Fluctuations in velocity due to atmospheric turbulence cause local displacements between frames, rather than uniform global shifts [18]. As a result, 3D operations are more effective than 2D ones in extracting and enhancing features. Specifically, non-rigid registration is achieved using deformable 3D convolutions [45]. After local spatial alignment among frames, the aligned features from each layer of the pyramid are enhanced using 3D Swin Transformers [42].

In summary, our main contributions can be summarised as follows:

- We propose a novel framework, DeTurb, for restoring long-range videos affected by spatiotemporal distortions due to atmospheric turbulence. With comparable inference speed, DeTurb significantly outperforms the state of the art in terms of video quality.
- DeTurb mitigates geometric distortion using a non-rigid registration module, and then enhances edges and texture details with a feature fusion module.
- The non-rigid registration module estimates the flow of random perturbations and moving objects via a UNet-like architecture, in which each scale performs deformable 3D convolutions.
- The feature fusion module combines features of registered frames with 3D Swin transformers arranged in a UNet-like architecture. This aims to enhance both local and global details for better visualization.

# 2 Existing methods

Atmospheric turbulence, usually resulting from temporal variations occurring near the ground, is typically anisoplanatic for large field-of-view objects, exhibiting spatial variations that complicate the correction process. Given these complexities, learning-based techniques, specifically those involving deep learning, have become increasingly effective in mitigating these effects. Unlike traditional methods (e.g., [1,3,23,52]), these techniques rely on the processing power of neural networks to predict and correct distortions, leading to more adaptable and robust solutions.

Convolutional Neural Networks (CNNs) have been at the forefront of this effort, as they are well-suited for image processing tasks due to their ability to extract hierarchical features from images. CNNs can effectively learn to recognize and mitigate the effects of turbulence, thereby significantly enhancing image clarity [13, 26, 44]. Another innovative approach is Generative Adversarial Networks (GANs), which employ a dual-network architecture comprising a generator and a discriminator that work in tandem to produce highly refined outputs from severely degraded inputs. GANs are capable of generating clear, high-resolution images from those distorted by atmospheric conditions. For example, ATVR-GAN [11] integrates a Recurrent Neural Network (RNN) into the GAN's generator, while in [24], phase disturbance reduction is performed in the Fourier domain, which is then used as a condition for the GAN. LTT-GAN [33] applies style transfer via GAN to restore faces degraded by atmospheric turbulence.



Fig. 1. Diagram of the proposed DeTurb. Top-row: end-to-end framework comprising the non-rigid registration and the feature fusion modules. Bottom-row: block diagrams of 3D Swin transformer encoder and decoder blocks.

More recently, transformer-based methods have shown promising results in handling atmospheric turbulence by employing different attention mechanisms to process image patches. This enables precise correction of atmospheric distortions across various scales within an image. These methods have discovered long-term dependencies in data and demonstrated great potential in this field [32,46,49,50]. TurbNet [32] extracts features through a transformer UNet-like architecture and utilizes physics-inspired downstream methods to reconstruct clean images. The Swin transformer is employed for estimating turbulent flow in [46]. TMT [49] utilizes vision transformers to remove blur, while ASF-Transformer [50] integrates spatial-aware and frequency-aware transformer blocks into a UNet framework.

Diffusion models (DMs) are also of interest in these applications. Two recent methods, [35] and [38], focus on restoring faces degraded by atmospheric turbulence using the Denoising Diffusion Probabilistic Model (DDPM), initially proposed by [16]. In [20], a physics-based simulator is directly integrated into the training process of a restoration model. Furthermore, [40] proposes a conditional diffusion model under a variational inference framework for more generic images. DMs serve as generative priors for blind restoration in [9], which exploits a degradation model involving tilt and blur, akin to the concept in TMT [49] that utilizes two concatenated modules for tilt removal and deblurring.

# 3 Proposed method

The proposed framework is shown in Fig. 1, where the non-rigid registration module (described in Sec. 3.1) is concatenated with the feature fusion module (described in Sec. 3.2).

### 3.1 Non-rigid registration module

As mentioned earlier, several methods address multiple frame inputs of the video using frame or feature alignment processes. Among these, we were inspired by [49], where a UNet architecture with depth-wise 3D convolutions is used to estimate flow across multiple frames. We improve this process with deformable 3D convolutions. A key contribution is that, in atmospheric turbulent environments, objects exhibit visual distortions within small ranges of pixel displacement, appearing randomly in all directions. The use of deformable 3D convolutions provides flexibility in capturing the shapes of the distorted objects across different scales of the UNet in space and time, enabling the extraction of appropriate features from the distorted scenes. After obtaining multi-scale flows each frame in the input group is warped to the current frame from coarse to fine displacements, as shown in Fig. 1.

Deformable 3D convolutions. The transition from the deformable 2D convolutions [10] to 3D ones involves extending the adaptability of convolution operations to the third dimension. Deformable 3D convolutions are defined as Eq. 1,

$$y(p_0) = \sum_{p_n \in \mathcal{G}} w(p_n) \cdot x(p_0 + p_n + \Delta p_n), \tag{1}$$

where  $p_0$  denotes a location in the output feature map y,  $p_n$  represents the *n*-th location in the convolution sampling grid  $\mathcal{G} \in \{(-1, -1, -1), (-1, -1, 0), ..., (1, 1, 0), (1, 1, 1)\}$  for a grid size of  $3 \times 3 \times 3$ , w is the convolution weight, and x is the input feature map.  $\Delta p_n$  is the learnable offset for the *n*-th location, introducing adaptability to the convolution operation [45]. Consequently, the convolutional kernel's receptive field can adjust in response to alterations in the input feature map's shape, thereby accommodating changes in the dimensions and scales of the identified patterns.

Architecture settings. The basic architecture is a 3D-UNet with a depth of 4. The number of depths of the non-rigid registration relies on the local displacement among frames due to atmospheric turbulence. The deformable 3D convolutions (Def3DConv) are used in depth 2<sup>nd</sup>-4<sup>th</sup> as listed in Table 1 (left). Each convolution is integrated with ReLU activation functions. Deformable convolutions still rely on a predefined kernel size, but it adds the ability to dynamically adjust the sampling locations within that predefined kernel grid during the convolution process. This adjustment allows them to adapt better to local variations in the input data. Following [49] to further enhance the perceptual field and improve the detection of features across varying scales, the kernel sizes are strategically set in the encoder. The first level uses a larger kernel size of 7, maintained for the second level to capture broader features and anomalies caused by turbulence before reducing to size 5 and 3 in the third and fourth levels, respectively. This arrangement allows the network to capture and process various distortion patterns, from broad to more localised disturbances. For decoder, we use 3D transposed convolutions (3DTranspConv) and the kernel size of 3 for all levels.

**Table 1.** DeTurb network configuration. Input size of  $H \times W \times N$ . 3DConv is a 3D convolution block, Def3DConv is a deformable 3D convolution block, and 3DTransp-Conv is a transposed convolution block. 3DSwinBlock\_Enc is a 3D Swin Transformer Encoder block and 3DSwinBlock Dec is a 3D Swin Transformer Decoder block.

Non-rigid registrati	Feature fusion module			
Layer (kernel size)	Out dimension	Layer (kernel size) Out dimension		
$3DConv+MaxPool (7 \times 7)$	$\frac{H}{2} \times \frac{W}{2} \times 64$	$2 \times 3 \text{DConv} (4 \times 4)  \frac{H}{2} \times \frac{W}{2} \times 32$		
Def3DConv+MaxPool $(7 \times 7)$	$\frac{\tilde{H}}{4} \times \frac{\tilde{W}}{4} \times 256$	$3DSwinBlock\_Enc \frac{\tilde{H}}{4} \times \frac{\tilde{W}}{4} \times 64$		
Def3DConv+MaxPool (5×5)	$\frac{\tilde{H}}{8} \times \frac{\tilde{W}}{8} \times 256$	$3DSwinBlock Enc \frac{\tilde{H}}{8} \times \frac{\tilde{W}}{8} \times 128$		
Def3DConv $(3 \times 3)$	$\frac{\ddot{H}}{8} \times \frac{\ddot{W}}{8} \times 512$	3DSwinBlock_Enc $\frac{H}{16} \times \frac{W}{16} \times 256$		
3DTranspConv+3DConv (3×3)	$\frac{H}{4} \times \frac{W}{4} \times (256 \to 2N)$	3DSwinBlock_Enc $\frac{\dot{H}}{32} \times \frac{\dot{W}}{32} \times 512$		
3DTranspConv+3DConv (3×3)	$\frac{\hat{H}}{2} \times \frac{\hat{W}}{2} \times (128 \rightarrow 2N)$	$3DSwinBlock_Dec \frac{H}{16} \times \frac{W}{16} \times 256$		
3DTranspConv+3DConv (3×3)	$\tilde{H} \times \tilde{W} \times (64 \to 2N)$	3DSwinBlock_Dec $\frac{\dot{H}}{8} \times \frac{\dot{W}}{8} \times 128$		
$2\text{DWarp}_k, k \in \{0, 1, 2\}$	$\frac{H}{2k} \times \frac{W}{2k} \times 3N$	$3DSwinBlock\_Dec \frac{H}{4} \times \frac{W}{4} \times 64$		
$3DConv(7\times7)$	$\tilde{H} \times \tilde{W} \times 3N$	$3 \text{DTranspConv}$ $\tilde{H} \times \tilde{W} \times 32$		
		3DConv $(1 \times 1)$ $H \times W \times 3$		

#### 3.2 Feature fusion module

Similar to many methods proposed for video processing [2,29,37], we process data through multiscale feature learning using a UNet-like architecture, as shown in Fig. 1 and parameters are listed in Table 1 (right). The process begins with input initial feature extraction through two blocks of 3D convolutions, which prepare the data by highlighting essential features for subsequent layers. The 3D Swin Transformer [42] is used as its efficiency in modelling complex dependencies. It utilises shifted window mechanisms to handle the input data's non-uniformity. The processed data is then channelled through a UNet-like architecture with a depth of 4. This approach ensures that the enhancement of local areas is related to the semantic information of those areas and their surroundings. This multi-scale processing is crucial for restoring fine textures and edges in distorted frames, enhancing the model's capability to effectively address a range of distortion scales introduced by atmospheric conditions. The last 3D convolution layer converts features to RGB output.

3D Swin Transformer. The 3D Swin transformers [4, 39, 42] extend the Swin transformer architecture to three dimensions, adapting it to better understand volumetric or sequential data by incorporating the temporal dimension. This offers significant advancements over traditional models in handling the complexities of video reconstruction. The model introduces a hierarchical structure to process data at multiple scales, capturing detailed spatial-temporal features. Our 3D Swin transformer uses a shifted window-based self-attention mechanism across three dimensions as used in [4,42], effectively capturing dynamic changes over time.

As shown in Fig. 1 bottom-row, the 3D Swin Transformer Encoder block uses 3D convolutions to merge and downsample features across multiple channels. These features are then split to undergo processing by one 3D ResNet block and one 3D Swin transformer block, with the final output being the sum of these two blocks' outputs. In each 3D Swin Transformer Decoder block, a 3D transposed convolution operator is used to combine and upsample feature maps.

Another benefit of the 3D Swin Transformer blocks is that in both the encoder and decoder, cyclic shifts are applied to the input of multi-head self-attention to improve interaction between adjacent and non-adjacent tokens in three directions. It works by shifting the input tokens cyclically by a specified number of units s in each dimension, where s is typically set to half the window size. This shifting rearranges the tokens so that those from adjacent windows in the original configuration may end up in the same window post-shift. This process enables the model to compute similarities and interactions between tokens that were initially in neighboring windows, which helps overcome the limitation of the original window-based self-attention that only computes interactions within the same window. The cyclic shifting thus facilitates the capture of broader contextual information across adjacent windows, enhancing the network's ability to learn long-distance dependency information. However, it also leads to an increase in the number of windows and variability in window sizes, which is managed through a window-masking mechanism to ensure that only relevant token similarities are considered. More details about the 3D Swin transformer block can be found in [4].

#### 3.3 Loss Functions

Two loss functions are employed similar to [49]: Charbonnier Loss and Edge Loss. In the distortion mitigating module, atmospheric turbulence effects can create outliers in the pixel-wise loss. To address this, Charbonnier loss is used, as it combines the benefits of both  $\ell_1$  and  $\ell_2$  losses, effectively handling outliers better [6]. Defined by Eq. 2,

$$L_{\text{Char}}(x,y) = \sqrt{(x-y)^2 + \epsilon^2},$$
(2)

where x and y represent the predicted and true values, and  $\epsilon$  is a small constant (e.g., 1e-3) to ensure numerical stability, this loss function effectively balances the error distribution. It provides a smooth gradient even when small errors are present, which is particularly beneficial for handling the subtle but critical differences in turbulence-affected images, where precision in error correction is essential. To ensure sharp results, the loss  $L_{\rm Edge}$  is added after training for a certain number of iterations (300k in this paper). The edge loss  $L_{\rm Edge}$  is defined as described in Eq. 3,

$$L_{\rm Edge}(x,y) = \lambda L_{\rm Char}(x - g((g(x)^{\downarrow 2})^{\uparrow 2}), y - g((g(y)^{\downarrow 2})^{\uparrow 2})),$$
(3)

where g is a Gaussian filter,  $g = g_h^T \times g_h$ , where  $g_h = [.05, .25, .4, .25, .05]$ .  $\lambda$  represents a small gain (set to 0.05 in this paper), and  $(\cdot)^{\downarrow 2}$  and  $(\cdot)^{\uparrow 2}$  denote

downsampling and upsampling by 2, respectively. Bilinear filtering is used for both upsampling and downsampling.

## 4 Datasets

#### 4.1 Synthetic dataset

As ground truth for atmospheric turbulence mitigation is not available, we generate synthetic distorted sequences from clean ones. We employed the method based on P2S transform, proposed in [31], to generate atmospheric turbulence distortions, as it has proved to be efficient for use as training data.

For static scenes, the dataset is sourced from the Place Dataset [51], where 9,017 images were randomly selected. These images serve as the basis for simulation, where each generates 50 corresponding turbulence-impacted images and their distortion-free counterparts, resulting in a total of 9,017 pairs of static scene sequences (total 450,850 image pairs).

For dynamic scenes, the dataset is enriched with video content from multiple sources to ensure variability and complexity, mimicking real-world conditions more accurately. The primary sources for these dynamic scenes are the Sports Video in the Wild (SVW [36]) dataset, the ground truth videos from the TSR-WGAN [22] project, and the Video Dataset of Perceived Visual Enhancements (VDPVE [14]). This integration forms a comprehensive collection of 6,495 video pairs (total 2,749,582 image pairs).

The dataset is split into training and testing subsets to facilitate practical training and evaluation. The static images are divided into 7,499 pairs for training and 1,518 pairs for validation. The dynamic videos are similarly split, allocating 4,700 videos for training and 1,795 for testing, maintaining a cap of 120 frames per video in the testing set to ensure uniformity in evaluation conditions.

#### 4.2 Real datasets

Videos with real atmospheric turbulence are used to evaluate the performance and generalization capabilities of the proposed method. This includes two datasets: the OTIS [15] dataset and the CLEAR [2] dataset. The OTIS dataset contains 16 static scenes with ground truth. The CLEAR dataset includes 11 dynamic scenes with significant motion and 8 static scenes with minimal or no motion. This dataset comes with *pseudo* ground truth generated using complex waveletbased image fusion. Collectively, these datasets encompass a broad spectrum of turbulence conditions.

#### 5 Results and discussions

#### 5.1 Experiment settings

We trained the two modules separately. Initially, the non-rigid registration module was trained for 400K iterations using 12 frames of input video, a patch size of 128, and a batch size of 2. This phase focused on establishing a robust base for angular correction before proceeding to more complex tasks. Once this module's training was solidified, the feature fusion module was trained for 600K iterations under the same batch size, patch size, and learning rate conditions, ensuring consistency in training dynamics. Both models employ the Adam optimizer, known for its efficiency in handling sparse gradients on noisy problems, in conjunction with a Cosine Annealing scheduler. This scheduler adjusts the learning rate from an initial value of  $2 \times 10^{-4}$  to  $1 \times 10^{-6}$ , promoting a gradual and controlled optimization process.

We evaluated and compared our method with four state-of-the-art models designed to mitigate atmospheric turbulence in long-range imaging: AT-Net [43], TurbNet [32], TSRWGAN [22], TMT [49], and DATUM [48]. Additionally, we included the state-of-the-art video restoration models, BasicVSR++ [5], STA-SUNet [29], and VRT [27]. These models were retrained with our synthetic datasets and tested on both synthetic and real-world datasets. All comparative models were trained using similar data augmentation strategies as employed for the TMT, ensuring comparable conditions and fair performance evaluation.

#### 5.2 Performance on synthetic datasets

With ground truth available, we perform an objective assessment using PSNR. SSIM, and LPIPS measurements. Table 2 presents the average results, calculated by first averaging the results of all frames within each scene, and then averaging these scene results. This approach ensures that the results are not biased toward longer videos. The results clearly show that our method outperforms other existing methods for all metrics, highlighting that the deformable property facilitates a more flexible convolution operation, effectively modeling and correcting geometric distortions caused by atmospheric variations. Additionally, 3D Swin transformers demonstrate their effectiveness in deblurring and visual quality enhancement. Interestingly, BasicVSR++ delivers good results in PSNR and SSIM, but not in LPIPS. BasicVSR++ also exploits deformable convolutions, but applies optical flows for large motion beforehand. This should benefit dynamic scenes, which we plan to investigate further in future work. Example subjective results of the synthetic videos are shown in Fig. 2, where our visual results are sharper, and the zoomed-in areas of the straight lines reveal cleaner and clearer restored lines than other methods. This comparison obviously aligns with the objective results.

#### 5.3 Performance on real dataset

For real atmospheric turbulence, we evaluated our proposed DeTurb and compared it with existing methods using both reference and no-reference metrics. The static scenes of the OTIS dataset come with ground truth, while the scenes of the CLEAR dataset come with pseudo ground truth; hence, objective assessment is possible. We employed NIQE scores [34] as a no-reference evaluation. NIQE is based on a quality-aware set of statistical features derived from a straightforward

Methods	St	atic Sce	nes	Dynamic Scenes			
	$ PSNR\uparrow$	SSIM $\uparrow$	$\mathrm{LPIPS}\downarrow$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\rm LPIPS\downarrow$	
AT-Net [43]	20.86	0.603	0.518	20.21	0.587	0.488	
TurbNet [32]	21.40	0.637	0.431	21.35	0.635	0.433	
STA-SUNet [29]	23.58	0.712	0.352	23.45	0.706	0.348	
BasicVSR++[5]	26.10	0.792	0.279	26.14	0.790	0.280	
TSRWGAN [22]	24.99	0.763	0.249	25.05	0.760	0.251	
VRT [27]	25.85	0.782	0.218	25.78	0.766	0.266	
TMT [49]	25.94	0.795	0.202	26.09	0.767	0.264	
DATUM [48]	<u>27.00</u>	0.787	0.198	27.35	0.819	0.250	
DeTurb (ours)	27.17	0.827	0.170	27.44	0.815	0.242	

 Table 2. Performance comparison on static and dynamic scenes using a synthetic dataset. Bold and underline indicate the best and second-best results, respectively.

 Table 3. Performance comparison on real atmospheric turbulence scenes. Bold and underline indicate the best and second-best results, respectively.

Methods	Static Scenes				Dynamic Scenes				
Modifious	PSNR ·	$\uparrow$ SSIM $\uparrow$	LPIPS .	$\downarrow$ NIQE $\downarrow$	PSNR ·	$\uparrow$ SSIM $\uparrow$	LPIPS ↓	, NIQE ↓	
Distorted inputs	-	-	-	26.48	-	-	-	28.07	
AT-Net [43]	15.65	0.593	0.483	16.37	18.04	0.773	0.472	27.56	
TurbNet [32]	15.08	0.684	0.453	26.11	19.38	0.736	0.411	27.23	
STA-SUNet [29]	20.74	0.698	0.433	26.94	22.66	0.784	0.358	26.68	
BasicVSR++[5]	21.66	0.754	0.232	26.72	26.84	0.845	0.206	25.03	
TSRWGAN [22]	16.87	0.710	0.228	26.72	24.73	0.809	0.130	24.88	
VRT [27]	15.25	0.620	0.338	28.46	25.26	0.818	0.133	25.63	
TMT [49]	15.17	0.611	0.329	29.21	25.59	0.822	0.122	27.98	
DATUM [48]	24.65	0.827	0.204	$\underline{25.01}$	26.46	0.838	0.096	22.72	
DeTurb (ours)	24.53	0.841	0.128	24.34	26.99	0.847	0.088	22.77	



Fig. 2. Subjective results of a synthetic scene. The bottom of each picture shows a magnified version of the straight lines.

yet effective natural scene statistic (NSS) model in the spatial domain. These features are extracted from a collection of natural, undistorted images. Lower scores indicate better perceptual quality.

Table 3 shows the average results of these scenes. Comparing the two existing methods that performed best on the synthetic data, BasicVSR++ and TMT, our method achieves better scores across all metrics, particularly in LPIPS. For the no-reference metric NIQE, our method achieves approximately 9% and 16% higher scores, respectively. Fig. 3 and Fig. 4 illustrate examples of the restored scenes, with the last column showing the pseudo ground truth from CLEAR. Our method produces results with clearer edges and more readable text compared to BasicVSR++ and TMT. Although it cannot restore textures or text as well as CLEAR, our method achieves smoother edges than CLEAR. Overall, DeTurb outperforms other learning-based methods.

Fig. 5 shows examples of varying levels of turbulence. When the atmospheric turbulence is low, DeTurb produces very sharp edges, achieving straight lines and clear curves, particularly at high contrast, such as text on road signs. However, strong distortions are more difficult to recover, as demonstrated in the last row where light rays propagate through the medium from a distance. Although the results appear sharper, random geometric distortion is still present in the results of all methods.

### 5.4 Ablation study

The results of the ablation study are shown in Table 4. First, we investigated the influence of deformable 3D convolutions by replacing them with depth-wise 3D convolutions, as used in TMT [49]. All metrics indicated a performance reduction, particularly for LPIPS. Next, we tested the impact of the 3D Swin transformation by replacing it with 2D Swin transformers [30]. Although the overall performance was reduced, it was not as significant as the decrease observed when replacing deformable 3D convolutions with depth-wise 3D convolutions. This highlights the importance of addressing spatiotemporal distortions caused by atmospheric turbulence. This significance is further confirmed by the substantial reduction in restoration quality when the non-rigid registration module is removed, which has a greater impact than removing the feature fusion module from the pipeline.

For clearer visualization, Fig. 6 shows the y-t plane images, illustrating how a specific line or point at a particular x position within an image evolves over time across video frames. This provides a view of the stabilization achieved by turbulence mitigation models. We selected the middle point of the width. The left image is from the original distorted scene, showing clear vertical streaks and inconsistencies. The feature fusion model alone shows improvement, with noticeably reduced streaking and clearer continuity of lines, but some random geometric distortions are still present, as seen in the middle figure. The entire pipeline, including the non-rigid registration, exhibits the highest level of clarity and consistency, effectively mirroring the original structure of the scene with superior stabilization of vertical elements.

**Table 4.** Ablation study showing the effects of replacing or removing certain modules in the proposed framework. The results are averaged from both synthetic and real scenes, except for NIQE, which is computed only on the real scenes.

Methods	Static Scenes				Dynamic Scenes			
	$ PSNR\uparrow$	SSIM $\uparrow$	$\mathrm{LPIPS}\downarrow$	NIQE $\downarrow$	$ PSNR\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	. NIQE $\downarrow$
$\rm \overline{Deform3D \rightarrow DW~Conv~3D}$	24.25	0.772	0.249	25.73	27.05	0.789	0.231	25.08
$3D \text{ Swin} \rightarrow 2D \text{ Swin}$	25.15	<u>0.800</u>	0.188	25.75	27.13	0.816	0.193	23.72
wo Non-rigid	23.88	0.761	0.261	26.02	26.87	0.795	0.246	25.12
wo Fusion	24.07	0.778	0.258	25.70	26.53	0.800	0.228	24.89
DeTurb	25.85	0.834	0.149	24.34	27.22	0.831	0.165	22.77



Fig. 3. Subjective results of a real distorted scene. The left column shows a static scene, while the other columns depict dynamic scenes. From top to bottom, the rows display the distorted input, BasicVSR++ results, TMT results, our DeTurb results, and CLEAR results.

### DeTurb 13



Fig. 4. Subjective results (similar to Fig. 11 in the DATUM paper [48]).



Fig. 5. Subjective results of a zoomed-in real scene with varying amounts of atmospheric turbulence distortion. The top to bottom rows show cases ranging from easy to difficult to restore.



Fig. 6. Example y-t planes of static 'Man' scene restored without and with the non-rigid registration module

#### 5.5 Computing Budget

We include an assessment of the inference-time computing budget for each model. This analysis is conducted on a single NVIDIA 4090 GPU, providing a standardized basis for comparing the computational demands of each model. Table 5 summarizes the computing resources and time required for each model. Although none of the learning-based methods currently meet real-time requirements, they are significantly faster than conventional model-based methods such as SID and CLEAR.

**Table 5.** Inference time computing budget, measured per frame on a  $256 \times 256$  resolution image (average speed calculated from five test trials).

Methods	# parameters (M)	FLOPs/frame (G)	speed (s)
SID [52]	-	-	132.32
CLEAR [2]	-	-	25.37
TurbNet [32]	26.60	190.69	5.79
STA-SUNet [29]	21.82	-	2.04
BasicVSR++[5]	9.76	127.30	1.20
TSRWGAN [22]	46.28	2,836	2.58
VRT [27]	18.32	7,759	7.35
TMT [49]	23.92	1,304	2.37
DeTurb (ours)	58.79	1,975	2.55

# 6 Conclusion

This paper introduces DeTurb, a novel framework for atmospheric turbulence reduction. The proposed framework consists of two modules: non-rigid registration and feature fusion. The first module uses deformable 3D convolutions to estimate flow and mitigate random geometric distortion across several distorted frames. The second module, utilizing a UNet-like architecture of 3D Swin transformers, further sharpens and enhances the details of the current frame. Experimental results demonstrate that DeTurb outperforms existing methods specifically designed for atmospheric turbulence problems as well as methods proposed for video restoration and enhancement. However, there is still room for improvement in situations of strong atmospheric turbulence.

Acknowledgments. This work was supported by UKRI MyWorld Strength in Places Programme (SIPF00006/1).

# References

 Anantrasirichai, N., Achim, A., Kingsbury, N., Bull, D.: Atmospheric turbulence mitigation using complex wavelet-based fusion. Image Processing, IEEE Transactions on 22(6), 2398–2408 (2013)

- Anantrasirichai, N.: Atmospheric turbulence removal with complex-valued convolutional neural network. Pattern Recognition Letters 171, 69–75 (2023). https://doi.org/https://doi.org/10.1016/j.patrec.2023.05.017
- Boehrer, N., Nieuwenhuizen, R.P.J., Dijk, J.: Turbulence mitigation in imagery including moving objects from a static event camera. Optical Engineering 60(5), 1 19 (2021). https://doi.org/10.1117/1.OE.60.5.053101, https://doi.org/10.1117/1.OE.60.5.053101
- Cai, Y., Long, Y., Han, Z., et al.: Swin Unet3D: a three-dimensional medical image segmentation network combining vision transformer and convolution. BMC Medical Informatics and Decision Making 23, 33 (2023). https://doi.org/10.1186/s12911-023-02129-z
- Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video superresolution with enhanced propagation and alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5972–5981 (June 2022)
- Charbonnier, P., Blanc-Feraud, L., Aubert, G., Barlaud, M.: Deterministic edgepreserving regularization in computed imaging. IEEE Transactions on Image Processing 6(2), 298–311 (1997). https://doi.org/10.1109/83.551699
- 7. Chen, G., Gao, Z., Wang, Q., Luo, Q.: Blind de-convolution of images degraded by atmospheric turbulence. Applied Soft Computing 89, 106131 (2020). https://doi.org/https://doi.org/10.1016/j.asoc.2020.106131, https://www.sciencedirect.com/science/article/pii/S1568494620300715
- Cheng, J., Zhu, W., Li, J., Xu, G., Chen, X., Yao, C.: Restoration of atmospheric turbulence-degraded short-exposure image based on convolution neural network. Photonics 10(6) (2023). https://doi.org/10.3390/photonics10060666, https://www.mdpi.com/2304-6732/10/6/6666
- Chung, H., Kim, J., Kim, S., Ye, J.: Parallel diffusion models of operator and image for blind inverse problems. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6059–6069 (jun 2023). https://doi.org/10.1109/CVPR52729.2023.00587
- Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (Oct 2017)
- 11. Ettedgui, B., Yitzhaky, Y.: Atmospheric turbulence degraded video restoration with recurrent GAN (ATVR-GAN). Sensors **23**(21) https://doi.org/10.3390/s23218815, (2023).https://www.mdpi.com/1424-8220/23/21/8815
- 12. Fried, D.L.: getting lucky Probability of a short-exposure imthrough turbulence\*. J. Opt. Soc. Am. **68**(12), 1651 age 1658(Dec 1978). https://doi.org/10.1364/JOSA.68.001651, https://opg.optica.org/abstract.cfm?URI=josa-68-12-1651
- 13. Gao, J., Anantrasirichai, N., Bull, D.: Atmospheric turbulence removal using convolutional neural network (2019)
- Gao, Y., Cao, Y., Kou, T., Sun, W., Dong, Y., Liu, X., Min, X., Zhai, G.: Vdpve: Vqa dataset for perceptual video enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 1474–1483 (June 2023)
- 15. Gilles, J., Ferrante, N.B.: Open turbulent image set (otis). Pattern Recognition Letters 86. 38 - 41(2017),https://www.sciencedirect.com/science/article/pii/S0167865516303750

- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems. vol. 33, pp. 6840–6851 (2020)
- Hu, D., Anantrasirichai, N.: Object recognition in atmospheric turbulence scenes. In: 2023 31st European Signal Processing Conference (EUSIPCO). pp. 561–565 (2023)
- Huebner, C.S.: Compensating image degradation due to atmospheric turbulence in anisoplanatic conditions. Mobile Multimedia/Image Processing, Security, and Applications (2009)
- Hunt, B.R., Iler, A.L., Bailey, C.A., Rucci, M.A.: Synthesis of atmospheric turbulence point spread functions by sparse and redundant representations. Optical Engineering 57(2), 024101 (2018). https://doi.org/10.1117/1.OE.57.2.024101, https://doi.org/10.1117/1.OE.57.2.024101
- Jaiswal, A., Zhang, X., Chan, S.H., Wang, Z.: Physics-driven turbulence image restoration with stochastic refinement. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 12170–12181 (October 2023)
- Jiang, W., Boominathan, V., Veeraraghavan, A.: Nert: Implicit neural representations for unsupervised atmospheric turbulence mitigation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 4236–4243 (June 2023)
- 22. Jin, D., Chen, Y., Lu, Y., Chen, J., Wang, P., Liu, Z., Guo, S., Bai, X.: Neutralizing the impact of atmospheric turbulence on complex scene imaging via deep learning. Nature Machine Intelligence 3(10), 876–884 (Oct 2021). https://doi.org/https://doi.org/10.1038/s42256-021-00392-1
- Kelmelis, E.J., Kozacik, S.T., Paolini, A.L.: Practical considerations for real-time turbulence mitigation in long-range imagery. Optical Engineering 56, 1 – 12 (2017). https://doi.org/10.1117/1.OE.56.7.071506, http://dx.doi.org/10.1117/1.OE.56.7.071506
- 24. Leihong, Z., Zhixiang, B., Hualong, Y., Zhaorui, W., Kaimin, W., Dawei, Z.: Restoration of single pixel imaging in atmospheric turbulence by fourier filter and cgan. Applied Physics B 127(3) (Mar 2021). https://doi.org/https://doi.org/10.1007/s00340-021-07596-8
- 25. Li, D.: Suppressing atmospheric turbulent motion in video through trajectory smoothing. Signal Processing (2009)
- 26. Li, J., Zhang, M., Wang, D., Wu, S., Zhan, Y.: Joint atmospheric turbulence detection and adaptive demodulation technique using the cnn for the oam-fso communication. Opt. Express 26(8), 10494–10508 (Apr 2018). https://doi.org/10.1364/OE.26.010494, https://opg.optica.org/oe/abstract.cfm?URI=oe-26-8-10494
- Liang, J., Cao, J., Fan, Y., Zhang, K., Ranjan, R., Li, Y., Timofte, R., Van Gool, L.: VRT: A video restoration transformer. IEEE Transactions on Image Processing 33, 2171–2182 (2024). https://doi.org/10.1109/TIP.2024.3372454
- Lin, R., Anantrasirichai, N., Huang, G., Lin, J., Sun, Q., Malyugina, A., Bull, D.R.: BVI-RLV: A fully registered dataset and benchmarks for low-light video enhancement. arXiv preprint arXiv:2402.01970 (2024)
- Lin, R., Anantrasirichai, N., Malyugina, A., Bull, D.: A spatio-temporal aligned sunet model for low-light video enhancement. In: IEEE International Conference on Image Processing (2024)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10012– 10022 (October 2021)

- Mao, Z., Chimitt, N., Chan, S.H.: Accelerating atmospheric turbulence simulation via learned phase-to-space transform. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14759–14768 (October 2021)
- Mao, Z., Jaiswal, A., Wang, Z., Chan, S.H.: Single frame atmospheric turbulence mitigation: A benchmark study and a new physics-inspired transformer model. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) Computer Vision – ECCV 2022. pp. 430–446. Springer Nature Switzerland, Cham (2022)
- 33. Mei, K., Patel, V.M.: LTT-GAN: Looking through turbulence by inverting gans. IEEE Journal of Selected Topics in Signal Processing 17(3), 587–598 (2023). https://doi.org/10.1109/JSTSP.2023.3238552
- Mittal, A., Soundararajan, R., Bovik, A.C.: Making a "completely blind" image quality analyzer. IEEE Signal Processing Letters 20(3), 209–212 (2013). https://doi.org/10.1109/LSP.2012.2227726
- Nair, N.G., Mei, K., Patel, V.M.: At-ddpm: Restoring faces degraded by atmospheric turbulence using denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 3434–3443 (January 2023)
- 36. Safdarnejad, S.M., Liu, X., Udpa, L., Andrus, B., Wood, J., Craven, D.: Sports videos in the wild (svw): A video dataset for sports analysis. In: 2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG). vol. 1, pp. 1–7 (2015). https://doi.org/10.1109/FG.2015.7163105
- Shang, W., Ren, D., Yang, Y., Zhang, H., Ma, K., Zuo, W.: Joint video multi-frame interpolation and deblurring under unknown exposure time. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13935–13944 (June 2023)
- Suin, M., Nair, N.G., Lau, C.P., Patel, V.M., Chellappa, R.: Diffuse and restore: A region-adaptive diffusion model for identity-preserving blind face restoration. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 6343–6352 (January 2024)
- 39. Tang, Y., Yang, D., Li, W., Roth, H.R., Landman, B., Xu, D., Nath, V., Hatamizadeh, A.: Self-supervised pre-training of swin transformers for 3d medical image analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20730–20740 (June 2022)
- Wang, X., López-Tapia, S., Katsaggelos, A.K.: Atmospheric turbulence correction via variational deep diffusion. In: 2023 IEEE 6th International Conference on Multimedia Information Processing and Retrieval (MIPR). pp. 1–4 (2023). https://doi.org/10.1109/MIPR59079.2023.00022
- 41. Wang, X., Chan, K.C., Yu, K., Dong, C., Loy, C.C.: EDVR: Video restoration with enhanced deformable convolutional networks. In: IEEE/CVF CVPR (June 2019)
- Yang, Y.Q., Guo, Y.X., Xiong, J.Y., Liu, Y., Pan, H., Wang, P.S., Tong, X., Guo, B.: Swin3d: A pretrained transformer backbone for 3d indoor scene understanding (2023)
- 43. Yasarla, R., Patel, V.M.: Learning to restore images degraded by atmospheric turbulence using uncertainty. In: 2021 IEEE International Conference on Image Processing (ICIP). pp. 1694–1698 (2021). https://doi.org/10.1109/ICIP42928.2021.9506614
- 44. Yasarla, R., Patel, V.M.: CNN-based restoration of a single face image degraded by atmospheric turbulence. IEEE Transactions on Biometrics, Behavior, and Identity Science 4(2), 222–233 (2022). https://doi.org/10.1109/TBIOM.2022.3169697

- 18 Z. Zou and N. Anantrasirichai
- Ying, X., Wang, L., Wang, Y., Sheng, W., An, W., Guo, Y.: Deformable 3d convolution for video super-resolution. IEEE Signal Processing Letters 27, 1500–1504 (2020). https://doi.org/10.1109/LSP.2020.3013518
- Zhang, M., Yousif, M.Z., Yu, L., Lim, H.C.: A Swin-transformer-based model for efficient compression of turbulent flow data. Physics of Fluids 35(8), 085108 (08 2023). https://doi.org/10.1063/5.0160755, https://doi.org/10.1063/5.0160755
- 47. Zhang, S., Wu, Y., Zhao, J., Wang, J.: Astronomical image restoration through atmosphere turbulence by lucky imaging. In: Zhang, T. (ed.) Third International Conference on Digital Image Processing (ICDIP 2011). vol. 8009, p. 80090B. International Society for Optics and Photonics, SPIE (2011). https://doi.org/10.1117/12.896183, https://doi.org/10.1117/12.896183
- Zhang, X., Chimitt, N., Chi, Y., Mao, Z., Chan, S.H.: Spatio-temporal turbulence mitigation: A translational perspective. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2024)
- Zhang, X., Mao, Z., Chimitt, N., Chan, S.H.: Imaging through the atmosphere using turbulence mitigation transformer. IEEE Transactions on Computational Imaging 10, 115–128 (2024). https://doi.org/10.1109/TCI.2024.3354421
- 50. Zhang, Z., Zhao, B., Chen, Y., Wang, Z., Wang, D., Sun, J., Zhang, J., Xu, Z., Li, X.: ASF-Transformer: neutralizing the impact of atmospheric turbulence on optical imaging through alternating learning in the spatial and frequency domains. Opt. Express **31**(22), 37128–37141 (Oct 2023). https://doi.org/10.1364/OE.503131, https://opg.optica.org/oe/abstract.cfm?URI=oe-31-22-37128
- 51. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence 40(6), 1452–1464 (2018). https://doi.org/10.1109/TPAMI.2017.2723009
- Zhu, X., Milanfar, P.: Removing atmospheric turbulence via space-invariant deconvolution. IEEE Transactions on Pattern Analysis and Machine Intelligence 35(1), 157–170 (2013). https://doi.org/10.1109/TPAMI.2012.82
- 53. Çaliskan, T., Arica, N.: Atmospheric turbulence mitigation using optical flow. In: 2014 22nd International Conference on Pattern Recognition. pp. 883–888 (2014). https://doi.org/10.1109/ICPR.2014.162